Real-time tone mapping

An evaluation of color-accurate methods for luminance compression

Sebastian Dille,
Arnulph Fuhrmann,
Gregor Fischer

TH Köln,
Betzdorferstr. 2,
D-50679 Köln
http://www.th-koeln.de

Abstract Recent advances in real-time rendering enable virtual production pipelines in a broad range of industries. These pipelines are based on a fast, latency-free handling in addition to the accurate appearance of results. This requires the use of high dynamic range rendering for photo-realistic results and a tone mapping operator for a matched display.

However, color-accurate tone mapping of real-time rendering has not yet been the focus of most publications. Tone mapping furthermore tends to cause image appearance effects, which causes changes in the perception of colors and therefore a discrepancy between displayed and real-world products.

To compensate for these changes in appearance, seven tone mapping operators with different approaches of color correction are evaluated in this paper. They are implemented as fragment shaders to facilitate real-time processing. In addition, their performance and subjective accuracy is measured.

As a result, the filmic tone mapping operator as a global sigmoid-based operator performs best and can be recommended for the usage in color-critical real-time applications.

1 Introduction

Conventional displays are technically limited in displaying naturally possible luminance values. The resulting dynamic range is comparably low and therefore referred to as LDR - low dynamic range. On the other hand, the latest developments in rendering algorithms enable physically-correct lighting simulations in real-time with up to 60 frames per second (FPS). By implementing real-life light sources and reflectance properties, scenes with a high dynamic range are created, requiring HDR-rendering. This powerful way of image synthesis goes far beyond the scope of pure entertainment. In various industrial sectors, HDR-rendering is used to design realistic models of future products. This so called virtual prototyping allows specific development steps to be taken in a very early stage without the need of a physical prototype. Changes can be made quickly and with little expenses.

Nevertheless, there is one crucial condition: the appearance of the image needs to match with reality as good as possible. There are two aspects affecting the image appearance. First, the rendering itself can be further enhanced to behave as realistic as possible. Second, the way of displaying the resulting images on a conventional display can be optimized. This paper focusses on the latter.

The procedure of compressing the luminance values for LDR displays is called tone mapping. There are several approaches to tone mapping with different intents, some emphasize preserving the details, others creating the most "natural" look etc..

Though luminance compression works great, it leads to a second problem: If the luminance of the image is changed, so is the color appearance.

Colors are perceptions created in the human brain. Therefore, colors are not absolute. The appearance of a colored surface to the viewer depends on different correlates, one of them is the surface luminance. If these color appearance correlates are changed, different color appearance effects occur, causing the viewer to perceive a different color.

2 Related work

Tone mapping itself is an established discipline. A few methods to handle high dynamic ranges existed already for analogue photography and postproduction like Adam's "dodging and burning" technique [1]. Since the beginning of digital image processing, the number of different tone mapping operators has steadily grown.

Current tone mapping operators (TMOs) can be separated into four groups. The first two groups of TMOs both work spatially: While global TMOs like Larson and Rushmeier's approach [11] only take overall image parameters into consideration, local TMOs compress the luminance of a pixel depending on its surrounding, which is used e.g. in Rheinhard's photographic TMO [19].

A third group is formed by operators working in the frequency domain. Tomasi and Manduchi for example have introduced an algorithm, where low and high frequencies are compressed differently to preserve details [20].

Lastly, Fattal's method [6] is an example of a fourth kind of TMO, which compresses luminance values based on the gradients of an image.

A good overview about these four groups is given by Reinhard et al. [18].

2.1 Tone mapping for video

While most tone mapping operators focus on compressing single images, several solutions for tone mapping of a series of images or video tone mapping respectively have been presented.

Kiser et al. have introduced a simple TMO based on histogram adaption for HDR video [9]. It is combined with a flicker-reduction solution and can process a 1080p resolution video at 30 FPS. However, it does not address color correction nor was it tested for higher frame rates.
A different approach has been presented by Boitard et al. [4]. It mainly compensates for luminance inconsistency between different frames of an image sequence and can be added to any TMO. To eliminate high luminance changes, it requires the maximum key value of the video sequence, which can only be computed during a preprocessing step. Therefore, it is not suitable for real-time applications.

2.2 Tone mapping evaluation

Since every TMO has its own characteristics resulting in specific strengths and weaknesses, it becomes inevitable to compare different approaches in order to find the most suitable one for a certain application.
Yoshida et al. were one of the first to compare tone-mapped images with corresponding real-world scenes [21]. They evaluated the naturalness of the reproduction among other things, though they did not compare the accuracy of color reproduction.
Ledda et al. also compared tone-mapping operators for still images, using a large number of test subjects to generate reliable data [12]. They also took real-world references into account by displaying the corresponding HDR image on an HDR-capable display.
Petit et al. have evaluated four candidates for video tone mapping, thereby focusing on overall quality and reproduction accuracy [15]. Lastly, a much wider evaluation in terms of tested operators was taken by Eilertsen et al. with an additional evaluation of image artifacts like noise and ghosting [5]. In contrast to [21] and [12], these two approaches do not take the real-world reference into consideration. Furthermore, though both approaches focus on video tone mapping, they do not distinguish between operators, which omit preprocessing, and those that do not. Hence, the results cannot be adopted directly for real-time applications.
In our evaluation we combine a performance analysis of different existing TMOs with a qualitative user study. This way we shall find a TMO fast enough for real-time purposes which, on top of that, simultaneously reproduces color as realistically as possible. Additionally, our evaluation focuses on tone mapping for indoor visualization purposes, which are based on physically correct rendering. The operators thus have to excel for medium dynamic ranges, whereas very large differences in luminance levels usually don't occur and can be neglected.

3 Tone mapping operators

Out of the existing field of TMOs, six candidates seemed promising with regard to both their color processing ability and execution speed, and were evaluated. They are described in the following sections. Additionally, a seventh one is explained in section 3.7, which has been developed to test a new combination of tone mapping and color correction.

3.1 Calibrated image appearance reproduction

Reinhard et al. have developed a model for the accurate color rendering of an image with respect to the viewing conditions [17]. It is called "Calibrated image appearance reproduction", abbreviated with CIAR in the following.
It uses a photoreceptor response function but with additional influence of the size of the pupil and the bleaching effect. By integrating viewing conditions like the display white point, the surrounding white point and maximum and minimum brightness levels respectively, an ideal reproduction of contrast and colors is intended.

3.2 Color appearance in high dynamic range imaging

A different approach is proposed by Reinhard in "Color appearance in high dynamic range imaging" [2].This method is a combination of a color appearance model with a tone mapping operator.
By applying the CAM, the image is first adjusted to match the viewing conditions. Afterwards, the luminance is reset to the original value in case of any CAM-related altering, before a tone mapping operator compresses the image.
The proposed method is designed to work with any combination of CAM and TMO. For this evaluation, the CIECAM02 algorithm [13] was chosen as well as the local variant of Reinhards' photographic TMO [19].
The algorithm also incorporates the "dodging-and-burning"-technique by using the surrounding luminance of the pixel, determined by a gaussian filter.
Below, this operator is called "CAMTMO".

3.3 Color correction for tone reproduction

A third operator is based on the method described by Pouli et.al. in "Color Correction for Tone Reproduction" [16]. It combines the photographic tone mapping operator from Reinhard et al. [19] with a color correction procedure of the resulting image. Here, this operator is referred to as "ColorCorrectedTMO".
The luminance compression is also based on a photoreceptor response function. A special characteristic of the photographic tone mapper is the use of a digital version of the "dodging-and-burning" technique [1] to further expose especially dark areas or darken especially bright ones respectively.
The color corrected version now adds an adjustment of the chroma of the image to counter a potential desaturation caused by the luminance compression, while keeping the compressed luminance intact.
For this reason, the original and the tone mapped image are first converted to CIE-LCh-coordinates. Second, the hue angle h of the original image is copied to the compressed one.

3.4 Filmic tone mapping

The filmic tone mapping operator was introduced by John Hable [7]. It is a global tone mapping operator, which imitates the response curve of a film negative material. It is

designed to be capable of reproducing soft transitions in bright image areas instead of hard clipping, and at the same time keeping saturated colors in dark areas.
For this reason, it uses several given parameters to get the final color for a given input value. These parameters are listed in table 1.

Table 1: One set of parameters for the film response curve proposed by John Hable

Name	Equivalent	Value
sStrength	shoulder strength	0.22
linStrength	linear strength	0.30
linAngle	linear angle	0.10
toeStrength	toe strength	0.20
toeNumerator	toe numerator	0.01
toeDenominator	toe denominator	0.30
toeAngle	toe angle	$\frac{toeNumerator}{toeDenominator}$
white	linear white point	11.2

3.5 iCAM06

The purpose of the iCAM algorithm is the accurate prediction of the attributes of the human visual system for the widest range of images as possible. An updated version from 2006 is used for this comparison [10].
As a CAM, iCAM06 in general is similar to CIECAM02. It incorporates a chromatic adaption, luminance compression based on photoreceptor behaviour and an adjustment of the appearance correlates.
Different to CIECAM02, iCAM06 uses a bilateral filter to split the image into a base layer and a detail layer, which are processed differently. This way, the details in the image shall be preserved.
Also, in iCAM06 the IPT color space is used to adjust the appearance correlates whereas in CIECAM02 the use of the HPE color space is intended.

3.6 Local laplacian filters

A completely different approach was introduced by Paris, Hasinoff and Kautz [14, 3]. It is based on local Laplacian filters and therefore referred to as ”LocalLaplacianTMO” in this paper.
As a first step, the corresponding Laplacian pyramid is constructed out of the image. Now each coefficient of each pyramid level is compared with its corresponding Gaussian coefficient. If the difference is greater than a predetermined threshold, an edge is present at the given location. In case the difference is smaller, the pixel represents details. Thirdly, for every Laplacian coefficient, the corresponding pixel region of the original image is determined. This region is then remapped, with details being treated differently than edges.

For tone mapping, the resulting luminance values are additionally treated with a simple gamma correction. Lastly, the image is reconstructed from the Laplacian coefficients.

3.7 Chroma-invariant tone mapping

The new approach was developed based on the following principles:

- The dynamic range shall be compressed by a photoreceptor model, without any additional effects like blending or similar.
- To compensate for any appearance effects, the color appearance correlates of the tone mapped image shall be set to the original ones.

These principles are strongly related to color correction for tone reproduction [16]. However, both the luminance compression and the subsequent color correction differ significantly.
To reduce the dynamic range, a few changes were made to Reinhard's photographic TMO, yielding the following formula:

$$L_{compressed}(x, y) = white \frac{luminance(x, y)^n}{1 + y_{Surround}^n}. \quad (1)$$

Here, *white* is a user defined upper luminance limit. Any greater value is set to white. With $y_{Surround}$ the surrounding luminance of the pixel is also taken into account. Additionally, the exponent n was added to the formula to control the contrast adjustment.
For the color correction, both images, original and compressed, are transferred to the HSL color space to directly manipulate the hue and saturation. Next, the compressed hue angle and saturation are set to the uncompressed values without further adaptation and the resulting image is transferred back to the RGB color space.
This new variant is referred to as "ChromaInvariantTMO".

4 Implementation

All tested tone mapping operators are implemented as GLSL-shaders. Thus, they process each pixel individually. Though this takes advantage of the full processing power of a parallel working GPU, global image parameters like the maximum or minimum luminance cannot be calculated in one single shader nor can the image be filtered sufficiently using only one shader.
Using render-to-texture- and stream-reduction techniques, this problem can be solved. With multiple calls of a preprocessing-shader, the image is first reduced to one single pixel, which contains the maximum luminance in the red channel, the minimum luminance in the green channel and the log-average luminance in the blue channel. The frame buffer containing the pixel is then used as a second input texture to the tone mapping shader. The achieved rendering speeds are summarized in table 4.

5 Evaluation

5.1 Apparatus

Table 2 and 3 are showing the platforms used to analyse the performance and the perceived quality respectively.

Table 2: Platform used for performance measurement

Component	Data
CPU	AMD FX(tm)-8320 Octacore 3,5 GHz
RAM	16 GB
GPU	GeForce GTX 650 128-bit GDDR5 1024MB
OS	Ubuntu 14.04 64 Bit

Table 3: Platform used for user study

Component	Specification
CPU	Intel i5-4440 Quadcore 3.1Ghz
RAM	8 GB
GPU	GeForce GTX 750Ti 128-bit GDDR5 2048MB
OS	Windows 8.1 64 Bit
Display	Dell Ultrasharp 2412M

5.2 Performance

The execution speed of a given program is an important indicator to judge how well a program is capable of fulfilling real-time operations. In computer graphics, current real-time applications are executed with 75, 90 or even 120 FPS, although the minimum requirement is usually defined by a frame rate of 60 frames per second or 16 ms per frame respectively. This shall also be the goal for the evaluated shaders.
To analyse the performance, the exact times before and after the render call were taken. By subtracting both measurements, one gets the amount of time for transferring the data from CPU to GPU, executing the renderer and applying the shader. Since the transfer rate and rendering process are the same for every shader in the study, the difference in time relates directly to the performance of a specific shader.
Additionally, the execution times for one frame were averaged over a total time of one minute for every shader.

5.3 User study

Color appearance is highly subjective [8]. For this reason, the reproduction quality cannot be fully evaluated without a user study. From the introduced set of tone-mapping

operators, the four most promising candidates were preselected to be evaluated further. These are specifically the 'CIAR' algorithm, the 'chroma-invariant' TMO, the 'Filmic' TMO and the 'ColorCorrected' TMO.
The chosen operators were tested via a pairwise comparison similar to Ledda et al. [12], as it proofs to be easier for the subjects than ranking all operators at once, and it omits the effect of fatigue. A group of 12 people would individually compare two operators with one another and select the better one. To test every operator against each other, the test subjects would select a total of six sets with two operators and answer the same five questions for each set:

- Which operator preserves more details ?
- Which operator shows more static artifacts ?
- Which operator shows more dynamic artifacts ?
- Which operator produces a better reproduction of the shown ColorChecker ?
- Which operator creates an overall more realistic image ?

For the artefact-related questions, the quality could additionally be rated using a scale from zero points for the heaviest perceived disturbance up to four points for a flawless result.
A score system was developed to rank the operators based on each persons decisions. Whenever one operator outperformed the second one, it got two points, the losing one zero. In case of a tie, both operators got one point. Regarding the artefact related questions, the scale points decided the score.
For the study, a specific test program has been developed:

Figure 1: The program's user interface showing the live rendering (left), the reference (top right) and the control buttons (bottom right).

Figure 1 shows its interface. In the left part of the program window, a live rendering produced by a given tone mapping operator is displayed. The scene consists of two

differently illuminated planes, each textured with an image of a ColorChecker. Through rendering, this image is first produced as an HDR image and then tone mapped. The HDR-skydome acts as background and can be replaced by an all-grey environment. On the right, the same texture image is displayed as a reference, not rendered and not tonemapped. By pressing one of the buttons underneath the reference, the active set can be chosen. With the radio button in the second line, the operators can be switched, while the second radio button toggles between the two different backgrounds. To evaluate dynamic artifacts, the test subjects could move the camera interactively using the mouse. This could also be used to examine different parts of the scene as well as to inspect details in the scene more closely by changing the focal length.
The whole study took place in a grey-painted room lit by two halogen-metal vapour lamps, with a spectrum close to D65. That way, the influence of changing lighting conditions could be reduced to a minimum.

6 Results

The described procedures yield the following results:

6.1 Performance

Table 4 shows the average render performance:

Table 4: Shader performance, optimized

Shader	ø Runtime
StandardTMO	1.4 ms
ColorCorrectedTMO	2.3 ms
CAMTMO	2.6 ms
iCam06	2.4 ms
CIAR	2.7 ms
FilmicTMO	2.5 ms
ChromaInvariantTMO	2.6 ms
LocalLaplacianTMO	4.2 ms

It is clearly visible that with optimization, every operator can fulfil the real-time condition. Also, the individual render times are very similar. Only the approach based on local laplacian filtering turns out to be slower, caused by the calculation of the laplacian pyramids. But with slightly over 4 ms per frame, it also runs in real-time.

6.2 Reproduction quality

Regarding the perceived quality, the tested operators show a varying performance. The new approach is slightly preferred in terms of luminance compression, where the filmic operator shows the weakest results. While the average score for the visibility of details is similar for all four operators (see figure 2), great differences appear e.g. regarding static artifacts.

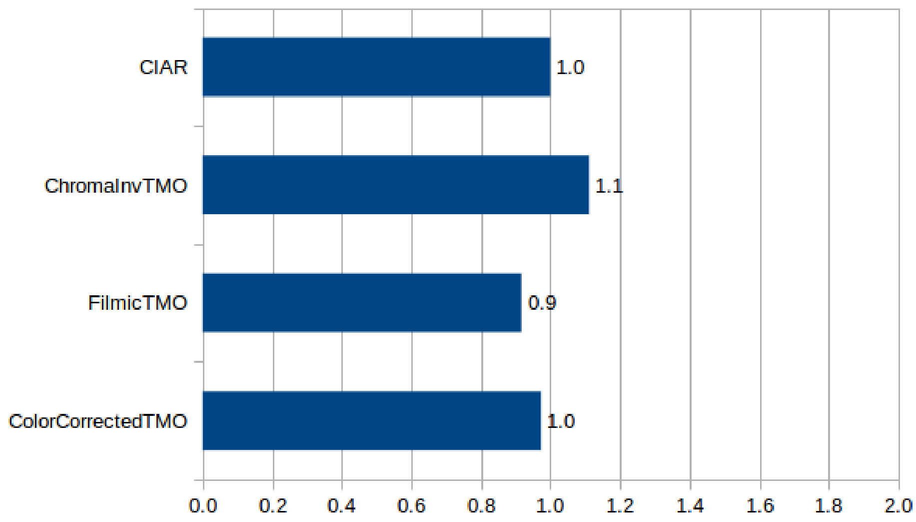

Figure 2: Average score for the reproduction of details

Here, the filmic operator clearly scores best with nearly no visible artifacts, followed with a statistically significant distance ($p < 5\%$) by the CIAR method, while the other operators score significantly lower. Figure 3 shows the corresponding results.

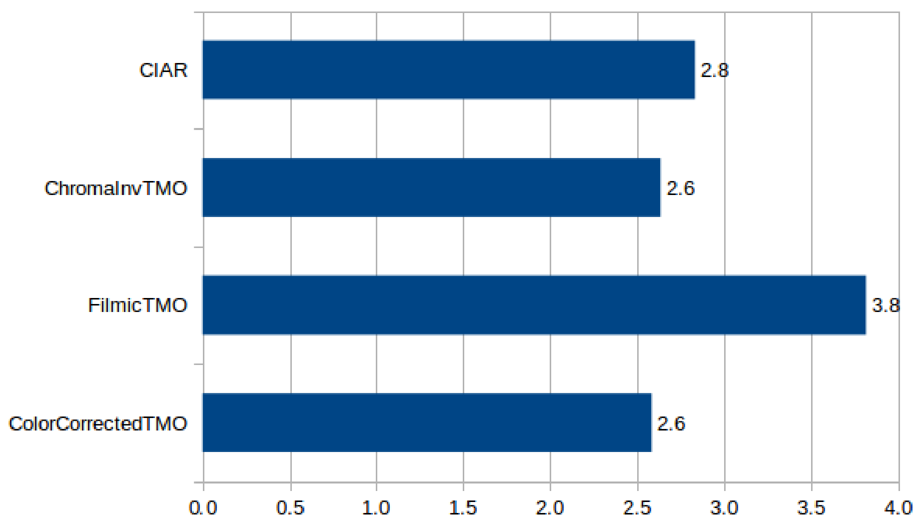

Figure 3: Average score for the impact of static artefacts

The same applies for dynamic artifacts, as shown in figure 4.
Again, the filmic operator scores best. However, the fact that every operator gets a comparably good score with at least 81% of the maximum confirms no significant differences between the single operators.

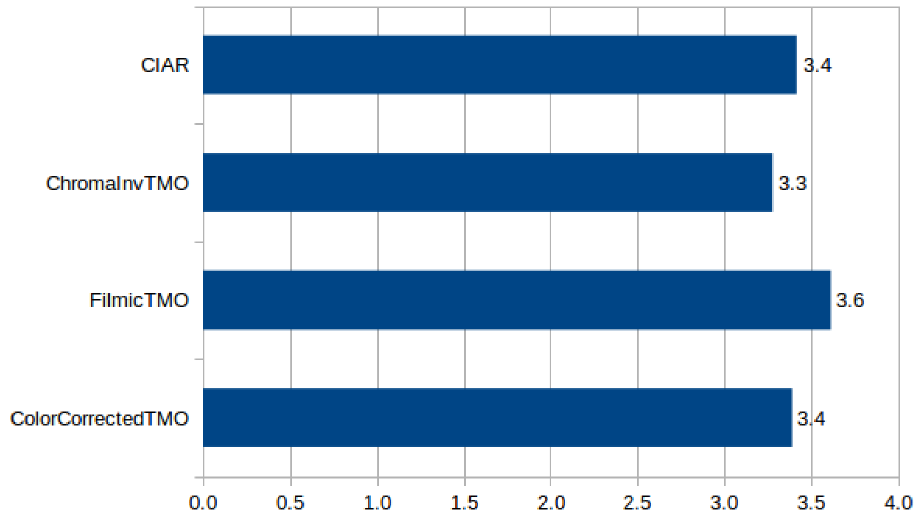

Figure 4: Average score for the impact of dynamic artefacts

A very large variety in score was found for the reproduction quality of color. Figure 5 displays the results:

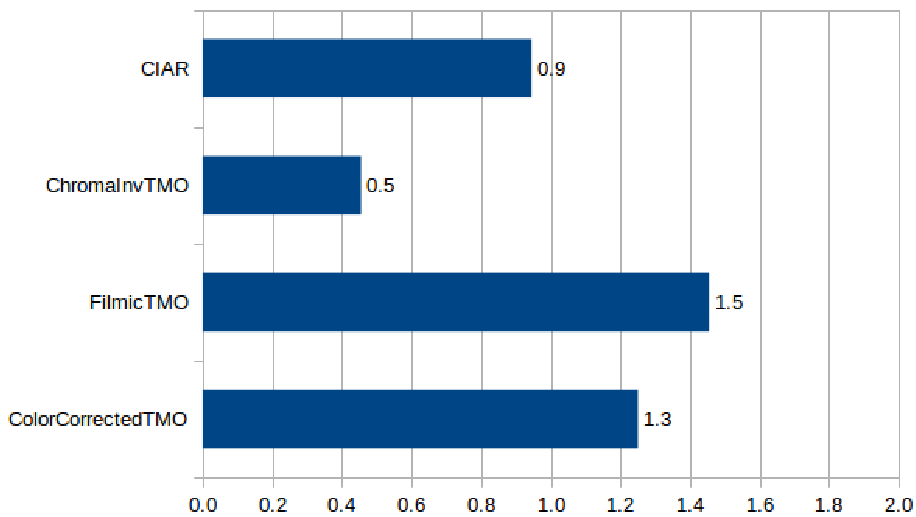

Figure 5: Average score for the reproduction of colour

The filmic operator results best also for this criterion and achieves 72% of the maximum score or approximately 1.45 points. It is followed by the color corrected approach with 62%. With a percentage of 22%, the new approach falls significantly behind ($p < 5\%$). The results for the natural appearance, or perceived authenticity, are similar. Once again, the filmic operator achieves over 70% and performs best, followed by the color corrected approach. Though this time the difference is bigger, as shown in figure 6. The weakest candidate concerning a realistic appearance is the CIAR algorithm.

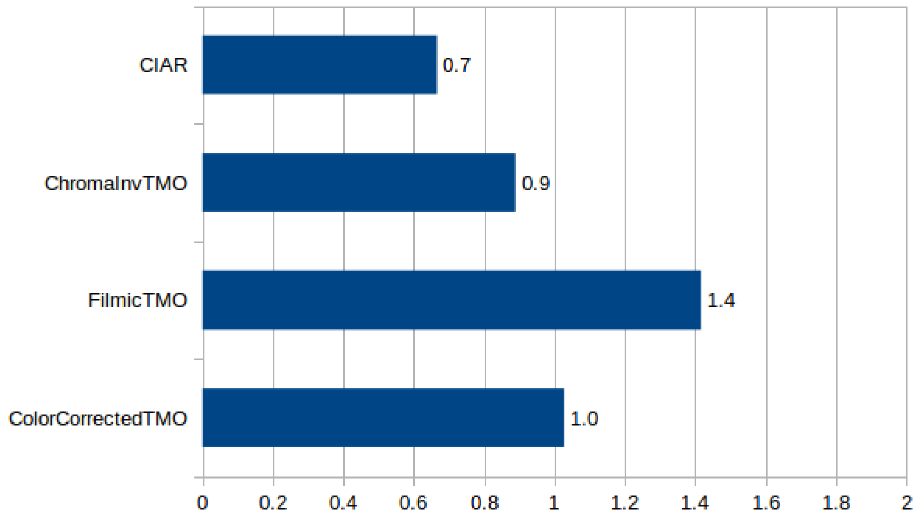

Figure 6: Average score for the perceived authenticity

By combining all five criteria, the following overall score can be calculated (see figure 7):

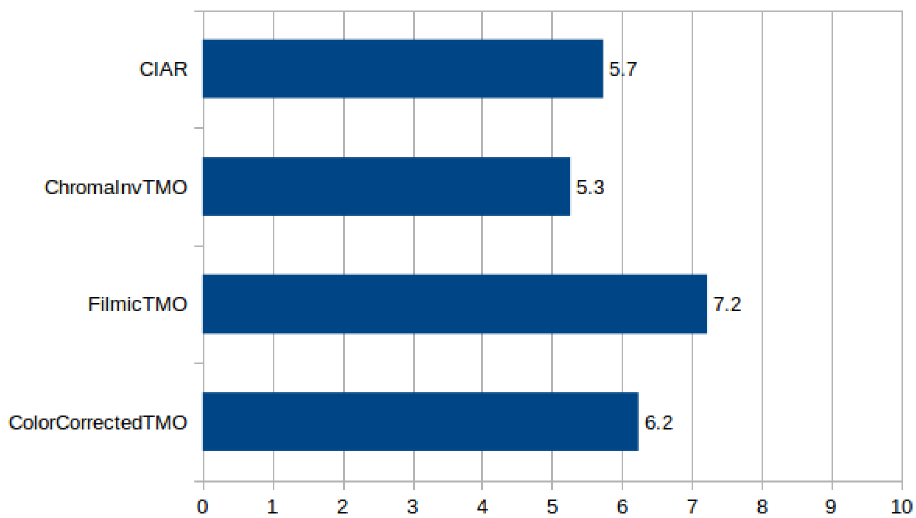

Figure 7: Average overall score

With 72%, the filmic operator produces the best results. As second best candidate, the color corrected TMO achieves a score of 62%, closely followed by CIAR with 57%. The new approach shows the weakest results and gets an overall score of 5.3 points or 53% respectively.
It is noticeable that although the last three operators get results close to each other, the lead of the filmic operator compared to CIAR and the chroma-invariant approach is significant again ($p < 1\%$ for both).

7 Conclusions and future work

In this paper, seven different tone mapping operators with color correction capabilities have been introduced. Their performance was measured and the quality of their color reproduction was compared.
All in all, the following conclusions can be drawn:

- All operators have an overall good performance. Interactive applications with real-time requirements can be implemented using any of those operators, with a difference in running time of less than 1.5ms per frame between the best and weakest candidate.

- Concerning the subjective quality, the operators show greater differences. Here, the Filmic TMO clearly performs best. It can therefore be recommended for future use in real-time rendering.

Although this study focuses on medium dynamic range scenarios, the obtained results should be reviewed for very high dynamic ranges in future work. Since global TMOs in general do not work well for extreme contrast ratios, the results may be different then.

References

[1] Ansel Adams and Robert Baker. *The print*. The Ansel Adams photography series. Bulfinch, 1995.

[2] Ahmet Oguz Akyüz and Erik Reinhard. "Color appearance in high-dynamic-range imaging". In: *Journal of Electronic Imaging* 15.3 (July 2006).

[3] Mathieu Aubry et al. "Fast Local Laplacian Filters : Theory and Applications". In: *ACM Trans. on Graphics* VV (2014), p. 15.

[4] Ronan Boitard et al. "Temporal coherency for video tone mapping". In: *Proc. SPIE 8499, Applications of Digital Image Processing XXXV*. Vol. 33. 2012.

[5] Gabriel Eilertsen et al. "Evaluation of tone mapping operators for HDR-video". In: *Computer Graphics Forum* 32.7 (2013), pp. 275–284.

[6] Raanan Fattal, Dani Lischinski, and Michael Werman. "Gradient domain high dynamic range compression". In: *ACM Transactions on Graphics* 21 (2002).

[7] John Hable. "Uncharted 2: HDR Lighting". In: *Game Developers Conference*. 2010, p. 56.

[8] R. W. G. Hunt. *The Reproduction of Colour*. Vol. 37. The Wiley-IS&T Series in Imaging Science and Technology. Wiley, 1969, p. 114.

[9] Chris Kiser et al. "Real time automated tone mapping system for HDR video". In: *Proc. of the IEEE Int. Conference on Image Processing, IEEE*. 2012.

[10] Jiangtao Kuang, Garrett M. Johnson, and Mark D. Fairchild. "iCAM06: A refined image appearance model for HDR image rendering". In: *Journal of Visual Communication and Image Representation* 18.5 (Oct. 2007), pp. 406–414.

[11] Gregory Ward Larson, Holly Rushmeier, and Christine Piatko. “A visibility matching tone reproduction operator for high dynamic range scenes”. In: *IEEE Transactions on Visualization and Computer Graphics* 3 (1997), pp. 291–306.

[12] Patrick Ledda et al. “Evaluation of Tone Mapping Operators Using a High Dynamic Range Display”. In: *ACM Trans. Graph.* 24.3 (July 2005), pp. 640–648. ISSN: 0730-0301.

[13] M. R. Luo and C. Li. “CIECAM02 and Its Recent Developments”. In: *Advanced Color Image Processing and Analysis*. 2013, pp. 19–58.

[14] S Paris, Sw Hasinoff, and J Kautz. “Local Laplacian filters: edge-aware image processing with a Laplacian pyramid”. In: *ACM Trans. Graph.* (2011).

[15] Josselin Petit and Rafał K. Mantiuk. “Assessment of video tone-mapping: Are cameras’ S-shaped tone-curves good enough?” In: *Journal of Visual Communication and Image Representation* 24.7 (Oct. 2013), pp. 1020–1030.

[16] Tania Pouli, Alessandro Artusi, and Francesco Banterle. “Color Correction for Tone Reproduction”. In: *Color and Imaging Conference* (2013), pp. 215–220.

[17] Erik Reinhard et al. “Calibrated image appearance reproduction”. In: *ACM Transactions on Graphics* 31.6 (2012), p. 1.

[18] Erik Reinhard et al. *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting (The Morgan Kaufmann Series in Computer Graphics)*. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2005.

[19] Erik Reinhard et al. “Photographic tone reproduction for digital images”. In: *ACM Transactions on Graphics* 21 (2002).

[20] C. Tomasi and R. Manduchi. “Bilateral filtering for gray and color images”. In: *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)* (1998).

[21] Akiko Yoshida et al. “Perceptual Evaluation of Tone Mapping Operators with Real-World Sceness”. In: *Human Vision and Electronic Imaging X, IS & T/SPIE’s 17th Annual Symposium on Electronic Imaging (2005)*. Ed. by Bernice E. Rogowitz, Thrasyvoulos N. Pappas, and Scott J. Daly. Vol. 5666. SPIE Proceedings Series. San Jose, USA: SPIE, Jan. 2005, pp. 192–203.