

Chromatic Weibull Tone Mapping for Underwater Image Enhancement

Chloe A. Game^{1,*}, Michael B. Thompson² and Graham D. Finlayson³

¹ University of East Anglia, UK; chloe.game@uea.ac.uk

² Gardline Ltd., UK; michael.thompson@gardline.com

³ University of East Anglia, UK; g.finlayson@uea.ac.uk

* Corresponding author: chloe.game@uea.ac.uk

Abstract

Image enhancement is often used to alleviate the low contrast, blurring and colour reduction effects, common in underwater imagery. Tone Mapping, a particularly simple yet elegant enhancement technique improves image quality by modifying image histograms to a more desirable tonal distribution. In previous work, we presented our novel chromaticity-preserving algorithm, Weibull Tone Mapping (WTM), that can simplify custom tonal manipulations, to increase conspicuousness of image features. In this paper, we present a natural non-chromaticity preserving counterpart, in which the WTM tone map is applied to all colour channels (R,G,B). We demonstrate, as before, that user's prefer WTM to unenhanced images. However, contrary to prior work, our non-chromaticity preserving WTM is less preferred to custom tonal manipulations. Thus, how we map the colour aspect of images (given a brightness only tonal adjustment) has a significant impact on users' subjective preference judgements.

Keywords: *Underwater image enhancement, Tone mapping, Histogram Specification, Weibull Distribution*

INTRODUCTION

Exploring underwater environments with imagery has a historic past, with the first underwater photographs taken in the 19th century (Ruppé and Barstad 2013: 22). Although a long-established tool, and technology continues to advance, imaging science still struggles with increased wavelength absorption and scattering with water depth, resulting in poor quality images, that may be dark, lacking contrast, blurry and de-saturated. Extracting biological information, by annotation and enumeration of image features (seafloor organisms), therefore remains hindered.

Improving image quality, through post-processing, can offer a solution (Schettini and Corchs 2010). Tone mapping, an intuitive and often simple enhancement method, offers a fast approach to suppress unwanted lighting effects and improve the conspicuity of image features. It involves the use of a transform function, or tone curve, that alters image pixel values to render an image more pleasing and of better quality. Certain distributions of pixel values have been linked to image quality (Zuiderveld 1994; Game et al. 2020). Information theory dictates that a flatter distribution contains more information or entropy – requiring more bits to encode (Shannon 1948). In a visual sense, this translates to increased contrast, with maximum contrast achieved when all bins in an image histogram contain equal values, as is the case with Histogram Equalization (HE). However, the slopes of tone curves – that when applied to an image, lead histograms to be uniform - are often very high or low, causing over-enhancement of contrast and contouring artefacts, and loss of details. In Contrast Limited Histogram Equalization (CLHE) (Zuiderveld 1994), the slope of the tone curve is bounded between a minimum and maximum value, thus weakening the compression and stretching of pixel values. In Figure 1a we show an input image with low image contrast. After enhancing the image by applying a HE and CLHE tone curve, we can there is a notable increase in image contrast in (b) and (c). However,

in the case of HE, the resulting image is oversaturated and noisy, whereas CLHE enhancement is more restricted, suitable and preferred.



Figure 1: Comparison of an unenhanced image (a) followed by its HE (b) and CLHE (c) tone-mapped output. CLHE min and max slope thresholds were set to 2.56/256 and 0.768/256, respectively.

Underwater image histograms typically possess a bell-shaped distribution, tapering to zero at the ends. Forcing histograms in to a uniform ‘flat’ shape, may therefore be inappropriate, resulting in abnormal brightness or colour changes. Instead, tone mapping operations that preserve the natural characteristics of underwater images, such as the dark pixels, whilst enhancing details within the spot illumination would be better suited to this application. This was the motivation for using a Weibull distribution (WD) to model the brightness statistics of underwater images (Game et al. 2020).

The WD is a smooth, unimodal distribution, parameterized by two parameters that control the location of the peak and spread of the distribution. It is not unlike the Rayleigh distribution (RD), that has been used to enhance underwater images previously (Eustice et al. 2002), however it is more general and we found it was more applicable to our data.

In Game et al. (2020) we described tonal manipulations made by end-users to enhance visibility of specific seabed features for their analyses and demonstrated that they can be modelled and in fact simplified by the WD, using our Weibull Tone Mapping (WTM) algorithm. Significantly, the user and the WD approximations are carried out in the brightness domain. Thus, how a brightness adjustment can be applied to a colour image needs to be considered. In Game et al. (2020), the image brightnesses were modified such that the image chromaticities were preserved.

The starting point for this paper was our recent realization that the user manipulations themselves were not chromaticity preserving. Rather, the output images were calculated by applying the same tone map – that mapped the input brightness image to the corresponding output - to each of the R, G and B channels. Candidly, the discrepancy between how WTM images and the users own custom adjustments were colour rendered was due to our incomplete understanding of the imaging platform we used (GIMP). Thus, in this paper, we re-run our preference experiment. We now consider image preferences for an original image, the user’s own adjusted image and the WTM solution. Where, in the latter two cases the respective tone curves are calculated in a brightness channel and then applied to all three channels (non-chromaticity preserving).

BACKGROUND

As defined in (Game et al. 2020), Weibull Tone Mapping (WTM) involves approximating the brightness histograms (sum normalized to 1) of an input and output (enhanced) image, with a Weibull Distribution (WD), and then solves for the curve that maps the former distribution to the latter, see (Game et al. 2020) for detailed algorithmic steps. In Figure 2 we show an example of this. An input histogram and its user-adjusted output, are presented as dashed lines in 2a. & c, with the best matching WD to these

target distributions shown as a solid line. The probability density function (PDF) of the WD is given in closed form by

$$PDF_W = \frac{k}{\lambda} \left(\frac{v}{\lambda}\right)^{k-1} e^{-(v/\lambda)^k}, \quad v \geq 0, \lambda > 0, k > 0 \quad (1)$$

where v is the brightness value, k is the shape parameter and λ is the scale parameter. The two parameters relate to peak position (k) and spread (λ), and can thus be ‘broadly’ understood to control contrast and brightness. Due to the high similarity of the WD to the input-output brightnesses, we can see that the target (user) tone map in 2b. - that transforms the input histogram to the enhanced output- is closely modelled by the derived WTM curve, but it is a little smoother.




Figure 2: Example of the WTM method. WTM approximations of an input PDF and custom (output) PDF are shown in a) and c), respectively. The user tone map and derived WTM tone map are shown in b).

Given that WTM, and indeed other enhancement algorithms, are designed to improve appearance of imagery for an end-user, it is important to validate their performance in this context. A typical and effective method is to consider the preferences of end-users in a paired comparison experiment. Thurstone (1927) theorized that although paired comparisons are subjective and preference may be inconsistent across pairs, a frequent preference response often exists and can be described statistically. Frequency matrices of preference votes are simply converted to proportional values and sampled from a normal distribution, giving standard Z-scores. As demonstrated in Game et al. (2020), this allows the stimuli (image variants in this case) to be ranked and a preferred type to be identified by end-users. In this case, it was found that a WTM simplification of user tonal adjustments, is preferred for the interpretation of underwater imagery.

Making colour tone-mapped images

In the original study (Game et al. 2020), after tone mapping with the WTM curve, they calculated an RGB colour image by multiplying, the red, green and blue pixel values, by the same scalar - a ratio of the output over the input brightnesses. Specifically, it is calculated as

$$L_{out}(x, y) = L_{in}(x, y) \left(\frac{L_{out}(x, y)}{L_{in}(x, y)} \right) \quad (2)$$

where $L_{in}(x, y)$ and $L_{out}(x, y)$ represent an input brightness image and a WTM-mapped brightness image, and $\underline{I}_{in}(x, y)$ and $\underline{I}_{out}(x, y)$ denote the input and output 3 channel RGB images (underscoring denotes a vector function). Note that we define brightness here as $\max(R, G, B)$, i.e. $L_{in}(x, y) = \max(\underline{I}_{in}(x, y))$. By construction, the output colour image, following Equation 2 (henceforth referred to as ‘original WTM’), will have the same chromaticities as the input.

Previously, we assessed the usefulness of tone-adjusted images in a preference experiment; comparing a user tone mapped image against the WTM simplification, finding the latter to be preferred over the former. However, in revisiting our work we realized that the user adjustments (carried out in GIMP) did not preserve chromaticity, yet we were comparing these against the WTM approximation (where the chromaticity was preserved). We found that GIMP calculated the output image as

$$\underline{I}_{out}(x, y) = t(\underline{I}_{in}(x, y)) \quad (3)$$

where the function $t()$ represents the user tone curve, mapping input brightness to an output. In Figure 3a we see an unenhanced input image of the seabed, followed by its custom output (user tone-mapped) in b. In c), it is clear that applying the WTM tone map, as in Equation 3 (hereby referred to as ‘matched WTM’), produces a more similar image to b than the original chromaticity preserving WTM in d. Visually, we found that this trend was followed for the remaining images in the dataset i.e. matched tone renditions looked more similar to the user’s own tone-adjusted images than the original chromaticity preserved variant. Additionally, following matched WTM, the colours of the image are boosted. In this paper, we therefore evaluate whether ‘original WTM’ versus ‘matched WTM’ has a material result on user preference of tone-mapped images.




Figure 3: Comparison of an unenhanced input image (a), a custom enhanced image (b), matched WTM (c) and original WTM (d).

PSYCHOPHYSICAL EVALUATION

Under ISO standard 3664:2009 conditions (ISO 2009), in a darkened room, 6 Gardline analysts- 3 of which created the custom enhanced dataset- were presented with uniquely randomized pairs of images. For each pair, they were asked to ‘Choose the image that best allows identification of the habitat (class) therein, or no preference if the images are equivalent’. Preference votes were assigned to 3 options; Image 1, Image 2 or No Preference. To minimize the duration of the experiment, 2 random subsets (one for each experiment repeat) of 18 images, 3 per habitat class (total=6), were selected. This kept the number of pairwise comparisons manageable, since the number of image variants, $n = 3$, would result in $\frac{n}{2(n-1)} = 3$ pair-wise comparisons per image, which is then viewed twice (left-right order switched). The number of pairwise comparisons, per experimental repeat, was therefore $18 \times 2 \times 3 = 108$. Preference votes for each image pair were processed by awarding a score of 1 to the chosen image and 0 to the other. If no preference was selected, each image was given a score of 0.5.

Votes for each analyst, in each experimental sitting, were converted to 3x3 frequency matrices, of which the score at $[i, j]$ represents the frequency of votes in which image variant i was preferred over variant j , across the 18 images. Thurstone's Law of Comparative Judgments, or Thurstone's Case V (Thurstone 1927), was used to convert frequency matrices to z-score (standard score) matrices.

Our results are interesting. In contradiction to the original WTM, we find for matched WTM - when chromaticity is not preserved- that users prefer their own manipulations to the WTM simplification (Figure 4). This demonstrates that how the image colour component is mapped, when based on a brightness tonal adjustment, can significantly influence the preference assessments of end-users. However, irrespective of this, both the custom and WTM adjustments are significantly preferred over doing nothing, for each WTM case. Therefore, given an unenhanced input image, end-users find a WTM enhancement useful for the extraction of biological data from imagery, regardless of the colour rendering method.

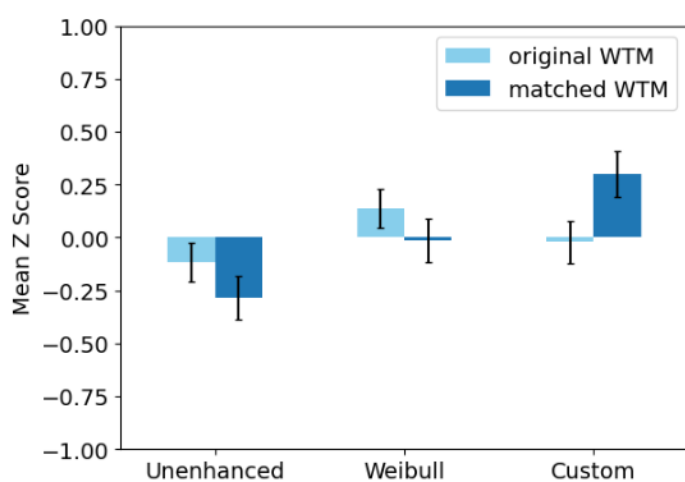


Figure 4: Preference Experiment Scores.

CONCLUSION

In this paper, we introduced a chromatic modification of our Weibull Tone Mapping (WTM) algorithm. We find that, given an input-output image pair; application of the WTM tone map to all colour channels (R, G and B) provides a more colour-rich image and a better approximation of user tonal adjustments, than WTM in its original form. Interestingly, the method of colour mapping when using WTM – non-chromaticity preserving vs chromaticity preserving – alters the preference of end users. It therefore seems prudent to consider how colours are rendered, when developing and assessing tone-mapping algorithms, as they can have a significant influence on quality perception by end-users. For WTM specifically, an enhancement that preserves chromaticity is preferred for the identification of marine habitats in underwater images.

ACKNOWLEDGEMENTS

This work was supported by the Natural Environment and the Engineering and Physical Sciences Research Councils [grant numbers NE/N012070/1 and EP/S028730/1]. The authors also thank the support of Gardline Limited and the Environmental Reporting team for their contribution to this research.

REFERENCES

- Eustice, R., O. Pizarro, H. Singh, and J. Howland. 2002. UWIT: underwater image toolbox for optical image processing and mosaicking in MATLAB. *Proceedings of the 2002 International Symposium on Underwater Technology* (Cat. No.02EX556), 19-19 April 2002.
- Game, C. A., M. B. Thompson, and G. D. Finlayson. 2020. Weibull Tone Mapping for Underwater Imagery. *Color and Imaging Conference 2020* (28): 156-161.
- ISO. 2009. *Graphic technology and photography—Viewing conditions ISO3664:2009*. Geneva.
- Ruppé, C. V., and J. F. Barstad. 2013. *International handbook of underwater archaeology*. Springer Science & Business Media.
- Schettini, R., and S. Corchs. 2010. Underwater Image Processing: State of the Art of Restoration and Image Enhancement Methods. *EURASIP Journal on Advances in Signal Processing* 2010 (1): 746052.
- Shannon, C. E. 1948. A mathematical theory of communication. *The Bell System Technical Journal* 27 (3): 379-423.
- Thurstone, L. L. 1927. A law of comparative judgment. *Psychological Review* 34 (4): 273-286.
- Zuiderveld, K. 1994. Contrast limited adaptive histogram equalization. *Graphics gems*: 474-485.