

Simplifying Tone Curves for Image Enhancement

James Bennett and Graham Finlayson

School of Computing Sciences, University of East Anglia; Norwich, UK.

Abstract

A single tone curve which is used to globally remap the brightness of each pixel in an image is one of the simplest ways to enhance an image. Tone curves might be the result of individual user edits or from algorithmic processing including in-camera processing pipelines. The precise shape of the tone curve is not strongly constrained other than it is usually limited to increasing functions of brightness. In this paper we constrain the shape further and define a simple tone adjustment, mathematically, to be a tone curve that has either no or one inflexion point. It follows that a complex tone curve is one with more than one inflexion point, visually making the curve appear ‘wiggly’. Empirically, complex tone curves do not seem to be used very often. For any given tone curve we show how the closest simple approximation can be efficiently found. We apply our approximation method to the MIT-Adobe FiveK dataset which comprises 5000 images that are manually tone-edited by 5 experts. For all 25,000 edited images - where some of the tone adjustments are complex - we find that they are all well-approximated by simple tone curve adjustments.

Introduction

Tone curves map input values to corresponding output values by a continuous, almost always increasing, function and are used for many purposes. A tone curve that describes the mapping of the real-world scene radiance to measured pixel values is sometimes called a camera response function [1] and is often a linear map. The inverses of these functions are also used to recover the scene radiance from an image. Tone curves are applied to map a large dynamic range of brightnesses that are measured at image capture to the lower dynamic range of a display [2]. Also, images are generally encoded after the application of a gamma function (a compressive tone-curve function) because the target display inverts the gamma when an image is displayed. However, the predominant use of tone curves - and the one of interest in this paper - is image enhancement. Simply, a tone curve is applied to the brightnesses of an input image which maps the brightness to output counterparts. This tone-mapped output image is - in some sense - enhanced.

As discussed in [3], an enhancement is only suitable according to its purpose and audience, thus there is no single optimal enhancement. However, regarding post-processing for photographic purposes, the hope and expectation are that the enhanced image is preferred compared to the original image. Many photographers routinely edit the photographs they capture to reproduce the tones of the image in a more pleasing manner. Significantly, while a tone curve can be described explicitly (by drawing a tone curve), tone adjustments are often implicit: users enhance images through manual slider-based adjustment, pre-made ‘filters’, or automatic enhancement tools. Often the effect of these adjustments can be modelled as the application of a tone curve. Perhaps unsurpris-

ingly, [4] has shown that the ‘best’ tone-mapped image for one person may not lead to the most preferred image for another.

Tone maps can be applied to all colour channels in an image [5], to the individual colour channels [6, 7] or to a brightness channel only [8, 9]. In this last approach, while the adjustment takes place on a brightness channel, the effect is transferred to the colour image. For example, we might map the brightness channel of an input image to an enhanced output brightness channel. Then, per pixel, the ratio of the output over input brightnesses is applied to each RGB. The result of this operation is to make an output image that has the same chromaticities as the input but where the brightnesses have been correctly reproduced (according to the enhancement). Alternatively, we might map an input image to the CIELAB colour space [10], enhance the lightness channel, L^* , and map back to RGB. It is this workflow we use in this paper.

There are many techniques employed for enhancing images [11, 12]. The focus of this work is a global brightness remapping. Given an input scalar brightness image $i(x,y)$ where the pixels are indexed using their location x and y , the enhanced output scalar image,

$$o(x,y) = f(i(x,y)), \quad (1)$$

where f is the mapping function. Note also that the inverse function f^{-1} will recover the original brightness from the enhanced counterpart. When the same function is applied to all pixels in the image, the enhancement is said to be *global*. When different functions are applied to different pixels then the mapping would be *local*. Almost always, f is an increasing function of brightness.

An example of tone mapping is shown in Figure 1. An input image is shown on the left which produces the image on the right that has been enhanced with the tone curve shown in the middle. The brightness image, $i(x,y)$, is the L^* channel in the CIELAB colour space. The curve depicts the enhancement f and is read that a pixel with a value of 0.3 in the input brightness image (marked as a ‘dot’ in the graph in the Figure) would be mapped to a value of roughly 0.7 in the output image.

The tone curve in Figure 1 is very simple, showing effectively a single control point at (0.3,0.7). Generally, tone curves map the entire input brightness range to the same output range, here the interval [0,1]. The three points (0,0), (0.3,0.7) and (1,1) are interpolated with a piecewise cubic spline such that the gradient at the control point is continuous to form the smooth curve, passing through the control point as seen in Figure 1.

In many image manipulation tools, a user might select several control points and adjust them independently. A second example is shown in Figure 2 with 5 control points and the same interpolation. Here, the tone curve is ‘wiggly’ and in that sense complex. Informally, one of the questions we consider in this paper is whether complex (wiggly) tone curves are ever chosen by users to enhance their images.

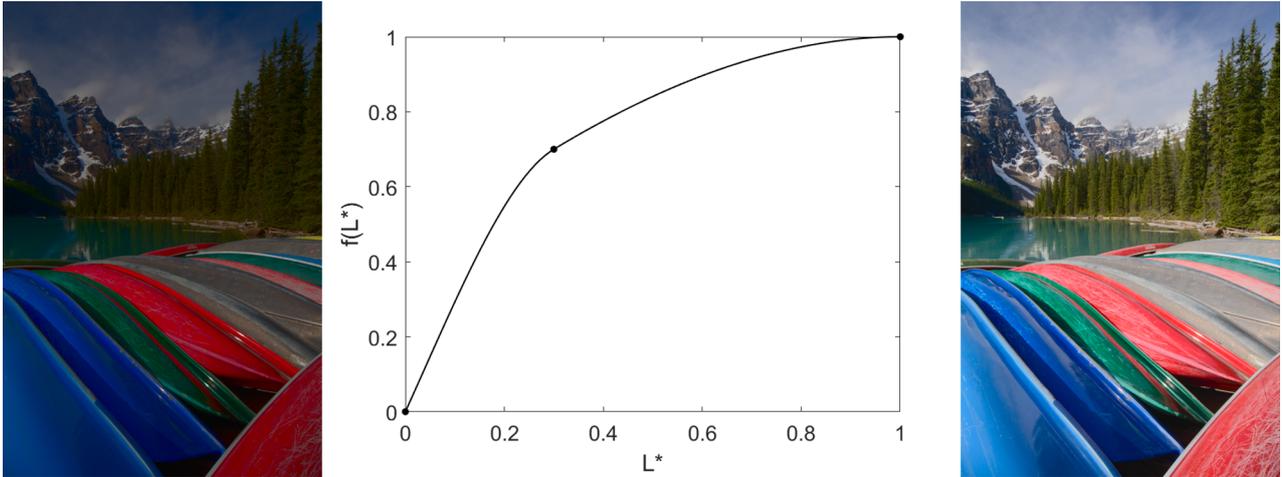


Figure 1. Left, middle and right show respectively, an input image, a tone curve that maps the input brightness to output counterparts and the result of applying the tone curve. The image is number 635 Expert E's a^* and b^* channels.

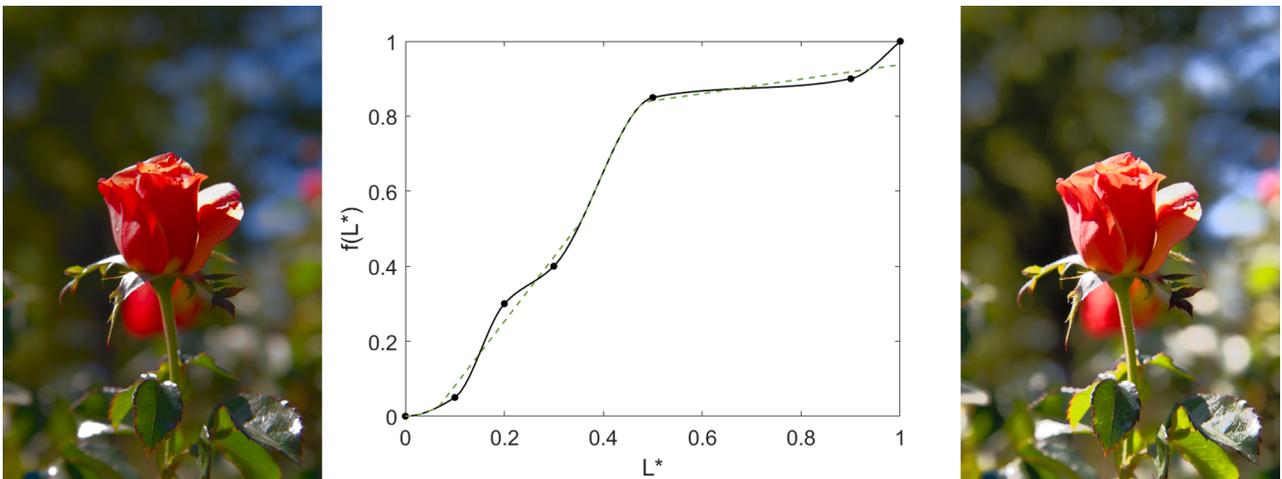


Figure 2. Left, middle and right show respectively, an input image, a tone curve that maps the input brightness to output counterparts and the result of applying the tone curve. The image is number 216 Expert A's a^* and b^* channels. The green dashed line shows the simple curve approximation and is discussed later.

We propose that simple tone curves have the property that they have 0 or 1 inflexion point. For monotonically increasing tone curves (which account for almost all tone curves applied to images), this means the gradient of a tone curve might be monotonically increasing (e.g. $o = i^2$) or it might increase to the inflexion point and then decrease as in an S-shaped curve. Or, the gradient might decrease monotonically (e.g. $o = i^{0.5}$) or decrease then increase at the inflexion point (an inverse S-shaped curve).

In this paper, we look at the tone manipulations made by 5 expert photo retouchers for the same set of 5000 images as made available in the MIT-Adobe FiveK dataset [8] (hereafter referred to as FiveK). Each individual manipulation can be the result of the direct manipulation of a tone curve and adjustments of sliders including, but not limited to, exposure and contrast. For **all** 5000 images we found that the effective tone adjustment that these experts make **is** simple in the sense that there exists a simple - single inflexion point - tone adjustment that results in a visually indistinguishable output image.

This begs the question as to whether photo-processing tools should allow (or at least make it easy) for users to make non-simple tone adjustments. It could be that they are allowing their users to make adjustments that will not be preferred or, perhaps, as preferred as a simple adjustment. Equally, would automated tone-mappers that produced only simple curves deliver preferred outputs (compared to the more wiggly tone curves that can currently be delivered by automated methods).

In the *Background* section, prior work related to the FiveK dataset is discussed. The *Method* section discusses the ground-truth data we use (where each expert's edit can be interpreted as a single tone curve adjustment) and presents our algorithm for approximating any tone curve by a simpler proxy. Experimental results are presented in the *Results* section which is followed by a short *Discussion* section. The paper finishes with a *Conclusion* section.

Background

In this work, we will use the FiveK dataset [8] which comprises 5000 photographs from a range of scenes including the natural world, people, built environment, and man-made objects. Each photograph has been retouched by five experts, resulting in 25,000 image pairs. Because each expert has edited according to their preference, the different renditions range from being similar to markedly different from one another. It is generally proposed that each expert’s individual edits can be well-approximated by a single global tone curve.

The FiveK dataset has been used in designing automatic image enhancement algorithms [13, 14] which are constrained to be global enhancements in [15] and further to be human interpretable curves in [16]. When a tone mapping (or its parameters) is the algorithmic output [5, 6, 7, 17] (as opposed to the image itself) a tone curve can be determined from a thumbnail of an image and then applied to the full-resolution image to produce an enhanced output, reducing computational demand [18].

Curiously, the premise behind the FiveK dataset - that there is a global tone curve to be learnt - turns out only to apply somewhat approximately in practice (at least for some images). In Figure 3 we plot input (unedited) versus output (expert-adjusted) L^* lightness values [10] for two image pairs. For the example in the left plot, the relationship between input and output is clearly almost global.

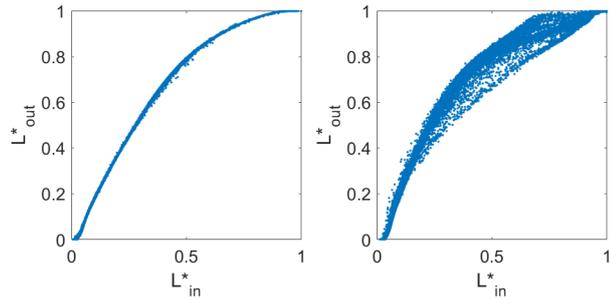


Figure 3. Scatter plot of L^* values of the input image against the expert’s rendition. Left shows image 294 against Expert D’s rendition where the relationship between input and output is well-modelled by a global function. Right shows image 281 against Expert D’s rendition where the relation is not well-modelled by a global curve.

A second example is shown in the right plot of Figure 3. While there is a single global trend the plot has a variety of outputs for every single input and this may indicate that some local processing has been applied (such as highlight recovery) to the image and/or the colour rendition has changed (e.g. a saturation boost). Typically, colour changes and spatial processing cannot be modelled by a tone curve operating on L^* .

Of course, any edits made by a photographer will be made with respect to a particular colour space representation. Should we expect the input-lightness to adjusted-lightness plot to result in a one-to-one global curve? After all, if an image is tone mapped, say where the brightnesses were equal to $(R+G+B)/3$ (as opposed to Luminance) then the resulting lightness to lightness plot would not be global. From informally looking at this issue we found the L^* channel often led to good global tone curves and other colour encodings led to curves that were less global. This said, the exact

colour space transforms used in mapping input to output images are not explicitly disclosed in the FiveK dataset.

Method

Dataset

The FiveK dataset images can be downloaded individually or as an Adobe Lightroom catalogue. We have 5000 input images and 5 output counterparts for each input (the same image manipulated by 5 experts). Adopting the procedure of [17] (also used by [6, 15]), the collection *Input/Input with Daylight WhiteBalance minus 1.5* is used as the input images and image pairs are exported from Adobe Lightroom in ProPhotoRGB. Each image is then converted into the CIELAB colourspace. As discussed earlier, we will model the tone mapping from input to output brightnesses by considering corresponding pixel values in the L^* channel only.

Generating a global ground truth tone curve

As shown in Figure 3, the relationship between input and output L^* images may or may not be well-modelled as a global tone curve. However, our analysis of whether simple or complex tone curves are used to enhance images is only applicable to images where there is a global tone curve. Let $\underline{I}(x,y)$, $\underline{P}(x,y)$ and $\underline{P}^G(x,y)$ denote respectively an input image, an expert adjusted output and an approximation thereof. For the purposes of this paper, the relationship between these three images is summarised as:

$$\begin{aligned} \underline{I} &= [L_I^* a_I^* b_I^*]^\top \\ \underline{P} &= [L_P^* a_P^* b_P^*]^\top \\ \underline{P}^G &= [f(L_I^*) a_P^* b_P^*]^\top \end{aligned} \quad (2)$$

where, for a single given pixel, \underline{I} , \underline{P} and \underline{P}^G are 3-vectors, dropping the spatial dependence on (x,y) . The function f denotes a tone curve that approximately maps L_I^* to L_P^* . Here and throughout the superscript $^\top$ denotes the transpose operator. We found that for defining f using histogram matching, the images $\underline{P}(x,y)$ and $\underline{P}^G(x,y)$ were almost always visually indistinguishable. Figure 4 shows an example of an input image $\underline{I}(x,y)$, the expert adjusted output $\underline{P}(x,y)$ and our global approximation $\underline{P}^G(x,y)$.

The tone-curve $f()$ is defined by k uniformly spaced quantiles $[0, 1/(k-1), 2/(k-1), \dots, 1]^\top$ relating the L^* image for the input and expert edited images [3]. The quantile-to-quantile plot can then be used to interpolate any input control points [19]. Thus, for the 100 uniformly spaced input L^* values that we use to define tonality $[0, 1/99, 2/99, \dots, 1]^\top$, we calculate the outputs of the tone curve as the 100-vector $\mathbf{f} = [f(0), f(1/99), f(2/99), \dots, f(1)]^\top$. That is, we assume the equivalence:

$$\mathbf{f} \equiv f() \quad (3)$$

For the rest of this paper we assume that the ground truth tone-mapping output for the FiveK dataset are the images $\underline{P}^G(x,y)$.

The key concern of this paper is to consider to what extent tone curves applied to images are or are not simple (see next section for definition). Thus, we will seek to find an approximate tone curve, $\hat{\mathbf{f}}$ such that

$$\hat{\mathbf{f}} \approx \mathbf{f}, \text{ where the shape of } \hat{\mathbf{f}} \text{ is constrained to be simple} \quad (4)$$



Figure 4. Left, middle and right show respectively, an input image, $I(x,y)$; the expert's rendition, $P(x,y)$; and the global approximation, $P^G(x,y)$. In this example, image 1925 and Expert C are shown.

Simple Tone Curves

Tone curves f are surjective increasing functions - that is, if $b_j \geq b_i$ then $f(b_j) \geq f(b_i)$ - where the input and output values are (typically) defined on the interval $[0,1]$. As we discussed in the introduction, we have an expectation that tone curves are additionally *simple* in some sense: we do not expect the tone curve to be wiggly. See, Figure 2 in the introduction. Wiggly here, intuitively, means that there are only so many 'turns' in the tone curve: the more it looks like a series of steps the more wiggly it is. In part, we feel that wiggly tone curves should not be used because, from the user's viewpoint, they seem to be difficult to define (and although they can be defined they may be difficult to replicate).

We define simple tone curves to be increasing functions that have no or only a single inflexion point. In terms of the derivatives of the tone curves, this means they are increasing to the inflexion point and then decreasing (or the converse) or have a strictly increasing or decreasing derivative. For a given tone curve represented by a vector \mathbf{f} - which may have more than one inflexion point - we would like to find an approximate tone-curve $\hat{\mathbf{f}}$ that is simple and as close as it can be to \mathbf{f} . Clearly, to do this will involve solving an optimisation and part of this optimisation will, perform, involve the calculation of discrete derivatives.

Let us define a square matrix \mathbf{D} to calculate the derivative (here discrete differences) of a vector:

$$\mathbf{D} \triangleq \begin{bmatrix} -1 & 1 & 0 & \cdots & 0 & 0 & 0 \\ -1 & 1 & 0 & \cdots & 0 & 0 & 0 \\ 0 & -1 & 1 & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -1 & 1 & 0 \\ 0 & 0 & 0 & \cdots & 0 & -1 & 1 \end{bmatrix}, \quad (5)$$

Notice that the first two rows are the same. This is because we need to make a statement about what the derivatives are at the boundary of the domain. By replicating the first two rows, we are assuming that the gradient is constant here (we adopt the so-called homogeneous Neumann boundary conditions). Given that our tone curves are represented by 100 component vectors the matrix \mathbf{D} is 100×100 . Using the subscript i to index the i th term in a vector, we see that,

$$\begin{aligned} [\mathbf{D}\mathbf{f}]_i &= \mathbf{f}_i - \mathbf{f}_{i-1}, \text{ for } i \in [2, 100] \\ [\mathbf{D}\mathbf{f}]_1 &= [\mathbf{D}\mathbf{f}]_2 \end{aligned} \quad (6)$$

Our concept of inflexion point is defined in terms of increasing and decreasing first derivatives. Equivalently, a positive sec-

ond derivative indicates an increasing gradient and a negative second derivative indicates a decreasing gradient. A second derivative is simply the derivative of the first derivative. Our second derivative operator is denoted \mathbf{D}^2 and is defined:

$$\mathbf{D}^2\mathbf{f} \triangleq \mathbf{D}(\mathbf{D}\mathbf{f}) \quad (7)$$

We are going to solve for the best simple approximate tone curve $\hat{\mathbf{f}}_i$ for 4 cases, $i = 1, 2, 3, 4$. Then the best tone curve overall will be the one that is closest to the original curve:

$$\arg \min_{\hat{\mathbf{f}} \in \mathbf{S}} \|\hat{\mathbf{f}} - \mathbf{f}\|, \mathbf{S} = \{\hat{\mathbf{f}}_1, \hat{\mathbf{f}}_2, \hat{\mathbf{f}}_3, \hat{\mathbf{f}}_4\} \quad (8)$$

Each of our 4 minimisations is formulated as the problem of finding the closest approximation to a given tone curve subject to the constraints that the tone curve is between 0 and 1 and also that it is increasing. Thus, the first part of each optimisation is written:

$$\mathbf{J}_i = \arg \min_{\hat{\mathbf{f}}_i} \|\hat{\mathbf{f}}_i - \mathbf{f}\| \text{ s.t. } \mathbf{0} \leq \hat{\mathbf{f}}_i \leq \mathbf{1}, \mathbf{D}\hat{\mathbf{f}}_i \geq \mathbf{0}. \quad (9)$$

Then the constraints on the second derivative are considered to yield the 4 minimisations. Case 1: gradient increasing (no inflexion point)

$$\arg \min_{\hat{\mathbf{f}}_1} \mathbf{J}_1, \mathbf{D}^2\hat{\mathbf{f}}_1 \geq \mathbf{0} \quad (10)$$

Case 2: gradient decreasing (no inflexion point)

$$\arg \min_{\hat{\mathbf{f}}_2} \mathbf{J}_2, \mathbf{D}^2\hat{\mathbf{f}}_2 \leq \mathbf{0} \quad (11)$$

Now let us consider the case where there is an inflexion point. The inflexion point λ , is an integer in $[2, 99]$ (any point in the tone curve other than the first and last). Let us use $\mathbf{f}_{[1, \dots, \lambda]}$ to denote the first λ terms in the vector and $\mathbf{f}_{[\lambda+1, \dots, 100]}$ to denote the remaining vector components.

Case 3: gradient increasing then decreasing (one inflexion point)

$$\arg \min_{\hat{\mathbf{f}}_3, \lambda \in [2, 99]} \mathbf{J}_3, [\mathbf{D}^2\hat{\mathbf{f}}_3]_{[1, \dots, \lambda]} \geq \mathbf{0}, [\mathbf{D}^2\hat{\mathbf{f}}_3]_{[\lambda+1, \dots, 100]} \leq \mathbf{0}. \quad (12)$$

Case 4: gradient decreasing then increasing (one inflexion point)

$$\arg \min_{\hat{\mathbf{f}}_4, \lambda \in [2, 99]} \mathbf{J}_4, [\mathbf{D}^2\hat{\mathbf{f}}_4]_{[1, \dots, \lambda]} \leq \mathbf{0}, [\mathbf{D}^2\hat{\mathbf{f}}_4]_{[\lambda+1, \dots, 100]} \geq \mathbf{0}. \quad (13)$$

Each of these objective functions characterises the best approximate tone curve $\hat{\mathbf{f}}$ as the one that minimises a mean error (sum of squares) norm subject to linear inequalities. Thus, each minimisation is precisely a quadratic program which can be solved efficiently and for which a global optimum is found [20]. For a single tone curve we need to carry out Case 1 and Case 2 minimisations and then, for every inflexion point λ , 98 variants of Case 3 and Case 4. Over all of these 198 minimisations (for each of which a global optimum is found) we choose the $\hat{\mathbf{f}}$ that is closest to \mathbf{f} , Equation 8.

When we apply the solved for tone curve, $\hat{\mathbf{f}}$, to an input image, $I(x,y)$, we generate an output $\hat{P}(x,y)$. Remembering that we can interchange a continuous tone-curve f by its vector representation \mathbf{f} and adopting the notation of Equation 2, per pixel,

$$\hat{P} = [\hat{f}(L_I^*) a_p^* b_p^*]^\top \quad (14)$$

which should be close to the ground-truth, $\hat{P}(x,y) \approx P^G(x,y)$. We examine the extent to which this is true in the next section.

Results

For each input image $I(x,y)$ we have the actual output image $P(x,y)$ which is the result of an expert’s image edits. However, this pair of images are not exactly related to one another by a global tone curve and we would like there to be such a relationship to test our method. Thus, for ground truth images, we use $P^G(x,y)$ formed using Equation 2.

Of course, we should only use this approach if, visually, $P^G(x,y)$ looks like $P(x,y)$. To test this, for all 25,000 images (5000 input images adjusted for preference by 5 experts) we calculate the mean CIELAB colour difference, ΔE , between the actual adjusted output $P(x,y)$ and our ground truth $P^G(x,y)$ as the Euclidean distance of their $L^*a^*b^*$ coordinates; the results are summarised in Table 1. We find 50% of our data has a mean ΔE of less than 0.36 and even the 0.99 quantile is only 1.88 (visually not noticeable in images [21, 22]). The maximum mean ΔE is 10 but this is an outlier; all of the remaining 24,999 images have a ΔE less than 4. Henceforth we use $P^G(x,y)$ as our ground truth.

Table 1: Quantile mean ΔE between $P(x,y)$ and $P^G(x,y)$ (25,000 images)

Quantile	0.5	0.9	0.95	0.99	1
Mean ΔE	0.355	0.598	0.928	1.88	10.1

Now we have pairs of input and output ground truth images where the Lightness channel of the output image is precisely a tone curve from the input lightness channel. We can now consider the extent to which these tone curves are simple or complex. For every input-output pair, we have the actual groundtruth tone curve \mathbf{f} that is used to map input to output and we can calculate the approximate simple curve, $\hat{\mathbf{f}}$ using Equation 8. We use this simple (zero or no inflexion point) tone curves to generate an image $\hat{P}(x,y)$ that is an approximation to the ground truth.

Let us consider again the ‘wiggly’ example from the *Introduction* shown in Figure 2. The green dashed line shows the simple tone curve obtained by our method. The simple curve is close to the original wiggly curve but without the wiggles. The images produced by the simple green dashed line and the original wiggly black line have a mean ΔE difference of 1.94.

In Figure 5, we show some qualitative visual results. In the left column, we show 4 input images $I(x,y)$ (before adjustment by the expert). We show the expert ground truth rendition, $P^G(x,y)$, in the second column and the simple tone curve approximation, $\hat{P}(x,y)$, in the third column. Clearly, the corresponding images in columns 2 and 3 are very similar to one another and both are preferred to the input. In the right column we plot, in blue, the ground-truth tone-curve and in dashed red the best simple tone-curve approximation.

Evidently, as we move from top to bottom in the figure the tone curves become more complex (and are less well-approximated by our simpler curve). In fact, these images were explicitly chosen to represent the results. In terms of the mean ΔE between ground truth and simple tone-mapped images, from top to bottom the images shown represent respectively the 0.5, 0.95, 0.99 and 1 quantile ΔE . The corresponding mean ΔE (calculated between the simple image output and the ground truth) are given in Table 2

Table 2: Error statistics of the images shown in Figure 5.

Image	Mean ΔE	Quantile mean ΔE
D4754	0.0116	0.5
B4737	0.135	0.95
C245	0.768	0.99
E4122	4.00	1

Now, we consider how well a simple tone curve can approximate the ground truth for the whole FiveK dataset. For each of the 5 experts, over their 5000 edits, we calculate the mean ΔE between the simple tone curve approximation and the ground truth. We calculate the errors for the quantiles 0.5, 0.9, 0.95, 0.99 and 1 and these are tabulated in Table 3. In all cases, at the 0.99 quantile, the mean ΔE to the ground truth is less than or close to 1 indicating that the ground truth and simple tone-mapped images are visually indistinguishable. Even the maximal errors are modest. There is only one image for the 5th expert (labelled E) where the approximation is 4. Even in this case, the images are visually very similar (and are likely not appear to be visually significantly different [21, 22]). Not reported here, we find no preference difference for images rendered by simple or complex curves (even for the hardest approximation cases).

Table 3: Quantiles of the mean ΔE per expert.

Quantile	A	B	C	D	E
0.5	0.0159	0.0157	0.0169	0.00694	0.00799
0.9	0.178	0.131	0.218	0.0957	0.0744
0.95	0.321	0.202	0.382	0.181	0.141
0.99	1.09	0.457	0.971	0.534	0.638
1	3.84	2.01	3.26	2.04	4.00

The images edited by each expert can be quite different from one another. Thus, per image, we calculate the median ΔE for the 5 adjusted outputs and then calculate the quantile errors for these median adjustments. The results of this experiment are reported in Table 4. The maximum median error here is less than 1. Thus, on average (for 5 experts) a simple tone curve well-approximates (renders an image visually indistinguishable from the ground truth) the actual tone-adjustments made to images.

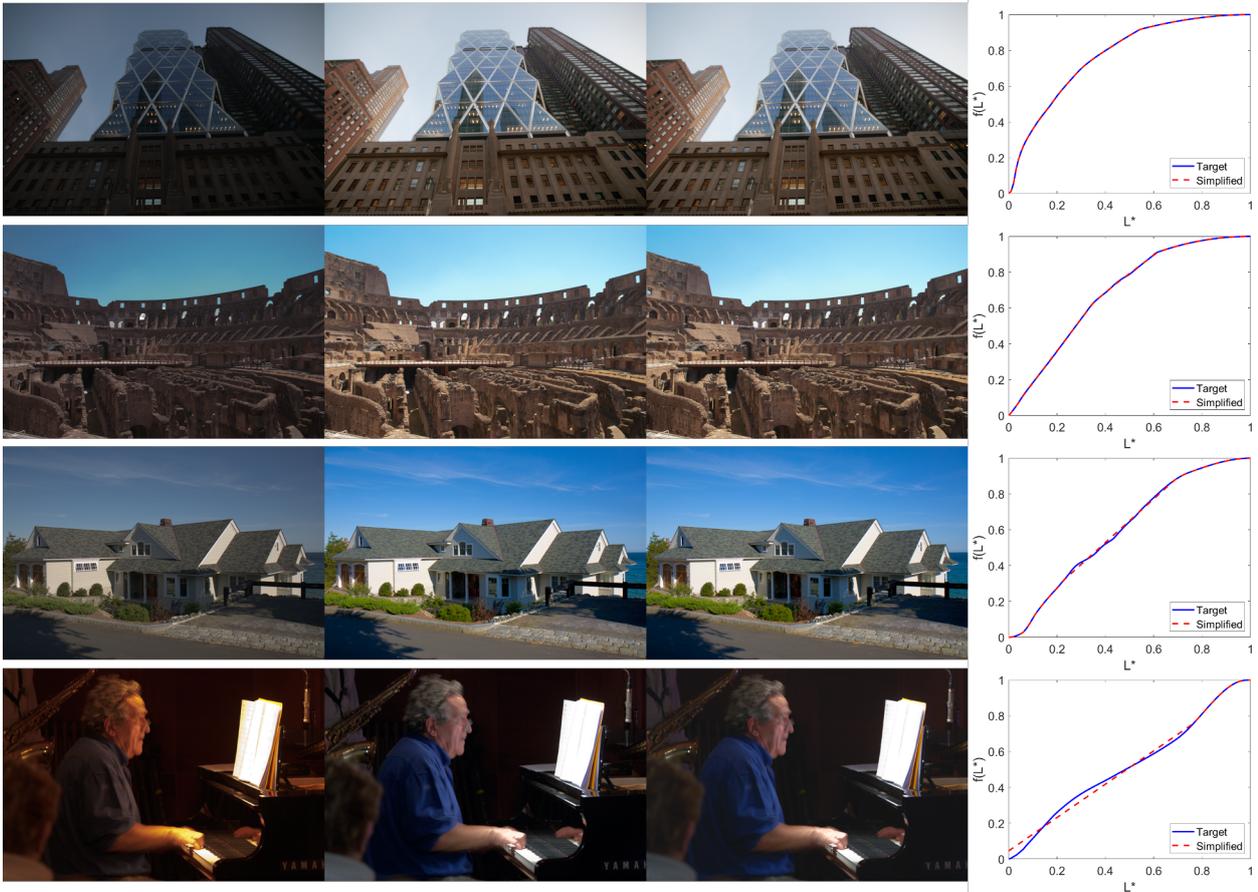


Figure 5. Example visual results. First column shows the input image $I(x,y)$, next is the ground truth image $P^G(x,y)$, then the approximated simply enhanced image $\hat{P}(x,y)$ and right shows the two tone curves that gave these enhancements.

Table 4: Quantiles of the ΔE values for the median adjustment.

Quantile	0.5	0.9	0.95	0.99	1
ΔE	0.0105	0.0675	0.107	0.277	0.958

Discussion

For a very large dataset of enhanced images, the tone adjustments made are either simple tone curves (defined in this paper as having no or a single inflexion point) or can be well-approximated by simple curves. This could be a significant result for tone adjustment in general. First, many photo-editing tools give the user the freedom to make non-simple tone adjustments. Is this wise if professional photographers never make these kinds of adjustments? Perhaps wiggly tone curves should only be available as an ‘advanced’ option. Second, if we assume tone manipulations should be simple then this constraint should be incorporated into automated tone mapping algorithms. Existing prior art like Contrast Limited Histogram Equalisation [23, 24] can produce quite wiggly curves.

Conclusion

Tone mapping is a very powerful technique for image enhancement and is a key part of the toolset in photo editing software as well as being implemented in every photographic camera pipeline. In this paper, we considered whether tone mappings made by users are simple or complex. Simple tone curves were defined to be monotonically increasing curves that either have no or one inflexion point. Conversely, complex curves are those that cannot be simply defined. A computational method is presented to find the best simple tone curve that approximates a complex tone mapping. Experiments conducted on a large set of 25,000 manually tone adjust images found that the tone adjustments made were either simple or that they could be well-approximated by a simple curve adjustment.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council and AgriFoRwArdS CDT [EP/S023917/1]. Graham Finlayson was funded by EPSRC Grant EP/S028730/1 and he also gratefully acknowledges the support of York University, Toronto, Canada (where he was a VISTA distinguished visiting scholar, funded by the Canada First Research Excellence Fund and York University).

References

- [1] M. Grossberg and S. Nayar, "Determining the camera response from images: What is knowable?" *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, pp. 1455–1467, 2003.
- [2] Y. Salih, A. S. Malik, N. Saad *et al.*, "Tone mapping of HDR images: A review," in *International Conference on Intelligent and Advanced Systems*, 2012, pp. 368–373.
- [3] R. C. Gonzalez and R. E. Woods, *Digital Image Processing*, 3rd ed. New York: Pearson, 2008.
- [4] S. B. Kang, A. Kapoor, and D. Lischinski, "Personalization of image enhancement," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2010, pp. 1799–1806.
- [5] C. Li, C.-L. Guo, S. Zhou, Q. Ai, R. Feng, and C. C. Loy, "Flexi-Curve: Flexible piecewise curves estimation for photo retouching," in *IEEE Conference on Computer Vision and Pattern Recognition Workshop*, 2023, pp. 1092–1101.
- [6] S. Bianco, C. Cusano, F. Piccoli, and R. Schettini, "Personalized image enhancement using neural spline color transforms," *IEEE Transactions on Image Processing*, vol. 29, pp. 6223–6236, 2020.
- [7] A. Mustafa, P. Hanji, and R. Mantiuk, "Distilling style from image pairs for global forward and inverse tone mapping," in *ACM SIGGRAPH European Conference on Visual Media Production*, 2022, pp. 1–10.
- [8] V. Bychkovsky, S. Paris, E. Chan, and F. Durand, "Learning photographic global tonal adjustment with a database of input/output image pairs," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2011, pp. 97–104.
- [9] R. Montulet, A. Briassouli, and N. Maastricht, "Deep learning for robust end-to-end tone mapping," in *British Machine Vision Conference*, 2019, pp. 167.1–167.10.
- [10] "Colorimetry — Part 4: CIE 1976 L*a*b* colour space," CIE International Commission on Illumination, Standard, Jun. 2019.
- [11] D. C. C. Wang, A. H. Vagnucci, and C. C. Li, "Digital image enhancement: A survey," *Computer Vision, Graphics, and Image Processing*, vol. 24, pp. 363–381, 1983.
- [12] Y. Qi, Z. Yang, W. Sun, M. Lou, J. Lian, W. Zhao, X. Deng, and Y. Ma, "A comprehensive overview of image enhancement techniques," *Archives of Computational Methods in Engineering*, vol. 29, pp. 583–607, 2021.
- [13] H.-U. Kim, Y. J. Koh, and C.-S. Kim, "PieNet: Personalized image enhancement," in *European Conference on Computer Vision*, 2020, pp. 374–390.
- [14] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, "Unpaired image-to-image translation using cycle-consistent adversarial networks," in *IEEE International Conference on Computer Vision*, 2017, pp. 2223–2232.
- [15] Y. Liu, J. He, X. Chen, Z. Zhang, H. Zhao, C. Dong, and Y. Qiao, "Very lightweight photo retouching network with conditional sequential modulation," *IEEE Transactions on Multimedia*, vol. 24, 2022.
- [16] S. Moran, S. McDonagh, and G. Slabaugh, "CURL: Neural curve layers for global image enhancement," in *IEEE International Conference on Pattern Recognition*, 2021, pp. 9796–9803.
- [17] Y. Hu, H. He, C. Xu, B. Wang, and S. Lin, "Exposure: A white-box photo post-processing framework," *ACM Transactions on Graphics*, vol. 37, pp. 1–17, 2018.
- [18] M. Gharbi, J. Chen, J. T. Barron, S. W. Hasinoff, and F. Durand, "Deep bilateral learning for real-time image enhancement," *ACM Transactions on Graphics*, vol. 36, pp. 1–12, 2017.
- [19] W. S. Cleveland, *Visualizing Data*. Hobart Press, 1999.
- [20] P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization*. Academic Press Inc., 1999.
- [21] G. W. Meyer, "Reproducing and synthesizing colour in computer graphics," *Displays*, vol. 10, pp. 161–170, 1989.
- [22] M. Stokes, M. D. Fairchild, and R. S. Berns, "Precision requirements for digital color reproduction," *ACM Transactions on Graphics*, vol. 11, pp. 406–422, 1992.
- [23] V. Chesnokov, "Image enhancement methods and apparatus therefor," U.S. Patent 7 302 110, 2007.
- [24] 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, pp. 355–368, 1987.

Author Biography

James Bennett is a PhD student at the University of East Anglia in the Colour & Imaging Lab. He is a member of the AgriFoRwArdS CDT and is working in collaboration with Antobot on imaging in agriculture. Professor Graham Finlayson is the Director of the Colour & Imaging Lab. He has published over 200 conference papers, over 75 journal papers and is the inventor of 30+ patents. His interests span perception, colour image processing and physics-based computer vision.