

Weibull Tone Mapping for Underwater Imagery

Chloe A. Game, University of East Anglia, Norwich, UK; Michael B. Thompson, Gardline Ltd, Great Yarmouth, UK; Graham Finlayson, University of East Anglia, Norwich, UK

Abstract

Imagery is a preferred tool for environmental surveys within marine environments, particularly in deeper waters, as it is non-destructive compared to traditional sampling methods. However, underwater illumination effects limit its use by causing extremely varied and inconsistent image quality. Therefore, it is often necessary to pre-process images to improve visibility of image features and textures, and standardize their appearance. Tone mapping is a simple and effective technique to improve contrast and manipulate the brightness distributions of images. Ideally, such tone mapping would be automated, however we found that existing techniques are inferior when compared to custom manipulations by image annotators (biologists).

Our own work begins with the observation that these user-defined tonal manipulations are quite variable, though on average, are fairly smooth, gentle waving operations. To predict user-defined tone maps we found it sufficed to approximate the brightness distributions of input and user adjusted images by Weibull distributions and then solve for the tone curve which matched these distributions from input to output. Experiments demonstrate that our Weibull Tone Mapping (WTM) method is strongly preferred over traditional automated tone mappers and weakly preferred over the users' own tonal adjustments.

Index Terms - Underwater image enhancement, Tone mapping, Histogram Specification, Weibull Distribution, Contrast Limited Histogram Equalization

Introduction

Underwater optical imaging is challenging, particularly with respect to illumination. Increasing light attenuation with depth, due to increased wavelength absorption and scattering, can result in colour reduction, low contrast and blurring effects in images. This attenuation drives the requirement for strong artificial lighting on camera platforms which causes non-uniform illumination and shadows within images [1]. Image lighting patterns are also inconsistent, through use of multiple imaging platforms and lighting adjustments to limit interest of fish shoals; which can obstruct seafloor imaging and impede investigation. As a result, images are varied in quality, with image features such as seafloor dwelling organisms, often poorly visible. Features may also be irregular in appearance, reducing correspondence between images. These issues are a severe hindrance to both manual and automated annotation tasks [2].

Manipulation of image histograms, or brightness distributions, through tone mapping, can be effective in suppressing unwanted lighting effects and enhance appearance and/or the visibility of image features. Algorithms that operate in this way are often fast and simple, requiring no *a priori* knowledge on the imaging environment, such as depth field, water quality, or distance between a camera and a target. Yet few have been successfully applied in the underwater image domain [3, 4]. For a detailed review of underwater image processing, see [5].

Tone mapping is illustrated in Figure 1. In (a) we show an input image which lacks details and appears *flat*. In (c), we plot a

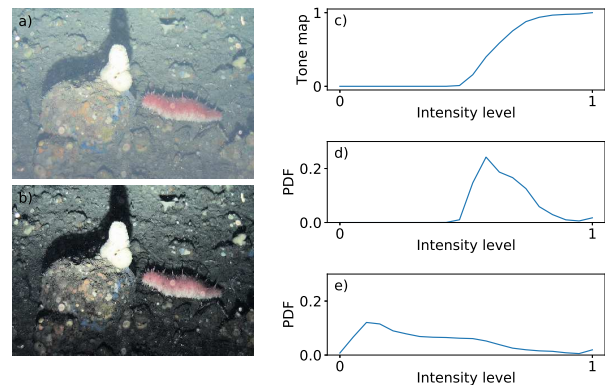


Figure 1: Tone mapping example: a) input image, b) output image, c) tone map, d) input brightness distribution and e) output brightness distribution

simple tone curve to map input to output brightnesses, generating a modified output image (b) from (a). Clearly, there are more image details after tone adjustment. Note that for this example and for the remainder of the paper, adjustments are made for the brightness signal only, with colour saturation and hue kept the same.

Often tone mapping is framed as a problem of mapping an input brightness distribution to a target output distribution. In Figure 1, we show the brightness histogram (d) of the unenhanced input image (a) and in (e) the target brightness distribution and corresponding image in (b). Notice that the distribution shown in (e) is more uniform, or flatter, than (d). Intuitively, the flatter the brightness distribution, the more conspicuous image details will be, since the whole range of brightness values is used (and almost equally). In an information theoretic sense, a flatter distribution has more information or higher entropy; it requires more bits to encode [6]. One might therefore think that the goal of all tone mapping would be to map the input brightness histogram to a perfectly uniform counterpart. Indeed, this is exactly what, possibly the oldest, image enhancement method called Histogram Equalisation (HE) does. However, the tone curve that generates a perfectly equalised (uniform) histogram often has ranges where the curve has very high or low slopes. High slopes generally correlate with the appearance of artefacts such as contouring (e.g. false regions in the sky) or unnatural contrast, as is evident in image (b); the contrast seems too high. Alternatively, low slopes can result in the loss of important visual details.

A powerful modification of HE, Contrast Limited Histogram Equalization (CLHE), performs more limited equalization generating a tone curve with a bounded slope, i.e. never too large or small. CLHE proceeds in three steps: (1) the input brightness histogram is approximated, (2) we solve for the tone curve that equalises the approximated histogram and (3) the same tone curve is applied to the input image. Although a significant step forward and widely deployed, CLHE too has problems.




Figure 2: Enhancement example: a) Unenhanced image, b) HE image and c) CLHE image (max slope = 2/256)

Chief amongst these is that a tone curve bounded by a minimum and maximum slope is stepped and *wiggly*; technically its second derivative can cross the x-axis many times. Not only can this introduce banding artefacts in images, it is a behaviour that we do not mimic when we manually adjust the tonality of an image.

In this paper, we focus on tone mapping problems in underwater imagery, with the long-term aim of developing automated image enhancement tools. As comprehension of desirable underwater image enhancements, for analytical purposes, is in its infancy, we began by investigating images processed by end-users in the field *i.e.* *biologists*, who manually manipulated input images so that details sought for analytical purposes were more conspicuous and did not contain artefacts. While the users could choose fairly arbitrary tone maps, we found that those selected were broadly, fairly smooth increasing functions that gently waved; some exhibiting linearly stretched S-shaped patterns. This observation led us to develop a new method for the adjustment of underwater images which we call Weibull Tone Mapping, or WTM.

A Weibull distribution (WD) is a smooth and highly generalisable function with a single peak and varying width, parameterised by two numbers. We found that by fitting a WD to brightness histograms of input and user-adjusted output images and solving for the tone map that mapped these distributions to each other, it resulted in an enhancement similar to those created by users. However, Weibull approximations of user tone curves were typically smoother and even simpler in shape. This begs the question of whether our approximate tone map works as well as the user's own tonal adjustment.

To test the efficacy of WTM we therefore ran preference experiments. Pairs of images were drawn from unenhanced input images, and those adjusted by CLHE, users and WTM. Image pairs were shown to underwater image analysts who were tasked to select the image which best shows details they seek for analysis. A Thurstonian analysis reveals that our WTM is clearly preferred over CLHE and the original image and weakly preferred over users' own adjustments.

Importantly, we draw the readers attention to the fact that the WTM method is not *constructive*. Rather, just now, to apply our method we need a user to make an adjustment which we then improve upon. However, WTM is clearly a way of helping users reach a better final endpoint, more quickly *i.e.* WTM generally creates an output image that appears similar or a little improved. We are aware that by choosing a WD, we are enforcing uni-modality on the underlying histogram, however, remarkably our preference tests show that people prefer this restriction. We are now investigating how the output Weibull distribution can be predicted given information inferred from the input image. Our preliminary work indicates that this inference cannot be based on the input image brightness alone but, rather, will require spatial

image analysis.

Background

An image histogram, \underline{h} , is a vector containing the frequency of pixels assigned to any given value, or bin. This can be used to either visualize pixel values in each of the colour channels of an image, $I(x, y, z)$, where $z = [R, G, B]$, or in the case of a brightness histogram, represents the intensity or *brightness* of each pixel in a single-channel greyscale image, $I(x, y)$. Note that for this paper, all greyscale brightness images are extracted from $\max(I(x, y, z))$, to correspond with software used in Experiment I. For a typical 8-bit encoded image, there are 256 possible values; integers in the interval $[0, 255]$. However for simpler explanation, we map the image values to the interval $[0, 1]$ *i.e.* we divide each pixel value by 255. It follows that if \underline{h} is a vector of 256 values, then h_j , where j is in $[1, 256]$, encodes the frequency that the j th brightness - which is calculated as $((j - 1)/255)$ - appears in the image.

The arguments we set forth below hold for any number of bins in the histogram, or intensity levels, and are denoted by the integer L . Also, it is often useful if \underline{h} is normalised, by dividing the raw frequencies by the sum of the histogram. In the continuous domain, a normalised histogram such as this is called a probability density function (PDF). Henceforth, we assume all histograms sum to 1.

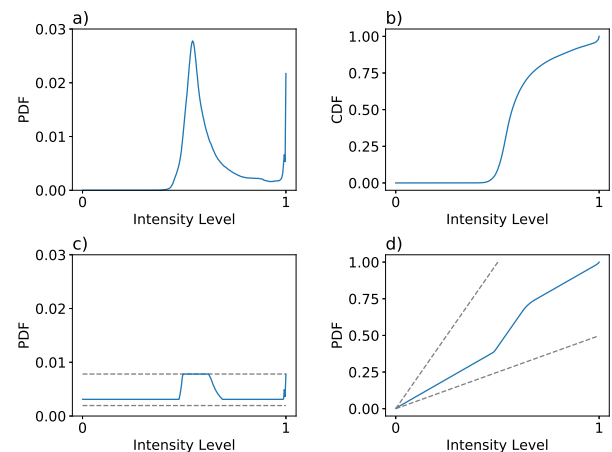


Figure 3: Tone mapping methods HE & CLHE: a) HE PDF, a) HE tone curve (CDF), c) CLHE PDF constrained by upper and lower slope bounds of $2/L$ and $0.5/L$ respectively, where $L = 256$ bins, and d) CLHE tone curve. Dashed lines in c) and d) depict upper and lower slope bounds.

A relatively simple approach to enhance an image, by modifying the tonality and contrast, is to map the brightness distribution \underline{h} to *match* a target distribution \underline{h}' . A popular enhancement

method that operates in this way is histogram equalization (HE). Here \underline{h}' is a uniform or equalised histogram. In HE, the cumulative distribution function (CDF) of \underline{h} maps the input brightnesses so that the resulting image brightness histogram is uniform [7]. The HE tone map, \underline{c} , is calculated as

$$c_i = \sum_{j=1}^i h_j \quad (1)$$

where $0 \leq c_i \leq 1$, and $i=1, 2, \dots, L$, for a histogram of L bins.

A computationally cheap and revertible method, HE can improve segmentation and identification of features in an image, as well as standardize the appearance of features under different illumination [8]. However, it can often generate an undesirable output image that is over enhanced, with significant and abnormal brightness changes, and where any background noise is amplified. In Figure 2, an input image (a) is histogram equalised (b). Figure 3a. shows the corresponding input image histogram and (b) the tone curve that results from HE. Clearly, the tone curve has regions where its slope is both very small and very large. These are precisely the main conditions where HE can produce poor results. In this example, the small slope in the bright pixel range has resulted in the loss of details in the fish (*Sebastidae sp.*), in the HE image.

Contrast Limited Histogram Equalisation is a modification of HE, in which limits can be placed on the slopes [3, 4]. A detailed summary can be found in [9]. Given that the tone map for HE is simply the cumulative distribution of the normalised histogram, or equivalently the integral of the PDF, the slope of the histogram is therefore directly related to the relative frequency of the histogram; since differentiation of the CDF returns the PDF. The intuition behind CLHE therefore, is to find an approximation of the input histogram, in which relative frequencies across all bins are bounded. This in turn, bounds the minimum and maximum slopes of the curve. Figure 3c. demonstrates the CLHE modification to the input brightness distribution in (a), with slopes limited to be above 0.5 and below 2, (d) shows the corresponding tone curve. The enhancement provided by CLHE to Figure 2a., is shown in Figure 2c. This image is a significant improvement on the input; the image does not look over-enhanced, there are no new artefacts and, compared to the original, details are far more visible.

As well as simplifying the input distribution to find a tone curve with desirable properties, such as a bounded slope in CLHE, it has also been argued that the target distribution should also have particular properties i.e. it should not always be a uniform histogram. For example, it has been suggested that a Rayleigh distribution (RD) is a favourable target distribution for underwater images [10], see Figure 4. This is a continuous probability distribution for positive-valued random variables, often resembling a bell-shape, and has been frequently enforced within a variant of the CLHE algorithm for underwater image enhancement [10, 11, 12, 13]. The Rayleigh PDF and CDF is given by

$$PDF_{Rayleigh} = \frac{x}{a^2} e^{-x^2/(2a^2)}, \quad x \geq 0, \quad a > 0, \quad (2)$$

$$CDF_{Rayleigh} = 1 - e^{-x^2/(2a^2)}, \quad x \in (0, \infty), \quad (3)$$

where x is the brightness value and a the scale parameter.

For underwater images, particularly those in deeper waters, it is important to preserve the fact that it is dark. If the target distribution does not tail toward zero, as Rayleigh does, then the

processed image will be over-enhanced including dark noisy pixels may become apparent. As underwater images are often dominated by a strong 'spot-light' in a region of interest and floating particulates, there are often many extremely bright pixels which should be reduced in intensity. Mapping to the Rayleigh distribution preserves the darkness of pixels that should not be enhanced as well as bringing back details that are compressed within the spot illumination.

Given these factors and its popularity with underwater image enhancement, we explored the use of the RD as an automated tool. However, we found that matching the histograms of brightness images to a RD did not provide compelling improvements or mimic the adjustments made by end-users.

Weibull Tone Mapping (WTM)

The Weibull probability distribution (WD) exhibits similarities to the RD and as one of the contributions of this paper, we propose it has desirable properties that make it a good target distribution for underwater images. Although not yet used for this purpose, to our knowledge, it has been demonstrated that the WD can explain the contrast statistics of natural images [14, 15] and is correlated with our own perception of natural images [16]. In its two-parameter form, the Weibull PDF and CDF are denoted as

$$PDF_{Weibull} = \frac{k}{\lambda} \left(\frac{x}{\lambda}\right)^{k-1} e^{-(x/\lambda)^k}, \quad (4)$$

$$x \geq 0, \quad \lambda > 0, \quad k > 0,$$

$$CDF_{Weibull} = 1 - e^{-(x/\lambda)^k}, \quad x \in (0, \infty), \quad (5)$$

where x is the brightness value, k is the shape parameter and λ is the scale parameter. Weibull distributions are smooth and unimodal, with k and λ accounting for the peak position and the spread of the distribution. In comparison to the RD, the WD can exhibit a greater variety of properties, see Figure 4. For example, it is able to approximate characteristics of Rayleigh, when $\lambda = \lambda/\sqrt{2}$ & $k = 2$, as well as Exponential, when $\lambda = 1/\lambda$ & $k = 1$




Figure 4: The Rayleigh distribution ($a = 4$) and 3 variations of the Weibull distribution.

Our Weibull Tone Mapping (WTM) method proceeds in three steps; 1) we calculate the best Weibull approximations to an input image and the corresponding user-enhanced output image, 2) we calculate the tone map that matches the input WD to the corresponding output WD and 3) we apply the calculated tone curve to the original input image; in so-doing we approximate the user-adjusted output image.



Figure 5: WTM example: a) Input image, b) Output image by biologist and c) WTM approximation of Output image (b).

Approximating a brightness distribution using the Weibull function

To find a WD that best *matches* a brightness distribution, Kullback-Leibler divergence (KLD) can be used. Also known as information divergence or relative entropy, KLD determines the difference between two probability distributions, and can be calculated as follows

$$KLD = \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right) \quad (6)$$

where P is an input brightness distribution, Q denotes a target distribution, here a WD. KLD provides a measure of the amount of information in P that allows discrimination of P and Q . If the two distributions are highly similar then the KLD will be low, with 0 reached only when $P=Q$. Remembering that the WD is parameterised by k (broadly, peak position) and λ (broadly, slope), we can therefore search for parameter pairs that create a WD that best fits a given input and output distribution, by those that achieve the lowest KLD. In our study we test the parameter pairs of λ in the interval $[0.1:0.1:3]$ and k in $[0.1:0.1:15]$. This resulted in the comparison of each brightness distribution to 4500 possible Weibull models.

Calculating the Weibull Tone Map

It is well known that histogram matching - finding the curve that maps an input distribution to a target output distribution - can be implemented as a forward and inverse HE step. Let us have two images, $I(i, j)$ and $O(i, j)$, corresponding to an input and a user tone-mapped output. Let $h_I(x)$ and $h_O(x)$ denote the Weibull input and output distributions, that approximate the brightness distributions of $I(i, j)$ and $O(i, j)$, and $C_I(x)$ and $C_O(x)$ their corresponding cumulative distributions. Remembering that the tone map that implements histogram equalisation is defined by the CDF of the images' brightness distribution, it follows that the Weibull tone-mapped image $\hat{I}_O(i, j)$ is therefore calculated as $\hat{I}_O(i, j) = C_O^{-1}(C_I(I(i, j)))$. The brightness distributions of $\hat{I}_O(i, j)$ and $I_O(x, y)$ are the same.

It is important to note, that in the discrete image domain we cannot in fact carry out exact HE. In order to obtain a uniform histogram, some of the brightnesses in the input image, with value v , might be mapped to either v_1 or v_2 in the output image. In tone mapping, this is not possible, as every input maps to each output uniquely. However, for almost all images this detail is insignificant. Also noteworthy, we process only the brightness channel. If a pixel at R , G and B , are intensities in $[0,1]$, our brightness channel, w , is equal to $w = \max(R, G, B)$. After WTM, each w is mapped to an output counterpart w' . Correspondingly, the new RGB is set to be equal to $R \frac{w'}{w}$, $G \frac{w'}{w}$ and $B \frac{w'}{w}$. Choosing the maximum as the definition of brightness has the advantage

that the output RGB image is also, and always, in $[0,1]$. Therefore, values will not map out of the display range, see[17]. It also ensured our tone-mapping methodology complied with software used in Experiment I.

In Figure 5, we show a worked example to illustrate the mechanics of our tone mapping approach. We show that our WTM method applied to image (a) results in image (c) that is almost indistinguishable from the target output image in (b). In Figure 6a, we see that the Weibull distribution closely *matches* the input brightness distribution of image 5a. and in 6c, we see that it also well approximates the target brightness distribution of image 5b. The respective tone maps to transform image 5a. to 5b & c. can be seen in Figure 6b. Here we demonstrate the close approximation of the target tone map by WTM, accounting for extremely similar output images. The tonal adjustments in this case produce a slightly darker output image. However, the visibility of fine morphological features that can aid annotation, is improved, such as *osicles* on the white sponge (likely *Mycale lingua*).

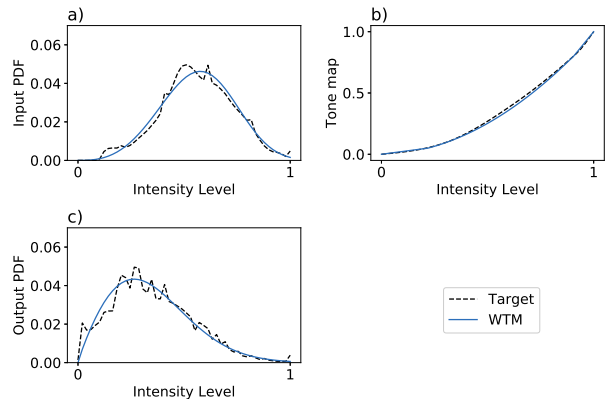


Figure 6: WTM method showing a) Input brightness PDFs, of a target output image and its WTM approximation, b) Tone maps used to adjust input PDFs and c) Output PDFs.

Experiment I: User tone mapping

This experiment involved the development of custom-per-image tone maps by analysts, to determine the type of enhancements that will aid their annotation efforts. For this, our collaborator, Gardline Ltd., provided a large set of underwater survey images. From this, 60 RGB (.jpeg) images, of size 3236x4320x3 pixels, were randomly selected. This selection contained images of 6 broad habitat classes (10 of each), representing the breadth of biological and physical features expected in the Gardline dataset. Our data set is intentionally small as we later test preference using pairwise comparisons, a time consuming process.

For these 60 images we asked three image analysts to, man-

ually, tonally adjust the images so that details required to annotate the content of the image (i.e. the habitat) are made as conspicuous as possible. Tone curves were created using an open source ‘GNU image manipulation programme’ (GIMP), and had to be conventional i.e. strictly increasing functions of brightness. Each analyst performed tone mapping on a unique randomization of the data-set under ISO standard 3664:2009 conditions [18]; sitting approximately 70 cm from the display in a neutrally painted and darkened room. On average, they each needed ≈ 90 minutes to complete the experiment.

The shape of tone curves created by analysts was variable within this study, however a large proportion were mildly *wavy* and linear patterns, with some appearing as stretched S-shaped curves; indicating more gentle contrast enhancement. Some curves were also more exponential-like in shape, causing some compression of the mid-tones. Figure 7. shows some example tonal adjustments performed by users.

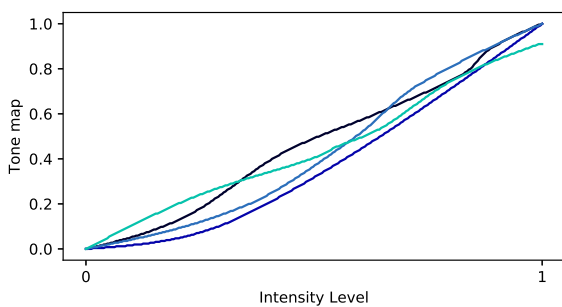


Figure 7: Some examples of user tone maps.

Experiment II: Quantative Analysis

Using our Weibull Tone Mapping method, we approximated the user tone mappings from Experiment I, henceforth defined as ‘Custom’. We found that WTM can closely mimic the brightness distributions created by end-users with a low average KLD of 0.12, but is often smoother and simpler in shape, see Figure 6. Indeed there are other smooth functions that could likely fit the data. However, we chose Weibull because it is well-known and it’s two parameters intuitively control the peak position and width of the fitted histogram.

The similarity of resulting Weibull images, that approximate their Custom counterparts, was assessed quantitatively using Universal Quality index (UQI) [19], Peak Signal to Noise Ratio (PSNR) and Mean Squared Error (MSE). This showed images to be highly similar, achieving an average UQI of 0.98, PSNR of 31.84 and MSE of 85.25.

Experiment III: Psychophysical Evaluation

Given the good performance of WTM in Experiment II, we undertook a pair-wise comparison experiment to determine to what extent analysts prefer - for their purpose of image interpretation - our WTM adjustment compared with 1) an existing enhancement algorithm CLHE, 2) their own custom enhancements from Experiment I and 3) the original images that are unenhanced.

Following the same experimental environment as Experiment I, a team of 10 analysts at Gardline, 3 of which created the custom images, were presented with, uniquely randomized, pairs of images. For each, they were asked to ‘Choose the image that best allows identification of the habitat (class) therein, or no preference if the images are equivalent’. Analysts could thus choose one of three options for each image; Image 1, Image 2 or No

Preference.

For each of the original images we have $n = 4$ variants, therefore there are $n/2(n - 1) = 6$ pair-wise comparisons. Since each pair of images is viewed twice, left-right order switched, the total number of comparisons is $60 \times 2 \times 6 = 720$ pairs; too many to compare in a reasonable time frame. We therefore chose a random selection of 18 images, 3 per habitat class (total=6), for the pairwise experiment. Each analyst therefore considered $6 \times 2 \times 18 = 216$ pairwise comparisons. A week later we repeated the experiment with a second random subset (of the same size), from the original 60 images. On average, each analyst took approximately ≈ 30 minutes to complete the experiment.

For each pair, the image chosen by an analyst was given a score of 1 and the other a score of 0. If no preference was selected, each image was given a score of 0, to remove the no-preference vote. Splitting the votes between the two pairs diluted the test result (signal-to-noise ratio); by assuming that within these comparisons analysts would respond randomly, a significant preference would be missed [20]. More detailed analysis of the pairwise data will be carried out in future work.

The data for each analyst, in each experimental sitting, was converted to a 4×4 frequency matrix, of which the value at $[i, j]$ represents the frequency of votes in which enhancement type i was preferred over type j across the 18 images. Each frequency matrix was then converted to a z-score (standard score) matrix using Thurstone’s Law of Comparative Judgments, or Thurstone’s Case (5) [21].

The expert-informed enhancements were considered, on average, somewhat more desirable than unenhanced input images and those enhanced with CLHE, see Figure 8. Although more tailored and preferable, the development of custom-per-image enhancements is time-consuming and therefore in-practical. Attempts must be made to develop an automated approach, that at the very least performs equivalently. Clear from these results is that, although popular, CLHE is not up to the task. However, our WTM shows good relative performance.

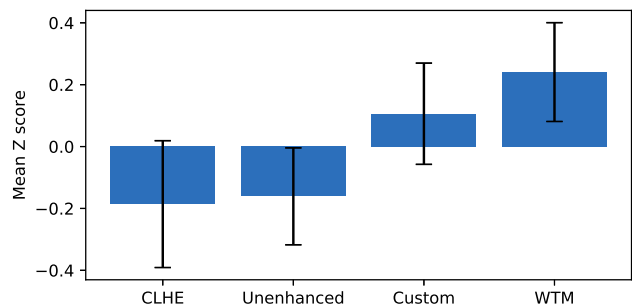


Figure 8: Mean Z scores for each image variant

Indeed, our WTM method performed best out of the pairwise comparisons. Figure 8 shows that end-users significantly prefer WTM images, over the original image and those enhanced with CLHE, to conduct their analysis. In Figure 9 we see a comparison of an unenhanced input image and its a corresponding WTM and CLHE enhancement. It is apparent that the CLHE enhancement is unsuitable here, over-enhancing the brightest pixels, emphasizing the spotlight, and resulting in a loss of details. The WTM enhancement however, although more subtle in this case, offers a more appropriate enhancement, improving visibility and conspicuousness of the community therein. Figure 8 also shows that analysts even weakly prefer WTM images over their own enhancements, see Figure 8. This is not unexpected given the good approximation of the custom images by WTM.




Figure 9: Image comparison of a) an unenhanced input image and the output image due to enhancement by b) WTM ($\lambda = 1, \kappa = 2$) and c) CLHE (Max slope = $3.2/256$, Min slope = $2.56e - 03/256$).

Conclusion

In this paper, we developed a method for tonally adjusting underwater images called Weibull Tone Mapping (WTM). Given an input and a user tone-mapped image, WTM provides an output image derived from the input-output pair. The method works by approximating the input and output distribution, with a Weibull PDF, and calculating the tone curve that maps between these distributions. This WTM map, is then applied to the input image to generate a new output.

The WTM method is designed, by construction, to approximate user adjustments but result in smoother and simpler tone maps. Experiments validate our method, showing that when compared to automatic image adjustments our method is strongly preferred. Compared with users' own adjustments, there is a weaker preference signature in WTM's favour.

Finally, we note that WTM, at this stage, is 'existential' in nature. Given user adjustment it provides even better tone maps. We are currently investigating how a WTM can be derived automatically from images.

References

[1] J. Horgan & D. Toal, Computer Vision Applications in the Navigation of Unmanned Underwater Vehicles, Underwater Vehicles, A. V. Inzartsev (Ed.), 2009, pg. 195-214.

[2] T. Schoening, M. Bergmann & T. Nattkemper, Investigation of hidden parameters influencing the automated object detection in images from the deep seafloor of the HAUSGARTEN observatory, Proc. Oceans, pg. 1-5. (2012).

[3] S. Pizer et. al., Adaptive histogram equalization and its variations, Computer vision, graphics, and image processing, 39, 3, pg. 355-368. (1987).

[4] K. Zuiderveld, Contrast Limited Adaptive Histogram Equalization, Graphics Gems, Academic Press, 1994, pg. 474-485.

[5] R. Schettini & S. Corchs, Underwater image processing: state of the art of restoration and image enhancement methods, EURASIP Journal on Advances in Signal Processing, 2010, 1, pg. 1-14. (2010).

[6] C. E. Shannon, A mathematical theory of communication, The Bell System Technical Journal, 27, 3, pg. 379-423. (1948).

[7] R. C. Gonzalez and R. E. Woods, Digital Image Processing, 2nd ed, Englewood Cliffs, NJ: Prentice-Hall, 2002, pg. 91-94.

[8] S. Shan, W. Gao, B. Cao & D. Zhao, Illumination normalization for robust face recognition against varying lighting conditions, Proc. 2003 IEEE International SOI Conference, pg. 157-164. (2003).

[9] J. McVey & G. Finlayson, Least-Squares Optimal Contrast Limited Histogram Equalisation, CIC 2019, pg. 256-261. (2019).

[10] R. Eustice, O. Pizarro, H. Singh & J. Howland, UWIT: Underwater Image Toolbox for optical image processing and mosaicking in MATLAB, Proc. International Symposium on Underwater Technol-

ogy, pg. 141-145. (2002).

[11] C. Town, A. Marshall & N. Sethasathien, Manta Matcher: automated photographic identification of manta rays using keypoint features, Ecology and evolution, 3, 7, pg. 1902-1914. (2013).

[12] A. S. A. Ghani & N. A. M. Isa, Underwater image quality enhancement through integrated color model with Rayleigh distribution, Applied Soft Computing, 27, pg. 219-230 (2015).

[13] J. Ma, X. Fan, S. X. Yang, X. Zhang & X. Zhu, Contrast limited adaptive histogram equalization-based fusion in YIQ and HSI color spaces for underwater image enhancement, International Journal of Pattern Recognition and Artificial Intelligence, 32, 7, 1854018 (26 pages). (2018).

[14] J. Geusebroek, A. W.M. Smeulders et al., A physical explanation for natural image statistics, Proc. The 2nd International Workshop on Texture Analysis and Synthesis, pg. 47-52. (2002).

[15] V. Yanulevskaya, J. Geusebroek et. al., Significance of the Weibull Distribution and its Sub-models in Natural Image Statistics, Proc. VISAPP, pg. 355-362. (2009).

[16] S. H. Scholte, S. Ghebreab, L. Waldorp, A. W. M. Smeulders, V. A. F. Lamme, Brain responses strongly correlate with Weibull image statistics when processing natural images, Journal of Vision, 9, 4, pg. 1-15. (2009).

[17] R. Sobol, Improving the Retinex algorithm for rendering wide dynamic range photographs, Proc. SPIE 4662, Human Vision and Electronic Imaging VII, (30 May 2002). (2002).

[18] ISO3664:2009, Graphic technology and photography-Viewing conditions, Geneva: ISO, 44 pages. (2009).

[19] Z. Wang & A. C. Bovik, A universal image quality index. IEEE signal processing letters, 9, 3, pg.81-84. (2002).

[20] H. R. Moskowitz, J. Beckley, and A. V. A. Resurreccion. 2006. Sensory And Consumer Research In Food Product Design And Development. Ames, Iowa: Blackwell, pg.226.

[21] L. L. Thurstone, A law of comparative judgment, Psychological review, 34, 4, pg. 273-286. (1927).

Acknowledgements

The authors gratefully acknowledge the support of EPSRC grant P007190, the NEXUSS CDT and Gardline Limited. We also thank the Gardline biologists for their contribution to this research.

Author Biography

Chloe Amanda Game is a PhD student in Computing Sciences at the University of East Anglia, focusing on underwater imaging. **Michael Barry Thompson** is an industry marine environmental scientist who has worked for Gardline Limited, a leading international marine survey contractor, since 2009. **Graham Finlayson** is a Professor of Computing Sciences at the University of East Anglia and leads the Colour Imaging research group.