https://doi.org/10.2352/issn.2694-118X.2020.LIM-19

Linear Histogram Adjustment for Image Enhancement

Jake McVey; University of East Anglia; Norwich, England.

Abstract

Tone curves are one of the simplest techniques for image enhancement. Specified as a function, a tone curve is a transformation that maps pixel levels of an input image to new output levels. Tone curves are the basis of many contrast enhancement algorithms, including Contrast Limited Histogram Equalisation (CLHE), which derives a tone curve from a modification of the image histogram. While these methods can provide good enhancement, they are generally non-linear.

In this paper we show the surprising result that a tone curve generated by the non-linear CLHE method (and HE) can be calculated by applying a linear transform to the histogram of the input image. Experiments validate our method.

INTRODUCTION

Many existing contrast enhancement techniques obtain a tone curve from the histogram associated with an image [1]. When applied to an image a tone curve can be simply thought of as a mapping from input to output intensities. In the context of contrast enhancement, a good tone curve is one that brings out the details of the image without introducing artefacts.

Histogram Equalisation (HE) is a well known method that obtains a tone curve from the input image histogram. To formalise the histogram definition, consider a typical greyscale image $\mathbf{I}$, with pixel values $\mathbf{I}(x,y) \in [0,255]/255$. The histogram representation for this image is a 256-vector $\mathbf{h}$, where each element, $\mathbf{h_k}$, counts the occurrences of the kth pixel value in the image. The histogram is generally normalised such that it sums to 1 (a PDF, or probability density function).

The HE algorithm seeks to find an image with maximum entropy. As is well known (but see [2] for review) this is achieved when the cumulative sum of the image histogram is used as the tone curve. Note then that there should be a direct relationship between the value of the histogram bins and the steepness of the tone curve. In fact, since we obtain the curve as a cumulative histogram sum, the derivative of the tone curve is indeed the histogram itself. This means that when the bins of a histogram grow large, the slope of the tone curve at that point will be very steep. Conversely, small values in the histogram result in shallow curves at those points. We illustrate this relationship in figure 1. Note the peak count in the histogram corresponds to the steepest part of the tone curve.

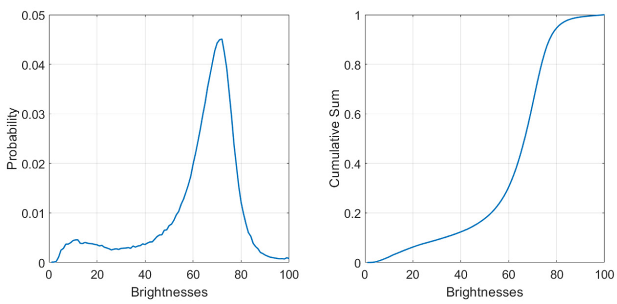

Figure 1: Left, histogram of an input image histogram. Right, tone curve obtained as the cumulative sum of the histogram.

Performing histogram equalisation - i.e. transferring an image with a tone curve defined to be the cumulative histogram of the image, often results in an unpleasant output image. From an entropy point of view the image has more information, but the image also often has unnatural contrast and false contours amongst other artefacts. The reason for this is when the bins of the histogram are too large or too small, the slope of the tone curve at those points maps values of the image to, respectively, a very wide or narrow range. Various methods have been developed that constrain the level of contrast enhancement through modification of bins in the histogram [3, 4, 5, 6, 7]. Recalling the relationship between histogram bin values and slope of the tone curve, Contrast Limited Histogram Equalisation (CLHE) [8] directly reduces the level of contrast in the output image by enforcing an upper-bound on the histogram bins.

Other methods in recent literature have recast histogram modification as an optimisation problem [9, 10, 11, 12], defining an objective function with penalty terms to enforce some characteristics on the tone curve. These characteristics - like the relationship between tone curve slope and level of enhancement - are designed to ensure the output image presents with a certain quality.

In figure 2a we show a Kodak test image. In 2b and 2c we show the outputs of applying two tone curves generated using HE and CLHE (where for CLHE the max and min slopes of the histogram are set to 2 and 0.5 respectively). Both these output images show a strong increase in contrast. But, the images are too punchy and neither are better than the original. In Figure 2d we show the output of the method we develop here. Specifically our tone curve is computed by:

$$\mathbf{t} = \mathbf{Mh} \tag{1}$$

where $\mathbf{M}$ is a 256×256 linear transform. For this example this simple linear tone curve generation function clearly outperforms CLHE and HE.

In section 2 we discuss in more detail relevant recent literature. Section 3 outlines our method. Experimental results are reported in section 4, and the paper is concluded in section 5.

Background

In Histogram Equalisation the tone curve, $\mathbf{t}$, - that maps an input image to an output image that has a uniform histogram - is simply the cumulative sum of the input histogram (normalised to sum to 1). As explained in the introduction, in this paper, we seek to develop a tone curve generation algorithm that can be written as $\mathbf{t} = \mathbf{Mh}$. We can neatly formalise the cumulative sum as a linear transformation of the input image histogram, $\mathbf{h}$, by the matrix $\mathbf{M}$, such that $\mathbf{Mh} = \mathbf{t}$. In this case $\mathbf{M}$ is a lower triangular matrix, with 1's on the bottom half, and 0's elsewhere:

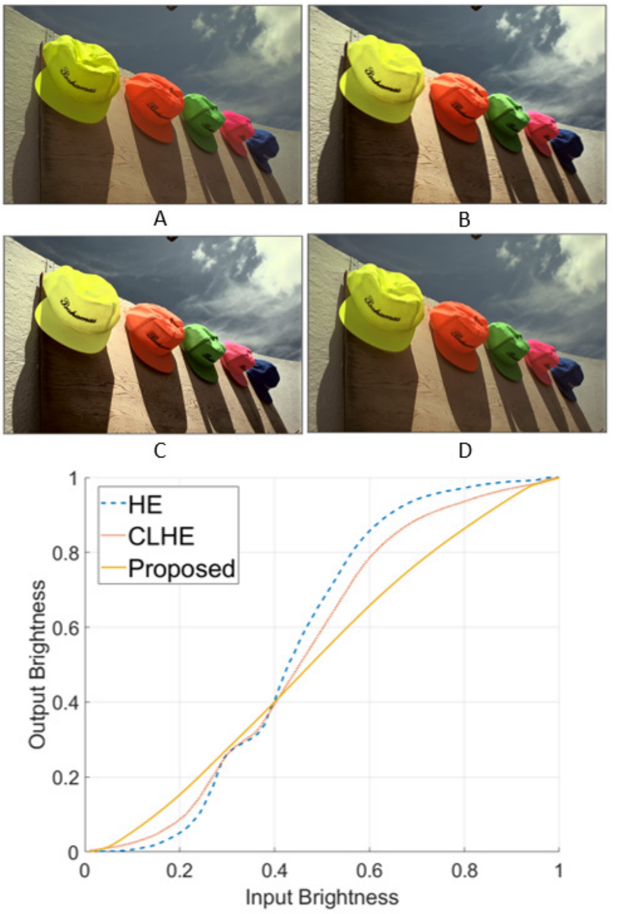

Figure 2: Tone curves (bottom) applied to an image. A) Original. B) CLHE. C) HE. D) Proposed method. Brightness calculated as $(R+G+B)/3$. Tone curve applied to each RGB channel independently.

$$\mathbf{M}_{HE} = \begin{bmatrix} 1 & 0 & \cdots & 0 & 0 \\ 1 & 1 & \cdots & 0 & 0 \\ 1 & 1 & \ddots & 0 & 0 \\ \vdots & \vdots & \cdots & 1 & 0 \\ 1 & 1 & \cdots & 1 & 1 \end{bmatrix}$$

where *HE* denotes *Histogram Equalisation*. Of course as illustrated in figure 2, Histogram Equalisation often produces disappointing results. Yet, it is interesting that this simple algorithm can be written in the linear framework. Perhaps the linear formalism idea is a good one but a better linear transform (than implied by Histogram Equalisation) can be found.

As illustrated in figure 2 the CLHE method allows us to find a tone curve where the maximum and minimum slope of the tone curve can be limited. The precise details of CLHE are not important. But, let us give the intuition. Suppose $\mathbf{h}$ is a histogram (that is normalised and sums to 1). Assuming again there are 256 bins and we are thinking of the image having brightness values in the interval $[0,1]$, then $\mathbf{h}$ is a 256 component vector and each component is in $[0,1]$ and the sum of $\mathbf{h}$ is 1. Since the tone curve $\mathbf{t}$ is the cumulative sum of $\mathbf{h}$ it follows that the slope of the bins is equal to $(\mathbf{h}_k)/(1/256)$ i.e. the derivative of the cumulative histogram $\mathbf{t}$ is the histogram itself and we divide by the step change (in this case $1/256$).

If we wish the max slope to be 2 then this means that $(\mathbf{h}_k)/(1/256) \leq 2$ or equivalently that $\mathbf{h}_k \leq 2/256$ for all k bins of the histogram. Here $2/256$ is sometimes called the "clip limit". Clearly if we calculate a histogram as $min(\mathbf{h}, 2/256)$, the cumulative sum of this new derived histogram meets the max slope limit. But, the clipping operation means the clipped histogram does not sum to 1. Thus the clipped values need to be redistributed (added back into the histogram). This means we may need to clip again. Indeed, CLHE works by iteratively clipping, redistributing, and clipping again until the resulting histogram meets the clip limit and sums to 1. A similar algorithm is adopted for enforcing minimum slope limits. Moreover, there are a variety of ways to redistribute the counts in the histograms that have been clipped.

We refer the reader to [8] for a detailed explanation of how CLHE works. For our purposes here the idea of clipping is a powerful one which we return to in the next section.

Importantly, CLHE was extended in [9] to impose extra constraints on the tone curve including smoothness and that the output images have "good" whites and blacks. This method in turn was reformulated in [12] as a single global least-squares optimisation. The method we develop in the next section is general - and can approximate any method that generates a tone-curve from a histogram. But, for the purposes of this paper we apply our method to [12].

Method

Equation 1 in the introduction encapsulates our method. Assuming $\mathbf{h}$ is a histogram that sums to 1 then $\mathbf{t} = \mathbf{Mh}$. However, we allow ourselves a little latitude in the definition of $\mathbf{h}$ and our actual formalism is equal to

$$\mathbf{t} = \mathbf{Mf(h)} \tag{2}$$

where $\mathbf{f}$ is a function that can pre-process the histogram. In the current paper we pre-process the input histogram by running CLHE for one iteration - i.e. we clip to the max value, a renormalise (without redistributing) so the resulting histogram sums to 1.

$$\mathbf{g}(\mathbf{h}) = min(\mathbf{h}, max\_val) \tag{3a}$$
$$\mathbf{f}(\mathbf{g}) = \mathbf{g}(\mathbf{h})/sum(\mathbf{g}(\mathbf{h})). \tag{3b}$$

We do this simply make the input histograms robust to large image areas that have the same value (i.e. those that cause the high slopes in the tone curve).

Now suppose we have N histograms denoted $\mathbf{h^i}$ $(i = 1,2,\cdots,N)$ we calculate the corresponding N tone curves $\mathbf{t^i}$ using [12] (a closed-form extension of CLHE which produce stable preferred results). Our hypothesis is that there exists a 256×256 matrix $\mathbf{M}$ such that $\mathbf{Mf(h^i)} \approx \mathbf{t^i}$.

To test this hypothesis we solve the following regression

$$\mathbf{M} = \underset{\mathbf{M}}{min}\left(\sum_{j=1}^{N} ||\mathbf{Mh}^j - \mathbf{t}^j||_2^2 + \lambda ||\mathbf{M}||_2^2\right) \tag{4}$$

Here λ is a user defined penalty term that regularises the expression. When $\lambda = 0$ the expression resolves to ordinary least squares. The purpose of this penalty term is to guide the prediction to a more 'stable' solution. In the context of our mapping problem, an unstable prediction is one that maps pixel values in the image outside of the expected range (e.g. mapping values of an sRGB image outside [0, 1], remember we are representing image values as lying between 0 and 1), or a curve that is not strictly increasing.

From a mathematical perspective regularisation is a method to deal with sets of equations which we can solve (or approximately solve here) but where the equation solving process is unstable. Relevant to the problem at hand, the instability follows from the fact that the histograms $\mathbf{f}(\mathbf{h^i})$ taken together do not form (span) 256 dimensional space. The regularisation term in Equation 4 penalises the norm of the recovered matrix, in effect simultaneously bounding the variance of the solved-for tone curves.

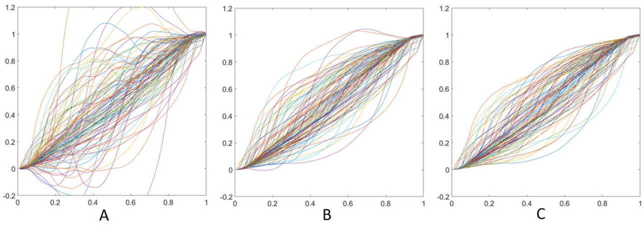

Figure 3: Plot of $\mathbf{M}$, calculated with Equation 5 with various different λ values. A) $\lambda = 0$. B) $\lambda = 0.01$ C) $\lambda = 1.5$.

The closed form solution to Equation 4 is

$$\mathbf{M} = (\mathbf{A}^T\mathbf{A} + \gamma\mathbf{I})^{-1}\mathbf{A}^T\mathbf{B} \tag{5}$$

where $\mathbf{A}$ and $\mathbf{B}$ are both $N \times 256$ matrices (where the rows are respectively corresponding histograms and tone curves), and $\mathbf{I}$ is the 256×256 identity matrix. In Figure 3 we show 3 examples of $\mathbf{M}$ calculated using 3 different values for λ using Equation 5. Notice that $\mathbf{M}$ - with an ill-fitting λ - in 3a and 3b resolves to (what we previously described as) an unstable solution. The result in 3c is clearly much more suitable. It is evident that - without any explicit enforcement on our part - the solved-for curves are naturally smooth.

The optimal Tikhonov regularisation parameter in Equation 5 was found via grid search on validation data.

Experiments

We evaluate the success of our model by comparing the *closeness* of images enhanced with the proposed method and with those generated by the target enhancement algorithm, in this case [12]. Closeness in this instance will be quantitatively measured using CIELAB ΔE* 76 [13]. That is, each image in the test data set is enhanced with the tone curve obtained from [12] and our method independently. Then we convert both enhanced images to the CIELAB color space, and finally obtain a per pixel error using ΔE*. Given this error image we can calculate mean, median, and 99 percentile error statistics for each image.

We test our method on two datasets: the well-known Kodak dataset [14], and from a randomly selected sub-set of images in the ImageNet dataset [15]. To build the proposed model we used 20,000 images from the ImageNet dataset as a training set (N in Equation 4), and 5,000 images for testing. Images from the Kodak dataset were not used to build the model, but all 24 images from [14] were used as an independent test set (accessed on 01-06-19).

The mean and standard deviation for these measures across all test images are calculated and shown in the Table 1. In Figure 4 we show some images for comparison. Left is the original unenhanced Kodak image. The image enhanced by [12] is shown in the middle. Our simple enhancement - where the tone curve applied is calculated as a simple linear transform of the input image's histogram - is shown right. The algorithm [12] produces pleasing results and our linear transform algorithm provides visually indistinguishable results, validating our method.

Results and Analysis

The target ΔE* value for image matching is not universally accepted. In [16] it is argued that a mean ΔE* of less than 2.15 in a complex image is required for close perceptual uniformity, while in [17] it is reported a ΔE* of 5 is sufficient. In Table 1, we show that our method is able to satisfy both benchmarks.

	Mean ΔE*	Median ΔE*	99 pt. ΔE*
ImageNet	1.69 (± 1.68)	1.61 (± 0.95)	3.37 (± 1.49)
Kodak	1.79 (± 0.57)	1.64 (± 0.68)	3.9 (± 1.02)

Table 1: Mean (± standard deviation) of the mean, median, and 99-percentile of ΔE* measured between images enhanced with [12] and the shown models over shown data sets.

In Figure 4 we show several images from the Kodak data set. The middle images enhanced with [12] all present with pleasing contrast, and are better than the input images on the left. On the right we show the same images enhanced with curves obtained by our proposed method. The ΔE* for each pair of enhanced images is also presented. The images in this Figure teach us that the suggested target ΔE* of 3 is sensible, since even the 'most-different' image in the set - the red door - only begins to show slight differences upon close observation.

We make two final remarks. First, that the optimisation framework [12] will generally be a preferred solution for achieving naturally enhanced images. But, the optimisation uses Quadratic Programming (QP) that is computationally expensive and therefore infeasible for many applications. Furthermore, dimensionality reduction techniques - like singular value decomposition (SVD) - can be introduced into the proposed linear framework to reduce the computational complexity (and perhaps improve accuracy) of our method. Second we note that the proposed method, that is, a linear regression model for approximation, was chosen for for it's simplicity.

A potential advantage of our method is that it will be easier to analyse how the algorithm works when the input is perturbed. If the input histogram $\mathbf{f}(\mathbf{h})$ is perturbed by a random bin adjustment ε the tone curve, perforce, will be the sum of the two i.e. $\mathbf{Mf}(\mathbf{h}) + \mathbf{M}\varepsilon$.

Conclusion

Tone curves are often used in image processing to enhance images: the output image is a tone adjustment of the input image. The first tone-curve adjustment algorithm was Histogram Equalisation - where the tone curve is the cumulative sum of the input histogram and which, unfortunately, often provides over enhanced outputs - and this method has been further developed to, more reliably, produce preferred image outputs [3, 4, 5, 6, 7].

In this paper we demonstrate the interesting fact that we can closely approximate Histogram Equalisation based methods which generate tone curves from the histogram of an image, by a linear matrix computation. Specifically, the tone curve we propose applying to an image is found by applying a linear transform to the input histogram. Further, this transform can be found by regularised regression.

We conduct experiments to test the effectiveness of our approximation and present the results quantitatively using the CIELAB ΔE* difference metric, and visually with enhanced images for comparison. Almost always (for 10s of thousands of image) our simple linear algorithm delivers enhanced images that are indistinguishable from those delivered by more complex non linear counterparts.

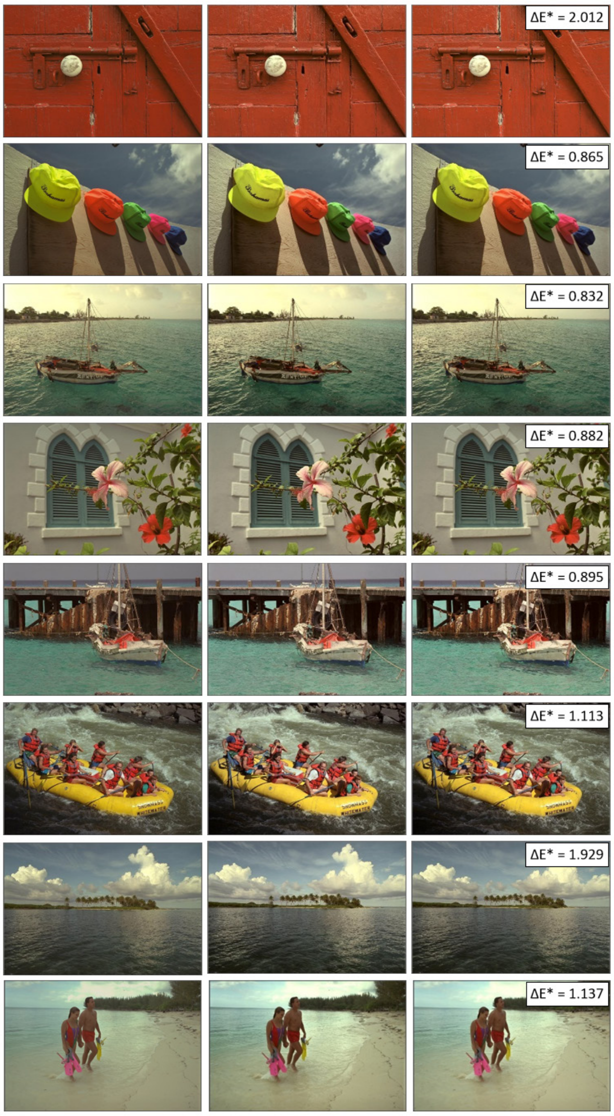

Figure 4: *For each image in the set:* First, original image. Middle, image enhanced with [12]. Right, image enhanced with proposed method.

Acknowledgements

The author is grateful for the support from Spectral Edge Ltd, and to Prof. Graham Finlayson for helpful discussions and guidance during this work.

References

[1] A. S. Parihar and O. P. Verma, “Contrast enhancement using entropy-based dynamic sub-histogram equalisation,” *IET Image Processing*, vol. 10, no. 11, pp. 799–808, 2016.

[2] R. C. Gonzalez and R. E. Woods, *Digital image processing*. Prentice-Hall, 2. ed. ed., 2002.

[3] Yeong-Taeg Kim, “Contrast enhancement using brightness preserving bi-histogram equalization,” *IEEE Transactions on Consumer Electronics*, vol. 43, pp. 1–8, Feb 1997.

[4] Soong-Der Chen and A. R. Ramli, “Minimum mean brightness error bi-histogram equalization in contrast enhancement,” *IEEE Transactions on Consumer Electronics*, vol. 49, pp. 1310–1319, Nov 2003.

[5] A. Paul, P. Bhattacharya, S. P. Maity, and B. K. Bhattacharyya, “Plateau limit-based tri-histogram equalisation for image enhancement,” *IET Image Processing*, vol. 12, no. 9, pp. 1617–1625, 2018.

[6] Y.-R. Lai, P.-C. Tsai, C.-Y. Yao, and S.-J. Ruan, “Improved local histogram equalization with gradient-based weighting process for edge preservation,” *Multimedia Tools Appl.*, vol. 76, pp. 1585–1613, Jan. 2017.

[7] S. Poddar, S. Tewary, D. Sharma, V. Karar, A. Ghosh, and S. K. Pal, “Non-parametric modified histogram equalisation for contrast enhancement,” *IET Image Processing*, vol. 7, pp. 641–652, October 2013.

[8] S. M. Pizer, E. P. Amburn, J. D. Austin, R. Cromartie, A. Geselowitz, T. Greer, B. ter Haar Romeny, J. B. Zimmerman, and K. Zuiderveld, “Adaptive histogram equalization and its variations,” *Computer vision, graphics, and image processing*, vol. 39, no. 3, pp. 355–368, 1987.

[9] T. Arici, S. Dikbas, and Y. Altunbasak, “A histogram modification framework and its application for image contrast enhancement,” *IEEE Transactions on Image Processing*, vol. 18, pp. 1921–1935, Sep. 2009.

[10] D. Kim and C. Kim, “Contrast enhancement using combined 1-d and 2-d histogram-based techniques,” *IEEE Signal Processing Letters*, vol. 24, pp. 804–808, June 2017.

[11] J. Shin and R. Park, “Histogram-based locality-preserving contrast enhancement,” *IEEE Signal Processing Letters*, vol. 22, pp. 1293–1296, Sep. 2015.

[12] J. McVey and G. Finlayson, “Least-squares optimal contrast limited histogram equalisation,” in *Color and Imaging Conference*, vol. 2019, pp. 256–261, Society for Imaging Science and Technology, 2019.

[13] A. R. Robertson, “The cie 1976 color-difference formulae,” *Color Research & Application*, vol. 2, no. 1, pp. 7–11, 1977.

[14] “Kodak lossless true color image suite.” http://r0k.us/graphics/kodak/. Accessed: 01-06-19.

[15] “Imagenet image database.” http://image-net.org/index. Accessed: 01-06-19.

[16] M. Stokes, M. D. Fairchild, and R. S. Berns, “Precision requirements for digital color reproduction,” *ACM Trans. Graph.*, vol. 11, pp. 406–422, Oct. 1992.

[17] G. W. Meyer, “Reproducing and synthesizing colour in computer graphics,” *Displays*, vol. 10, no. 3, pp. 161–170, 1989.

Author Biography

Jake McVey is a PhD student at the University of East Anglia (UEA). His work is focussed on colour image processing and image preference.