Optimal Global Approximation to Spatially Varying Tone Mapping Operators

Jakkarin Singnoo, Graham D. Finlayson, School of Computing Sciences, University of East Anglia, Norwich, UK

Abstract

Compared with spatially-varying tone-mapping operators, global tone maps have the advantage that the input is mapped to an output image without introducing spatial artifacts common to spatially-varying tone-mapping operators (e.g. halos and intensity inversions). However some local detail can be compressed (visually lost). In this work, we propose a global tone-mapping operator that optimally, in a sum of least-squares sense, approximates spatially-varying tone-mapping operators.

Our method is based on a modification of the simple but elegant constrained optimization technique called Pool-Adjacent-Violators-Algorithm (PAVA). In a second step, we show how any lost local detail can be brought back through copying, in an edge sensitive manner, detail from the original input (an approach already developed in the literature).

Our new global tone-curve approach has a specific advantage: we show it suffices to learn the tone-curve by processing a small thumbnail and then produce the final output by applying the tone-curve to the full resolution input. Not only does processing on thumbnails deliver excellent results we can, using this approach, significantly increase the speed of tone-mapping operators.

To evaluate our method we carried out a paired comparison psychophysical experiment. Preference scores resulting from the experiment show that in general the perceived quality of our proposed operator is similar (equally preferred) to a range of spatially-varying tone-mapping operators.

Introduction

HDR images are formed by blending together multiple exposures [3] and have the advantage that the true range of intensities in the original scene are recorded (which is not the case for typical camera images where highlights and shadows are often clipped). However, typical display devices (monitor or even worse printer) cannot output the large acquired dynamic range. In order to visualize HDR contents on these devices, an operator that performs dynamic range compression is needed. That is, an operation called "Tone-Mapping Operator" (TMO) is used. Various tone-mapping operators have been developed in recent years [19]. In general, there are two types of TMO: one is the *global* (spatially-uniform) and the other is *local* (spatially-varying).

Global TMOs (G-TMOs) are non-linear *surjective* functions that map an input HDR image to the output LDR (low dynamic range) image for display. The function is typically parameterised by simple image statistics drawn from the input (e.g. mean and quantiles) and the dynamic range of the output display device. Once the G-TMO function is defined, every pixel of the image is mapped globally (independent from surrounding pixels in the image). Global TMOs often work well and are also simple and fast.The solid blue line in Figure 2. is an example of the G-TMO mapping input log-scale values to an output linear display range.

However, by their very nature G-TMOs compress or expand the input signal. If the slope of the G-TMO function is less than 1 then detail is compressed in the output images. Such compression often happens in highlight areas of an image and if it does the output image appears flat. Often G-TMOs produce images where contrast is lacking.

Spatially-varying TMOs (SV-TMOs) on the other hand, take into account the spatial context when they adjust pixel intensity values. In other word, the parameters of the non-linear function change at each pixel according to the local features extracted from the neighboring pixels. Often this leads to improved local contrast and images that are preferred by observers: simply, they look more appealing. On the downside, SV-TMO algorithms are far more complicated than global tone-curves. Moreover, spatially-varying operators such as Ashikhmin's [1] and Retinex also often introduce strong halo artifacts in the area around high contrast edges.

Our work in this paper is motivated by the idea of finding the optimal global function that best approximates the spatially-varying operator. We seek to build a processing workflow that retains the simplicity of a G-TMO but keeps the preferred local detail that is characteristic of SV-TMOs.

In [16], the tone-curves map input to output (for non optimal) approximations to SV-TMOs are analysed statistically. To a first approximation all curves studied could be modelled by a 5-parameter sigmoid. based on this statistical analysis, they force all G-TMOs to be modelled by this sigmoidal function. In [8], a G-TMO is applied with some additional pre-and post-processing that leads to local detail being preserved (a strategy we also use in this paper). The G-TMO also has a constrained shape (a simple brightness and gamma adjustment).

In our work we seek to find an *optimal* G-TMO. Specifically, given an input image and the output of an arbitrary SV-TMO, we find the G-TMO that approximates the output in a least-squares sense. We make no prior assumption about the shape of the curve. We have found that our agnosticism is important to achieving the best image outputs. indeed we found that our G-TMOs are not, in general well modelled by a sigmoidal curve or brightness+gamma adjustment.

We realised from the outset that mapping an input to an output image using a G-TMO could be cast as a quadratic programming (QP). Yet, QP is a rather general and computationally expensive procedure. Thus, instead, we propose using the simpler (and for this problem equivalent) Pool-Adjacent-Violaters-algorithm (PAVA). PAVA optimally finds the monotonically increasing function that maps scale inputs to scale outputs. We also (very slightly) modify PAVA so the procedure runs very rapidly and for the problem at hand (tone-mapping of images) also produces pleasing results (the default PAVA can result in a loss of detail).

Figure 7 shows the outputs of SV-TMOs (left panels) and our G-TMO approximation (center column). It is clear a G-TMO can produce a very good approximation to a spatially varying algorithm. The images in the rightmost column are the outputs from the G-TMO after applying a detail recovery procedure.

To quantify how well a G-TMO fits the output of a SV-TMO we could look at the % RMSE difference. Typically this is a very small percentage (less than 5%). But, actually, the outputs from both operators are going to be viewed by a human observer. Thus we ran a preference experiment to see whether observers judged our G-TMO outputs to be equally acceptable as the SV-TMO images. Our preference experiments (which evaluate 3 leading SV-TMOs) found that our G-TMO produced equally preferred results.

This is an encouraging result. Our G-TMO vehicle produces its outputs at a fraction of the computational cost and provably does not suffer from spatial artifacts.

Related works are considered in the next Section. Then, our Optimal Global Tone-Mapping Operator is proposed. In order to evaluate image preference, a psychophysical experiment is conducted in Experiment. In the last section, we conclude the work.

Background

The goal of SV-TMO is to reduce the dynamic range of the HDR images while preserving subtle details i.e. preserving local contrast. However, it is well-known that SV-TMOs often result in halo artifacts: intensity inversions near high contrast edges. There is a trade-off between the dynamic range compression and the image rendition: A strong, spatially-varying, dynamic compression leads to halo artifacts while a weaker has no halos (or they are much reduced) but the detail in the output can be more muted (compared with the original). This tradeoff is discussed in [24, 11, 18, 1, 5, 8] and in general reviews in [4, 19].

We would like to use a G-TMO but to 'bring back' any local contrast detail that might be lost. One way to do this is to form an input smoothed image using an edge-sensitive Bilateral Filter [22].

Bilateral Filtering

Bilateral filtering adds photometric distance enforcement to traditional filtering proximate pixels are averaged together in proportion to their similarity to a central pixel. In *cross bilteral filtering* [9, 17] the distance enforcement is according to a second refernce image. Using cross-bilateral filtering we can decompose an image into a low frequency component where edges are preserved - called base layer - and a high frequency component called detail layer defined to be the original minus the base. Base and detail layers are calculated:

$$L_{base} = BF(reference, target) \tag{1}$$
$$L_{detail} = L_{in} - L_{base} \tag{2}$$

where both *reference* and *target* images are in log-space.

The cross-bilateral filter uses the $reference$ image to determine the photometric weighting where the actual smoothing is carried out on the *target* image. If the reference is the same as the target image the cross bilateral filter is the same as the conventional bilateral filter.

In previous tone-mapping research [8], a bilateral filter was used to decompose an HDR image into two-scale layers: an HDR base-layer (because large edges are preserved) and a small-scale LDR detail-layer. Then the base-layer is tone-mapped to LDR using a G-TMO e.g. Tumblin et al. [23]. Finally, the detail is added back onto the output image. Mathematically, this method is summarised as:

$$L_{base} = BF(L_{in}, L_{in}) \tag{3}$$
$$L_{base'} = G-TMO(L_{base}) \tag{4}$$
$$L_{out} = L_{base'} + L_{detail} \tag{5}$$

L_{in} denotes a brightness image (e.g. average of R, G and B channels) in logarithmic domain of the input HDR, $L_{base'}$ is the LDR tone-mapped image of the L_{base} using a G-TMO. The final step is to reconstruct the output image L_{out} by adding L_{detail} to the tone-mapped base layer $L_{base'}$.

Color Processing

Most TMOs are applied only to the brightness image. Full colour output images are calculated according to [20].

$$C_{out} = \left(\frac{C_{in}}{L_{in}}\right) L_{out} \tag{6}$$

C denotes one of the color channels (red, green, or blue). L_{in} and L_{out} denote the luminance before and after tone-mapping, respectively. All values are given in a linear domain (no gamma-correction is applied).

Later [24] observed that the resulting images of their TMO using Equation 6 are often over-saturated due to a stronger contrast compression. To solve this problem, they simply introduce a color saturation factor s to the Equation 6. The color components C_{out} of the tone-mapped image are now computed as:

$$C_{out} = \left(\frac{C_{in}}{L_{in}}\right)^{s} L_{out} \tag{7}$$

The Optimal Approximation

Our model of the optimal G-TMO approximation of a SV-TMO consists of three components: (1) tone-curve optimisation (the main contribution of this paper), (2) detail recovery, and (3) color reconstruction. Figure 1 illustrates our processing workflow. The main tone-mapping processing is summarised as:

$$L_{out} = DR(G-TMO(L_{in})) \tag{8}$$

$G-TMO()$ is the global tone-curve applied to the input image, which can be either derived from a down-sampling or full-size image. $DR()$ denotes the detail recovery process (see below). The color image can be then reconstructed using Equation 6. An additional unsharp masking step is available to tweak preferred contrast (not necessary for the experiments we report later).

Tone-curve Optimization

Let us now consider the function G-TMO. The optimal global tone-curve applied to an image is not an arbitrary function. Rather, it should be monotonically increasing both to avoid intensity inversions and to allow image manipulations to be undone. Almost all curve adjustments made to images (brightening, contrast changes and gamma) are monotonically increasing functions.

Given an HDR image and its spatially-varying tone-mapped LDR image, we want to find a 1-D surjective and monotonically increasing function that best maps HDR to LDR. We point out that we have a choice of how to *encode* the HDR brightness image. Throughout this paper we will choose to represent our input

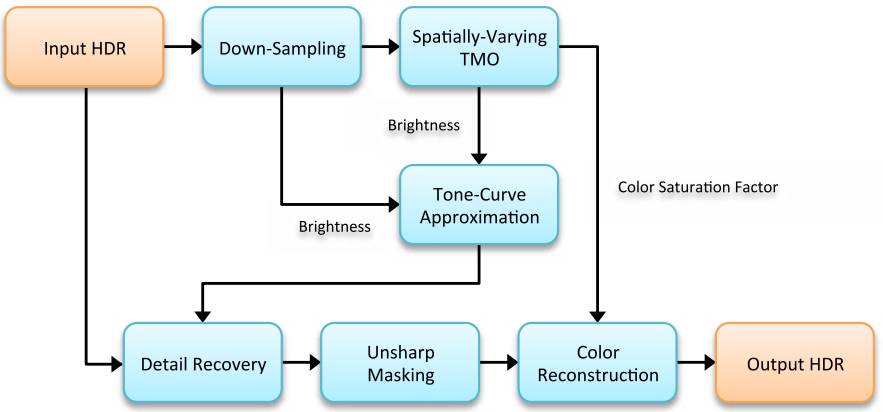

Figure 1. *Image processing of the optimal approximation operator.*

data in the *log* domain. We do this because relative differences are most meaningful to human observers (we have approximately a *log* visual response) and, practically, most tone mappers take the *log* of the input image as input.

Pool-Adjacent-Violators-Algorithm (PAVA)

To monotonically constrain the shape of the function (and minimise squared residual error), Quadratic Programming (QP) might be used. However, QP is a complex computational procedure and is not something easily implementable in (say) a digital camera. A simpler (but equally optimal) solution can be found using the Pool-Adjacent-Violators-Algorithm (PAVA) [2]. To our knowledge, this paper is the first application of PAVA to image reproduction.

PAVA is a simple iterative algorithm for solving monotonic (either increasing or decreasing) regression problems. It pools values together until an optimal solution is found.

Let us assume we have a two-dimensional data $\{(X_i, Y_i)\}_{i=1}^{n}$ where X_i is in ascending order. We seek a monotonically increasing function $\hat{m}()$ that minimizes

$$\sum_{i=1}^{n}(Y_i - \hat{m}(X_i))^2 \tag{9}$$

subject to

$$\hat{m}(X_{(1)}) \leq \hat{m}(X_{(2)}) \leq ... \leq \hat{m}(X_{(n)}). \tag{10}$$

The PAVA algorithm (originally presented in [14]), which optimally solves this problem, works as follows:

1. Sort the data according to X (Y is reordered according to the X) [X_i is in monotonically increasing order but Y_i may not be]$\{\hat{m}(X_{(i)})\}_{i=1}^{n}\}$.
2. Starting from the leftmost of the function $Y_{(1)}$ move to the right and stop if the pair (Y_i, Y_{i+1}) violates the monotonicity constraint we seek: $Y_i > Y_{i+1}$.
3. Pool $Y_{(i)}$ and the adjacent $Y_{(i+1)}$ together and replacing them both by their average, $Y^*_{(i)} = Y^*_{(i+1)} = (Y_{(i)} + Y_{(i+1)})/2$.
4. Next check that $Y_{(i-1)} \leq Y^*_i$. If not, pool $\{Y_{(i-1)}, Y_{(i)}, Y_{(i+1)}\}$ into one average. Continue to pool to the left unless the monotonicity requirement is satisfied.
5. Proceed to the right and keep repeating from step 2 until an monotonically increasing solution is derived.

By the way the algorithm works the computational complexity of PAVA is $O(n^2)$ in the worst case (very expensive if every pixel has a unique intensity!). The worst case for PAVA happens when, in step 4, we have to search back to our first data point. In practice (for our application) the worst case is never encountered (if a G-TMO is a good approximation, the expected complexity of PAVA can be shown to low).

Reducing the Complexity of PAVA with Fixed Quantization Levels

Suppose we have $n+1$ quantization levels of X : $q_n, q_n - 1, \cdots, 0$. If the minimum *log*-value is M then let $q_i = \frac{i}{n}M$. For each quantization level there may be many different output values. But, the complexity of PAVA is bounded by the n quantization levels (say 32, compared with the millions of pixels in the original image). We calculate PAVA only for $X_i = q_i$ and the corresponding output Y (a single quantization level can have many different output values). Of course our original HDR image is not quantized thus we must calculate the outputs by linear interpolation. For an arbitrary X (a brightness in the input HDR image whose brightness is between quantile levels u and $u+1$) we calculate the position of this brightness between the appropriate quantization levels:

$$\alpha = \frac{X - q_u}{q_{u+1} - q_u} \tag{11}$$

We assume that the output Y is the same linear combination of the outputs for these quantization levels:

$$\hat{m}(X) = (1-\alpha)\hat{m}(q_u) + \alpha\hat{m}(q_{u+1}) \tag{12}$$

Smoothed PAVA

One negative characteristic of PAVA is that it sensitive to outliers. In the presence of outliers, PAVA can produce long flat levels (step function). Such tone-curves (though least-squares optimal) will not lead to good looking images. The visual meaning of the flat part of a tone-curve is that a range of input values are all mapped to the same output values (with a potential loss

of detail as a result). To avoid this flattening behavior, we simply smooth the PAVA solution which mathematically equivalent to finding an optimal smooth curve as part of the general optimization [12]. But, *post*-smoothing has the advantage that it is carried out on a small number of regression points (e.g. for the 32 quantization levels) which leads to much faster processing than smoothing the set of actual mapping points. The robust loess estimation procedure is used to smooth our data.

An example of PAVA and its robust smoothed version is shown in Figure 2. As can be seen, there is a flat area occurred in the highlight of the original PAVA function which is smoothed out.

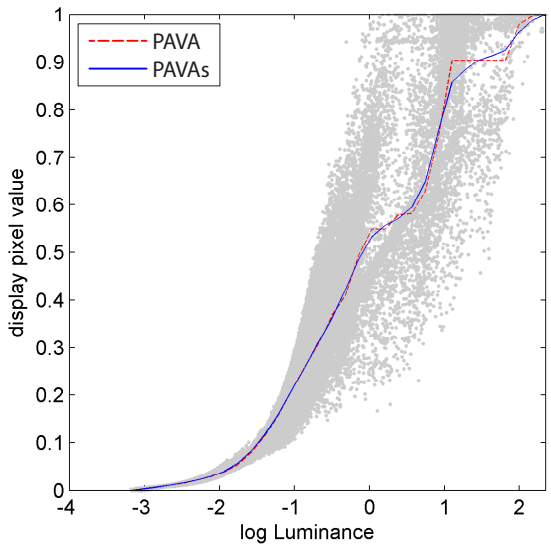

Figure 2. *Tone-curves resulted from PAVA and its robust smoothed version (PAVAs).*

Even though the basic PAVA solution is optimal in terms of Root Mean Square Error (RMSE), flat regions of tone-curve can produce poor image outputs. Significantly, we have found that the smoothed PAVA curve has almost the same RMSE. For the data fit shown in Figure 2, the optimal PAVA output captures 95.02% of the spatially varying TMOs output (RMSE of 0.0498) whereas the smoothed version (PAVAs) is almost as good 94.94% (RMSE of 0.0506). Similar smoothing results were found for all images we tested. Smoothed PAVA produces visually more pleasing output at the cost of a very small decrement in the data fit.

Discussion

Tone-curve manipulations tend to stretch contrast in some image areas and compress in others. When the derivative of the tone curve is less than 1 detail is being compressed and when it is greater than 1 there is an increase in contrast. Thus, while tone-mapped images generate by Smoothed PAVA are often similar to the spatially-varying outputs they can look very flat (typically in the highlight region area, the best global tone-curve has a <1 derivative). Simple unsharp masking (as shown in the workflow of Figure 1) can often (very slightly) ameliorate this problem.

However, unsharp masking does not always work. Spatially-varying TMOs such as Retinex manipulate images in a very local manner. Indeed, one of the problems of Retinex is that it can introduce artifacts such as halo around high contrast edges. Although, often problematic the processes that make halos also add a contrast boost (a 'punch') to local areas of images

Figure 3. *An example of a poor fit of our tone-curve approximation. Left, the reference image resulted from Retinex has higher local contrast compare to our mapping result on the right.*

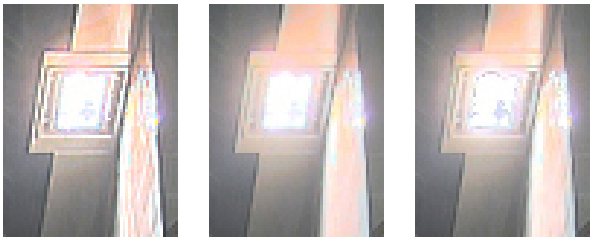

Figure 4. *Details recovered from the cross-bilateral of the proposed operator. Images were cropped from images shown in middle row of Figure 7.*

which is often preferred results. Of course, the more locally an image is processed, the less well a global tone-curve can approximate the outputs. Figure 3 gives an example of this poor *visual* fit. Arguably the Retinex output on the left has an almost 'hyper' realism. But, the G-TMO output on the right looks unnaturally flat.

Thus, in common, with previous works on G-TMOs, we must recover the detail that is missing in the G-TMO reproduction.

Detail Recovery

In [8] the detail is recovered using Equations (3) through (5). Here we adopt a slightly different approach using the cross bilateral filter. Denoting the ouptut of our PAVA method as L_{approx}, the final output of our tone-curve algorithm is:

$$L_{out} = L_{in} + BF(L_{in}, L_{approx} - L_{in}) \quad (13)$$

The advantage of (13) is that the $BF()$ is applied only to the gain. Equation (13) is a concise summary of the exact computation needed to bring back the detail. However, mathematically the approaches of (13) and (1) through (5) are similar). They are not precisely the same as L_{approx} is calculated given the original input image (not the input base layer). For our purposes, we found 13 gave slightly better results. Figure 4 gives an example of how detail recovery process can bring back the detail that is missing in the PAVA process.

Experiment

To test image preference of the optimal approximation operator against spatially-varying TMOs, a paired comparison psychophysical experiment [10] was conducted. The purpose of this experiment was to compare all the test SV-TMO images and to ensure that in general the G-TMO images derived from the smooth-PAVA optimization are comparable.

Five HDR images were used in our experiments: Figure 5 shows the image dataset. Tone-mapped images used in the experiment were generated by applying the optimal approximation op-

erator to tone-mapped images of 3 well-known spatially-varying TMOs. Table 1 lists TMOs that have been used in the experiment (note the labels we use to index the results in Table 2 and Figure 6). To avoid unfaithful implementation and bias from parameter selection problems, we use the images available from a previous study [7].

There are 6 tone-mapped images per HDR image (3 reference tone-mapped images, and 3 created using smooth-PAVA in the context of the processing workflow shown in Figure 1). Thus, there are 30 tone-mapped images in total. These images were evaluated by 21 observers (9 males and 12 females) with normal color vision, naive to the goal of the experiment. Images were viewed under controlled experimental conditions [15].

There were a total of 75 pairwise comparisons in the experiment (5 images * 15 pairs of algorithms). Participants were asked to make judgements of the TMOs based on overall appearance. For each pair, participants were instructed to observe the two tone-mapped images and select the one they preferred. Images were shown randomly on the left or the right of the screen. All algorithms were shown roughly an equal number times left and right. The whole procedure per participant took approximately 8-10 minutes. In order to evaluate the results, preference scores are generated using Thurstone's Law of Comparative Judgement Case V [21].

Figure 5. *5 well-known images used in the experiment. From left to right AtriumMorning, AtriumNight, HotelRoom, Memorial, and NapaValley. Image courtesy of Frédéric Drago, Paul Debevec, Simon Crone, and Spheron AG.*

Algorithm parameters

For PAVA, the number of quantization points that we use is 32. For bilateral filter parameters, we set the spatial closeness (σ_s) to 2% of the image size and the photometric similarity (σ_r) to 0.6*log* units since these two values performed consistently well for all test images. All processing is carried out in the brightness domain. The color output is created according to Equation 6.

Results and Discussions

The average preference scores of 5 test scenes from the 21 subjects was given in Figure 6 (the x-axis shows the operators. The actual scores are given in Table 2. The numbers in Table 2 are the number of times a particular algorithms is preferred. The last two columns compared the pooled algorithm performance: SV-TMO vs G-TMO. Using statistical assumptions (Thurstone's law of comparative judgement Case V) we can turn these raw preference numbers into a preference score with confidence intervals. See [13] for a full discussion of how this is done.

In Figure 6, The 95% confidence interval (error bars) are shown in a normalised preference score interval. The y-axis here (preference score) can be interpreted as a z-score: if one operator were strongly preferred it could, in theory have a score of 2 and if an operator was strongly unpreferred its normalised preference could be as low as -2. In cases where algorithms do not deliver strong preference, the preference scores tend to cluster around 0 on this normalised scale. If error bars do not overlap then one algorithm is better than another at the 95% confidence interval.

The results show that on average among six TMOs, TMO P (Reinhard's photographic tone reproduction operator) was most preferred but that the G-TMO $\hat{P}$ produced images which were not significantly different (in terms of preference). We can also clearly see that A (Ashikhmin) has the lowest score whereas its global approximating version $\hat{A}$ have a significant higher score (the score difference exceeds the confidence interval) and in fact ranked the third overall. This indicates that TMO which generate too much detail may be judged poorer than its global approximating version.

The *pooled* average scores for each category (S for spatially-varying and G for the proposed global approximation) are given on the right of the Figure. Since the two scores fall in to the same interval scale, it is reasonably to conclude that in general, the perceived quality of the optimal approximation operator is similar to its spatially-varying operators.

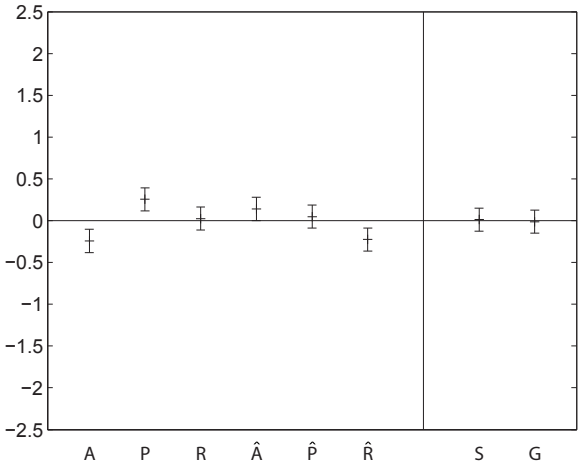

Figure 6. *Overall preference scores of 6 different TMOs (6 scores on the left). The average score of spatially-varying TMOs and the average score of the proposed operator are shown on the right.*

A	P	R	Â	P̂	R̂	S	G
213	315	267	291	273	216	795	780

The number of times each TMO is preferred over the others.

Conclusions

We demonstrate that many spatially varying TMOs can be visually approximated by the optimal global tone-curve approximation so long as care is taken to preserve local detail. A psychophysical experiment validates our method. Our experiment demonstrates that for three of the most widely used spatially-varying TMOs, their optimal global tone-curve approximation generates images that were equally preferred. This is a significant result as our global operator does not suffer from spatial ar-

Name	Label	Category
A Tone Mapping Algorithm for High Contrast Images [1]	A	S
Photographic Tone Reproduction [18]	P	
Retinex adapted to tone-mapping [6]	R	
Optimal Global Approximation to A	$\hat{A}$	G
Optimal Global Approximation to P	$\hat{P}$	
Optimal Global Approximation to R	$\hat{R}$	

Tone mapping operators used in the experiment together with their labels and categories used in the paper. For the category, S stands for spatially-varying TMOs, G stands for the optimal global approximation operator.

tifacts (such as halos) and can be implemented to run extremely rapidly. indeed the global tone-curve itself can be *learned* using only a small input thumbnail.

References

[1] M. Ashikhmin. A tone mapping algorithm for high contrast images. In S. Gibson and P. E. Debevec, editors, Rendering Techniques, volume 28 of ACM International Conference Proceeding Series. Eurographics Association, 2002.

[2] R. E. Barlow and H. D. Brunk. The isotonic regression problem and its dual. Journal of the American Statistical Associ- ation, 67(337):pp. 140147, 1972.

[3] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In Proceedings of the 24th annual conference on Computer graphics and interactive techniques, SIGGRAPH97, pages 369378, New York, NY, USA, 1997. ACM Press/Addison-Wesley Publishing Co.

[4] K. Devlin. A review of tone reproduction techniques. Technical Report CSTR-02-005, Department of Computer Sci- ence, University of Bristol, November 2002.

[5] J. M. Dicarlo and B. A. Wandell. Rendering high dynamic range images. volume 3965, pages 392401. SPIE, 2000.

[6] F. Drago, W. Martens, K. Myszkowski, and H.-P. Seidel. Perceptual evaluation of tone mapping operators with regard to similarity and preference. Research Report MPI-I-2002-4-002, Max-Planck-Institut für Informatik, Stuhlsatzenhausweg 85, 66123 Saarbrücken, Germany, August 2002.

[7] F. Drago, K. Myszkowski, H.-P. Seidel, and W.L. Martens. Gallery of reference tone mapped images. http://www.mpi-inf.mpg.de/resources/tmo/NewExperiment/TmoOverview.html, 2010. [Online; accessed 06-July-2011].

[8] F. Durand and J. Dorsey. Fast bilateral filtering for the display of high-dynamic-range images. In SIGGRAPH02: Proceedings of the 29th annual conference on Computer graphics and interactive techniques, pages 257266, New York, NY, USA, 2002. ACM.

[9] E. Eisemann and F. Durand. 2004. Flash photography enhancement via intrinsic relighting. In ACM SIGGRAPH 2004 Papers (SIGGRAPH '04), Joe Marks (Ed.). ACM, New York, NY, USA, 673-678.

[10] P. Engeldrum. Psychometric scaling: A toolkit for imaging systems development. 2000.

[11] R. Fattal, D. Lischinski, and M. Werman. Gradient domain high dynamic range compression. ACM Trans. Graph., 21:249256, July 2002.

[12] J. Friedman and R. Tibshirani. The monotone smoothing of scatterplots. 1984.

[13] P. Green and L. MacDonald. Colour engineering: achieving device independent colour. Wiley SID series in display technology. Wiley, 2002.

[14] W. Häerdle. Monotonic and unimodal smoothing. http://fedc.wiwi.hu-berlin.de/xplore/ebooks/html/anr/anrhtmlnode43.html. [On-line; accessed 06-July-2011].

[15] ISO 3664:2009. Graphic technology and photography Viewing conditions. ISO, Geneva, Switzerland.

[16] R. Mantiuk and H. peter Seidel. Modeling a generic tonemapping operator. Computer Graphics Forum, 27:699708, 2008.

[17] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama. 2004. Digital photography with flash and no-flash image pairs. ACM Trans. Graph. 23, 3 (August 2004), 664-672.

[18] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda. Photographic tone reproduction for digital images. In Proceedings of the 29th annual conference on Computer graphics and in- teractive techniques, SIGGRAPH02, pages 267276, New York, NY, USA, 2002. ACM.

[19] E. Reinhard, G. Ward, S. Pattanaik, and P. Debevec. High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting (The Morgan Kaufmann Series in Computer Graphics). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2005.

[20] C. Schlick. Quantization techniques for visualization of high dynamic range pictures. pages 720. Springer-Verlag, 1994.

[21] L. L. Thurstone. A law of comparative judgment. Psychological Review, 34(4):273286, 1927.

[22] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. pages 839846, 1998.

[23] J. Tumblin, J. K. Hodgins, and B. K. Guenter. Two methods for display of high contrast images. ACM Trans. Graph., 18(1):5694, 1999.

[24] J. Tumblin and G. Turk. LCIS: a boundary hierarchy for detail-preserving contrast reduction. In Proceedings of the 26th annual conference on Computer graphics and interactive techniques, SIGGRAPH 99, pages 8390, New York, NY, USA, 1999. ACM Press/Addison-Wesley Publishing Co.

Author Biography

Jakkarin Singnoo received his BSc in Imaging and Printing Technology from Chulalongkorn University, Thailand (2004) and his MSc in Advanced Computer Graphics and Applications from University of East Anglia, UK (2006). He is now a PhD student in Colour Group, School of Computing Sciences at University of East Anglia. His work has focused on HDR imaging including Tone Mapping Operators and Psychophysics.

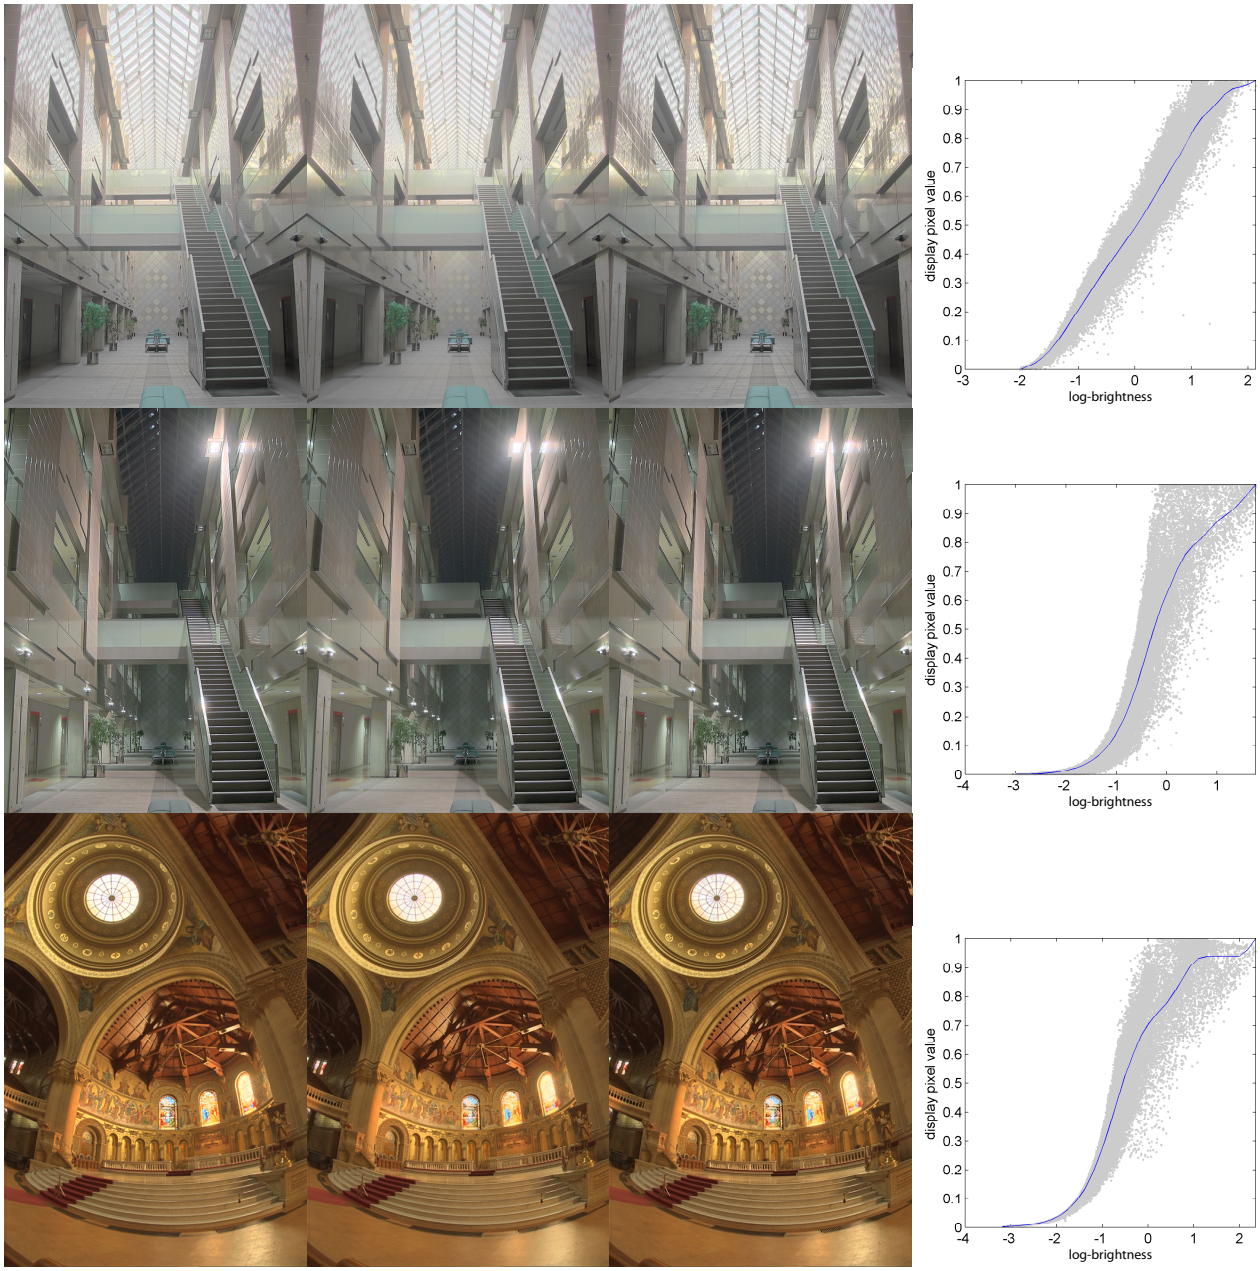

Figure 7. *Results of our approximation to the three tone-mapped images that we test. Left column, the reference tone-mapped image. Our approximation output before and after the detail recovery process are given in middle left and middle right columns, respectively. Right column, approximated tone-curve (blue) and the scatter plot showing the correlation between each pixel of the LDR and HDR brightness values. Note that: Although the regression is done in log-log space, we found that it is more intuitive to visualize it in log-linear space.*