Mixing Tone Mapping Operators on the GPU by Differential Zone Mapping Based on Psychophysical Experiments

Francesco Banterle[1], Alessandro Artusi[2], Elena Sikudova[3],
Patrick Ledda[4], Thomas Bashford-Rogers[5], Alan Chalmers[5],
and Marina Bloj[6]

[1]ISTI-CNR, Italy
[2]Universitat de Girona, Spain
[3]Comenius University, Bratislava, Slovakia
[4]MPC Ltd, UK
[5]University of Warwick, UK
[6]University of Bradford, UK

October 19, 2016

Abstract

In this paper, we present a new technique for displaying High Dynamic Range (HDR) images on Low Dynamic Range (LDR) displays in an efficient way on the GPU. The described process has three stages. First, the input image is segmented into luminance zones. Second, the tone mapping operator (TMO) that performs better in each zone is automatically selected. Finally, the resulting tone mapping (TM) outputs for each zone are merged, generating the final LDR output image. To establish the TMO that performs better in each luminance zone we conducted a preliminary psychophysical experiment using a set of HDR images and six different TMOs. We validated our composite technique on several (new) HDR images and conducted a further psychophysical experiment, using an HDR display as the reference, that establishes the advantages of our hybrid three-stage approach over a traditional individual TMO. Finally, we present a GPU version, which

is perceptually equal to the standard version but with much improved computational performance.

1 Introduction

The growth of the HDR imaging has led to the development of a variety of TMOs that attempt to simulate different aspects of the Human Visual System (HVS). The intent is to solve issues, such as preserving local contrast, avoiding halo artifacts, simulating visual adaptation, enhancement and expansion of the original dynamic range, contrast reduction, image detail preservation, etc. Figure 1 shows how two TMOs attempt to reproduce two different aspects such as the contrast and the detail: the left side of Figure 1 shows how the detail is well reproduced in the windows, but the contrast is not preserved. The right side of Figure 1 shows how the contrast is preserved in the bright (see the green squares) and dark areas (see the red squares) but the details are not well reproduced (see the windows).

At the time of writing, two main problems remain unsolved. First, choosing the proper TMO for a specific application can be difficult. Second, a comprehensive TMO that takes into account all aspects specified above has not yet been developed. Motivated by these two open issues, we present a new technique for retargeting HDR images into LDR display that takes into account the benefits of different existing TMOs and combines them, resulting in a more perceptually accurate image compared to the images obtained by existing TMOs. Figure 2 depicts our framework. We decomposed, based on the luminance characteristics, the input HDR image into zones and analyzed the benefits of six different TMOs for each luminance zone. The TMOs that perform better in each luminance zone were selected. This step, based on human observations, and on the specific application where the TMOs are required to be applied, is able to choose the most appropriate TMO for the specific luminance zone. Finally, the results of the selected TMOs are combined to reproduce the final output LDR image. We demonstrate the results of our hybrid approach on several HDR images. Using an HDR display, showing the original HDR image as reference, we conducted a psychophysical experiment that clearly shows the advantages of our hybrid approach over using a traditional single TMO.

This paper extends the previous work by Banterle et al. [12], with the following:

1. A comparison of previous Banterle et al.'s work [12] with the combin-

ing strategy proposed by Yeganeh and Wang [48]; see Section 5 and additional material[1].

2. An approximated GPU-friendly version of the whole algorithm that enables a speed-up of 100 times on average; see Section 6.

3. An evaluation of the GPU-friendly version; in order to estimate if the proposed approximations produce or not an identical result compared to the original algorithm; see Section 6.4 and additional material[2].

4. A GPU-friendly approximation of the local Photographic Tone Mapping Operator that works in real-time with HD, 4K, and 8K content; see Section 6.2.

The rest of the paper is organized as follows: Section 2 describes related work. Section 3 presents our technique. Section 4 presents the pshychophysical validation of our method, and Section 5 presents comparisons against the state-of-the-art. Section 6 shows how to approximate the algorithm for graphics hardware and evaluates its performance and quality. Finally, Section 7 concludes and presents future work.

2   Related Work

The concept of tone mapping in computer graphics was introduced by Tumblin and Rushmeier [43], who proposed a tone reproduction operator that preserves the apparent brightness of scene features. Subsequently, many TMOs have been proposed. It is not the purpose of this paper to give a complete overview on HDR imaging and TM techniques. For a full overview see [37, 11, 9, 10]. In this section, we will review only the approaches that are related to our work.

TMOs, that make use of segmentation to identify regions in the input image with common properties, are not new in the field of computer graphics. For example, the histogram adjustment proposed by Ward et al. [46], was one of the first operators presented that belongs to this group of algorithms. This operator divides an image into bins and equalizes its histogram. However, it does not make use of the spatial information in the segmented areas for further compression of the dynamic range.

---

[1]http://www.banterle.com/work/papers/ic_hybrid_tmo/additional_our_vs_attm.pdf

[2]http://www.banterle.com/work/papers/ic_hybrid_tmo/additional_cpu_vs_gpu.pdf

Figure 1: Example of two TMOs that take into account two different aspects. On the left side, the contrast in the bright (green squares) and dark areas (red square) is preserved but the detail in the windows is not well reproduced. On the right side, the detail in the windows is well reproduced (i.e., the global contrast is well reproduced), but the detail in the dark areas are not preserved. The original image is copyright of Paul Debevec.

Yee and Pattanaik [47] presented a model that makes use of the spatial information. The HDR image is segmented into bins in the logarithmic space, pixels are labeled into groups, and the local adaptation luminance is calculated for each group. Finally, the local adaptation luminance is used to compress the dynamic range of the input image.

Krawczyk et al. [29] proposed a perceptual approach that leverages the anchoring lightness perception theory [25]. Their operator segments the image in frameworks, using a k-means approach, and calculates for each framework a weight map. After this step, for each framework the anchoring is calculated, and the frameworks are merged together using the weight maps. We have used a faster and simpler segmentation approach and a simple binary map in the fusion step. It is true that our segmentation approach may generate superficial boundaries between segments but this problem is addressed by the use of our blending technique.

A user centered approach was investigated by Lischinski et al. [33]. In their method a soft segmentation is generated using rough brushes created by the user. Each brush determines the exposure that the user chose for the specific area. Then, the different exposures areas are merged to create the output image. The authors also proposed an automated initialization version of their approach, but user intervention is not completely eliminated due the interactive nature of their approach.

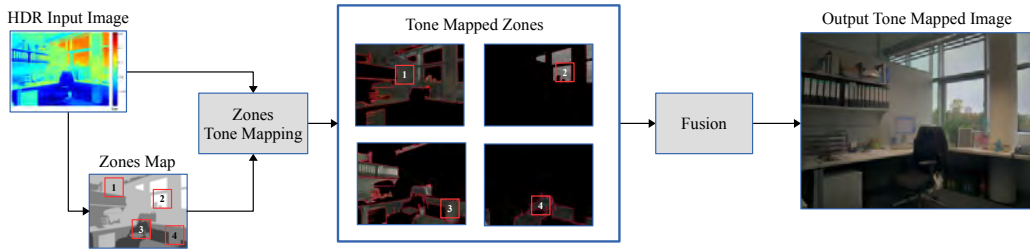Mertens et al. [34] presented an operator that can merge a stack of ex-

Figure 2: Schematic description of our framework. First, the HDR input image is segmented into luminance zones and a zone map is created. Second, the optimal TMO for each zone is applied on the whole input image generating different output images; i.e., one for each zone. Finally, the tone mapped images are merged and the final output is generated. For visualization purposes, we only show the specific zone tone mapped in the Tone Mapped Zones. The original image is copyright of Ahmet Oğuz Akyüz.

posures without the generation of an HDR image. This method, for each exposure, generates a weight map of pixel well-exposed, which can be seen as a soft segmentation. Exposures are merged using these weight maps and Laplacian pyramids blending is adopted for avoiding seams. In our case, the weight map is a simple binary map and can be easily generated and further acceleration of our approach is straightforward (see Section 3.3).

Cadik [16] presented a perceptual motivated approach that is based on combination of arbitrary global and local tone mapping operators. This technique must carefully select the local and global TMOs in order to reproduce high quality output images and this selection is not based on human observation.

Artusi et al. [7, 4] introduced the concept of *selective tonemapper* that is based on some drawbacks of visual attention approaches [6]. In this model, the local luminance adaptation is computed only in high frequency regions. The resulting luminance image is afterwards compressed for a LDR display by using a global tone mapping curve. This avoids the typical drawbacks of a merging step (i.e., seams) when tone mapping an HDR image with two different TMOs (local and global), and greatly reducing computations.

Banterle et al. [12] proposed a novel method for combining different tone mapping operators based on human observation information for selecting different TMOs. They performed formal psychophysical validation using an HDR monitor in order to capture data for an automatic technique for blending. Yeganeh and Wang [48] proposed an extension of the SSIM index, TMQI, for comparing tone mapped images and HDR images. This is based on psychophysical experiments exploiting Amazon Mechanical Turk. Yeganeh

5

and Wang also proposed to blend different TMOs based on the index using Laplacian pyramids, in a very similar fashion to Banterle et al.'s work [12]. The blending is driven by quality index values of tone mapped images; these are computed per pixel for each TMO.

Recently, Boitard et al. [15] presented a video tone mapping operator that preserves the temporal brightness coherency locally. First, the operator segments each frame of a video into zones using the frame log-luminance histogram with constant boundaries. Second, a key value for each zone in the current frame is computed. Third, Boitard et al.'s algorithm [14] is applied for each zone, and blending is applied to avoid artifacts at zones' boundaries. Authors showed that this method can preserve both spatial and temporal local brightness coherencies and they validated it with a subjective evaluation. Note that this method, as Krawczyk et al.'s one [29], is not spatial, and a spatial grouping is important to help experiment participants to evaluate zones without the need to take into account isolated few pixels.

In this work, we also present a novel GPU implementation of Reinhard et al.'s operator [38]. GPU acceleration of TMOs have started in the last decade proposing general frameworks [5, 4] or TMO-specific implementation [27, 39, 2]. Compared to the previous work, we present a TMO-specific implementation where its strength resides in the fact that only a single GPU pass is required in order to tone map an image making the algorithm memoryless. This is thanks to the use of an efficient in terms of memory and computational speed bilateral filter [13].

The solution proposed in this paper, does not need user intervention, takes into account spatial information, and it is validated via a psychophysical study where the original HDR image is shown on an HDR display.

3 Algorithm

Our Framework, as illustrated in Figure 2, can be separated in three straightforward steps. First, we segment the input HDR image into luminance zones and a zone map is generated. Second, for each zone we choose a TMO and apply it on the whole input image generating different output images; i.e., one for each zone. Finally, we merge the resulting outputs for each zone. In the following subsections, each step is described in detail.

3.1 Zone Map Generation

We generate the zone map, Z, using the luminance information contained in the original HDR input image. First, we assign to a pixel its closest order of

magnitude value as

$$Z(\mathbf{x}) = \left\lceil \log_{10} L(\mathbf{x}) \right\rceil, \tag{1}$$

where $L(\mathbf{x})$ is the luminance value of the image at pixel $\mathbf{x}$, and $[\cdot]$ is the rounding operator. However, the resulting Z can have very small isolated areas, even a few pixels, that are difficult to evaluate during our psychophysical experiments; see Section 3.2. Therefore, based on the display spatial resolution used in our experiments, we decided to merge small isolated areas that size is below 5% of the image size. From the smallest isolated area to largest one (below 5% area threshold), the merging function assigns a small isolated area to its spatially closest neighbor area with the closest dynamic range. This process stops when there are no more small isolated areas in Z; see Figure 3 (b). Connected components algorithm [26] is employed to determine isolated small areas and neighbor queries.

Note that more sophisticated approaches can be employed. For example, we tested SuperPixels [1]; see Figure 3 (c). The effort to use a more sophisticated approach when generating Z may produce better quality zones. However, in our experiments, we found out that extra computational efforts are not going to impact on the large areas but only at the boundaries of the zones; see Figure 3 (b) and (c).
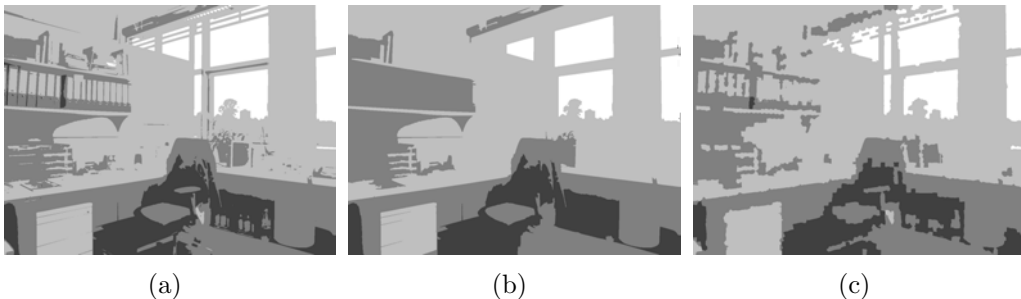

(a) (b) (c)

Figure 3: A zone map output from an HDR image: (a) a zone map with noisy areas as result of our thresholding approach; (b) the zone map in (a) after merging and removing small isolated regions; (c) a zone map generated using SuperPixels [1]. The original image is copyright of Ahmet Oğuz Akyüz.

Due to the rigid selection of the luminance threshold, our method may generate superficial boundaries between areas. However, these are removed afterwards in the fusion step; see Section 3.3. Keeping this step simple enables a future straightforward implementation on the GPU.

3.2 Tone Mapper Selection

Our main goal is to show that is possible to select, based on human observations, a number of existing TMOs and use a combination of their results to reproduce an output image that takes into account all the advantageous aspects of single TMOs. In order to identify which TMO performs better in a specific luminance zone, we conducted a series of psychophysical experiments.

To prove the concept above and be able to generalize its validity, we chose a subset of TMOs from literature. These TMOs were selected based on the previous work by Ledda et al. [32]. Therefore, in our work we investigated the following six TMOs: the algorithm proposed by Tumblin et al. [42]; the histogram model presented by Ward et al. [46]; the time-dependent visual adaptation model [35]; the photographic model (local operator) [38]; the bilateral filtering in the context of TM [22]; and finally, the model presented by Drago et al. [21]. The TMOs are implemented as provided in [37], using the default parameters for each operator. In the case of the operator presented by Pattanaik et al. [35], which takes into account the time dependency adaptation of the HVS, we used only the final frame of the adaptation process. We have chosen these operators based on their quality performances, usability, and popularity. We also limited the number of TMOs tested based on the fact that the purpose of this paper is to prove the concept and not to have a complete comparison of the qualities of the all existing TMOs in the computer graphics community. The iCAM method [24] was originally included but some of the generated images had a violet tinge that did not entirely disappear even adopting the solution suggested by the authors [23].

Note that using a single TMO and varying its parameter values for each zone is not a suitable approach. The parameters are difficult to manually choose. In fact, this process can be very time consuming, because TMOs can have more than 3-4 parameters making the space exploration a tedious task for a user. On the other hand, an optimization process can be employed to find TMO parameters, but it can be computationally expensive. For example, generating the tone mapped images for all zones in Figure 4 requires 254 seconds using TMQI [48] as function goal of the optimization process and Drago et al. operator as TMO. Moreover, an optimization process can fail to determine parameters that are suitable for merging different zones into a single picture; see Figure 4. In fact, it can get stuck in local minima.

3.2.1 Set-up

Once a luminance zone map is obtained, using the method described in subsection 3.1, the performance of the six TMOs in these luminance zones were
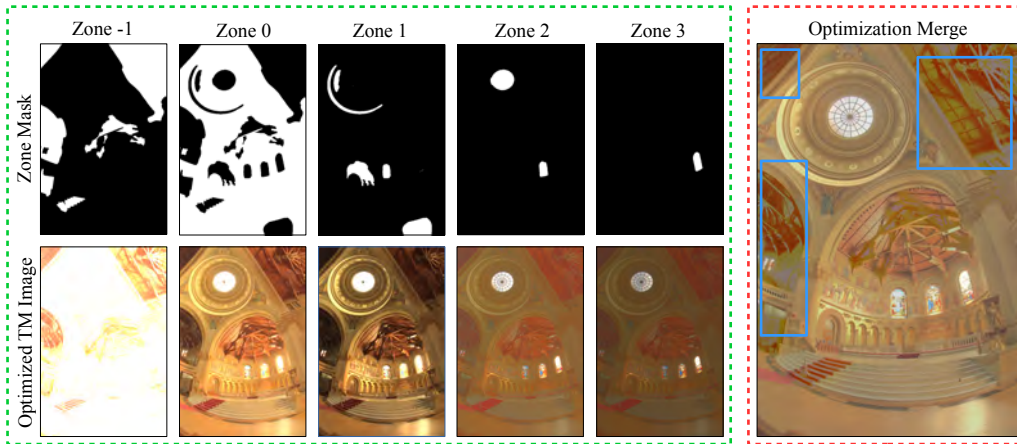
Figure 4: An example of an optimization process (using TMQI as maximization goal) to determine for each zone the parameters of a Drago et al. operators. In the green box, the tone mapped images (bottom) for each corresponding zone (top). In the red box, the result of merging all tone mapped images using Laplacian pyramids and each zone mask as weight (top in the green box) for the corresponding tone mapped image (bottom in the green box). Note that the merging result can produce posterized areas (blue boxes). This is because a TMO may be not suitable for certain dynamic range zones, and the optimization process may struggle in finding high quality results.

compared psychophysically. We performed this test using a Dolby DR37-P HDR display [19] and two LDR LCD displays - which are identical to the front panel of the HDR display. All have a resolution of 1920 × 1080 pixels.

In these experiments, we used the set-up shown in Figure 6. The reference HDR image was presented on the HDR display (center). On the LDR display on the right, all of the LDR images obtained with the six TMOs together with the luminance zone map were displayed. By selecting one of the six images on this monitor, a full size version could be seen on the LDR display on the left. The specific luminance zone under examination was outlined in red on the luminance zone map and on both the HDR and full size LDR images. This outlining ensured that the observer was aware of the particular zone currently being evaluated. The participants could take as much time as they required to complete the task.

We used 15 HDR images with the characteristics reported in Table 3.2.1 and shown in Figure 5. The HDR images were reproduced with physically correct luminances following the method suggested in Akyüz et al. [3]. Also in our case, the highlights that exceed the peak luminance of the HDR display

were clipped around 3000 cd/m^2. Moreover, due to the flanking displays, a veiling effect can appear on the reproduced images. This problem was taken into account. A pilot experiment demonstrated that making use of a proper distance between each LDR and the HDR displays (half meter) avoided this problem.
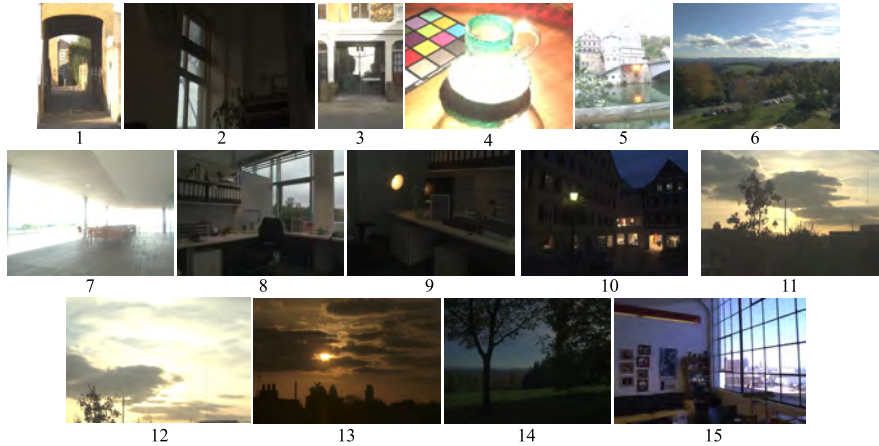

Figure 5: Images used in the psychophysical experiment. The number under the image corresponds to the numbers in the Table 3.2.1. Images 1-3 and 11-13 are copyright of Patrick Ledda. Images 4-10 and 14 are copyright of Ahmet Oğuz Akyüz. Image 15 is copyright of Fredo Durand.

The luminance zones available on the HDR display cover the range from 10^{-2} to 3×10^3 cd/m^2. The images, used in our experiments (see Table 3.2.1), cover the range from 10^{-2} to 10^3 cd/m^2 as shown in the Table 3.2.2. The participants were asked to select the TMO that was considered to reproduce the results as close as possible to the HDR reference image in a specified luminance zone. A group of 24 naïve participants between 20 and 60 years old with normal or corrected to normal vision took part in this experiment. The display was placed in a dark room minimizing the effects of ambient light. Prior to the start of the experiment, each participant was given five minutes to adapt to the environment.

3.2.2 Data Analysis

In our task, observers were asked to indicate for a given image and zone which of the six TMOs was closer to the reference. This includes not only luminance contrast, but also fine details reproduction and color information. Since the observers did not rank all the algorithms we do not have ordinal data but rather categorical data. This allows us to use one-variable χ^2 test to establish if the frequency distribution that we obtain (i.e, the numbers

10

HDR Image	Resolution	Dynamic Range
1	1080×721	2.16
2	1080×1650	3.50
3	1080×723	2.68
4	1080×1440	5.33
5	1080×810	4.47
6	1080×1440	3.39
7	1080×1440	3.38
8	1080×1440	3.93
9	1080×1440	5.46
10	1080×1440	5.90
11	1080×1619	2.48
12	1080×1588	2.83
13	1080×1621	3.27
14	1080×1440	3.27
15	1080×1646	4.32

Table 1: Images used in the psychophysical experiment for the evaluation of the TMOs in the specific luminance zone. The dynamic range is expressed as order of magnitude (logarithm base with base ten of the dynamic range).

on a row in Table 3.2.2) is significantly different from that which we would expect if participants showed no difference (all scores would be $100/6 = 16.6$ %).

Zone	TMO 1	TMO 2	TMO 3	TMO 4	TMO 5	TMO 6	χ^2
3	**28%**	22%	23%	1%	20%	6%	78.23
2	23%	15%	**31%**	1%	20%	10%	98.24
1	**36%**	6%	32%	1%	14%	11%	138.4
0	27%	6%	**33%**	1%	22%	11%	72.40
-1	**52%**	9%	20%	0%	13%	5%	74.94
-2	**68%**	0%	6%	0%	26%	0%	98.35

Table 2: Statistical results of the psychophysical experiment evaluating which TMO performs best in each luminance zone. The enumeration corresponds to the followings TMOs: 1 - Drago et al., 2 - Pattanaik et al., 3 - Reinhard et al., 4 - Bilateral Filtering, 5 - Ward et al., 6 - Tumblin et al. Drago et al. operator was ranked the best in the luminance zones -2, -1, 1 and 3; and Reinhard et al. operator in the luminance zones 0 and 2.

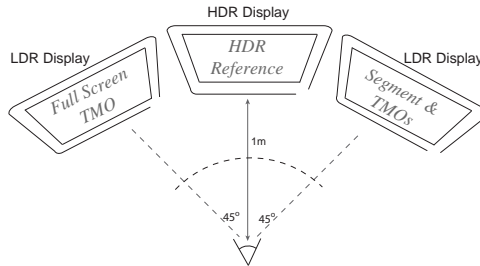Each row corresponds to a luminance zone and each column to a TMO.

Figure 6: Set-up used in the psychophysical experiment evaluating which TMO performs best in each luminance zone. An HDR display is used as reference (center). The LDR display on the right is used to display the six TMOs used in the experiments and the luminance zone under examination. The LDR on the left is used to visualize full screen the TMO that the observer selected at each time in order to perform his/her evaluation.

The number in each cell indicates the percentage of times that a particular algorithm was chosen for a given image zone. The last column of the table reports the corresponding value of the test statistic χ^2 for each image zone. In all cases $p < 0.0001$ and the degree of freedom $df = 5$ (*num. of categories -* 1). In our case the number of categories is equal to the number of TMOs used in the experiment. Our results indicate that for all image zones observers chose an algorithm over others to be closer to the reference.

In Table 3.2.2 the TMO that was most often selected for each zone is also highlighted in bold. The Drago et al. [21] operator was ranked the best operator in the luminance zones -2, -1, 1 and 3; and the Reinhard et al. [38] operator in the luminance zones 0 and 2. Several TMOs evaluation works have been published lately [20], [30], [32], [49], [17], [3] and [31]. The main difference of these works, in comparison with our experiment, is that the evaluation was done on the whole input images and not on a particular luminance zone. This can lead to a different evaluation since the subject, in these other experiments, can be influenced by the overall quality of the input image. In our case this is less evident since the subject is evaluating the TMO in a specific luminance zone without paying attention to the other areas of the input image. Our results confirm, either partially or completely, previous TMOs evaluation works [49], [20], [17]. However, the bilateral filter in our experiment was not performing well, as judged by our subjects, and this is partially in contradiction with the study presented by [31] and [3]. This can be explained by the fact that the original bilateral filtering presented by [22] was modified in the works of [31] and [30] resulting in an improvement of the overall image contrast. In addition, [31] and [30] agree with our result that

12

Reinhard et al. [38] was always well ranked.

Concerning the work of Akyuz et al. [3] we have noticed that the number of images used is 33% less than what has been used in our experiment and this may lead to different results. Also Akyuz et al. [3] is in contradiction not only with our results, but also with the results of the others such as [49], [20], [17], [31], and [30], that confirm, either partially or completely, our results.

3.3 Fusion Results

Following the results of the previous Section we concluded that the Drago et al. [21] and the Reinhard et al. [38] operators would be utilized in the fusion step. Once the TMO to be used in each zone has been selected, the different TMOs are applied to the relevant zone and the outputs are blended to create the whole image. The blending operation is necessary to eliminate discontinuities in the luminance, that may occur at the edges of the zones within the map. A naïve blending of TMOs, is a simple multiplication of a weight map times the TMO with a consequent accumulation. This method creates artifacts such as seams and discontinuity at edges. To avoid these problems we blend the various TMOs using Laplacian blending, similar to that used in [34]. Note that the use of Laplacian blending helps masking contouring, which may appear when applying spatial blending without the need to use more sophisticated approaches for segmentation; i.e., soft segmentation [40]. This allows us to keep a straightforward approach to implement that can be implemented on the GPU; see Section 6.
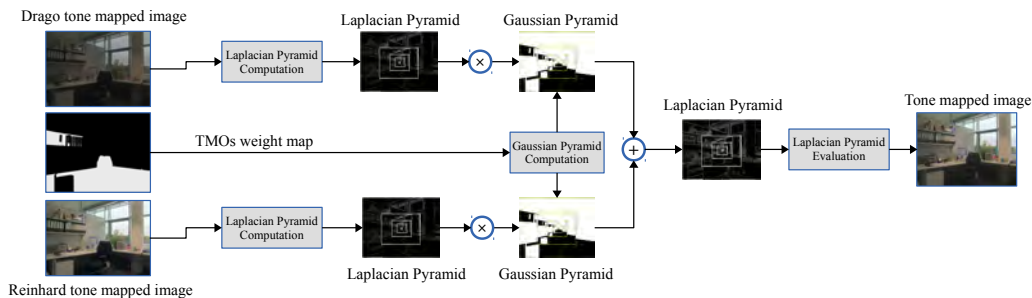

Figure 7: Graphical representation of the fusion step of our Framework. The original image is copyright of Ahmet Oğuz Akyüz.

Figure 7 shows the graphical description of the fusion step. First, a Laplacian pyramid of the best performing TMOs is computed. Second, the Laplacian pyramid is multiplied by the Gaussian pyramid of the weight map associated with each TMO. Third, these results are summed and the fusion

pyramid is generated. Finally, from the fusion pyramid the output image is reconstructed (*Fused Result*).

Based on the fact that the subjects' evaluation of the TMOs in the specific luminance zone, is equivalent to give the same weight to all the pixels of the zone, we decided to use a simple Gaussian filtered binary map. We also investigated if a simple linear combination fusion scheme was able to give comparable results to the one obtained with the proposed fusion scheme. As is shown in Figure 8, several artifacts are visible in the output obtained with the simple fusion scheme (left image).

| (a) | (b) | (c) |

Figure 8: Comparison of the results obtained with two different fusion schemes: (a) the TMOs weight map; (b) a result using the simple fusion scheme; (c) a result using the proposed fusion scheme. Artifacts are visible on the image obtained with the simple fusion scheme (b). The original image is copyright of Max Lyons.

3.4   Limitations

Our method can reproduce details of different dynamic range zones with high quality. However, our proof of concept is tailored for the dynamic range available in the HDR display used in the psychophysical experiments.

Regarding efficiency, our algorithm needs to tone map an image a number of times equal to the number of tone mappers used for the luminance zones. In our proof of concept this is equal to two; i.e., two tone mappers are used Drago et al. [21] and Reinhard et al. [38]. Moreover, the results need to be combined using Laplacian blending.

HDR Image	Resolution	Dynamic Range
1	1016×760	4.34
2	874×644	5.82
3	1080×1626	5.13
4	1080×1626	4.69
5	906×928	4.33
6	769×1025	5.00
7	1080×718	4.08
8	684×912	5.09
9	768×1024	3.86
10	1023×767	4.53
11	768×512	5.40
12	1080×1624	4.08
13	769×1025	2.45

Table 3: Images used in the psychophysical experiment for the evaluation of our TMO. The dynamic range is expressed as order of magnitude and it is covering the same dynamic range as the images used in the first experiment.

4 Results: Psychophysical Validation

In this Section we present the experimental results obtained from testing our hybrid approach.

4.1 Set-up

We tested and psychophysically validated our hybrid approach on several HDR images, using the same HDR and LDR displays as in the previous experiment. Again the HDR display was used to visualize the HDR image as the reference. We conducted a pair-wise comparison test between the two best TMOs (Drago et al. [21] and Reinhard et al.[38]) and our hybrid operator. Figure 9 shows the set-up for this experiment; the reference HDR image is displayed on the HDR display (center), and the two TMOs that are being compared are displayed on the two LDR displays. This is a similar approach to evaluating TMOs as used in [32].

We tested 13 images, different from the ones used in the previous experiment, resulting in 39 possible pair comparisons for each participant (3 possible pairs × 13 images). A group of 21 naïve participants between 20 and 60 years old with normal or corrected to normal vision took part in this experiment. Environmental conditions were the same as the previous TMO selection experiment. The characteristics of the 13 HDR images, used in this
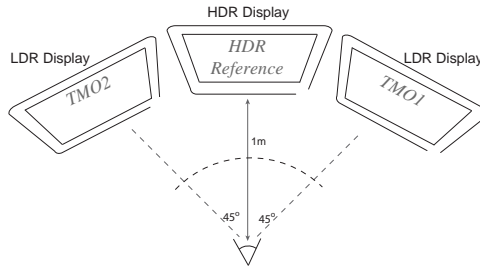
Figure 9: Set-up used in the validation experiment done for evaluating our hybrid TMO versus the two best TMOs of the previous psychophysical experiment. An HDR display is used as reference (center). Two LDR displays (right and left) are used for pairwise comparison.

Figure 10: Images used in the psychophysical experiment; the number under the image corresponds to the numbers in Table 4. Image 1 is copyright of Karol Myszkowski. Images 2 and 5 are copyright of ILM Ltd. Images 3-4 are copyright of Mark Fairchild. Image 6 and 13 are copyright of Raanan Fattal. Image 7 is copyright of Jack Tumblin. Image 8 is copyright of Francesco Banterle. Image 9 is copyright of Greg Ward. Image 10 is copyright of Max Lyons. Image 11 is copyright of Paul Debevec. Image 12 is copyright of Patrick Ledda.

experiment, are reported in Table 4 and shown in Figure 10.

4.2 Data Analysis

In order to determine which TMO performs best, observers were asked to select, for each pair of images, the TMO that was considered to reproduce the result perceptually closer to the HDR reference image. We used a *Multiple Comparison Range test*, for this paired-comparison experiment to determine if any pairwise difference was significant. This series of tests determines if inconsistencies in the scores between any of the two TMOs are present. This procedure, equivalent to Tukey's method used with ANOVA [18], is based on the range of the scores obtained by the three TMOs. In this test,

HDR Image	First	Second	Third
1	HYB 26	PHO 20	DRA 17
2	DRA 37	HYB 25	PHO 1
3	HYB 38	DRA 20	PHO 5
4	PHO 30	DRA 18	HYB 15
5	HYB 30	DRA 27	PHO 5
6	HYB 26	DRA 23	PHO 14
7	HYB 34	DRA 24	PHO 5
8	DRA 35	HYB 23	PHO 5
9	HYB 30	PHO 19	DRA 14
10	HYB 38	DRA 14	PHO 11
11	DRA 34	HYB 24	PHO 5
12	HYB 31	DRA 24	PHO 8
13	HYB 30	DRA 26	PHO 7
Overall	HYB 370	DRA 298	PHO 150

Table 4: Results of the pairwise validation test. Each row represents the results for each image used in the experiment. The row at bottom is the overall result. The columns represent the **First**, **Second** and **Third** TMO classified respectively. The number beside each operator represents the number of times it was chosen. The notation is the following: HYB is the Hybrid TMO presented in this paper, DRA is the Drago et al.'s operator, and PHO is the Reinhard et al.'s method or photographic operator.

the difference between the scores of two TMOs is considered significant if it is greater than a critical value, R [18]. In other words R represents the minimum value where the difference between the scores of two TMOs is consider non-significant.

We have chosen the significance level α equal to 0.05 and the critical values R is computed as:

$$R = \frac{1}{2} W_{t,\alpha} \sqrt{st} + \frac{1}{4} \tag{2}$$

where $W_{t,\alpha}$ is the upper significance point of the W_t distribution. At a significance level of 0.05 and for three TMOs ($t = 3$ the $W_{t,\alpha}$ is equal to 3.31; see Pearson [36], table 22). The value of s specifies the total number of possible cases in the experiment. Table 4.1 shows the results of the second experiment.

We performed the multiple comparison range test on the results of the second experiment, for each image individually ($s = num.\ operators$) and for the combined preference over all images ($s = num.\ operators \times num.\ images$). This test shows that overall the hybrid TMO performs significantly better than the other two TMOs used on their own. The critical value R is 48 when evaluated over all images. Since the difference in overall scores between each TMO is greater than R, the differences can be considered significant.

The test results show that the hybrid operator is not significantly worse for any of the test images than the other two operators. In the cases when Reinhard et al. or Dargo et al. operators appear to have higher scores the difference is less than the critical value of R, which for the case of single images is 14.

Figure 11 shows a compelling result where our method can reproduce both dark and bright regions. More results can be found in the additional material [3].

5   Results: Comparisons against the Adaptive TMQI Tone Mapping

In this Section, we compared our hybrid approach against the *adaptive TMQI tone mapping* (ATTM) [48] (Section IV.B). This is a strategy for blending different TMOs based on their TMQI quality index using Laplacian pyramids in similar fashion to Banterle et al.'s work [12].

---

[3]http://www.banterle.com/work/papers/ic_hybrid_tmo/img_htmo_val_second_experiment.pdf

| (a) | (b) | (c) | (d) |

Figure 11: A compelling example showing the efficacy of our approach: (a) Drago et al.'s method, (b) our method, (c) Reinhard et al.'s operator, and (d) the blending mask used by our method (white for Drago et al.'s method, and black for Reinhard et al.'s method).The original image is copyright of Max Lyons.

The main goal of this comparison is to determine *if the strategy to select a particular TMO for each region is effective* as the one using the TMQI. We want to stress out that we are not interested in having the best parameters of a particular TMO but only *how the blending is carried out.*

In order to understand if our method has similar blending performance, we implemented ATTM using only two TMOs, Drago et al. [21] and Reinhard et al. [38], which are the same used by our approach. Furthermore we used standard parameters of TMOs in our experiments for both our method and ATTM. This was done again to reduce complexity and focusing on our goal.
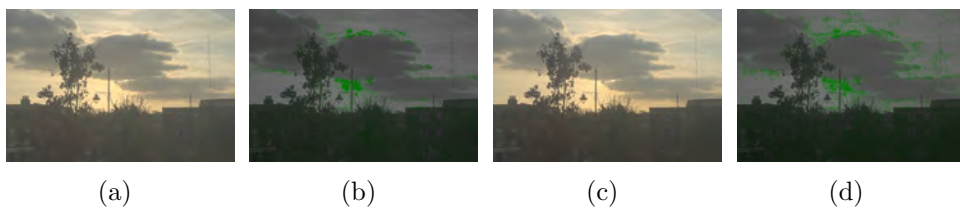

| (a) | (b) | (c) | (d) |

Figure 12: An example of DRIQAM [8] (note that green, blue and red respectively identify loss of visible contrast, amplification of invisible contrast and contrast reversal): (a) An image tone mapped using our method. (b) The amplification/loss/reversal map resulting after comparing (a) and the original HDR image. (c) The same image in (b) tone mapped with ATTM image tone mapped using our method. (d) The amplification/loss/reversal map resulting after comparing (c) and the original HDR image.

In our experiment, we tone mapped 28 HDR images Figure 10 and Figure 5 using our method and ATTM. We compared each image against the original HDR image using the dynamic range independent quality assessment metric (DRIQAM) [8]. This metric is capable of operating on an image pair where both images have arbitrary dynamic ranges which makes it suitable to compare HDR images and their tone mapped counterparts [8]. The result of the metric are three maps, determining if there is an amplification, loss or reversal of the contrast in the testing image compared to the reference image.

Figure 13 summarizes all results for the dataset and three maps (Please see additional material for full results). For each results (amplification, loss, and reversal), Figure 13 provides the percentage of pixels which may be detected by the human visual system with probability higher than $95\%(P(X > 0.95))$ (Please see Table.1 in the additional material[4]). From Figure 13, we can notice that both the tested methods, our method and ATTM, have similar performances on average with very small differences.

6 Algorithm on the GPU

This section presents the GPU implementation of our method. Typically, GPU implementations have a trade-off between quality and efficiency; i.e., speed and memory management. This is also valid in our case. In particular, we need to re-design the zone map and the local tone mapping operator implementations. In the next subsections, we describe the implementation details of these two steps.

6.1 Approximated Zones Map

To improve efficiency in the zone map calculations, in terms of speed, we have substituted the merging and removal of the small areas with a preprocessing filtering step. The idea underlying this design choice is based on the fact that to be merged areas are characterized by having small difference in luminance values. Therefore, to merge these areas, we can smooth these differences while keeping their original structure; i.e., edges. To achieve this, we can

---

[4]http://www.banterle.com/work/papers/ic_hybrid_tmo/additional_our_vs_attm.pdf
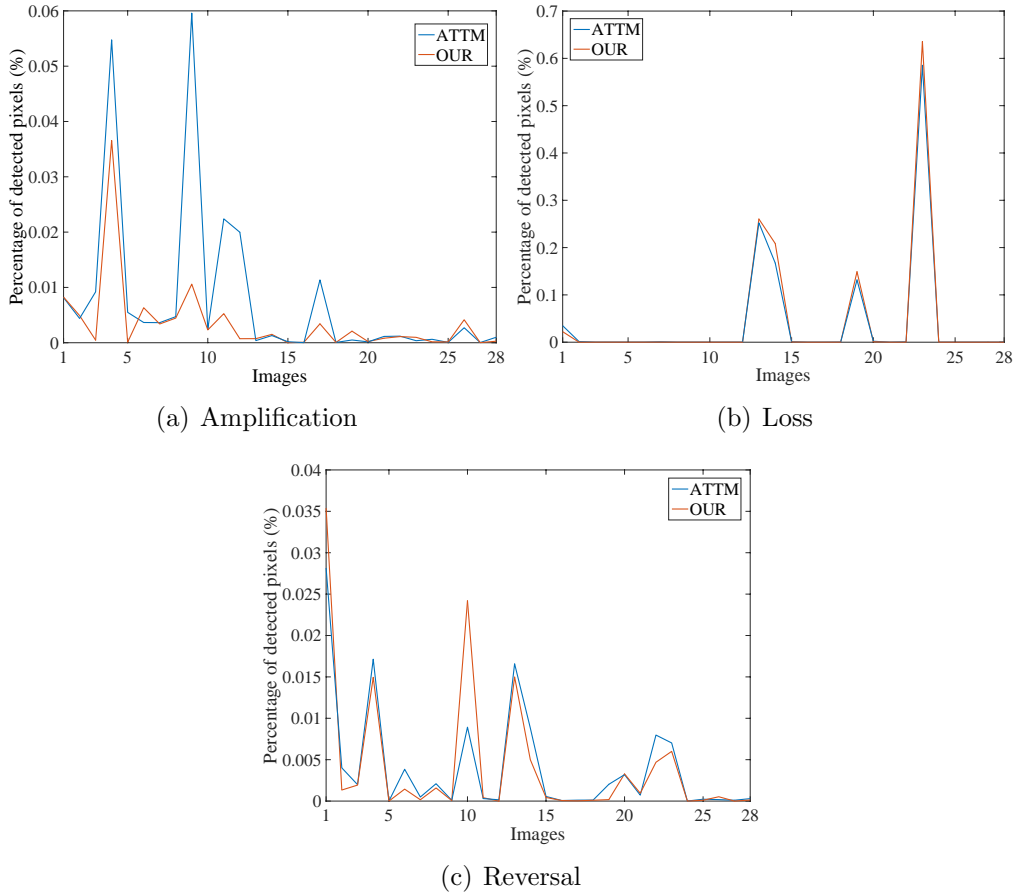
(a) Amplification

(b) Loss

(c) Reversal

Figure 13: Results of comparing our method and ATTM [48] using the DRIQAM [8]. We present results for amplification, loss, and reversal of contrast of the tone mapped image compared to the HDR image (reference). For each result (amplification, loss, and reversal), we provide the percentage of pixels which may be detected by the human visual system with probability higher than and $95\%(P(X > 0.95))$

apply the bilateral filter [41] for a few iterations. This filter is defined as

$$B(I, \mathbf{x}, \sigma_s, \sigma_r) = \frac{\sum_{\mathbf{y} \in \Omega(\mathbf{x})} I(\mathbf{y}) G_r(\|I(\mathbf{y}) - I(\mathbf{x})\|, \sigma_r) G_s(\|\mathbf{y} - \mathbf{x}\|, \sigma_s)}{N(I, \mathbf{x})} \quad (3)$$

$$N(I, \mathbf{x}) = \sum_{\mathbf{y} \in \Omega(\mathbf{x})} G_r(\|I(\mathbf{y}) - I(\mathbf{x})\|, \sigma_r) G_s(\|\mathbf{y} - \mathbf{x}\|, \sigma_s), \quad (4)$$

where B is a bilateral filter, I is the input image, $\Omega(\mathbf{x})$ is the set of pixels $\mathbf{y}$ in a window centered at $\mathbf{x}$, and G_r and G_s are the *range attenuation* and *spatial attenuation* functions, respectively. These are typically Gaussian functions with standard deviations σ_r and σ_s, respectively. Given this definition, we iteratively apply Equation 4 in the $\log_{10}$ domain (i.e., the same domain of segmentation). We set the filtering parameters as:

1. $\sigma_s = 1$; we want to propagate the filtering gradually over iterations as in a dilate/erode approach. This means a 5×5 wide filter.

2. $\sigma_r = 0.5$; we want to flatten values in order to belong to a dynamic range zone.

3. iterations $= \sqrt{0.05 \times w \times h}$, where w and h are, respectively, the width and height of the image. We want to iterate enough times in order to cover with our filter regions large as 5% of the image as in the original algorithm; see Section 3.1.

These parameters do not change across iterations. Figure 14 shows an example of our GPU-friendly zone map generation method. Note that the solution proposed maintains features as advanced segmentation method such as SuperPixels.

6.2 A Bilateral Filtering Based Local Reinhard et al.'s Operator

Before describing our novel GPU implementation of the local Reinhard et al. [38]'s operator, we provide a brief overview of its main steps (refer to the original paper for more details). This operator is defined as

$$L_d(\mathbf{x}) = \frac{L(\mathbf{x})}{1 + V_1(\mathbf{x}, s_m(\mathbf{x}))} \quad L(\mathbf{x}) = \frac{a}{\hat{L}_w} L_w(\mathbf{x}), \quad (5)$$

where L_w is the world luminance, a the key image (typically $a = 0.18$), $\hat{L}_w$ is the log-average luminance, and V_1 is the local adaptation, which is defined as
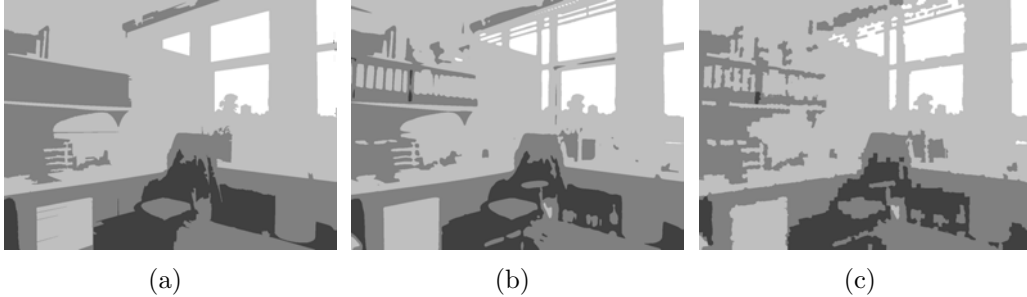
Figure 14: Approximated Zone map output from an HDR image, when compared with the CPU implementation of our method: (a) the result of our full algorithm. (b) The approximated GPU version with prefiltering. (c) Segmentation using SuperPixels [1]. Note that the approximated version preserves better edges in the image, and keeps features that only the SuperPixels based method can preserve. The original image is copyright of Ahmet Oğuz Akyüz.

Gaussian filtered L versions at different scales s ($s_{\max} = 8$ with an increasing factor of 1.6) with $\sigma_1 = 0.35s$. $s_m(\mathbf{x})$ is the largest possible neighborhood of pixel $\mathbf{x}$ and it is defined as

$$s_m(\mathbf{x}) = \arg\max_s |V(\mathbf{x}, s) < 0.05| \qquad V(\mathbf{x}, s) = \frac{V_2(\mathbf{x}, s) - V_1(\mathbf{x}, s)}{2^\Phi a/s^2 + V_1(\mathbf{x}, s)}, \qquad (6)$$

where V_2 is similar to V_1 with $\sigma_2 = 0.56s$, and Φ is a sharpening parameter; i.e., directly proportional to sharpness of edges. Given this definition of the operator, we remapped V_1 as a result of bilateral filtering; $V_1 \approx V_1^{\mathrm{B}}$. We propose the following remapping:

$$V_1(\mathbf{x}, s_m(x, y)) \approx V_1^{\mathrm{B}}(\mathbf{x}) = \phi^{-1}\Big(B\big(\phi\big(L(\mathbf{x})\big), 1.6, 0.05\big)\Big), \qquad (7)$$

where $\phi(x) = x/(2^\Phi a/s^2 + x)$. The main reasons for this remapping are:

1. $V_1^{\mathrm{B}}(\mathbf{x})$ is computed in a sigmoid domain, in order to be independent of absolute luminance levels as in the original formulation; see Equation 6.

2. $\sigma_s = 1.6$ is meant for maintaining fine details with a sufficient kernel size [38]. Note that a larger $\sigma_s = 1.6$ would lead to a smooth luminance adaptation, which does not preserve fine details.

3. $\sigma_r = 0.05$, this is a conversion of the *hard* test in $s_m(\mathbf{x})$ into a *smooth Gaussian* test.

23

Method	PSNR	SSIM	MS-SSIM
Our	**28.74**	**0.89**	**0.98**
Akyüz's method	28.52	0.87	**0.98**

Table 5: A quality comparison between our bilateral filtering based Reinhard et al.'s operator and Akyüz's method against the original operator. For PSNR, SSIM, and MS-SSIM a higher number is better.

6.3 Implementation Details

Regarding implementation details, we implemented the zone map generation, tone mapping, and pyramid blending using OpenGL. We greatly improved efficiency in terms of computational speed and memory by tone mapping images with GPU versions of Drago et al. [21] and local Reinhard et al. [38] on luminance channels only. In order to achieve real-time performances, we used an approximated version of the bilateral filter. In particular, we used the Banterle et al.'s bilateral filtering approximation [13], which is GPU friendly, easy to integrate, and it provides a very small memory footprint. Furthermore, it enables to compute the adaptation luminance in the same pass of the tone mapping avoiding to store it as in other algorithms [39, 2].

6.4 Results

We evaluated the time performance of our algorithm and compared with the original algorithm implemented on the CPU in MATLAB and a CPU optimized version. In these tests, we varied the size of Image 9 in Figure 5, which has 5 orders of magnitude and creates all possible zones. The machine, that we used in these tests, is an Intel Core i7 2.8Ghz equipped with 4GB of main memory and a NVIDIA GeForce GTX 480 with 2.5GB of memory under a 64-bit Windows 7 Professional. The results of our tests are shown in Figure 15. From this graph, our approximated version is 4 orders of magnitude faster (10,000 times) than the MATLAB original algorithm, and 3 orders of magnitude faster (1,000 times) than the optimized C++ version of the original algorithm. Moreover, the approximated GPU version can achieve nearly 45 frames per seconds (22ms) for HD content (1920 × 1080; i.e., 2 Megapixel).

We have also evaluated the GPU implementation results in terms of quality. In this case, we have compared the GPU vs. the CPU versions of our technique. An example is shown in Figure 16 (see additional material for all visual comparisons [5]). The GPU implementation result can define a more

---

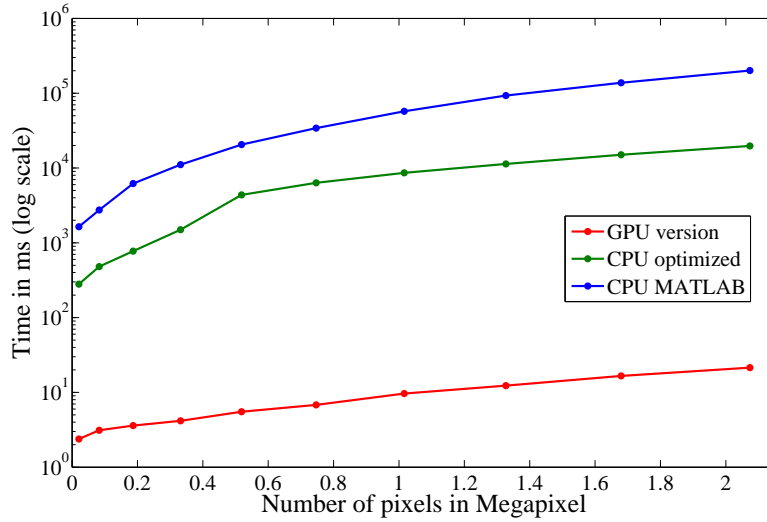[5]http://www.banterle.com/work/papers/ic_hybrid_tmo/additional_cpu_vs_

Figure 15: A timing comparison between our classic algorithm implemented in MATLAB, optimized in C++, and the GPU implementation. Note that the GPU version is 3 orders of magnitude faster than the optimized C++ and 4 orders of magnitude faster than the MATLAB version.

detailed zone map when compared with the CPU implementation; see the red square areas. On the other hand, the final look, of the two output images, may differ in the areas where the zone map is more detailed. This can be seen as a quality improvement, over the output of the original CPU implementation, where details may be better preserved.
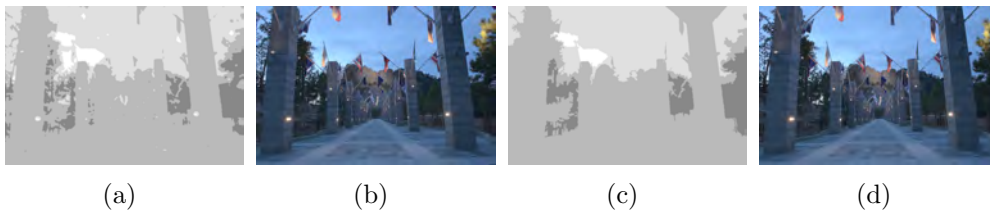

Figure 16: An example of differences between our GPU and CPU version of the algorithm on an image: (a) the zones map using our approximated algorithm on the GPU; (b) the tone mapped image using the GPU method; (c) the zones map using the original algorithm. Note that the red squares are the largest areas which differs from the GPU version; (d) the tone mapped image using the original algorithm.
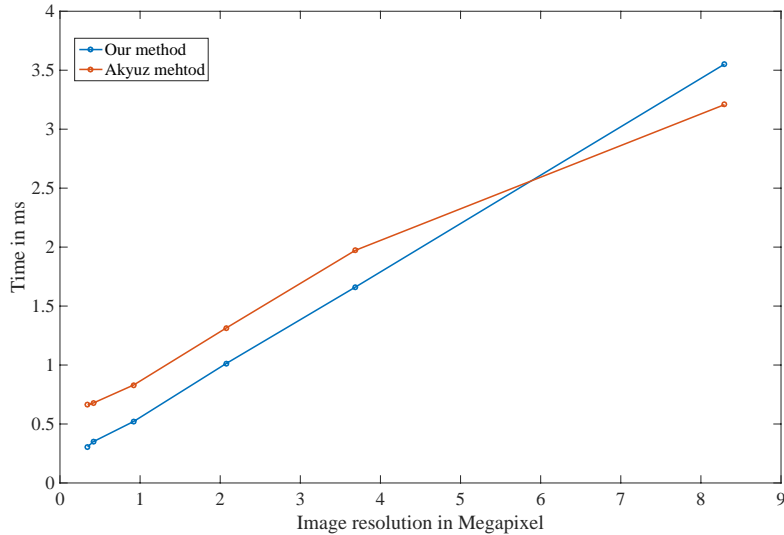

`gpu.pdf`

6.4.1 Evaluation

Figure 17: A timing comparison between our bilateral filtering based Reinhard et al.'s operator and Akyüz's method. Our method is faster for images lower than 4K resolution, but it starts to be a bit slower for 8K images.

(a)
(b)

Figure 18: An example of failure case of Akyüz's method [2]: (a) Akyüz's method, which shows some contouring artifacts (red and green boxes). This is due to the threshold test for selecting the best adaptation scale. (b) Our method, which does not have a threshold test, does not exhibit contouring (red and green boxes).

Timings. We evaluated the time performance of our algorithm and Akyüz [2], which is the state-of-the-art for accelerating the Reinhard et al.'s operator. In this test, the machine used was an Intel Xeon E5-2690 (2.9Ghz)

equipped with 32GB of main memory and a Nvidia TITAN GTX under a 64-bit Windows 7 Professional. The results of our tests are shown in Figure 17. From this graph, our method is slightly faster than Akyüz's one for images that have a resolution lower or equal to 4K (3840 × 2160), and it scales linearly with image resolution. However, Akyüz's method is faster for 8K images (7680 × 4320), and it scales sub-linearly with image resolution.

Quality. We compared, in terms of quality, our method and Akyüz's method against the original Reinhard et al.'s method using original Reinhard's source code and using the default parameters. Given that all images are tone mapped; i.e., LDR images, we used classic LDR quality metrics/indices such as PSNR, SSIM [45], and multi-scale SSIM (MS-SSIM)[44]. For this test, we used a dataset of 14 images; see the additional material [6]. Table 6.2 reports the average of the results for all 14 images; results for each image are in the additional material [7]. Although, our method performs slightly better, the results are very similar. However, our method does not produce contouring artifacts that Akyüz's method does in a few images; see Figure 18. These contouring artifacts are due to the fact that Akyüz's method has a hard test for determining the adaptation level for each pixel independently.

6.5   Extensions for Handling Video Tone Mapping

Flickering artifacts may be introduced when tone mapping HDR videos due to the computation of Z, thresholding and merging, and the use of static TMOs. In order to have a temporally coherent operator, both Z and the TMOs need to be temporally coherent. For the first task, we can use 3D iterative bilateral filtering across frames instead of iterative 2D bilateral filtering. For the second task, we can employ temporally coherent versions of the Drago et al. method [21] and Reinhard et al.'s operator [38]. This can be achieved by either using general frameworks [14] or smoothing statistics[28]. Some pilot experiments, using videos from the Stuttgart HDR video dataset [8], are available on youtube[9] or in the additional material [10].

---

[6]http://www.banterle.com/work/papers/ic_hybrid_tmo/additional_akyuz_vs_our.pdf

[7]http://www.banterle.com/work/papers/ic_hybrid_tmo/additional_akyuz_vs_our.pdf

[8]https://hdr-2014.hdm-stuttgart.de/

[9]https://youtu.be/7d1G2kc3XLA

[10]http://www.banterle.com/work/papers/ic_hybrid_tmo/video_tone_mapping_experiments.mp4

7   Conclusions and Future Work

We have presented a novel approach for tone mapping. Our work is an hybrid TMO which tackles the issue that existing TMOs do not achieve similar quality results on different luminance zones of an image. Combining existing TMOs in a selective fashion enables us to take advantage of just the best properties of existing TMOs and hence produce images that human observers more often rank as closer to a reference HDR image.

Such a hybrid approach helps to address two of the open issues in the area of HDR imaging discussed in Section 1. First, based on psychophysical experiment we can determine which existing TMOs are more suitable for a specific application. Second, the resulting hybrid TMO will be comprehensive that takes into account all the important aspects of the selected existing TMOs.

Our approach validated on numerous HDR images through a detailed psychophysical experiment, using an HDR display as the reference. We also confirmed the statistical significance of the performance of the hybrid compared to the two best single TMOs. We also compared our method against a similar recently proposed method [48]. From this comparisons, we found out that both methods have a similar behavior in selecting between Reinhard et al.'s operator [38] and Drago et al.'s operator [21] in the validated dynamic range on average. Given the simplicity of our approach, a look-up to a table, we can conclude that our method provides a faster selection.

Finally, we proposed a GPU-friendly version which is 100 times faster than the original optimized implementation on the CPU. This version can achieve a frame rate of roughly 45 frames per second on HD content, but it approximates the original algorithm for tone mapping. We evaluated our algorithm to determine if the introduced approximations were visible or not to the human eye. From our evaluation, we found out that our graphics hardware version of our method produces accurate results as the original algorithm.

Acknowledgements

References

[1] R. Achanta, A. Shaji, K. Smith, A. Lucchi, P. Fua, and Sabine Süsstrunk. Slic superpixels compared to state-of-the-art superpixel methods. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(11):2274–2282, 2012.

[2] Ahmet Oğuz Akyüz. High dynamic range imaging pipeline on the gpu. *Journal of Real-Time Image Processing*, pages 1–15, 2012.

[3] Ahmet Oğuz Akyüz, Roland Fleming, Bernhard E. Riecke, Erik Reinhard, and Heinrich H. Bülthoff. Do hdr displays support ldr content?: A psychophysical evaluation. *ACM Trans. Graph.*, 26(3), July 2007.

[4] Alessandro Artusi, Ahmet Oğuz Akyüz, Benjamin Roch, Despina Michael, Yiorgos Chrysanthou, and Alan Chalmers. Selective local tone mapping. In *IEEE International Conference on Image Processing*, September 2013.

[5] Alessandro Artusi, Jiří Bittner, Michael Wimmer, and Alexander Wilkie. Delivering interactivity to complex tone mapping operators. In Per Christensen and Daniel Cohen-Or, editors, *Rendering Techniques 2003 (Proceedings Eurographics Symposium on Rendering)*, pages 38–44. Eurographics, Eurographics Association, June 2003.

[6] Alessandro Artusi and Alan Chalmers. Visual attention in computer graphics. *Visual Perception: New Research*, page 289, 2008.

[7] Alessandro Artusi, Despina Michael, Benjamin Roch, Yiorgos Chrysanthou, and Alan Chalmers. A selective approach for tone mapping high dynamic range content. In *SIGGRAPH Asia 2011 Posters*, SA '11, pages 50:1–50:1, New York, NY, USA, 2011. ACM.

[8] Tunç O. Aydın, Rafal Mantiuk, Karol Myszkowski, and Hans-Peter Seidel. Dynamic range independent image quality assessment. In *ACM Transactions on Graphics (Proc. of SIGGRAPH)*, volume 27(3), 2008. Article 69.

[9] Francesco Banterle, Alessandro Artusi, Tunç O. Aydin, Piotr Didyk, Elmar Eisemann, Diego Gutierrez, Rafal Mantiuk, and Karol Myszkowski. Multidimensional image retargeting. In *SIGGRAPH Asia 2011 Courses*, SA '11, pages 15:1–15:612, New York, NY, USA, 2011. ACM.

[10] Francesco Banterle, Alessandro Artusi, Tunç O. Aydin, Piotr Didyk, Elmar Eisemann, Diego Gutierrez, Rafal Mantiuk, and Tobias Ritschel. Mapping images to target devices: spatial, temporal, stereo, tone, and color. In *Eurographics 2012 Tutorials*, EG '12. Eurographics Association, 2012.

[11] Francesco Banterle, Alessandro Artusi, Kurt Debattista, and Alan Chalmers. *Advanced High Dynamic Range Imaging: Theory and Practice.* AK Peters (CRC Press), Natick, MA, USA, 2011.

[12] Francesco Banterle, Alessandro Artusi, Elena Sikudova, Thomas Bashford-Rogers, Patrick Ledda, Marina Bloj, and Alan Chalmers. Dynamic range compression by differential zone mapping based on psychophysical experiments. In *Proceedings of the ACM Symposium on Applied Perception*, SAP '12, pages 39–46, New York, NY, USA, 2012. ACM.

[13] Francesco Banterle, Massimiliano Corsini, Paolo Cignoni, and Roberto Scopigno. A low-memory, straightforward and fast bilateral filter through subsampling in spatial domain. *Computer Graphics Forum*, 31(1):19–32, February 2012.

[14] Ronan Boitard, Kadi Bouatouch, Rémi Cozot, Dominique Thoreau, and Adrien Gruson. Temporal coherency for video tone mapping. In Alexander M. J. van Eijk, Christopher C. Davis, Stephen M. Hammel, and Arun K. Majumdar, editors, *Proc. SPIE 8499, Applications of Digital Image Processing XXXV*, pages 84990D–84990D–10, San Diego, October 2012.

[15] Ronan Boitard, Rémi Cozot, Dominique Thoreau, and Kadi Bouatouch. Zonal brightness coherency for video tone mapping. *Signal Processing: Image Communication*, 29(2):229–246, feb 2014.

[16] Martin Cadik. Perception motivated hybrid approach to tone mapping. In *Proceedings WSCG*, 2007.

[17] Martin Cadik, Michael Wimmer, Laszlo Neumann, and Alessandro Artusi. Image attributes and quality for evaluation of tone mapping operators. In *Proceedings of the 14th Pacific Conference on Computer*

Graphics and Applications, pages 35– 44, Taipei, Taiwan, 2006. National Taiwan University Press.

[18] Herbert A. David. *The Method of Paired Comparisons, 2nd ed.* Oxford University Press, 1988.

[19] Dolby. Dolby-DR37P. *http://www.dolby.com/promo/hdr/technology.html*, December 2008.

[20] Frédéric Drago, William L. Martens, Karol Myszkowski, and Hans-Peter Seidel. Perceptual evaluation of tone mapping operators. In *In Proceedings of the SIGGRAPH 2003 conference on Sketches and applications*, New York, NY, USA, 2003. ACM.

[21] Frédéric Drago, Karol Myszkowski, Thomas Annen, and Norishige Chiba. Adaptive logarithmic mapping for displaying high contrast scenes. Computer Graphics Forum (Proceedings of Eurographics '03), 2003.

[22] Frédo Durand and Julie Dorsey. Fast bilateral filtering for the display of high dynamic range image. In *Proceedings of SIGGRAPH 2002*, pages 257–265, 2002.

[23] Mark Fairchild. MEET ICAM. http://www.cis.rit.edu/mcsl/icam/, 2004.

[24] Mark Fairchild and Garrett Johnson. Meet icam: A next-generation color appearance model. In *IST SID Tenth Color Imaging Conference, 2002.*, 2002.

[25] Alan Gilchrist, Christos Kossyfidis, Frederick Bonato, Tiziano Agostini, Joseph Cataliotti, Xiaojun Li, Branka Spehar, Vidal Annan, and Elias Economou. An anchoring theory of lightness perception. *Psychological Review*, 106(4):795–834, October 1999.

[26] Rafael C. Gonzalez and Richard E. Woods. *Digital image processing.* Prentice-Hall, 2. ed. edition, 2002.

[27] Nolan Goodnight, Rui Wang, Cliff Woolley, and Greg Humphreys. Interactive time-dependent tone mapping using programmable graphics hardware. In *Rendering Techniques*, pages 26–37, 2003.

[28] Sing Bing Kang, Matthew Uyttendaele, Simon Winder, and Richard Szeliski. High dynamic range video. *ACM Trans. Graph.*, 22(3):319–325, July 2003.

[29] Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel. Lightness perception in tone reproduction for high dynamic range images. *Computer Graphics Forum*, 24(3), 2005.

[30] Jiangtao Kuang, Hiroshi Yamaguchi, Changmeng Liu, Garrett M. Johnson, and Mark D. Fairchild. Evaluating hdr rendering algorithms. *Color Image Conference*, pages 315–320, 2004.

[31] Jiangtao Kuang, Hiroshi Yamaguchi, Changmeng Liu, Garrett M. Johnson, and Mark D. Fairchild. Evaluating hdr rendering algorithms. *ACM Trans. Appl. Percept.*, 4(2):9, 2007.

[32] Patrick Ledda, Alan Chalmers, Tom Troscianko, and Helge Seetzen. Evaluation of tone mapping operators using a high dynamic range display. *ACM Trans. Graph.*, 24(3):640–648, 2005.

[33] Dani Lischinski, Zeev Farbman, Matt Uyttendaele, and Richard Szeliski. Interactive local adjustment of tonal values. *ACM Trans. Graph.*, 25(3):646–653, 2006.

[34] Tom Mertens, Jan Kautz, and F. Van Reeth. Exposure fusion. In *Proceedings of Pacific Graphics 2007*, pages 382–390. IEEE, 2007.

[35] Sumanta N. Pattanaik, Jack Tumblin, Hector Yee, and Donald P. Greenberg. Time-dependent visual adaptation for fast realistic display. In *Computer Graphics (SIGGRAPH '00 Proceedings)*, pages 47–54. Addison Wesley, July 2000.

[36] Egon S. Pearson and Herman O. Hartley. *Biometrika tables for statisticians 3rd ed., vol. 1*. Cambridge University Press, 1988.

[37] Erik Reinhard, Wolfgang Heidrich, Paul Debevec, Sumanta N. Pattanaik, Gregory Ward, and Karol Myszkowski. *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann, second edition edition, May 2010.

[38] Erik Reinhard, Michael Stark, Peter Shirley, and Jim Ferwerda. Photographic tone reproduction for digital images. In *Proceedings of SIGGRAPH 2002*, pages 267–276, 2002.

[39] Marco Slomp and Manuel Oliveira. Real-time photographic local tone reproduction using summed-area tables. In *CGI '08: Computer Graphics International 2008*, 2008.

[40] Yu-Wing Tai, Jiaya Jia, and Chi-Keung Tang. Soft color segmentation and its applications. *IEEE Trans. Pattern Anal. Mach. Intell.*, 29(9):1520–1537, September 2007.

[41] Carlo Tomasi and Roberto Manduchi. Bilateral filtering for gray and color images. In *ICCV '98: Proceedings of the Sixth International Conference on Computer Vision*, pages 839–847, Washington, DC, USA, 1998. IEEE Computer Society.

[42] Jack Tumblin, Jessica K. Hodgins, and Brian K. Guenter. Two methods for display of high contrast images. *ACM Trans. Graph.*, 18(1):56–94, 1999.

[43] Jack Tumblin and Holly E. Rushmeier. Tone reproduction for realistic images. *IEEE Computer Graphics and Applications*, 13(6):42–48, November 1993.

[44] Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, April 2004.

[45] Zhou Wang, Eero P. Simoncelli, and Alan C. Bovik. Multiscale structural similarity for image quality assessment. In *Conference Record of the Thirty-Seventh Asilomar Conference on Signals, Systems and Computers*, volume 2, pages 1398–1402, November 2003.

[46] Gregory Ward Larson, Holly Rushmeier, and Christine Piatko. A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes. *IEEE Transactions on Visualization and Computer Graphics*, 3(4):291–306, October 1997.

[47] Yangli Hector Yee and Sumanta Pattanaik. Segmentation and adaptive assimilation for detail-preserving display of high-dynamic range images. *The Visual Computer*, 19(7-8), 2003.

[48] Hojatollah Yeganeh and Zhou Wang. Objective quality assessment of tone mapped images. *IEEE Transactions on Image Processing*, 22(2):657–667, February 2013.

[49] Akiko Yoshida, Volker Blanz, Karol Myszkowski, and Hans-Peter Seidel. Perceptual evaluation of tone mapping operators with real-world scenes. In Bernice E. Rogowitz, Thrasyvoulos N. Pappas, and Scott J. Daly, editors, *Human Vision and Electronic Imaging X, IST SPIE's*

17th Annual Symposium on Electronic Imaging, volume 5666 of *SPIE Proceedings Series*, pages 192–203, San Jose, USA, January 2005. SPIE.