

Dynamic Range Reduction inspired by Photoreceptor Physiology

Erik Reinhard, *Member, IEEE*, and Kate Devlin

Abstract—A common task in computer graphics is the mapping of digital high dynamic range images to low dynamic range display devices such as monitors and printers. This task is similar to adaptation processes which occur in the human visual system. Physiological evidence suggests that adaptation already occurs in the photoreceptors, leading to a straightforward model that can be easily adapted for tone reproduction. The result is a fast and practical algorithm for general use with intuitive user parameters that control intensity, contrast and level of adaptation respectively.

Index Terms—tone reproduction, dynamic range reduction, photoreceptor physiology.

I. INTRODUCTION

THE real world shows a vast range of light intensities ranging from star-lit scenes to white snow in sun-light. Even within a single scene the range of luminances can span several orders of magnitude. This high dynamic range within a single scene can easily be computed with computer graphics techniques. They can also be captured using a composition of multiple photographs of the same scene with different exposures [1]. In the near future high dynamic range sensors will become generally available to directly capture high dynamic range images. Furthermore, the dynamic range of data captured with medical data acquisition techniques and scientific simulations may be arbitrarily high.

As a result, the availability of high dynamic range data will become much more commonplace than it is now. In contrast, the dynamic range of display devices is currently limited, and economically sensible high dynamic range display devices are not yet commonplace. This may change in the near future as recent research has already produced a high dynamic range display by combining LCD and LED technologies [2]. The dynamic range of printers, on the other hand, will remain low. The mismatch between high dynamic range data acquisition and high dynamic range display technology will therefore persist in one form or another.

This leads to the problem of how to display high dynamic range data on low dynamic range display devices, a problem which is generally termed tone mapping or tone reproduction [3], [4]. In principle this problem is simple: we need to turn an image with a large range of numbers into an image containing integers in the range of 0 to 255 such that we can display it on a printer or a monitor. This suggests linear scaling as a possible solution. However, this approach is flawed because

details in the light or dark areas of the image will be lost due to subsequent quantization, and the displayed image will therefore not be perceived the same as the scene that was photographed (Fig. 1).



Fig. 1. Linear scaling of HDR data (inset) will cause many details to be lost. Tone reproduction algorithms such as the technique described in this paper attempt to solve this issue, in this case recovering detail in both light and dark areas as well as all areas inbetween.

Tone reproduction algorithms therefore attempt to scale the high dynamic range data in such a way that the resulting displayable image has preserved certain characteristics of the input data, such as brightness, visibility, contrast or appearance. Algorithms can be classified into two broad categories: global and local operators. Global operators compress contrasts based on globally derived quantities, which may include for example the minimum and maximum luminance or the average luminance. In particular the log average luminance is often computed to anchor the computation. The compression algorithm then compresses pixel contrasts according to a non-linear function based on its luminance, as well as those global variables. No other information is used to modulate the compression curve [5]–[15].

The shape of the compression curve is what differentiates these global algorithms. While visual inspection of the compression curves, i.e. the functions that map high dynamic range luminances to display luminances, may lead to the suggestion that most of these algorithms are very similar in nature, we have found that small differences in their functional form may lead to substantial differences in visual appearance.

Global algorithms tend to be computationally efficient, but may have distinct disadvantages. In particular, loss of detail is often associated with global operators. The more recent algorithms tend to exhibit fewer artifacts than earlier attempts,

however.

A distinguishing feature of local operators is their use of neighboring pixels to derive the amount by which to compress a pixel [13], [16], [17]. Local operators may show haloing or ringing artifacts which indicate that although the principle may be valid, the calibration of these models is critical and is often not well understood.

Tone reproduction operators may also be classified based on whether they rely on models of human visual perception or on mathematical or engineering principles. Some tone reproduction operators use explicit perceptual models to control the operator [6], [8]–[11], [16]–[19], and in particular work on the assumption that local spatial interaction is a key feature in dynamic range reduction [16]. Other spatially varying operators have used bi- or tri-lateral filtering [20], [21] or compress the gradient of an image followed by numerical integration [22].

The human visual system (HVS) successfully and effortlessly overcomes dynamic range issues for a vast range of intensities by using various adaptation mechanisms. In addition to the photoreceptors (rods and cones), the retina contains additional types of cells, such as horizontal and amacrine cells providing lateral interconnectivity, and bipolar and ganglion cells giving distal connectivity [23]. Although this alone provides several loci where adaptation may occur, a key observation is that *all* cells in the HVS have a limited capability to produce graded potentials or spike trains. By definition this includes the very first cells in the chain of visual processing: the photoreceptors. Hence, dynamic range reduction must already occur in the rods and cones. Results from electro-physiology have confirmed this [24]–[27].

In this paper we adapt a computational model of photoreceptor behavior to help solve the tone reproduction problem. The aim of this work is to provide a new global tone reproduction operator that is fast and produces plausible results that are useful in practical settings such as high dynamic range photography. We believe that for a large range of images our method combines the speed of global tone reproduction operators with the ability to compress high dynamic range images as well as or better than current operators.

While our method is grounded in results obtained from electro-physiology, we do not present a full and complete model of photo-receptor behavior, because this would add unnecessary complexity to the model. The dynamic range of cells at various stages of visual processing may differ, and different adaptation mechanisms exist at different loci in the human visual system [23]. We therefore do not aim to present a complete model of the early stages of human vision, but focus on the first step of visual processing - the photoreceptors. Also, this step is only modeled to the extent that it allows the problem of tone reproduction to be addressed. The model of visual processing employed here should therefore not be seen as complete or even predictive for human visual perception.

Also, we deviate from this model in certain areas to increase the practical use of our algorithm. In particular, we have fitted the model with four user parameters which allow overall intensity, contrast, light- and chromatic adaptation to be independently controlled. However, we do show that initial

estimates may be computed for these parameters that provide results that in most cases require only very small adjustments.

II. ALGORITHM

Various mechanisms in the HVS mediate adaptation to lighting conditions. We specifically employ a model of photoreceptor adaptation, which can be described as the receptors' automatic adjustment to the general level of illumination [25], [28]. The potential V produced by cones as function of intensity I may be modeled by [29]:

$$\begin{aligned} V &= \frac{I}{I + \sigma(I_a)} V_{max} \\ \sigma(I_a) &= (f I_a)^m \end{aligned} \quad (1)$$

These equations are a subtle but important deviation from the more common Naka-Rushton equation [26] and are derived by Hood and colleagues for reasons mentioned in their paper [29]. The semi-saturation constant $\sigma(I_a)$ describes the state of long-term adaptation of the photo-receptor as function of adaptation level I_a . Both f and m are constants, but will be treated as user parameters in our adaptation of the model. Their values differ between studies, but for m it is found to lie between 0.2 and 0.9 [29]. The value of the multiplier f is not discussed further by Hood et. al., but we have found that setting $f = 1$ provides a useful initial estimate. The maximum incremental response elicited by I is given by V_{max} , which we set to 1. One reasonable assumption made for (1) is that the input signal is positive, so that the output V lies between 0 and 1.

The adaptation level I_a for a given photoreceptor can be thought of as a function of the light intensities that this photoreceptor has been exposed to in the recent past. If a sequence of frames were available, we could compute I_a by integration over time [12]. This approach may account for the state of adaptation under varying lighting conditions. However, even under stationary lighting conditions, saccadic eye movements as well as ocular light scatter cause each photoreceptor to be exposed to intensity fluctuations. The effect of saccades and light scattering may be modeled by computing I_a as a spatially weighted average [30].

Some previous tone reproduction operators that use similar compression curves, compute σ by spatial integration [13], [17]. However, if σ is based on a local average, then irrespective of the shape of the compression curve, ringing artifacts may occur [31]. By carefully controlling the spatial extent of σ , these artifacts may be minimized [13], [20]. We compare different choices of global and local adaptation levels I_a in Section IV.

In practice, we may assume that each photoreceptor is neither completely adapted to the intensity it is currently exposed to, nor is it adapted to the globally average scene intensity, but instead is a mixture of the two. Rather than compute an expensive spatially localized average for each pixel, we propose to interpolate between the pixel intensity and the average scene intensity. In the remainder of this paper we will use the term *light adaptation* for this interpolation.

Similarly, a small cluster of photoreceptors may be adapted to the spectrum of light it currently receives, or it may be

adapted to the dominant spectrum in the scene. We expect photoreceptors to be partially adapted to both. The level of chromatic adaptation may thus be computed by interpolating between the pixel's red, green and blue values and its luminance value. By making the adaptation level dependent on luminance only no chromatic adaptation will be applied, whereas keeping the three channels separate for each pixel achieves von Kries-style color correction [32].

We have found that for most images, keeping all channels fully dependent suffices, whereas using the pixel intensity itself rather than the scene average produces better compression. While for most images the setting of the interpolation weights is not critical, for the purpose of demonstrating the effect of different weights, we present an atypical result in Fig. 2. Our default settings would result in the image in the top right corner, which we deem overly compressed. In our opinion, the image on the middle right presents an attractive trade-off between detail visibility and contrast. The effect of manipulating the two interpolation weights is generally smaller because most images have a less pronounced color cast. Results shown in the remainder of this paper will have the two interpolation weights set to their default values, unless indicated otherwise.

Finally, we note that we could simulate the effect of time dependent adaptation for a still image by making the two interpolation weights functions of time and creating a sequence of images tonemapped with different interpolation weights. We illustrate this in Fig. 3, where both weights were incremented from 0 to 1 in steps of 0.2. Note that we simultaneously achieve adaptation to luminance levels as well as chromatic adaptation. The image on the right shows more detail in both the dark and light areas, while at the same time the yellow color cast is removed.

III. USER PARAMETERS

For certain applications it may be important to have a tone reproduction operator without any user parameters. Other applications may benefit from a small amount of user intervention, provided that the parameters are intuitive and that the number of parameters is small. We provide an intermediary solution by fitting the model with carefully chosen user parameters that may be adjusted within a sensible range of values. These parameters have an intuitive effect on the resulting images, so that parameter adjustment involves as little guesswork as possible.

In addition, we provide initial estimates of these parameters that produce plausible results for a wide variety of images. This benefits applications that require fully automatic tone reproduction, and also creates reasonable initial images that may be further modified by the user.

Two of the user parameters were introduced in the previous section. These are m and f , which control contrast and intensity respectively. In this section we discuss their useful range of operation, as well as reasonable initial estimates. We also provide further details for the parameters that govern the level of chromatic and light adaptation.

Although the constant m has been determined for specific experimental set-ups [29], we have found that its value may

successfully be adjusted based on the composition of the image. In particular, we make m dependent on whether the image is high- or low-key (i.e. overall light or dark). The key k can be estimated from the log average, log minimum and log maximum luminance (L_{av} , L_{min} and L_{max}) [14]:

$$k = (L_{max} - L_{av}) / (L_{max} - L_{min}) \quad (2)$$

with luminance specified as:

$$L = 0.2125 I_r + 0.7154 I_g + 0.0721 I_b. \quad (3)$$

We choose mapping the key k to the exponent m as follows:

$$m = 0.3 + 0.7k^{1.4}. \quad (4)$$

This mapping is based on extensive experimentation and also brings the exponent within the range of values reported by electro-physiological studies [29]. It was chosen for engineering purposes to make the algorithm practical for a wide variety of input data. By anchoring the exponent m to the log average luminance in relation to the log minimum and log maximum luminance, the model becomes robust in the sense that the input data does not need to be calibrated in particular units of measurement. This computation produces a reasonable initial estimate for m , although sometimes images may benefit from manual adjustment. We have found that the range of operation for this parameter should be limited to the range $[0.3, 1)$. Different values of m result in different shapes of the compression curve, as shown in Figure 4 for a range of values of m . This plot was created by tonemapping an image with a linear ramp between 10^{-3} and 10^3 . For this ramp, the exponent m would be initialized to a value of 0.493. This parameter has a direct effect on the slope of the curve, and thus trades contrast in medium-intensity regions for detail visibility in the dark and bright regions by becoming more or less "S"-shaped.

While the parameter f discussed above may be set to 1 as an appropriate initial value, we allow f to be varied by the user as well. Although it is possible to set f directly, the range of useful values is large and non-intuitive. We therefore replace the multiplier f by an exponential function:

$$f = \exp(-f') \quad (5)$$

By changing the parameter f' the overall intensity of the image may be altered; higher values will make the result lighter whereas lower values make the image darker. For most images, the range of useful values of this parameter is between -8 and 8 , with an initial estimate of 0 (such that $f = 1$ as indicated in the previous section). The tone curve follows a similar progression of shapes for different choices of f' as seen for parameter m in Figure 4. However, in practice the visual impact of this parameter is different from m and we therefore keep both parameters.

We apply (1) to each of the red, green and blue channels independently, because in the HVS different cone types do not interact. However, it may be desirable to remove strong color casts in the image, which can be achieved by computing the adaptation level I_a for each of the red, green and blue

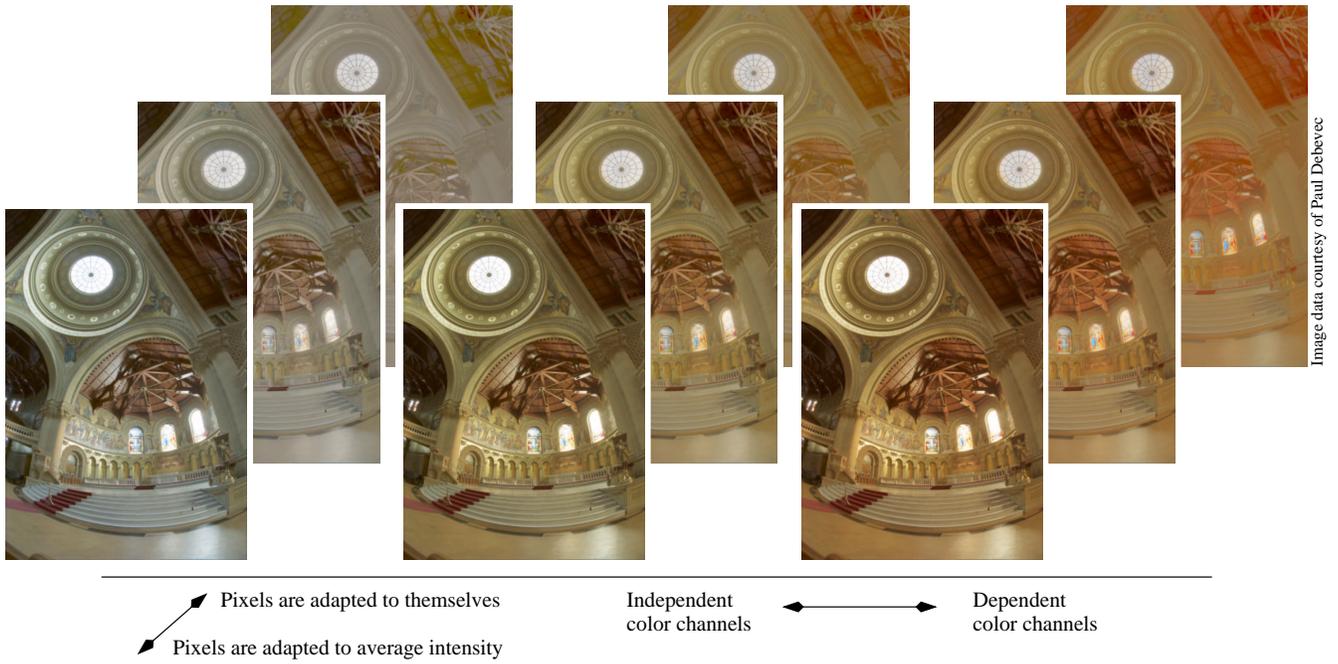


Fig. 2. Memorial church image, showing the effect of different methods to compute the adaptation luminance I_a . The values of the weights are varied from 0 to 0.5 to 1. This image is a particularly good example to show this effect because it has a strong yellow color cast. For most images, the setting of I_a is less critical and more benign.



Fig. 3. Outdoor scene taken at dawn and indoor scene taken around mid-day, both with simulated time dependent adaptation. On the left, the adaptation level I_a is computed from the average scene luminance, whereas on the right the adaptation level is computed using the independent red, green and blue components of each individual pixel.

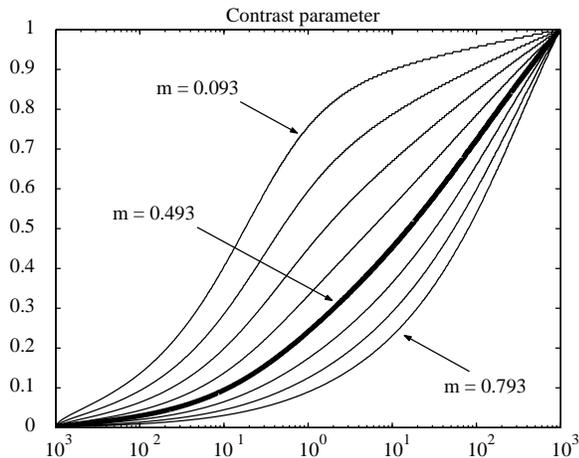


Fig. 4. Mapping of input luminances (horizontal) to display luminances (vertical) for different values of m .

channels as a weighted sum of the pixel's luminance L and the intensity value of the channel:

$$I_a = cI_{r|g|b} + (1 - c)L \quad (6)$$

The amount of color correction is now controlled by the weight factor c , which should be between 0 and 1. By setting this user parameter to 1, the red, green and blue color channels are treated independently and this achieves color correction in the spirit of a von Kries model [32]. By default, we do not apply chromatic adaptation by setting $c = 0$ so that the adaptation level is the same for all three color channels.

Similarly, in rare instances we would like to control whether the pixel adaptation is based on the pixel intensity itself, or on global averages:

$$I_a = aI_{r|g|b} + (1 - a)I_{r|g|b}^{av} \quad (7)$$

Here, we use a second weight a which interpolates between the

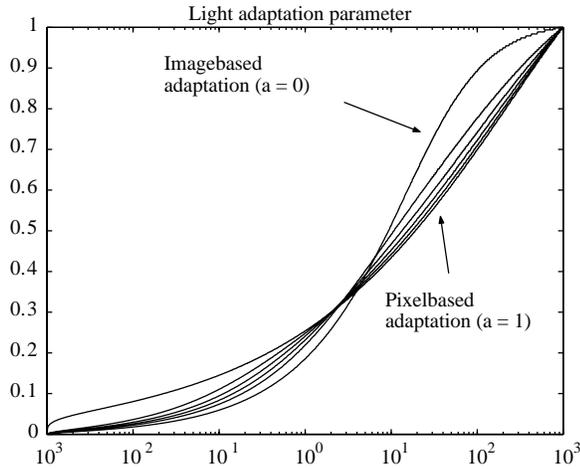


Fig. 5. Mapping of input luminances (horizontal) to display luminances (vertical) for different values of our light adaptation parameter a .

pixel intensity $I_{r|g|b}$ and the average channel intensity $I_{r|g|b}^{av}$. For a value of $a = 1$, adaptation is based on the pixel intensity, whereas for $a = 0$ the adaptation is global. This interpolation thus controls what we will refer to as light adaptation. Its impact on the compression curve is shown in Figure 5. As with m and f' , this parameter steers the shape of the compression curve. Although this family of curves does not span a very wide range of shapes, its visual impact can be considerable, as shown in Figure 2. By default, we set $a = 1$ to trade detail visibility for contrast. These interpolation schemes may be combined using tri-linear interpolation, with $I_{r|g|b}^{av}$ and L^{av} arithmetic averages (Fig. 2):

$$\begin{aligned} I_a^{local} &= cI_{r|g|b} + (1-c)L \\ I_a^{global} &= cI_{r|g|b}^{av} + (1-c)L^{av} \\ I_a &= aI_a^{local} + (1-a)I_a^{global} \end{aligned}$$

For reference, Table I shows all user parameters, their operating range as well as their initial estimates. In practice, manipulating m and f' allows sufficiently fine control over the appearance of the tone-mapped image. In rare cases c and a need minor adjustments too. All four parameters are set once per image, whereas I_a and V are computed per pixel and per channel. After normalization of V , which typically expands the range of values rather than compress them further, we set the display intensity to the photo-receptor output V , making this a simple and fast global operator. The normalization step was implemented by computing the minimum and maximum luminance in the image. The R, G and B channels are then individually scaled using $(I_{r|g|b} - L_{min}) / (L_{max} - L_{min})$ and clipped to 0 and 1. In summary, the source code of the full operator is given in Fig. 6.

IV. RESULTS

In this section we show the effect of manipulating the user parameters m , f' , a and c on visual appearance and compare our results with existing tone mapping operators in terms of visual quality as well as computation time.

```
double Cav[3];           // channel averages
double Lav;             // average luminance
double Llav;           // log average lum
double Lmin, Lmax;     // min and max lum

void tonemap (double ***rgb, // input/output image
             double f,      // overall intensity
             double m,      // contrast
             double a,      // adaptation
             double c)     // color correction
{
    int x, y, i;          // loop variables
    double L;            // pixel luminance
    double I_a;          // pixel adaptation
    double I_g, I_l;     // global and local

    f = exp(-f);
    m = (m > 0.) ? m : 0.3 + 0.7 * pow((log(Lmax) -
        Llav) / (log(Lmax) - log(Lmin)), 1.4);

    for (y = 0; y < height; y++)
        for (x = 0; x < width; x++) {
            L = luminance (rgb, x, y);
            for (i = 0; i < 3; i++) {
                I_l = c * rgb[y][x][i] + (1-c) * L
                I_g = c * Cav[i] + (1-c) * Lav;
                I_a = a * I_l + (1-a) * I_g;
                rgb[y][x][i] /= rgb[y][x][i] + pow(f * I_a, m);
            }
        }
    normalize (rgb, width, height);
}
```

Fig. 6. Source code. Note that user parameter m is computed from globally derived quantities unless the calling function specifies a value for m .

TABLE I
User parameters.

Parameter	Description	Initial value	Operating range
m	Contrast	$0.3 + 0.7k^{1.4}$	[0.3, 1.0]
f'	Intensity	0.0	[-8.0, 8.0]
c	Chromatic adaptation	0.0	[0.0, 1.0]
a	Light adaptation	1.0	[0.0, 1.0]

The images in Fig. 7 vary in the choice of parameters f' and m with the middle image using default settings for all parameters. Both f' and m may be modified beyond the range shown in this figure.

While the operator is global because m is computed from globally derived quantities, the method may be extended to a local operator by setting the adaptation level I_a to a local average of pixel intensities. Using their respective default parameter settings, we experimented with two such local operators, namely bilateral filtering [20] and adaptive gain control [33].

Fig. 8 shows that our global operator performs almost as well as bilateral filtering as applied to our operator. In our opinion, bilateral filtering causes no artifacts due to its ability to avoid filtering across high contrast edges. However, the advantage of applying bilateral filtering to our operator is relatively modest, judging by the visual difference between our global operator and our local operator using bilateral filtering. This observation does not necessarily extrapolate to other tone reproduction operators that may benefit from bilateral filtering.

The effect of applying adaptive gain control is more pronounced. While bilateral filtering applies a (Gaussian) operator in both the spatial as well as the intensity domain, adaptive gain control only filters pixels in the intensity domain [33]. We believe that for this reason adaptive gain control has a

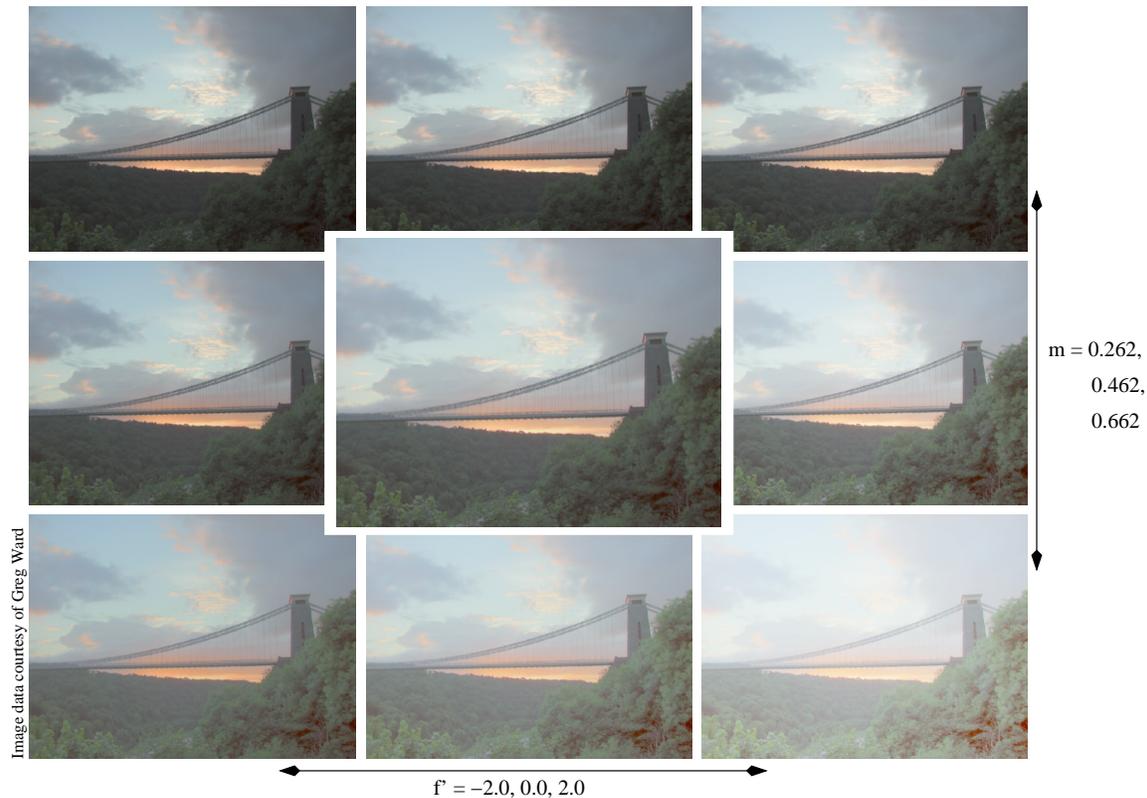


Fig. 7. Clifton suspension bridge showing the effect of varying user parameter m between ± 0.2 of the default value, and the parameter f' which was varied between -2.0 and 2.0 . The enlarged image in the middle was created using our default parameter settings.

somewhat cruder visual appearance. However, this approach has also increased the contrast of the image, which may be desirable for certain applications.

One of the more difficult images to tonemap is the 'desk' image, shown in Fig. 9. Because it is impossible to know how light or dark the image should be displayed to be faithful to the original photograph, we contacted the photographer to discuss and calibrate our result. For the desk image, as a general rule of thumb, the dark area underneath the desk should be quite dark, but some details are visible. In the real scene it was difficult to distinguish details of the bag on the left. The highlight in front of the book should appear bright. The text on the book should be visible, and the light bulb should be distinguishable from the lamp-shade. Note that this image has a fairly strong color cast, which we chose to remove by setting $c = 1$.

In the parking garage in Fig. 10, the foreground should be fairly dark with visible detail, and the area outside should be bright and is also showing detail. Timing results are given in Table II and were obtained using a 2.53 GHz Pentium 4 CPU.

For each of the algorithms in our comparison we manipulated the user parameters to show the details in both the light as well as the dark areas as well as possible. While this may not be in keeping with the intent of each of the operators, our aim is to provide a practical and useful operator. The fairest comparison possible is therefore one where the parameters for each method are optimized to produce the visually most pleasing results. This optimization is by its nature subjective. We applied gamma correction to all images afterwards ($\gamma = 1.6$). The methods we compare against are:

Logarithmic compression. In this algorithm we take the logarithm and apply a linear shift and scale operation to bring the data within displayable range. This operator is included because it is one of the most straightforward techniques that produces a baseline result against which all other operators may be compared.

Adaptive logarithmic mapping. This global tonemapping algorithm logarithmically compresses each pixel, but the base of the logarithm is chosen for each pixel separately according to a bias function [15].

Bi- and trilateral filtering. Here, we applied the bilateral filter as it was originally presented [20]. The method separates the image into a high dynamic range base layer and a low dynamic range detail layer with the aid of a bilateral filter which has the desirable property that the image is blurred without blurring across sharp edges. The base layer is then compressed, shifted and recombined with the detail layer. The two user parameters for this method are shifting and scaling of the base layer. We also compare against the trilateral filter which is an extension of bilateral filtering [21].

Histogram adjustment. Histogram adjustment is a fast and widely used technique which produces good results for a large class of images [10]. We did not include the optional veiling luminance, color sensitivity and visual acuity techniques to promote a fair comparison, but used the pcond program from the Radiance package [34] with no parameters.

Photographic tone-mapping. Photographic tone reproduction may be executed as a global as well as a local operator [13]. We used a parameter estimation technique to find the

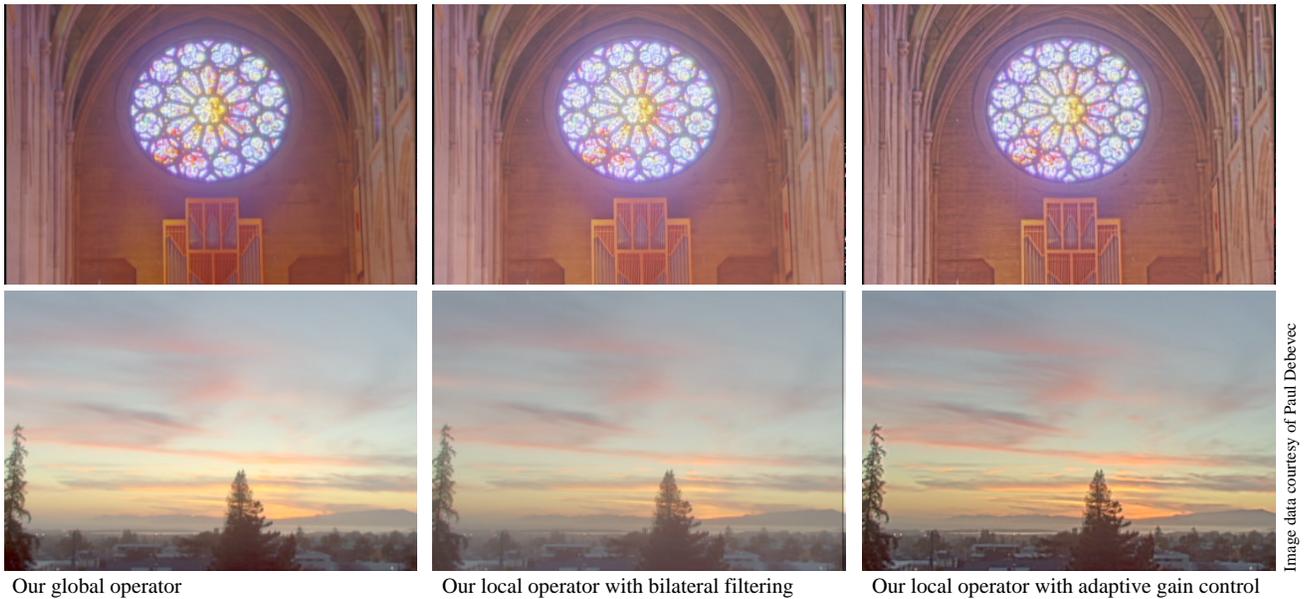


Fig. 8. Rosette and sunset images comparing local and global versions of our operator.

TABLE II

Computation times in seconds for the desk image (1536x1024 pixels) and garage image (748x492 pixels). See also Figs. 9 and 10.

Algorithm	Computation time	
	Desk	Garage
Ours	2.01	0.53
Ours + bilateral filtering	5.80	1.26
Ours + adaptive gain control	135.0	39.66
Photographic (local) [13]	28.86	7.47
Photographic (global) [13]	0.85	0.15
Ashikhmin's operator [17]	46.11	12.17
Bilateral filtering [20]	4.49	0.80
Trilateral filtering ^a [21]	488.0	110.8
Histogram adjustment [10]	1.78	0.43
Logarithmic compression	1.88	0.47
Adaptive logarithmic mapping [15]	1.64	0.40
Time dependent adaptation [12]	6.32	1.39
Revised Tumblin-Rushmeier [11]	2.41	0.59
Uniform rational quantization [7]	1.48	0.27

^aOptimizations as applied to the bilateral filter [20] could also be applied to the trilateral filter, which would reduce the computation time by at least an order of magnitude. Our code, based on the implementation made available by the authors, does not incorporate these optimizations.

appropriate settings for each image [14].

Ashikhmin's operator. This is a local operator based on human visual perception [17]. For each pixel a local adaptation level is computed in a manner similar to the local photographic tonemapping operator. There are no user parameters.

Time-dependent adaptation. This algorithm is a sigmoid using the original Naka/Rushton equation [26] with a fixed semi-saturation constant. Although the algorithm was originally presented to explore time-dependent adaptation [12], we have adapted the algorithm for still images with help from

the author. This algorithm assumes that the input is given in cd/m^2 . Because the units used for the images are unknown, this leaves two parameters to be set manually to convert the image roughly to SI units. It should be noted that for the work on time dependent adaptation, this operator was applied to sequences of low dynamic range images. Similar compression curves were also used for high dynamic range compression, but then the adaptation was local [33], not global as the results shown here. As such, the images shown for this operator are not directly comparable to the results obtained by Pattanaik et al [12], [33].

Revised Tumblin-Rushmeier. This global operator is essentially a power-law based on psychophysical data [11]. Like the previous method, the algorithm is calibrated in cd/m^2 . We linearly scaled the input data and normalized the output afterwards to produce what we believe the best possible results for this method in a practical setting.

Uniform rational quantization. This is another early operator which produces plausible results for many images [7]. The user parameter M was manipulated per image to produce reasonable results.

The images shown in Figs. 9 to 12 are fairly typical. In general, global operators tend to either appear washed-out or lose visible detail in the dark and/or light areas. Local operators tend to show the details better, but frequently do this at the cost of ringing or haloing artifacts. In our opinion, the method presented in this paper produces sensible results without obvious artifacts. It also allows strong color-casts to be removed should that be desirable.

With the exception of the iCAM color appearance model which addresses color fidelity in the context of high dynamic range data compression [35], the issue of color fidelity in tone reproduction has not received a great deal of attention. Many tone reproduction operators only compress the luminance channel, and apply the result to the three color channels in such a way that the color ratios before and after compression

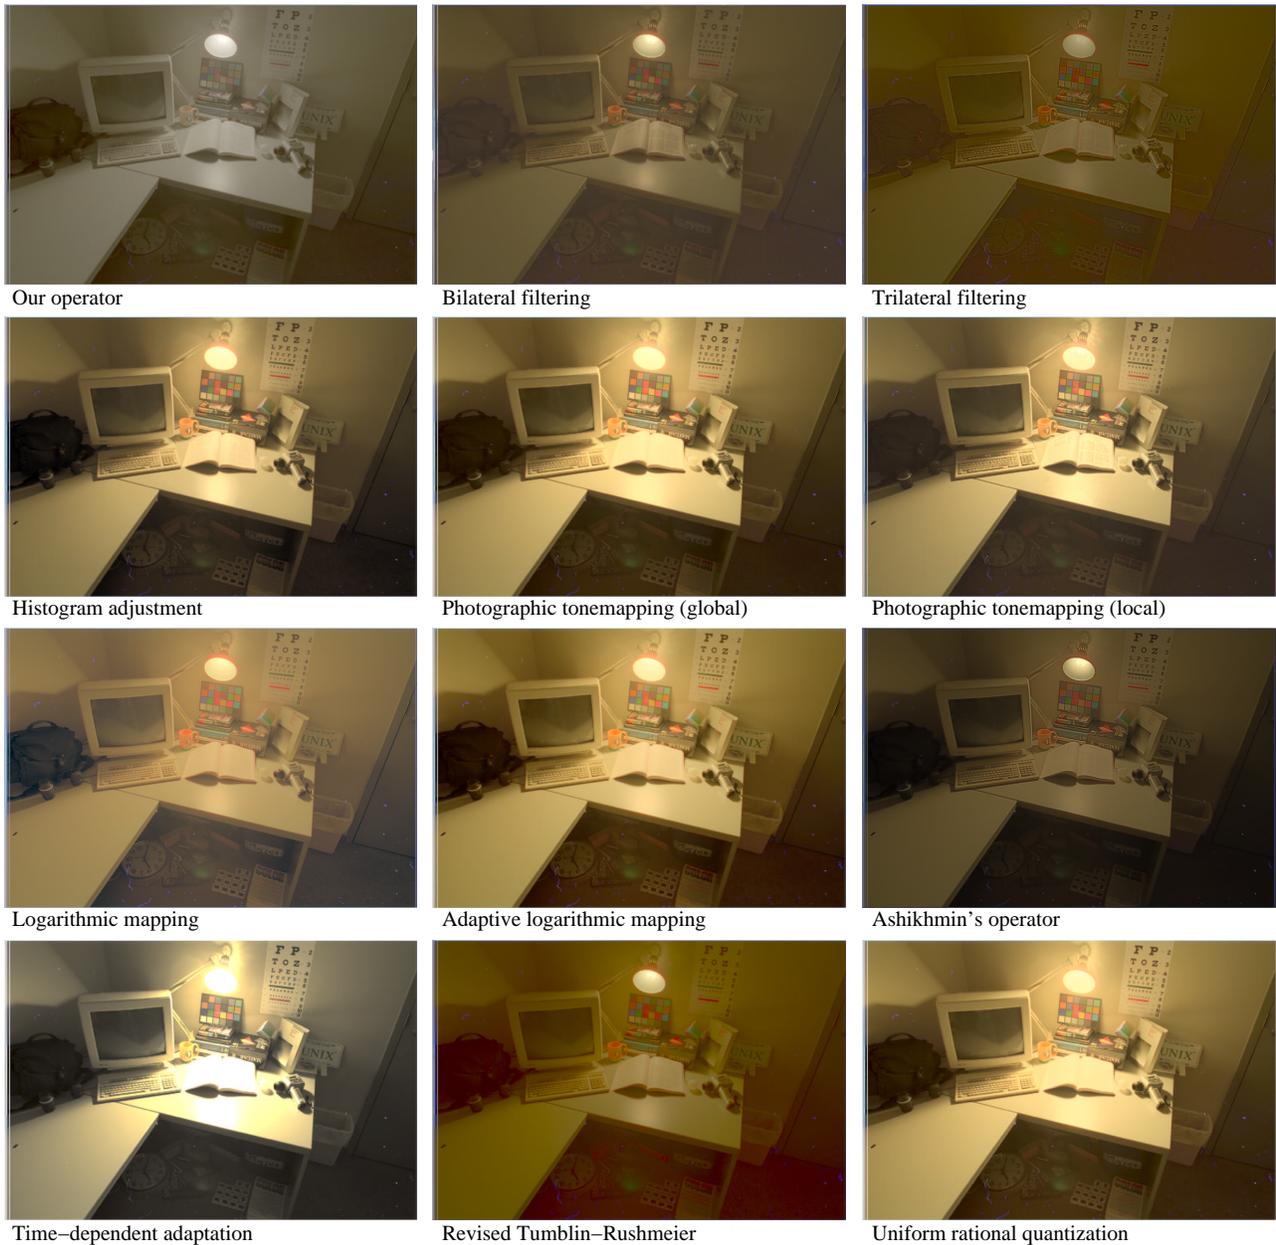


Fig. 9. Desk image.

are preserved [7]. Fattal et al [22] build upon this tradition by introducing an exponent s to control saturation. For the red channel, the display intensity R_d is then a function of the input intensity R_w , the pixel's luminance before and after compression (L_w and L_d respectively): $R_d = L_d (R_w/L_w)^s$. The green and blue channels are processed analogously. This is a reasonable first step but ignores the fact that color appearance varies with the overall intensity of the scene [36]. While our method does not address this issue either, in the absence of a satisfactory solution, we prefer to provide the user control over the amount of chromatic adaptation.

Our comparison is by no means exhaustive. There are many more research images available, as well as further tone reproduction techniques that we have not mentioned. However, we do believe that the above comparison is indicative of

the results one may expect from various tone reproduction operators, including the one presented in this paper.

While sigmoidal mapping functions are employed by others to describe aspects of vision [26], [29], and were later used in the field of tone reproduction [12], [33], [37] and color appearance modeling [38]–[40], we believe that its successful use in engineering applications strongly depends on the appropriate selection of tuning parameters. We have provided sufficient tools to shape the sigmoidal curve to suit most high dynamic range imagery.

An example response curve generated by our algorithm is compared with the global version of photographic tone mapping (using automatic parameter estimation) [14] in Figure 13. Our curve shows the response to a linear ramp using default parameter settings. While the shape of the curve

Fig. 10. *Parking garage.*

depends strongly on the average luminance in relation to the minimum and maximum luminance (compare with Figures 4 and 5), a general observation to be made is that our curve is generally shifted further to the right than the curve defined by photographic tone mapping. The curve also does not flatten out at the bright end if there are many bright pixels in the image, and therefore the average luminance is relatively high. This means that details in bright regions of bright scenes are preserved better than details in bright regions of dark scenes when the curve would in fact flatten out at the bright end. This desirable behavior is not as readily available with photographic tonemapping.

V. DISCUSSION

We envision tone reproduction operators to be used predominantly by photographers as well as in other artistic applications. It is therefore our aim to present a generally applicable operator that is fast and practical to use, and provides intuitive user parameters. We used findings from electro-physiology to motivate the design of our algorithm, but made engineering-based design decisions where appropriate. Experimentation with bilateral filtering and adaptive gain control techniques showed that the visual quality of our spatially varying operator is only marginally better than for our global operator. We therefore believe that for most practical applications, our fast global operator will suffice.

There are many criteria one might apply to compare qualitative results [41]–[43]. For instance, one could measure how

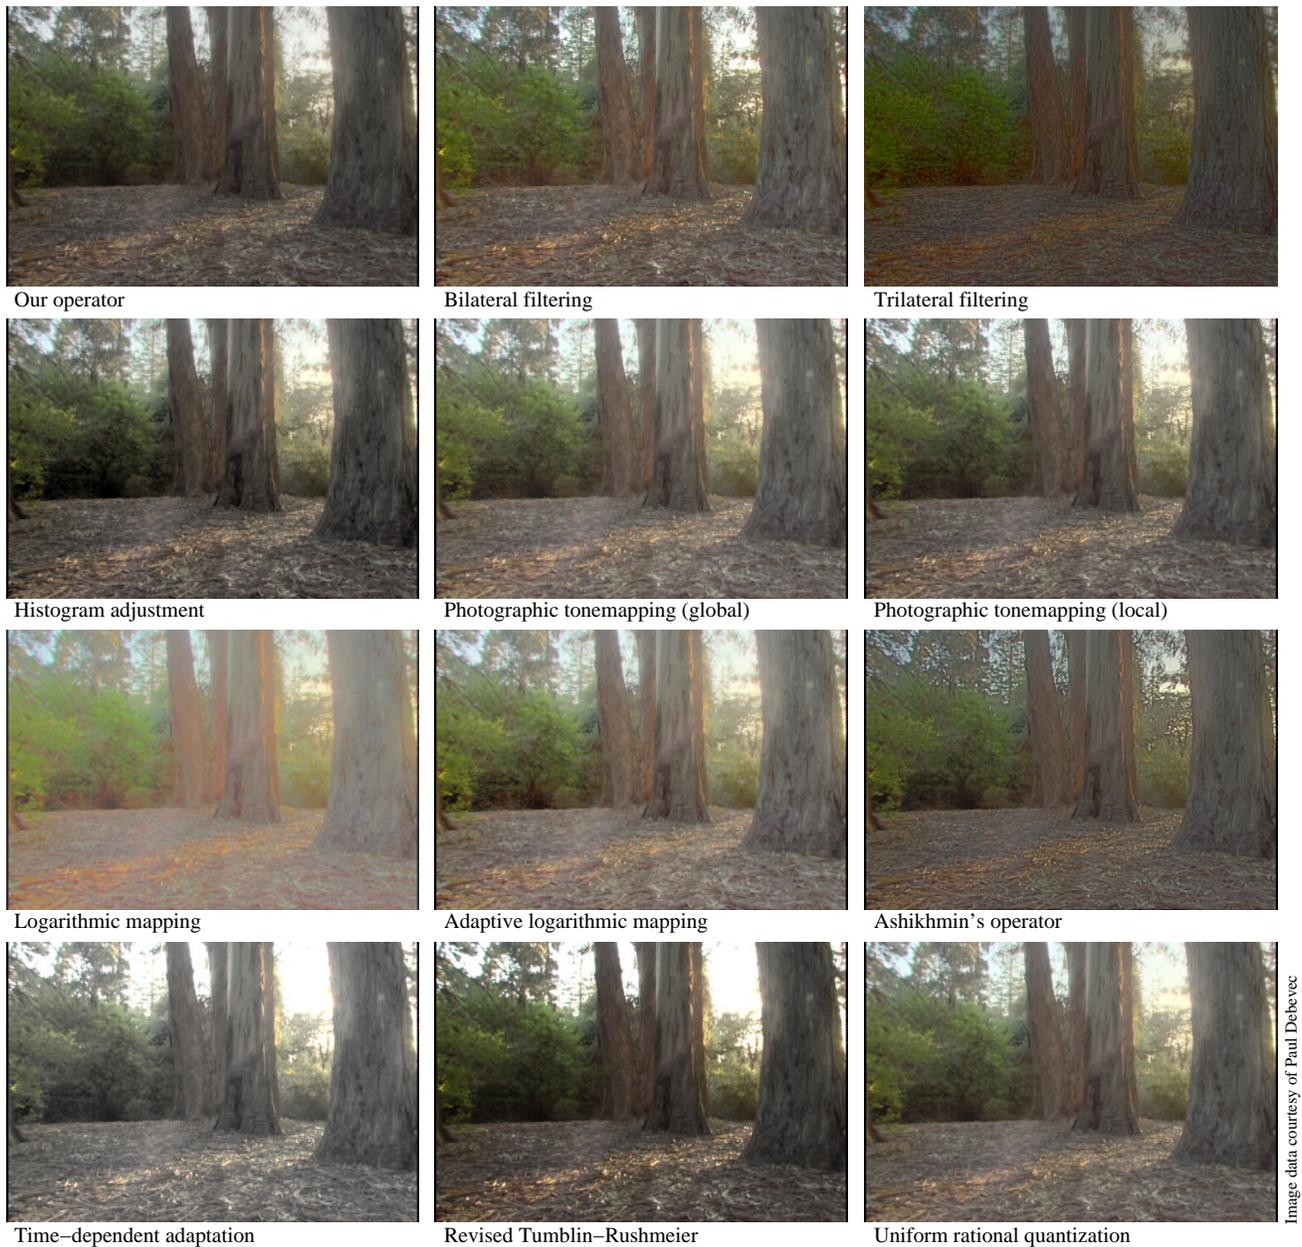


Fig. 11. Grove image.

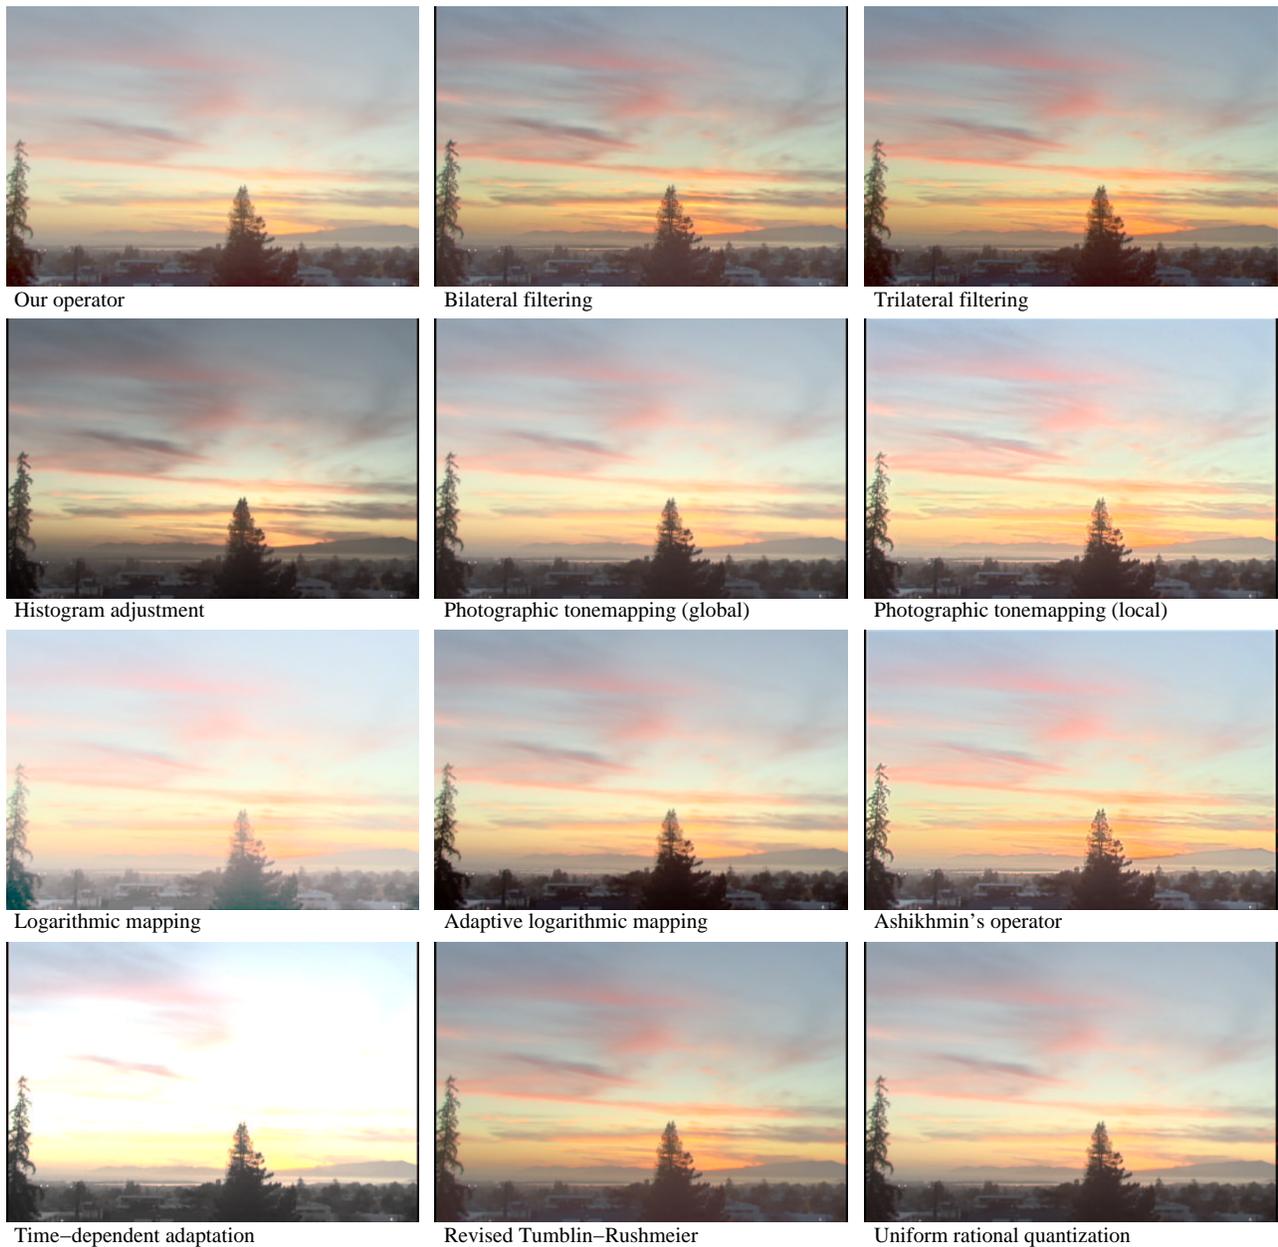
well details are preserved, or how well the method models certain aspects of the human visual system. These are all worthwhile criteria, but they also assume that tone reproduction operators will be used for specific applications, such as perhaps explaining visual phenomena. Validation of tone reproduction operators for specific tasks is a very necessary avenue of research that has yet to mature, although insight into this matter is beginning to accumulate [44]. For this reason, and because we aim for general applicability, we have used visual comparison to show qualitative results.

In the absence of straightforward validation techniques, judgment of operators is currently a matter of taste. In our opinion, the global operator presented in this paper produces visually appealing output for a wide variety of input. It shares speed-advantages with other global operators while

compressing images with a quality that in our opinion rivals local operators, albeit without any ringing artifacts. The method has four user parameters, each with sensible initial estimates that orthogonally control contrast, overall intensity, light- and chromatic adaptation, yielding a tone reproduction operator that is fast, easy to use and suitable for applications where plausible results is the main criterion for selection of a particular technique. We therefore believe that this algorithm will be a useful addition to the current collection of tone reproduction operators.

ACKNOWLEDGMENTS

We would like to thank Sumanta Pattanaik, Jim Ferwerda, Greg Ward and Paul Debevec for making their software and research images available. The memorial church, sun-

Fig. 12. *Sunset image.*

set, grove and rosette HDR data may be downloaded from <http://www.debevec.org/Research/HDR/>

REFERENCES

- [1] P. E. Debevec and J. Malik, "Recovering high dynamic range radiance maps from photographs," in *SIGGRAPH 97 Conference Proceedings*, ser. Annual Conference Series, August 1997, pp. 369–378.
- [2] H. Seetzen, L. A. Whitehead, and G. Ward, "A high dynamic range display using low and high resolution modulators," in *The Society for Information Display International Symposium*, May 2003.
- [3] K. Devlin, "A review of tone reproduction techniques," Computer Science, University of Bristol, Tech. Rep. CSTR-02-005, 2002.
- [4] J. DiCarlo and B. Wandell, "Rendering high dynamic range images," in *Proceedings of the SPIE Electronic Imaging 2000 conference*, vol. 3965, 2000, pp. 392–401.
- [5] N. J. Miller, P. Y. Ngai, and D. D. Miller, "The application of computer graphics in lighting design," *Journal of the IES*, vol. 14, pp. 6–26, October 1984.
- [6] J. Tumblin and H. Rushmeier, "Tone reproduction for computer generated images," *IEEE Computer Graphics and Applications*, vol. 13, no. 6, pp. 42–48, November 1993.
- [7] C. Schlick, "Quantization techniques for the visualization of high dynamic range pictures," in *Photorealistic Rendering Techniques*, P. Shirley, G. Sakas, and S. Müller, Eds. Springer-Verlag Berlin Heidelberg New York, 1994, pp. 7–20.
- [8] G. Ward, "A contrast-based scalefactor for luminance display," in *Graphics Gems IV*, P. Heckbert, Ed. Boston: Academic Press, 1994, pp. 415–421.
- [9] J. A. Ferwerda, S. Pattanaik, P. Shirley, and D. P. Greenberg, "A model of visual adaptation for realistic image synthesis," in *SIGGRAPH 96 Conference Proceedings*, August 1996, pp. 249–258.
- [10] G. Ward, H. Rushmeier, and C. Piatko, "A visibility matching tone reproduction operator for high dynamic range scenes," *IEEE Transactions on Visualization and Computer Graphics*, vol. 3, no. 4, 1997.
- [11] J. Tumblin, J. K. Hodgins, and B. K. Guenter, "Two methods for display of high contrast images," *ACM Transactions on Graphics*, vol. 18 (1), pp. 56–94, 1999.
- [12] S. N. Pattanaik, J. Tumblin, H. Yee, and D. P. Greenberg, "Time-

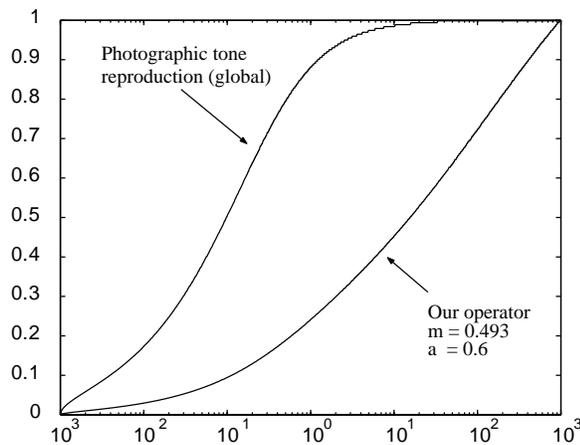


Fig. 13. Mapping of input luminances (horizontal) to display luminances (vertical) for a typical curve obtained by our operator compared with the sigmoid produced by photographic tone reproduction [13].

- dependent visual adaptation for fast realistic display,” in *SIGGRAPH 2000 Conference Proceedings*, July 2000, pp. 47–54.
- [13] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, “Photographic tone reproduction for digital images,” *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 267–276, 2002.
- [14] E. Reinhard, “Parameter estimation for photographic tone reproduction,” *Journal of Graphics Tools*, vol. 7, no. 1, pp. 45–51, 2003.
- [15] F. Drago, K. Myszkowski, T. Annen, and N. Chiba, “Adaptive logarithmic mapping for displaying high contrast scenes,” *Computer Graphics Forum*, vol. 22, no. 3, 2003.
- [16] S. N. Pattanaik, J. A. Ferwerda, M. D. Fairchild, and D. P. Greenberg, “A multiscale model of adaptation and spatial vision for realistic image display,” in *SIGGRAPH 98 Conference Proceedings*, July 1998, pp. 287–298.
- [17] M. Ashikhmin, “A tone mapping algorithm for high contrast images,” in *Proceedings of 13th Eurographics Workshop on Rendering*, 2002, pp. 145–155.
- [18] S. Upstill, “The realistic presentation of synthetic images: Image processing in computer graphics,” Ph.D. dissertation, University of California at Berkeley, August 1985.
- [19] F. Drago, W. L. Martens, K. Myszkowski, and N. Chiba, “Design of a tone mapping operator for high dynamic range images based upon psychophysical evaluation and preference mapping,” in *IS&T SPIE Electronic Imaging 2003. The human Vision and Electronic Imaging VIII Conference*, 2003.
- [20] F. Durand and J. Dorsey, “Fast bilateral filtering for the display of high-dynamic-range images,” *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 257–266, 2002.
- [21] P. Choudhury and J. Tumblin, “The trilateral filter for high contrast images and meshes,” in *Proceedings of the Eurographics Symposium on Rendering*, 2003, pp. 186–196.
- [22] R. Fattal, D. Lischinski, and M. Werman, “Gradient domain high dynamic range compression,” *ACM Transactions on Graphics*, vol. 21, no. 3, pp. 249–256, 2002.
- [23] J. E. Dowling, *The Retina: An Approachable Part of the Brain*. Harvard University Press, 1987.
- [24] D. A. Baylor and M. G. F. Fuortes, “Electrical responses of single cones in the retina of the turtle,” *J Physiol*, vol. 207, pp. 77–92, 1970.
- [25] J. Kleinschmidt and J. E. Dowling, “Intracellular recordings from gecko photoreceptors during light and dark adaptation,” *J gen Physiol*, vol. 66, pp. 617–648, 1975.
- [26] K. I. Naka and W. A. H. Rushton, “S-potentials from luminosity units in the retina of fish (cyprinidae),” *J Physiol*, vol. 185, pp. 587–599, 1966.
- [27] R. A. Normann and I. Perlman, “The effects of background illumination on the photoresponses of red and green cones,” *J Physiol*, vol. 286, pp. 491–507, 1979.
- [28] L. Spillmann and J. S. Werner, Eds., *Visual Perception: the Neurological Foundations*. San Diego, CA: Academic Press, 1990.
- [29] D. C. Hood, M. A. Finkelstein, and E. Buckingham, “Psychophysical tests of models of the response function,” *Vision Res*, vol. 19, pp. 401–406, 1979.

- [30] B. Baxter, H. Ravindra, and R. A. Normann, “Changes in lesion detectability caused by light adaptation in retinal photoreceptors,” *Invest Radiol*, vol. 17, pp. 394–401, 1982.
- [31] K. Chiu, M. Herf, P. Shirley, S. Swamy, C. Wang, and K. Zimmerman, “Spatially nonuniform scaling functions for high contrast images,” in *Proceedings of Graphics Interface '93*, May 1993, pp. 245–253.
- [32] G. Wyszecki and W. S. Stiles, *Color science: Concepts and methods, quantitative data and formulae*, 2nd ed. New York: John Wiley and Sons, 1982.
- [33] S. N. Pattanaik and H. Yee, “Adaptive gain control for high dynamic range image display,” in *Proceedings of Spring Conference in Computer Graphics (SCCG2002)*, Budmerice, Slovak Republic, 2002.
- [34] G. Ward Larson and R. A. Shakespeare, *Rendering with Radiance*. Morgan Kaufmann Publishers, 1998.
- [35] M. D. Fairchild and G. M. Johnson, “Meet iCAM: an image color appearance model,” in *IS&T/SID 10th Color Imaging Conference*, Scottsdale, 2002, pp. 33–38.
- [36] M. D. Fairchild, *Color appearance models*. Reading, MA: Addison-Wesley, 1998.
- [37] A. Chalmers, E. Reinhard, and T. Davis, Eds., *Practical Parallel Rendering*. A K Peters, 2002.
- [38] R. Hunt, *The reproduction of color*. England: Fountain Press, 1996, fifth edition.
- [39] CIE, “The CIE 1997 interim colour appearance model (simple version), CIECAM97s,” CIE Pub. 131, Vienna, Tech. Rep., 1998.
- [40] N. Moroney, M. D. Fairchild, R. W. G. Hunt, C. J. Li, M. R. Luo, and T. Newman, “The CIECAM02 color appearance model,” in *IS&T 10th Color Imaging Conference*, Scottsdale, 2002, pp. 23–27.
- [41] S. Daly, “The visible difference predictor: an algorithm for the assessment of image fidelity,” in *Digital Images and Human Vision*, A. B. Watson, Ed. MIT Press, 1993, pp. 179–206.
- [42] H. Rushmeier, G. Ward, C. Piatko, P. Sanders, and B. Rust, “Comparing real and synthetic images: Some ideas about metrics,” in *Proceedings of the Sixth Eurographics Workshop on Rendering*, Dublin, Ireland, June 1995, pp. 213–222.
- [43] A. Chalmers, S. Daly, A. McNamara, K. Myszkowski, and T. Trscianko, “Image quality metrics,” July 2000, siggraph course 44.
- [44] F. Drago, W. L. Martens, K. Myszkowski, and H.-P. Seidel, “Perceptual evaluation of tone mapping operators with regard to similarity and preference,” Max Plank Institut für Informatik, Tech. Rep. MPI-I-2002-4-002, 2002.



Erik Reinhard is assistant professor at the University of Central Florida and has an interest in the fields of visual perception and parallel graphics. He has a B.Sc. and a TWAIO diploma in computer science from Delft University of Technology and a Ph.D. in computer science from the University of Bristol. He was a post-doctoral researcher at the University of Utah. He is founder and co-editor-in-chief of *ACM Transactions on Applied Perception*.



Kate Devlin has a B.A. Hons. in archaeology and an M.Sc. degree in computer science from Queen's University, Belfast. She is currently a Ph.D. student in computer graphics at the University of Bristol. Her interest is in the perceptually accurate rendering and display of synthetic images, and its application to the field of archaeology.