

Color Correction for Tone Reproduction

Tania Poulil^{1,5}, Alessandro Artusi², Francesco Banterle³, Ahmet Oğuz Akyüz⁴, Hans-Peter Seidef and Erik Reinhard^{1,5}

¹Technicolor Research & Innovation, France, ²Department of Computer Science and Applied Mathematics, University of Girona, Spain, ³Visual Computing Lab ISTI-CNR, Italy, ⁴Department of Computer Engineering, Middle East Technical University, Turkey, ⁵Max-Planck-Institut für Informatik, Germany

Abstract

High dynamic range images require tone reproduction to match the range of values to the capabilities of the display. For computational reasons as well as absence of fully calibrated imagery, rudimentary color reproduction is often added as a post-processing step rather than integrated into the tone reproduction algorithm. However, in the general case this currently requires manual parameter tuning, although for some global tone reproduction operators, parameter settings can be inferred from the tone curve. We present a novel and fully automatic saturation correction technique, suitable for any tone reproduction operator, which exhibits better color reproduction than the state-of-the-art and we validate its comparative effectiveness through psychophysical experimentation.

Introduction

Recent advances in both capture and display technologies allow images of a much wider dynamic range to be photographed, manipulated and displayed, better capturing the light of natural scenes and giving artists unparalleled freedom. Unlike prevalent consumer imaging pipelines though, no high dynamic range (HDR) standard has yet emerged defining the precise range, format or encoding to be used. As such, HDR data often needs to be compressed for display on most current displays, a process known as tonemapping [15, 2].

The aim of this paper is to preserve the appearance and information content of the image as much as possible while ensuring that it can be displayed on the chosen display device. To achieve that, tonemapping algorithms typically operate on the luminance of the image with little to no consideration for the color information present, leading to noticeable changes in the color appearance of the image, as shown in Figure 1. Commonly, tone compressed images acquire an over-saturated appearance when only the luminance channel is processed [12, 18].

Image appearance models, which can be seen as tone reproduction operators with integrated color appearance management [7, 9, 16], aim to reproduce color appearance, but they are designed with calibrated applications in mind and often come at the cost of higher computational complexity due to spatially varying processing. Despite their accuracy, these factors can limit their general applicability.

Some solutions exist for correcting saturation mismatches after tone reproduction [12, 18]. This leads to computationally efficient correction, although we have observed that existing methods tend to create hue and luminance artefacts. Moreover, they require manual parameter selection which is strongly image and tone reproduction operator dependent. Recently, a psychophysi-

cal study was conducted for defining an automatic model to derive the parameters necessary for such corrections, but only allows parameters to be predicted when the tone compression or expansion function is global [12].

Instead, we propose a new approach for correcting saturation mismatches after dynamic range compression. We base our algorithm on insights from color science and on the observation that the amount of desaturation can be inferred from the non-linearity applied by the tone curve, irrespective of whether the tone reproduction operator was spatially varying or not. As such, our approach is parameter-free and agnostic to the operator used for mapping the dynamic range of the image or video. We find that our algorithm reproduces saturation significantly better than the current state-of-the-art.

Related Work

Differences in viewing conditions may result in significant mismatches in perceived color, which can be attributed to idiosyncrasies of the human visual system. To ensure that the appearance of a scene is correctly reproduced on a display, many issues will have to be taken into account, all broadly belonging to the field of color reproduction [8]. Image appearance models can be used to reproduce images as a human observer would see them under given viewing conditions [5, 16]. Such algorithms can be configured to yield calibrated color reproduction, and therefore do not require color post-processing. However, measurements of scene and display conditions are needed as inputs to image appearance models so that the human visual response can be accurately predicted. This requires specialist equipment such as photometers. These algorithms also tend to be computationally expensive, further limiting their use to offline processing.

Dynamic range mismatches between scenes and display devices are therefore typically handled by tone reproduction operators. In essence, most of these algorithms focus on one dimension of the color gamut, namely compression along the luminance direction [15, 2]. Appearance effects are often ignored, leading to images which may appear too saturated. This problem can be mitigated by combining tone reproduction and color appearance algorithms [1]. However, this solution still requires calibrated data and measured viewing conditions to drive the color appearance component.

A more common approach to saturation reproduction is to post-process the tone-mapped image, manually adjusting saturation to levels that appear plausible. Perhaps the most well-known technique for color correction involves the adjustment of color values by means of a power function, according to user parameter $p \in [0, 1]$ [18]. Given an original high dynamic range im-




Figure 1. The same HDR image was tonemapped with different operators (left - [10], right - [16]). The left tonemapped image is overly saturated, while the tonemapping algorithm used on the right has reduced the saturation too far. With our method, both images are automatically corrected to have a very similar appearance by considering their relation with the original HDR image. (Source image from Mark Fairchild’s HDR Survey)

age with input pixels $\mathbf{M}_o = (R_o, G_o, B_o)$ specified in some linear RGB color space, and its associated per-pixel luminances L_o , it is first tonemapped with an operator $f()$ that modifies the image’s luminances, $L_t = f(L_o)$. The color-corrected image \mathbf{M}_c is then produced with:

$$\mathbf{M}_c = \left(\frac{\mathbf{M}_o}{L_o} \right)^p L_t. \quad (1)$$

The primary drawback of this solution is that the selection of parameter p is both image and tone reproduction operator dependent. As this formulation may also introduce undesirable luminance shifts, an alternative adjustment was proposed¹ [12]:

$$\mathbf{M}_c = \left(\left(\frac{\mathbf{M}_o}{L_o} - 1.0 \right) p + 1.0 \right) L_t. \quad (2)$$

Although this equation is claimed to produce smaller luminance shifts, it may still create hue shifts [14]. Here, user parameter $p \in [0, 1]$ can be set manually with the same disadvantages as above.

Alternatively, the setting of p in either technique can be automated based on the slope of the tone curve at each luminance level [12]:

$$p = \frac{(1 + k_1)c^{k_2}}{1 + k_1 c^{k_2}} \quad (3)$$

where k_1 and k_2 are constants² and c is a factor indicating the amount of compression or expansion applied. This factor is calculated as the derivative of the tone curve:

$$c(\log(L_t)) = \frac{d}{d \log(L_o)} f(\log(L_o)). \quad (4)$$

We note that although in its original derivation $f()$ was a simple power function, it produces reasonable results as long as certain conditions are met, most important of which is that the operator needs to be global, i.e. spatially invariant. We view this as an important limitation, as local tone reproduction operators often allow better compression.

¹In the remainder of this paper, we will refer to Equation (1) as Schlick’s method, and Equation (2) as Mantiuk’s method.

²For Schlick’s correction: $k_1 = 1.6774$, $k_2 = 0.9925$. For Mantiuk’s correction: $k_1 = 2.3892$, $k_2 = 0.8552$ [12].

Hue and Saturation Correction

The aim of tonemapping is two-fold; images need to be processed so that their absolute luminance range is compressed, but pixel relations also need to be altered to maximize visible detail, therefore changing the contrast in the image. Changes to contrast and luminance, however, often lead to changes in the appearance of colors in the image and specifically in their saturation and hue. Thus, our algorithm is designed to correct the image’s appearance while minimizing luminance and contrast modifications.

Algorithm Overview

The input to the algorithm consists of two images given in linear RGB space: the tone-compressed image \mathbf{M}_t and the original, unprocessed HDR image \mathbf{M}_o as it contains the original saturation and hue values that we aim to reproduce. The goal of our algorithm is to modify \mathbf{M}_t such that it matches the color appearance of \mathbf{M}_o in terms of hue and saturation, while preserving luminance values from the tonemapped image \mathbf{M}_t . Note that matching the appearance of saturation requires active non-linear management of saturation values to account for the Hunt effect.

Since in most cases accurate radiometric data is not available for HDR images, luminance values computed from the images will be inherently inaccurate. As such we focus on contrast changes between the two input images and therefore normalize both \mathbf{M}_t and \mathbf{M}_o before converting them to XYZ tristimulus values. The image data is then transformed to IPT as this color space has better hue uniformity than CIE $L^*a^*b^*$ [4].

As we need separate access to lightness, hue and chroma, we then convert to a cylindrical color space akin to CIE $L^*C^*h^*$ [19]. This space is based on IPT and therefore we will refer to it as the *Ich* space, where *I* encodes lightness, *C* represents chroma and *h* is a measure of hue. The lightness channel *I* is not further processed, because this was the main purpose of the preceding tone reproduction operator. The hue in the tonemapped image h_t is subsequently set to the hue h_o of the original image, restoring any hue distortions that may have arisen due to gamut clipping during tone mapping.

The quantity that needs to be matched between high dynamic range and tonemapped images is saturation s . However, the aforementioned cylindrical color space produces chroma C . Nonetheless, we can adjust chroma on the basis of per-pixel saturation values computed on both images.




Figure 2. Comparisons between different variants of our algorithm, in particular comparing performance in CIE $L^*C^*h^*$ against the cylindrical version of IPT, termed *ICh*, paired with two different saturation formulations, namely $s = C/L$ and $s = C/\sqrt{C^2 + I^2}$ (substitute L for I in the case of IPT).

Appearance Parameters

After the input images are normalized and converted to *IPT*, chroma and hue parameters are computed for both images. To convert from *IPT* to a cylindrical color space *ICh* [19], we follow standard procedure and leave the I channel unchanged while setting hue h and chroma C as follows:

$$h = \tan^{-1}(P/T) \quad (5)$$

$$C = \sqrt{P^2 + T^2} \quad (6)$$

Saturation s is commonly computed as $s(C, I) = C/I$. Recently, however, an alternative formula was proposed that follows human perception more closely [11]:

$$s(C, I) = \frac{C}{\sqrt{C^2 + I^2}} \quad (7)$$

Note, however, that to our knowledge application of this formula in *ICh* is novel; its development was centered around CIE $L^*C^*h^*$. The merit of using this formulation is shown in Figure 2.

Saturation Correction

Tone reproduction typically maps luminance values in a non-linear manner. As a result, although the absolute luminance levels of the tonemapped image are likely to be lower than the original HDR scene if displayed on a conventional monitor, the relative luminance of many pixels will be increased compared to their surrounding pixels. To deal with this mismatch, we first scale the chroma of the tonemapped image. This step scales C_t to approximately what it would be if the original HDR image had been tonemapped in the *ICh* space:

$$C'_t = C_t \frac{I_o}{I_t} \quad (8)$$

Then, based on (7), we compute the ratio r between the saturation of the original and tonemapped image, albeit that we compute the latter using C'_t :

$$r = \frac{s(C_o, I_o)}{s(C'_t, I_t)} \quad (9)$$

This ratio is then applied to chroma C'_t as a second factor to find the chroma appropriate for the tonemapped image:

$$C_c = r C'_t = r \frac{I_o}{I_t} C_t \quad (10)$$

For convenience, in the following we will refer to the full adjustment factor as:

$$r' = r \frac{I_o}{I_t} \quad (11)$$

Finally, we reset the hue by copying values from the HDR images ($h_c = h_o$). Together with the corrected chroma C_c , it is combined with the lightness channel of the tonemapped image $I_c = I_t$ to produce the final corrected result, which can then be converted back to *RGB*.

Evaluation

To assess the performance of our algorithm, we compressed the dynamic range of many challenging scenes with different tonemapping operators. We then processed the results with our color correction method and compared our results against both the automatic and manual versions of Schlick's and Mantiuk's algorithms (Equations (1) and (2)) by means of psychophysical experimentation.

Tone Curve Estimation

For Schlick and Mantiuk's techniques we estimate the tone curve from the image pair directly so that Equation (3) can be applied to estimate p . If a global tone reproduction algorithm is used a one-to-one mapping between the original luminance L_o and the tonemapped luminance L_t can be obtained. For spatially varying tone mapping operators, many different input levels may be mapped to the same output level. To be able to infer a reasonable approximation for parameter p in the automatic Mantiuk and Schlick corrections, we compute the contrast factor c in (4) based on the average luminance level in L_o that corresponds to each luminance level L_t in the tonemapped image. To further enforce smoothness, this computation is carried out on a down-sampled version of the image and the resulting tone curve is filtered with a Gaussian filter kernel³.

In the following, we show the effect of our correction combined with several tone mapping solutions as well as side-by-side comparisons with other saturation correction techniques. The comparative performance of saturation reproduction is also assessed with a psychophysical experiment.

Results and Comparisons

The color correction method proposed in this paper is fully parameter-free and aims to be applicable irrespective of the type of processing that was applied to the image. The algorithm was implemented in MATLAB, running on an Apple Macbook Pro with an Intel Core 2 Duo processor running at 2.3 GHz. Although our current implementation is not optimized for performance, typical examples tested at resolutions of around 1MP were processed in approximately 5 seconds.

³Note that this approximation serves only for comparison purposes as the relation between p and c is only formally defined for global tonemapping operators.

Tonemapped input images



Corrected results



Pattanaik 2000

Reinhard 2002

Durand 2002

Drago 2003

Mantiuk 2006

Mantiuk 2008

Figure 3. The Memorial image was tonemapped using both global and spatially varying tone mapping operators. The tone mapped images (top) obtain very different appearances, which are corrected with our algorithm (bottom). Although the tone mapped images have different luminance and contrast distributions, our correction equalizes the color appearance between them. In particular, the different materials in the scene obtain a more natural appearance, notably the white marble of the stairs or the gold leaf on the walls.

Our method corrects the saturation in the image on a per-pixel basis. This ensures that even extreme changes in saturation due to tonemapping or any other manual or automatic image processing can be corrected. The quality of our algorithm is shown in Figures 3 and 4. Note that if both the high dynamic range image and the tonemapped image are individually normalized, the tone reproduction process does not universally reduce the image’s dynamic range. Instead, some pixels are reduced in level, whereas others are increased. As a result, some pixels require a commensurate decrease in saturation, while others need their saturation to be increased.

Figure 3 shows that one effect of our method is that material appearance can be correctly reproduced, irrespective of tone reproduction operator. The gold leaf on the wall still appears as gold for instance; an effect that is difficult to reproduce with other methods that tend to create more washed-out colors. Figure 4 demonstrates that existing methods tend to desaturate parts of the image that are both light and saturated, turning the yellow sign and the shop interior white in the top images, and the sky grey in the bottom images.

Psychophysical Evaluation

To assess saturation performance, we designed a 2-alternative forced-choice experiment (2AFC) whereby two identically tonemapped images were post-processed with different saturation correction algorithms and shown side-by-side on the display, underneath the high dynamic range input image as shown in Figure 5a. A SIM2 HDR47E S 4K was used, which can emit

up to 4000 cd/m^2 . To allow prolonged stable and calibrated use, we used a peak luminance of no more than 2500 cd/m^2 . The background of the stimuli was set to 18 cd/m^2 while the peak luminance for the tonemapped images was 100 cd/m^2 . The left and right 7 cm of the display were unused as we found luminance reproduction to be less accurate in those regions. The display was driven by an Apple Macbook Pro running Matlab using the Psychophysics toolbox extensions [3] and employing a custom OpenGL shader for driving the display in calibrated HDR mode.

A set of 8 HDR images were drawn from the HDR Photographic Survey [6] and were tonemapped with the global version of the photographic operator [17] and a spatially varying operator [10]. Subsequently, the images were post-processed with three different saturation correction algorithms: the proposed technique, as well as the automatic versions of the methods given in (1) and (2). A stimulus then consisted of the HDR image, below which two differently post-processed images were shown. Tone mapping operators were varied between stimuli, but not within stimuli. In each trial, the participant was asked to select the image which matched saturation best to the HDR image.

Before starting an experiment, participants were shown written instructions, followed by a short training session to familiarize participants with the difference between saturation and other appearance phenomena. General feedback was solicited after the experiment, which lasted on average 20 minutes.



Figure 4. Comparisons between our new algorithm and Schlick and Mantiuk’s automatic corrections. The two images were tone mapped with a spatially varying [10] and a global [17] operator and then processed with the three correction methods. The local variations of the spatially varying operator lead to very strong local desaturation when images are processed with Schlick’s and in particular Mantiuk’s correction formulae.

Experiment: Evaluation of automatic algorithms The task for the experiment was to match the impression of saturation between tonemapped color processed images and their HDR originals. Stimuli were created to compare our algorithm with the automated version of Schlick and Mantiuk’s algorithm using Li’s [10] and Reinhard’s [17] tone reproduction operators, leading to a total of 48 trials per participant to account for all paired comparisons. There were 18 participants in this experiment, who were between 23 and 53 years old, and all had normal or corrected-to-normal vision as well as normal color vision.

We used a multiple comparison range test to determine if any pairwise difference was significant. We have calculated the coefficient of consistency ξ per image and per tonemapping operator. For the photographic operator we find an average coefficient of consistency of $\xi = 0.78 \pm 0.1$ (mean and standard deviation). For Li’s operator we find $\xi = 0.85 \pm 0.08$. Thus, we have obtained overall high consistency, supporting the following findings.

Significance tests were calculated on the differences between the scores of pairs of color correction methods. These differences are considered significant if they are greater than a critical value R which is defined as:

$$R = \frac{1}{2}W_{t,\alpha}\sqrt{ut} + \frac{1}{4} \quad (12)$$

where $W_{t,\alpha}$ is the upper significance point for the W_t distribution, $t = 3$ is the number of compared methods, and u is the number of observations. At a significance level of $\alpha = 0.001$, $W_{t,\alpha}$ values

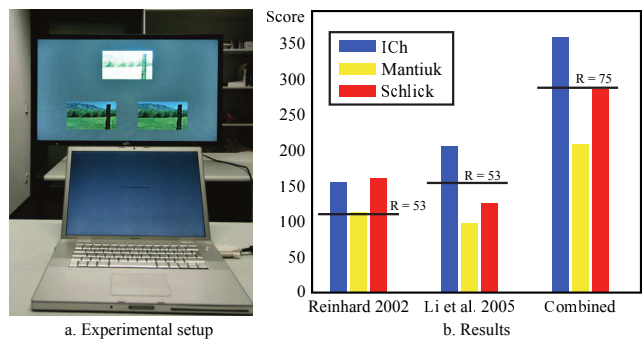


Figure 5. a. The setup used in our experiment. b. Results from our experiment, grouped by tone reproduction operator. Also indicated with a horizontal line is the difference with the longest bar in each group at which significance occurs.

is of 5.06, see Table 22 from [13]. Figure 5b shows the overall results of our experiment.

When we assessed the overall performance, for each tonemapping operator, over all images, we found statistical significance for Li’s tone reproduction operator at significance level $\alpha = 0.001$, with critical value critical value $R = 53$, given $u = 144$ for 18 participants \times 8 images. In this case our method was selected significantly more often. This is visualized in Figure 5b where we have drawn a horizontal line at a height 53 below the maximum score, noting that the bars for Schlick and Mantiuk’s

methods do not cross this line. For the photographic operator, we found no statistically significant differences.

We have observed that Li's operator on average requires stronger saturation correction than the photographic operator. It is therefore interesting to see that especially in the case of a local operator our saturation correction method performs particularly well. Moreover, for the photographic operator our algorithm performs on par with the current state-of-the-art. Although for the experimental evaluation only two tone mapping techniques were included, our experiments indicate that our findings generalize well to other operators, especially when the luminance channel is processed in a locally varying way.

We also computed scores for the two tone mapping operators combined. Here $R = 75$ as $u = 288$ (18 participants \times 8 images \times 2 tone reproduction operators). Overall, our method was selected significantly more often ($\alpha = 0.001$). In essence, this means that our algorithm matches the impression of saturation between tone mapped images and their HDR originals measurably better than the current state-of-the-art.

Conclusions

We developed a novel saturation correction algorithm for the purpose of removing the often over-saturated appearance of tonemapped images. Tone reproduction tends to be carried out on a luminance channel, while leaving chromaticities unaffected. As the appearance of saturation depends on relative luminance levels, ideally saturation should co-vary with luminance changes when applying tone reproduction operators. Nonetheless, it is possible to post-correct saturation mismatches given the input and the output images of a tone reproduction algorithm.

Our algorithm is based on recent insights into the design of perceptually linear color spaces as well as a recent formulation of saturation. This has led to an algorithm that with respect to the state-of-the-art better reproduces the color appearance of the HDR input images, while preserving the luminance compression applied by the tonemapping operator. We evaluated our algorithm and assessed its performance compared to the state-of-the-art with many challenging images as well as a psychophysical experiment. As the computational cost is similar to existing techniques, we believe that our algorithm is a good candidate for color post-processing of tone reproduction operators as well as manually processed images.

Acknowledgements

This work was partially supported by Ministry of Science and Innovation Subprogramme Ramon y Cajal RYC-2011-09372.

References

- [1] A. O. Akyüz and E. Reinhard. Color appearance in high-dynamic-range imaging. *Journal of Electronic Imaging*, 15(3):0330001, 2006.
- [2] F. Banterle, A. Artusi, K. Debattista, and A. Chalmers. *Advanced High Dynamic Range Imaging: Theory and Practice*. AK Peters (CRC Press), Natick, MA, USA, 2011.
- [3] D. H. Brainard. The Psychophysics toolbox. *Spatial Vision*, 10:433–436, 1997.
- [4] F. Ebner and M. D. Fairchild. Development and testing of a color space (IPT) with improved hue uniformity. In *Sixth*

Color Imaging Conference: Color Science, Systems and Applications, pages 8–13, 1998.

- [5] M. D. Fairchild. *Color appearance models*. John Wiley & Sons, Chichester, UK, 2nd edition, 2005.
- [6] M. D. Fairchild. The HDR photographic survey. In *Proceedings of the 15th IS&T/SID Color Imaging Conference*, pages 233–238, 2007.
- [7] M. D. Fairchild and G. M. Johnson. Meet iCAM: an image color appearance model. In *Proceedings of the 10th IS&T/SID Color Imaging Conference*, pages 33–38, 2002.
- [8] R. W. G. Hunt. *The reproduction of color*. Fountain Press, England, 1996. Fifth edition.
- [9] J. Kuang, G. M. Johnson, and M. D. Fairchild. iCAM06: A refined image appearance model for HDR image rendering. *Journal of Visual Communication and Image Representation*, 18(5):406–414, 2007.
- [10] Y. Li, L. Sharan, and E. Adelson. Compressing and expanding high dynamic range images with subband architectures. *ACM Transactions on Graphics*, 24(3):836–844, 2005.
- [11] E. Lübbecke. *Colours in the Mind - Colour Systems in Reality: A formula for colour saturation*. Books on Demand GmbH, 2008.
- [12] R. Mantiuk, R. Mantiuk, A. Tomaszewska, and W. Heidrich. Color correction for tone mapping. *Computer Graphics Forum*, 28(2):193–202, 2009.
- [13] E. S. Pearson and H. O. Hartley. *Biometrika tables for statisticians 3rd ed.*, volume 1. Cambridge University Press, 1988.
- [14] E. Reinhard. Tone reproduction and color appearance modeling: Two sides of the same coin? In *Proceedings of the 19th IS&T/SID Color Imaging Conference*, pages 171–176, 2011.
- [15] E. Reinhard, W. Heidrich, P. Debevec, S. Pattanaik, G. Ward, and K. Myszkowski. *High Dynamic Range Imaging: Acquisition, Display and Image-Based Lighting*. Morgan Kaufmann, 2nd edition, 2010.
- [16] E. Reinhard, T. Pouli, T. Kunkel, B. Long, A. Ballestad, and G. Damberg. Calibrated image appearance reproduction. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)*, 31(6):article 201, 2012.
- [17] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. *ACM Transactions on Graphics*, 21(3):267–276, 2002.
- [18] C. Schlick. Quantization techniques for visualization of high dynamic range pictures. *Proceedings of the 5th Eurographics Workshop on Rendering*, pages 7–18, 1994.
- [19] S. Shen. *Color Difference Formula and Uniform Color Space Modeling and Evaluation*. PhD thesis, Rochester Institute of Technology, 2009.

Author Biography

Tania Pouli is a Researcher at Technicolor Imaging Science Laboratory, working on high dynamic range imaging and color reproduction problems. She received her BSc in 2007 from the University of Bath and her PhD from the University of Bristol in 2011, where she worked with Erik Reinhard on image statistics and creative image editing. She is the lead author of the book “*Image Statistics in Visual Computing*”, published by CRC Press.