Faster Simple Tone Curves

James Bennett and Graham Finlayson, University of East Anglia, Norwich, UK.

Abstract

Tone curves, applied to images, are an effective way to enhance their visual appeal and to make details more conspicuous. Tone curves are widely used in image editing workflows and camera pipelines. Recently, it was proposed that tone curves should be simple. A simple tone curve has either 0 or 1 inflexion points (the curve shouldn't be wiggly). However, to find a simple approximation that is closest to a non-simple target curve, the prior-art method applies quadratic programming (QP) multiple times. And, such is the high computational cost of running QP that it will not always be practically possible to calculate simple tone curves (e.g. in a camera processing pipeline). In this paper, we show how the simple tone curve optimisation can be reformulated so it runs much more quickly. First, we equivalently enforce simplicity with respect to the derivative of tone curves. In the derivative - or density domain (the derivative of a tone curve is a histogram/density) - a simple curve has 0 or 1 extrema (a single maximum or a single minimum). Second, we show how we can optimise with respect to an extrema constraint using a method called the Pool Adjacent Violators algorithm. PAVA is, both from a computational complexity and run-time perspective, a much more efficient algorithm: we can find simple curves 300 times faster. For the MIT-Adobe FiveK dataset with 25,000 tone curves, the QP and PAVA optimisations produce visually similar results. Simple tone curves are closer to camera pipeline ready.

Introduction

Tone mapping is an effective way to enhance the visual appeal of images and make details more conspicuous, making them widely used in image editing workflows and camera pipelines. A tone curve – which is the function describing a mapping from input to output brightnesses – may have a fixed shape, be custom-defined per image, or be computed algorithmically based on the image's statistics or content. Camera pipelines have historically employed fixed tone curves, leading to a characteristic rendering style. However, in contemporary pipelines, tone curves are computed per-image. Compute time is important: tone mapping algorithms are generally very fast, such as with Contrast Limited Adaptive Histogram Equalisation (CLAHE), which derives local tone curves from local image brightness histograms [1].

In general, the shape of a tone curve is only weakly constrained. It is possible that the tone curve has a "wiggly" shape (CLAHE often generates wiggly curves). Previous work [2] has argued that tone curves should not be *wiggly*, but instead should be *simple*, defined to have at most one inflexion point. Analysis of the 25,000 expert tone adjustments in the MIT-Adobe FiveK dataset found that expert tone curves were inherently simple or could be closely approximated by simple counterparts, with no visually important difference in the resulting images.

In the top left of Figure 1, we show an input image (without tone mapping) from the MIT-Adobe FiveK dataset. In the top right, we plot two tone curves. The solid line is the wiggly tone curve adjustment that was made by a professional photographer, and the dashed line is its simple approximation. Respectively, bottom left and right show the tone adjusted output images

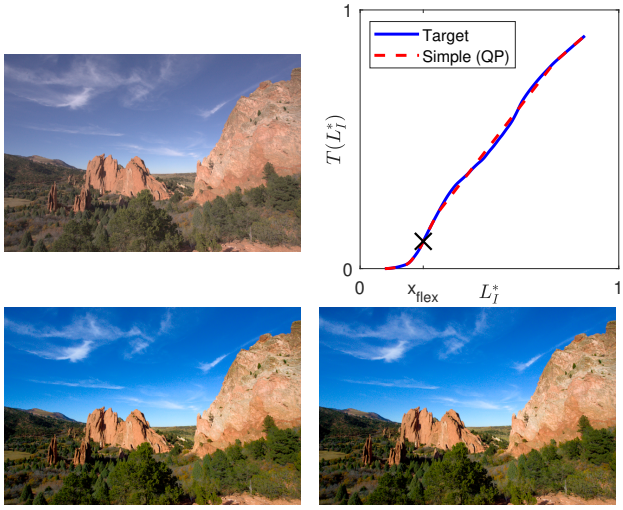


Figure 1: Top left, input image. Bottom left and right, image tone mapped with respectively the blue (wiggly) and dashed red (simple) curves, shown top right.

(photographer versus simple approximation). Evidently, tone-mapping has made a much more appealing image. Moreover, despite the two tone-curves being significantly different shape the output reproductions are similar.

If we denote brightness as $x \in [0,1]$ then a tone curve – an input to output brightness mapping – can be written as $y = f(x)$ where f has the property that it is an increasing function of brightness: $f(x_2) \geq f(x_1)$ when $x_2 > x_1$. Additionally, a simple tone curve is defined to be a tone curve that has 0 or 1 inflexion points. There are 4 cases to consider for simple curves: derivatives that are either strictly increasing or strictly decreasing (these curves have 0 inflexion points); and where the derivative is increasing-then-decreasing or decreasing-then-increasing (these curves have one inflexion point). For the 0 inflexion point case, the slope (or derivative) of the curve increases (or decreases) monotonically. Examples of 0 inflexion tone curves include $f(x) = x^2$ (increasing) and $g(x) = x^{0.5}$ (decreasing), which respectively have slopes/derivatives $f'(x) = 2x$ and $g'(x) = \frac{0.5}{x^{0.5}}$. Respectively, $f'(x_2) \geq f'(x_1)$ and $g'(x_2) \leq g'(x_1)$ when $x_2 > x_1$. In the 1-inflexion point case, the slope of a tone curve monotonically decreases to an inflexion brightness x_{flex} and then monotonically decreases (or vice-versa). An S-shaped curve is an example of a 1-inflexion point tone-curve. Returning to Figure 1, we see the slope of the simple tone curve increases and then decreases where the point of inflexion is at x_{flex} (see large X plotted on the graph top right).

In the prior art simple tone curves algorithm [2], one is given a non-simple tone curve and the closest – in a least-squares sense – simple approximation is found, with simplicity achieved through enforcing the linear equality and inequality constraints that account for the 4 cases of simple curves. This combination of objective function and constraints resists fast single pass solvers, instead formulating naturally as a quadratic program (QP). Although there are efficient implementations of QP, solving for a

simple tone curve is computationally laborious (in part, because QP has to be applied multiple times). Indeed, such is the required compute time that it is doubtful whether the method could be implemented in modern camera ISPs.

In this paper, we reformulate the problem so that we can use these single pass solvers, allowing simple tone curves to be found more quickly. The key observation that makes this possible is that inflexion points in tone curves are peaks in a density representation. Returning to the example in Figure 1, we plot the density of the blue photographer curve as a blue line in Figure 2. The wiggly tone curve has multiple extrema (maxima and minima). In fact, seeing the inflexion points as peaks in density space emphasises the number of inflexion points in the tone curve.

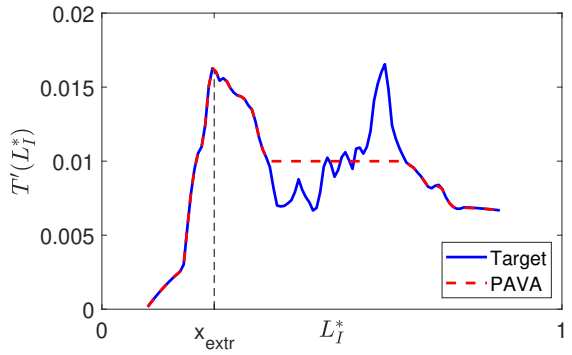


Figure 2: Derivative of the tone curves from Figure 1. The derivative of the complex and simple curves respectively have many (blue) or a single (red) peak.

Also shown in Figure 2, by the dashed red line, is the solution that our method obtains. It has a single peak corresponding with the leftmost extrema of the target density, and the curve is monotonic either side of the peak. Given the peak is a maxima, the solution is increasing up to it and decreasing thereafter. When converted back to tone curve space, this is a tone curve whose gradient increases up to an inflexion point and decreases beyond; an S-shaped curve (simple CASE 3).

The interesting question is how do we find the *best* single peak red curve given the blue curve. We will proceed in a similar manner as for the tone curve approximation. To find the best single peak approximation, we vary the peak location x_{extr} and solve (relative to the target) for the best increasing approximation up to the peak and then the best decreasing approximation after the peak. The solution with the best peak location is used.

Solving for the best increasing approximation is an isotonic regression problem which has a particularly efficient numerical algorithm that is suited to this task. Drawn from the numerical methods literature, the Pooled Adjacency Violators Algorithm (PAVA) [3] is an extremely rapid method for finding increasing (or decreasing) approximations. In fact for a given function $f(x)$ (say) $x \in [0,1]$ PAVA finds the best approximate function, $\hat{f}(x)$ that minimises $||f(x)-\hat{f}(x)||$ such that if $x_2 > x_1$ then $\hat{f}(x_2) \geq \hat{f}(x_1)$. For the case (in density representation) of finding a unimodal curve that increases to a peak brightness and then decreases (or the converse), it is a simple matter to apply PAVA twice.

In our experiments, we measure the speed improvement from using PAVA, finding it implements a simple curve calculation more than 300 times faster. Importantly, we also find these tone curves (the integral of the closest density found) produce good approximations to the photographer's own adjustment.

Next we review the prior-art QP method for calculating simple tone curves the Adobe-MIT FiveK image dataset used in our experiments. The new PAVA-based algorithm is presented in the method section. Following, we report on experiments comparing algorithm speed and tone-mapped outputs (QP versus PAVA tone mapped curves). The paper finishes with a short conclusion.

Background

Dataset and notation

The FiveK dataset [4] comprises 5,000 input images and the corresponding tonally adjusted outputs made by 5 photographers [2]. See, for example, [5, 6, 7, 8], for research that makes use of the FiveK dataset. For each tone curve adjustment in the dataset, let $\underline{I}$ denote an input image and $\underline{P}$ denote the expert modified output, expressed as a global tone curve adjustment. At each (x,y) location the image pixel colour is represented by a L^*, a^* and b^* CIELAB [9] triple.

$$\begin{aligned}\underline{I}(x,y) &= [L_I^* \; a_I^* \; b_I^*]^\top \\ \underline{P}(x,y) &= [L_P^* \; a_P^* \; b_P^*]^\top\end{aligned} \tag{1}$$

A tone curve maps input values to output values using a function that is continuous on the input domain. Where the same function is applied to every pixel of the image, the mapping is *global*. Mathematically, the relation between the brightnesses of the input, $\underline{I}$, and expert adjusted image, $\underline{P}$, is described by the tone mapping function T, according to equation,

$$L_P^* = T(L_I^*). \tag{2}$$

Throughout this work we are interested in finding a tone map $\hat{T}$ that approximates T, subject to additional constraints. When applied to the same input brightnesses, $\hat{T}(L_I^*) \approx L_P^*$, which yields an approximate adjusted image $\underline{\hat{P}}$, defined as

$$\underline{\hat{P}}(x,y) = [\hat{T}(L_I^*) \; a_P^* \; b_P^*]^\top, \tag{3}$$

noting the expert's a^* and b^* are retained.

All computations will be carried out in terms of a finite number of brightness levels i.e. in the discrete domain. Let us define an equivalence between a discrete tone-curve vector $\boldsymbol{t}$ and the corresponding continuous tone mapping function:

$$\boldsymbol{t} \equiv T \tag{4}$$

Given m uniformly spaced input tonal values in the domain $[0,1]$, written $\boldsymbol{b} = [0, 1/(m-1), 2/(m-1), ..., 1]^\top$. Evaluating the tone curve T from Equation (2) at these input tonal values, results in the m-vector $\boldsymbol{t} = [T(0), T(1/(m-1)), T(2/(m-1)), ..., T(1)]^\top$. We use t_i to denote the i-th component of $\boldsymbol{t}$. With an appropriate interpolation scheme, the vector pair $\boldsymbol{b}$, $\boldsymbol{t}$ defines the tone curve.

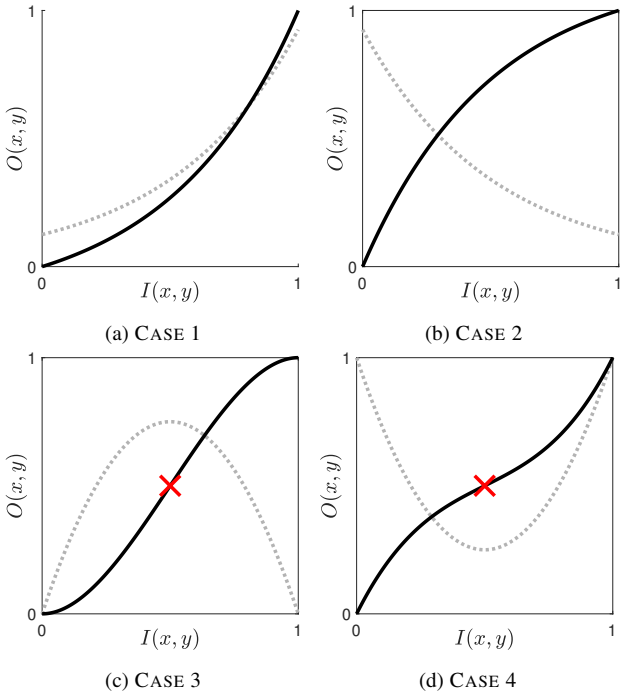
Simple Tone Curves

In the prior art, we seek the *simple* curve $\hat{\boldsymbol{t}}$ that best approximates a target curve $\boldsymbol{t}$.

$$\arg\min_{\hat{\boldsymbol{t}}} \|\hat{\boldsymbol{t}} - \boldsymbol{t}\|. \tag{5}$$

The curve $\hat{\boldsymbol{t}}$, to be simple, must adhere to a series of constraints. Firstly, the approximate curve must have the same start and end points as the target curve, ensuring it maps to the same dynamic range. Second, the curve must be monotonically increasing. The third constraint allows at most one inflexion point. There are four cases to consider, illustrated in Figure 3. CASE 1 and CASE 2 have zero inflexion points, with the former having

an increasing gradient over the whole curve and the latter a decreasing gradient. CASE 3 and CASE 4 each have one inflexion point, whose location can vary over the entire domain. A CASE 3 curve's gradient increases up to the inflexion point and decreases thereafter; CASE 4 is the converse.



(a) CASE 1 (b) CASE 2

(c) CASE 3 (d) CASE 4

Figure 3: Depiction of the four simple cases. Solid line shows tone curves, red cross marks the inflexion point. Dashed lines are derivatives of tone curves.

Given a non simple tone curve, in [2] it was shown how the closest simple approximation could be found. The problem was formulated as an *enumerative* least-squares optimisation, subject to linear equality and inequality constraints. By enumerative we mean that the best curve according to all 4 simplicity cases (and all possible inflexion point locations) are found and the best approximation overall is drawn from this set. Each individual optimisation is formulated as a QP [10, 11]. QP problems have the property that they return a unique global minima.

A leading algorithm for solving a QP is the interior-point method described in [12] which has a worst-case complexity of $O(n^3)$ [13, 14]. Because the simple curve optimisation requires a $O(n)$ search over the inflexion points, the overall complexity is $O(n^4)$. In practice, QP problems often converge to solutions more rapidly than this worst-case complexity would suggest.

The key insight that will make it possible to use PAVA (which we review in the methods section) – a fast linear time solver – is that, rather than optimising directly in tone curve space, we instead optimise in the space of the tone curve's derivative – a density representation of the tone curve. Whereas in tone curve space, the problem sat naturally as a QP and did not decompose cleanly for single pass solvers, in density space, it does. Derivative shape constraints of the simple tone curve become monotonicity constraints on the derivative. In Figure 3, the derivatives are shown by the grey dashed line. In density representation, a CASE 1 curve is a monotone increasing function. This is precisely the isotonic regression problem, which finds the least-squares optimal non-decreasing function and is solved in linear time by PAVA [3]. Previously, PAVA has been applied to tone mapping for the purpose of enforcing monotonicity [15]; here, we extend its application to solving for simple tone curves. CASE 2 is also entirely monotone, but decreasing, whilst CASE 3 and CASE 4 each have a single extrema which corresponds to the inflexion point of the tone curve, marked by the red cross. Importantly, these cases are monotone on either side of the extremum, meaning isotonic regression can be deployed.

Method

The prior art finds a simple tone curve by solving a QP in tone curve space. In this work, we instead optimise in the space of the tone curve's derivative – and the equivalence that makes this possible is elegant. The derivative of the tone curve is a density representation, which can be thought of as a histogram. Mathematically, we obtain $\boldsymbol{h}$ by computing the finite differences (analogous to the derivative) of the discrete tone curve vector $\boldsymbol{t}$,

$$h_i = t_i - t_{i-1}, \qquad \text{for } i = 2,3,\ldots,m, \tag{6}$$

which yields an n-vector, where $n = m-1$. To visualise this concept, see Figure 4. The tone curve spans the range t_1 to t_m and has the inflexion points marked. The density of this (obtained with Equation 6), is shown top right. Observe that increasing derivatives in tone curve space become increasing histogram counts in density space. Decreasing derivatives become decreasing histogram counts. We see that inflexion points in tone curve space are extrema (maxima and minima) in density space. Finally, we note that the area under the density distribution is equal to the range of the tone curve, put mathematically $\sum_i h_i = t_m - t_1$.

Of course, we require tone curves for mapping the brightnesses in images, and this is the re-integration of the density space solution. Suppose we have obtained the solution $\hat{\boldsymbol{h}}$ that best approximates the target $\boldsymbol{h}$ (which is the derivative of $\boldsymbol{t}$), the solution tone curve $\hat{\boldsymbol{t}}$ is recovered by integration:

$$\hat{t}_i = \begin{cases} t_1, & \textit{when } i = 1 \\ t_1 + \sum_{j=1}^{i-1} \hat{h}_j, & \textit{otherwise} \end{cases} \tag{7}$$

This is illustrated in the bottom of Figure 4. Left, shows the best single extrema approximation to the density, $\hat{\boldsymbol{h}}$; it is monotone increasing up to the peak and decreasing beyond. Right, plots the simple tone curve (the integral of the density, bottom left).

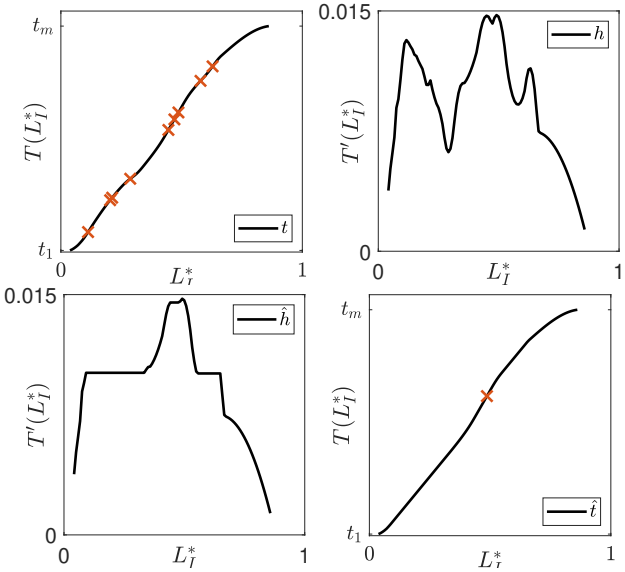


Figure 4: TL: wiggly tone curve whose inflexion points are marked 'x', TR: density representation, BL: simple density solution, BR: integration of BL gives simple tone curve.

Let us now, formally, define the **general** isotonic regression problem. In isotonic regression given $\boldsymbol{h} \in \mathbb{R}^n$, we seek

$$\min_{\hat{\boldsymbol{h}}} \|\hat{\boldsymbol{h}} - \boldsymbol{h}\|, \quad \text{s.t.} \quad \hat{h}_i \leq \hat{h}_{i+1} \quad \forall i. \tag{8}$$

PAVA, introduced by [3] and described in [16], solves this problem exactly and in linear time. Operationally, PAVA scans the observations from left to right, and whenever value lower than the one to the left – a violation of the non-decreasing constraint – pool the offending elements together and replace them with their mean. If the resulting pool now violates its left neighbour, merge again. Continue until no violations remain. The plot of $\hat{\boldsymbol{h}}$ in the bottom left of Figure 4 helps us understand this principle. The points that violated the monotonicity constraint have been pooled to a constant value (seen as the flat regions of the curve) which is the mean value of the violating points. Intuitively, this procedure enforces monotonicity (and, it turns out it also returns the optimal increasing curve in a least-squares sense).

Let us now use PAVA for the 4 cases of simple tone curves as they are represented in the density domain. CASE 1 requires the density histogram to be non-decreasing. PAVA is applied directly to $\boldsymbol{h}$ and the tone curve is recovered by re-integration. We denote the function that returns the closest monotonically increasing approximation to a histogram $\boldsymbol{h}$ as $P_F(\boldsymbol{h})$. For CASE 1 we write,

$$\hat{\boldsymbol{h}} = P_F(\boldsymbol{h}). \tag{9}$$

For CASE 2, a non-increasing histogram is required. We additionally define $P_R(\boldsymbol{h})$ which finds the least-squares optimal decreasing creasing function.

$$\hat{\boldsymbol{h}} = P_R(\boldsymbol{h}), \tag{10}$$

which is essentially the same as PAVA but now we index the histogram from right to left. A decreasing function left-to-right is increasing right-to-left.

CASES 3 and 4 each have one inflexion point, which corresponds to an extremum in derivative space – a peak for CASE 3 and a trough for CASE 4. Importantly, on either side of the extremum the histogram is monotone, so isotonic regression can be deployed on each side independently. The CASE 3 solution, for a given inflexion point position p, is obtained with,

$$\hat{\boldsymbol{h}} = C_3(\boldsymbol{h}, p) = \begin{bmatrix} P_F(\boldsymbol{h}_{[1..p]}) \\ P_R(\boldsymbol{h}_{[p+1..n]}) \end{bmatrix}, \tag{11}$$

where the bracketed subscript denotes a slice of the matrix, with $[a..b]$ denoting the integers from a to b inclusive. For CASE 4,

$$\hat{\boldsymbol{h}} = C_4(\boldsymbol{h}, p) = \begin{bmatrix} P_R(\boldsymbol{h}_{[1..p]}) \\ P_F(\boldsymbol{h}_{[p+1..n]}) \end{bmatrix}. \tag{12}$$

Now that we can generate candidate solutions for a given case and inflexion point, we address how the *best* candidate is selected. First, let us generate the set of **all** candidate solutions

$$\mathcal{H}_3 = \{C_3(\boldsymbol{h},2), C_3(\boldsymbol{h},3), ..., C_3(\boldsymbol{h},n-1)\} \tag{13}$$

$$\mathcal{H}_4 = \{C_4(\boldsymbol{h},2), C_4(\boldsymbol{h},3), ..., C_4(\boldsymbol{h},n-1)\} \tag{14}$$

$$\mathcal{H} = \mathcal{H}_3 \cup \mathcal{H}_4 \cup \{P_F(\boldsymbol{h})\} \cup \{P_R(\boldsymbol{h})\} \tag{15}$$

The set $\mathcal{H}$ contains the closest density (histogram) to a target that adheres to the 4 case constraints (including a varying peak position for cases 3 and 4). We now find the *best* member of this set according to two criteria. First, we find the closest approximation $\hat{\boldsymbol{h}}$ to $\boldsymbol{h}$, the closest histogram returned by our PAVA method, denoted **PAVA-H**,

$$\arg\min_{\hat{\boldsymbol{h}} \in \mathcal{H}} ||\hat{\boldsymbol{h}} - \boldsymbol{h}||. \tag{16}$$

Alternately, in **PAVA-C** we find the closest PAVA solution in tone curve space. First, let us define the set $\mathcal{C}$ where each member is the recovered tone curve of each histogram in $\mathcal{H}$.

$$\mathcal{C} = \{S(\hat{\boldsymbol{h}}) : \hat{\boldsymbol{h}} \in \mathcal{H}\}, \tag{17}$$

where $S(\hat{\boldsymbol{h}})$ denotes the computation in Equation 7, returning the solution tone curve $\hat{\boldsymbol{t}}$. For PAVA-C, the best tone curve approximation $\hat{\boldsymbol{t}}$ is returned by,

$$\arg\min_{\hat{\boldsymbol{t}} \in \mathcal{C}} ||\hat{\boldsymbol{t}} - \boldsymbol{t}||. \tag{18}$$

PAVA itself operates in linear time, however the inflexion point search evaluates $O(n)$ candidate positions. Thus, the overall algorithm is thus $O(n^2)$ (compared with $O(n^4)$ for the prior art QP method).

Results and Discussion

For the 25,000 tone curves in the FiveK dataset, the best simple tone curve approximation is found using three algorithms; QP uses the prior-art method, and PAVA-H and PAVA-C use the method presented in this paper with the density and tone curve space cost functions respectively, see Equations 16 and 18. Our experiments are conducted in MATLAB, using the function `quadprog` for the core of the QP with an interior point method and using `lsqisotonic` from the Statistics toolbox as the PAVA implementation. Both algorithms were timed from being passed a tone curve, to returning the solution tone curve (over all candidates), any setup or image rendering and evaluation is not considered. The tests were run on the same computer under the same conditions for fair comparison. These experiments assume brightnesses are quantised to 100 levels (a tone curve is defined by 100 values). There will be a variance in the compute time per tone curve. The PAVA optimal solution is found faster in some cases than others depending on the number of violators and merges that take place. QP is search-based and how quickly a minimum is found also depends on its input.

Compute Time

Table 1 reports the median, interquartile range (25 and 75 percentile), and maximum solve time for each method. The median solve time for QP is 3.54 seconds, compared to 0.011 seconds for both PAVA variants – a speed-up of over 300× at the median. Importantly, this advantage is not only confined to typical cases: the speed-up exceeds 100× across the entire distribution, including the worst case. PAVA-H and PAVA-C are essentially identical in solve time, confirming that evaluating candidates in tone curve space rather than density representation adds negligible cost.

Table 1: Compute time in seconds

Solver	Median	IQR	Max
QP	3.54	3.45 - 3.66	8.49
PAVA-H	0.0112	0.0086 - 0.0136	0.0661
PAVA-C	0.0114	0.0089 - 0.0138	0.0781



Figure 5: Example results for (top-bottom) B1850, C4537, A21. Left to right, images are, input, expert and the PAVA-H and PAVA-C simple approximations

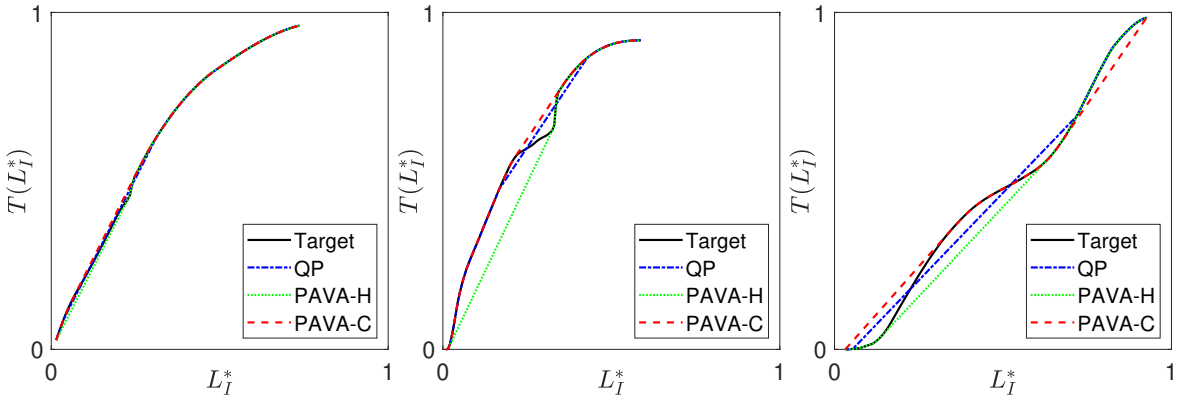


Figure 6: Tone curves for example results (left-right) B1850, C4537, A21 for examples in Figure 5

Image Quality

The MIT-Adobe FiveK dataset comprises 5000 input images and 25000 output images (5 tone-mapped outputs generated by 5 professional photographers). According to [17, 18], a mean ΔE below 3 indicates that an image pair is almost visually indistinguishable for complex scenes of the type found in the FiveK dataset. Table 2 reports selected quantiles. QP, as the globally optimal simple curve approximation, achieves the lowest ΔE across the entire distribution, as expected. Being the globally optimal curve approximation does not *guarantee* a small ΔE, but there is a strong positive correlation between residual error and ΔE; these results show QP is the best in terms of ΔE. The PAVA methods are close throughout: PAVA-H and PAVA-C track one another well at low ΔE values, with PAVA-C pulling ahead of PAVA-H in the tail. Impressively, 99.4% of QP images fall below $\Delta E = 1$. This compares to 98.8% for PAVA-C and 98.7% for PAVA-H,. Below the less stringent $\Delta E = 3$ threshold, all three methods perform similarly well, with QP, PAVA-C, and PAVA-H achieving 99.99%, 99.94%, and 99.84% respectively. For the overwhelming majority of images, the proposed method produces visually indistinguishable results from the QP baseline at a fraction of the computational cost.

Table 2: Quantiles for mean ΔE

	0.5	0.9	0.95	0.99	1
QP	0.0115	0.141	0.262	0.783	3.22
PAVA-H	0.0136	0.196	0.369	1.19	10.5
PAVA-C	0.0137	0.194	0.362	1.10	4.34

Figure 5 shows example results for three images from the FiveK dataset: B1850, C4537, and A21. For each, the input $\underline{I}$, expert $\underline{P}$, and the simple renditions $\hat{\underline{P}}$ for PAVA-H, and PAVA-C are shown, with corresponding tone curves in Figure 6.

B1850 is representative of the 98th percentile. The tone curve is well approximated by all three methods, with ΔE values of 1.23, 0.82, and 0.53 for PAVA-H, PAVA-C, and QP respectively. The rendered images are visually similar throughout, and the tone curves in Figure 6 are nearly indistinguishable.

C4537 is a more challenging case. QP finds the globally optimal simple approximation with $\Delta E = 2.09$, and PAVA-C is close, achieving $\Delta E = 2.49$. The renditions are similar and both broadly capture the photographer's intent, though with observable differences from the expert. PAVA-H, however, performs considerably worse at $\Delta E = 8.57$, the rendered image bearing visibly less resemblance to the target. In Figure 7 we show, in

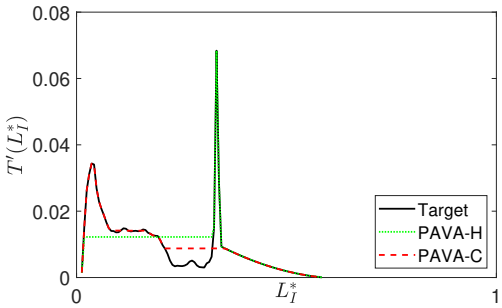


Figure 7: Tone curves for C4537 in density space.

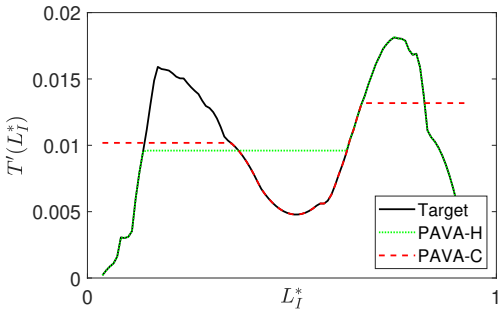


Figure 8: Tone curves for A21 in density space.

the density (derivative of tone curve) representation respectively the target (black), PAVA-H (green) and PAVA-C (red) approximations. It is evident that the PAVA-C method compared to the PAVA-H method has effectively fit to different density peaks.

A21 presents a different mode of failure, and one that affects all three methods. The target tone curve exhibits broad sweeping curvature with multiple inflexion points – a complex curve where a simple approximation struggles to capture the experts enhancement intent. This manifests in derivative space as two broad maxima with three extrema, as seen in Figure 8. Different cases are selected for the two PAVA methods, with PAVA-H placing the inflexion point at a maximum whilst PAVA-C places it at a minimum, resulting in ΔE values of 1.71 and 4.34 respectively. Interestingly, PAVA-C performs worse than PAVA-H here. Even QP, with $\Delta E = 2.68$ struggles to capture the photographer's intent.

Smoothness

A tone curve that has a large gradient discontinuity can introduce Mach band artefacts (false perceptual contours). For the QP solution this problem was mitigated by solving for a tone curve where, by construction, the solved-for tone curve did not have sharp gradient discontinuities. In the context of the current work, Friedman [19] has shown that smoothing the PAVA solution with a Gaussian kernel yields the optimal smooth isotonic regression. Due to the commutativity of convolution and differentiation, this smoothing can be applied equivalently in density or tone curve space, offering a natural route to C^∞ smooth simple tone curves. This result connects directly to the even simpler tone curve framework of [20], and extending the present method to incorporate Gaussian smoothing will be a natural direction for future work.

Conclusion

Previously, it was proposed that tone curves should be simple. Informally, a tone curve is simple if it is not too wiggly. This idea is captured mathematically by enforcing that a tone curve should have 0 or 1 inflexion point. In the prior art a method was developed that finds the best simple approximation to a given tone curve. The method involved multiple applications of quadratic programming (on order of the number of quantisation levels in a histogram representation of the image). While QP is a fast algorithm the requirement to apply it multiple times resulted in a method which could take several seconds to find a single tone curve.

The algorithm developed in this paper tackles this speed problem head on. First we reformulate the simple curve approximation problem so it happens in the density domain. If we differentiate a tone curve we effectively create a density (or histogram) distribution. In this representation a simple tone curve is now an increasing or decreasing density distribution or one where there is a single extrema (a single peak or trough). With respect to densities we show a particularly fast numerical method called the pooled adjacency violators algorithm (PAVA) can be used instead of QP, resulting in a 300x speed-up. A simple curve can now be found in less than 1/100th of a second. In experiments we show that our new PAVA-based simple curve algorithm generates tone-curve approximations almost the same as the original QP version.

Acknowledgments

This work was supported by the Engineering and Physical Sciences Research Council, EP/S023917/1 and EP/S028730/1.

References

[1] 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.

[2] J. Bennett and G. Finlayson, "Simple tone curves: theory and applications," *The Visual Computer*, vol. 42, no. 8, p. 302, May 2026.

[3] R. E. Barlow, D. J. Bartholomew, J. M. Bremner, and H. D. Brunk, *Statistical Inference under Order Restrictions*. John Wiley & Sons, 1972.

[4] 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.

[5] C. Li, C. L. Guo, S. Zhou, Q. Ai, R. Feng, and C. C. Loy, "FlexiCurve: 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] 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.

[9] "Colorimetry — Part 4: CIE 1976 L*a*b* colour space," CIE International Commission on Illumination, Standard, Jun. 2019.

[10] P. E. Gill, W. Murray, and M. H. Wright, *Practical Optimization*. Academic Press Inc., 1999.

[11] C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems*. Prentice-Hall, 1974.

[12] R. D. C. Monteiro and I. Adler, "Interior path following primal-dual algorithms. part II: Convex quadratic programming," *Mathematical Programming*, vol. 44, no. 1, pp. 43–66, May 1989.

[13] G. H. Golub and C. F. V. Loan, *Matrix Computations*, 3rd ed. The

John Hopkins University Press, 1996.

[14] J. Nocedal and S. J. Wright, *Numerical Optimization*, 2nd ed., ser. Springer Series in Operations Research and Financial Engineering. New York: Springer, 2006.

[15] J. Singnoo and G. D. Finlayson, "Optimal global approximation to spatially varying tone mapping operators," *Conference on Colour in Graphics, Imaging, and Vision*, vol. 6, no. 1, pp. 182–188, 2012.

[16] W. Härdle, *Applied nonparametric regression*. Cambridge university press, 1990.

[17] 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.

[18] M. K. E. Mahy, L. V. Eycken, and A. Oosterlinck, "Evaluation of uniform color spaces developed after the adoption of CIELAB and CIELUV," *Color Research & Application*, vol. 19, no. 2, 1994.

[19] J. Friedman and R. Tibshirani, "The Monotone Smoothing of Scatterplots," *Technometrics*, vol. 26, no. 3, pp. 243–250, 1984.

[20] J. Bennett and G. Finlayson, "Even simpler tone curves," in *London Imaging Meeting*, 2024, pp. 124–128.

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.