EUROGRAPHICS 2008 / G. Drettakis and R. Scopigno
(Guest Editors)

Volume 27 (2008), Number 2

Modeling a Generic Tone-mapping Operator

Rafał Mantiuk and Hans-Peter Seidel

Max-Plack-Insitut für Informatik
Saarbrücken, Germany

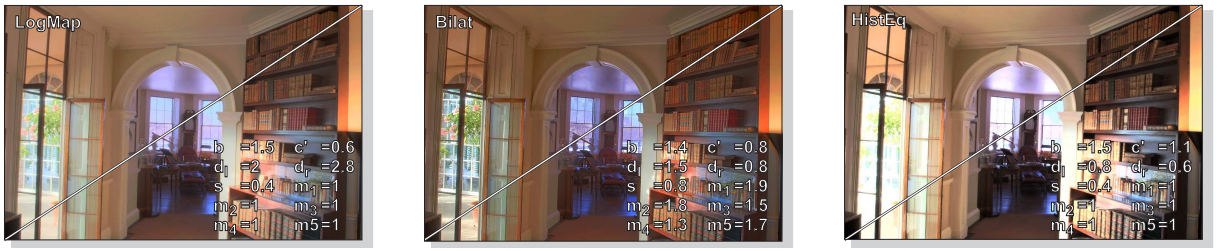

Figure 1: *The upper-left half of each image has been tone-mapped using three popular operators (see Table 1 for reference), while the lower-right half of each image is the result of the same generic tone-mapping operator (TMO). The parameters of the generic TMO may be adjusted to mimic a broad range of operators.*

Abstract

Although several new tone-mapping operators are proposed each year, there is no reliable method to validate their performance or to tell how different they are from one another. In order to analyze and understand the behavior of tone-mapping operators, we model their mechanisms by fitting a generic operator to an HDR image and its tone-mapped LDR rendering. We demonstrate that the majority of both global and local tone-mapping operators can be well approximated by computationally inexpensive image processing operations, such as a per-pixel tone curve, a modulation transfer function and color saturation adjustment. The results produced by such a generic tone-mapping algorithm are often visually indistinguishable from much more expensive algorithms, such as the bilateral filter. We show the usefulness of our generic tone-mapper in backward-compatible HDR image compression, the black-box analysis of existing tone mapping algorithms and the synthesis of new algorithms that are combination of existing operators.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation Display algorithms; I.4.2 [Image Processing and Computer Vision]: Enhancement Greyscale manipulation, sharpening and deblurring

1. Introduction

In recent years the problem of tone-mapping has attracted much attention and several dozens of tone mapping algorithms have been proposed. However, it is still disputable how to validate, analyze performance or quality, or simply benchmark tone mapping algorithms. The most recent operators tend to exhibit increased complexity and more sophisticated image processing algorithms, but it is not clear how much quality can be gained by this additional cost. What constitutes the improvement of the new tone mapping operators (TMOs) over the state-of-the-art? And finally, how different are tone mapping operators from each other?

In this paper we attempt to address some of these questions by modeling the processing that is performed inside a TMO. In our black-box approach, a generic tone-mapping algorithm is fit to a high-dynamic range (HDR) image and its tone-mapped low-dynamic range (LDR) rendering. We demonstrate that many tone-mapping operators, both global and local, can be satisfactorily approximated by computationally inexpensive image processing operations, such as a

per-pixel tone curve, modulation transfer function and color saturation adjustment. As shown in Figure , the results produced by our a generic tone-mapping algorithm are often visually indistinguishable from much more expensive algorithms, such as the bilateral filter.

We introduce a generic tone-mapping operator, that can closely approximate other operators in Section 3. Then, we show the usefulness of such a generic tone-mapper in the black-box analysis of existing tone mapping algorithms (Section 4), in backward-compatible HDR image compression (Section 5), and the synthesis of new algorithms, which are a combination of existing operators (Section 6).

2. Previous Work

This work is complementary to tone-mapping and inverse tone-mapping [RTS*07], in the sense that instead of computing LDR image from HDR or HDR image from LDR, we attempt to find an unknown tone-mapping operator. It also offers methods to objectively and quantitatively compare tone-mapping operators, a task that was usually the domain of subjective quality studies.

Typically, tone-mapping algorithms produce a display-ready (low-dynamic range) image based on a linear radiance map, often of high dynamic range (HDR). This challenging problem can be dated back to gray-scale photography or even to the Renaissance painters, as pointed out in [McC07]. Therefore, it is no surprise that some tone-mapping operators (TMO) borrow ideas from photography, such as dodging and burning [RSSF02], unsharp masking [DD02], or are inspired by arts where such techniques as counter-shading [KMS07] had been known for centuries. There have been also attempts to derive tone-mapping from subjective studies on human preference for contrast, brightness and color saturation [YMMS06]. What is common for these and many other TMOs, [WLRP97, DMAC03, MMS06, LFUS06], is that they aim to *maximize the subjective image quality* in terms of producing pleasing images. This goal is quite different from the other major TMO objective, proposed in computer graphics by Tumblin and Rushmeier [TR93], which aims at rendering a displayable image that would be perceived as close to the original scene as possible. This goal of *perceptual-match reproduction* was pursued by others, [FPSG96, PFFG98, PTYG00, KMS05], who employed the models of the human visual system to maintain the original scene appearance in tone-mapped images. These are only some of the several dozens of TMOs that have been proposed in the last 20 years. The most complete review of recent TMOs, although lacking those proposed after 2005, can be found in [RWPD05, chapters 6–8].

The large number of proposed TMOs motivated studies on their analysis, comparison and validation. Several subjective studies attempt to rank TMOs and analyze their effect on perceived image contrast, brightness, detail and artifact visibility, and overall quality and naturalness [DMMS02, KYJF04, LCTS05, CWNA06]. Such studies, due to tedious experimental procedures, focus on the results produced with default parameters, while ignoring the vast diversity of results that the same operators can produce when their parameters are varied. They also rarely make distinction between the TMOs that *maximized the subjective image quality* and aimed at *the perceptual-match reproduction*. Since natural-looking images are not always preferred images, it is quite expectable that a perceptual-match TMO will be ranked lower than a subjective-quality oriented TMO, while both tone-mappers can be equally suitable for their distinct applications. A larger number of factors and images could be considered given an objective metric. Smith et al. proposed an objective metric for measuring perceived global contrast and detail visibility change in tone mapped images [SKMS06] and compared several popular operators. In contrast to Smith et al.'s approach, we do not attempt to model perceived attributes, which are difficult to justify and assess even with extensive studies, but rather offer methods to quantitatively compare TMO characteristics and analyze their variability with different parameter settings and different image content.

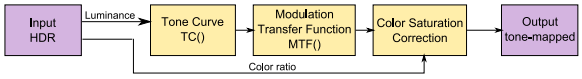

Figure 2: *Data-flow for the generic TMO model.*

3. Modeling generic TMO

Given a pair of images comprised of an HDR image and its tone-mapped LDR counterpart, we want to find a tone-mapping operator that is based on a simple model and is controlled with as few parameters as possible. Such a model, however, must be general enough to approximate a broad range of tone-mapping operators.

A typical approach to modeling complex systems, such as tone-mapping operators, usually involves graphical visualization of the system behavior, expert knowledge of the underlying processes and finally, a trial-and-error selection of the best model from possible candidates. Following this procedure we developed several models, from which we chose the one that gives the best fit and is stable. Our best performing model consists of three components: tone curve, modulation transfer function and color saturation correction, as illustrated in Figure 2. The result for a single color component (red, green or blue) is given by:

$$C_{TMO} = MTF(TC(L_{HDR})) \cdot R^s \qquad (1)$$

where $MTF()$ is the modulation transfer function, $TC()$ is a tone curve applied to each pixel separately, L_{HDR} is the luminance of the input HDR pixels, s is color saturation pa-

rameter and R is a color ratio, given by:

$$R = \frac{C_{HDR}}{L_{HDR}} \tag{2}$$

where C_{HDR} is a color component of the input HDR image. The above model lets us use a single tone curve for all color channels, and makes any necessary color adjustments using only the single parameter s. This approach is typically employed in tone-mapping [Sch94, TT99].

3.1. Tone curve

To find the most suitable model of a tone curve, we perform Principal Component Analysis (PCA) on a set of tone curves. These curves are the result of tone mapping 5 images using 12 different tone mapping operators. The first five significant principal components are shown in Figure 3. The first principal component clearly indicates that the S-shaped curve is prevalent in tone mapping, while the remaining principal components mostly modulate luminance compression in different parts of the tone-scale. There are also several practical reasons for an S-shaped tone curve: it distorts the contrast in the middle-gray range the least, which usually contains the most important part of the scene; typical images contain more pixels in the middle-gray range than in highlights and shadows; analog film has an S-shaped response; and photoreceptors are also found to have an S-shaped response. Low eigenvalues above the fifth principal components suggest that about 5 parameters should suffice to describe a tone curve.

We do not use actual principal components in our generic tone operator, as these are too abstract and difficult to analyze. Instead, we use a four-segment sigmoidal function:

$$TC(L_{HDR}) = \begin{cases} 0 & \text{if } L' \le b - d_l \\ \frac{1}{2}\, c \dfrac{L'-b}{1-a_l(L'-b)} + \frac{1}{2} & \text{if } b - d_l < L' \le b \\ \frac{1}{2}\, c \dfrac{L'-b}{1+a_h(L'-b)} + \frac{1}{2} & \text{if } b < L' \le b + d_h \\ 1 & \text{if } L' > b + d_h \end{cases} \tag{3}$$

where L' is the logarithm of luminance ($L' = log_{10}(L_{HDR})$), b is the image *brightness* adjustment parameter, c is the *contrast* parameter, and a_l, a_h decide on the contrast compression for shadows and highlights. Note that the *brightness* parameter is not an estimate of the perceived image brightness but a relative adjustment factor. For more intuitive control, the a_l, a_r parameters can be replaced by:

$$a_l = \frac{c \cdot d_l - 1}{d_l} \qquad a_h = \frac{c \cdot d_h - 1}{d_h} \tag{4}$$

where d_l is the lower midtone range and d_h is the higher midtone range. The curve is C^1-continuous everywhere except $b - d_l$ and $b + d_h$. To analyze actual contrast change, it is more convenient to express the contrast parameter c as the slope of an effective tone curve (displayed luminance instead of pixel values) on a log-log plot:

$$c' = \gamma\, c / log(10) \tag{5}$$

where γ is the display gamma, usually equal to 2.2. An intuitive illustration of the tone curve parameters is shown in Figure 4.

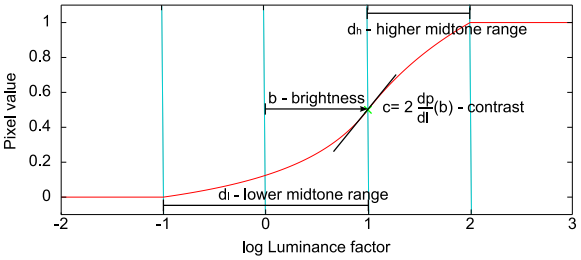

Figure 4: *Tone curve used in the generic TMO and its parameters.*

The tone function given in Equation 3 is one of multiple choices for tone curve parametrization. We choose it mostly because it offers intuitive parameters and can approximate both S-shaped sigmoidal tone curves used in tone-mapping and J-shaped gamma curves found in displays and cameras. The low number of parameters may result in imperfect fits for some global TMOs. We found, however, that such imperfect fits rarely result in visibly different results.

3.2. Spatial Modulation

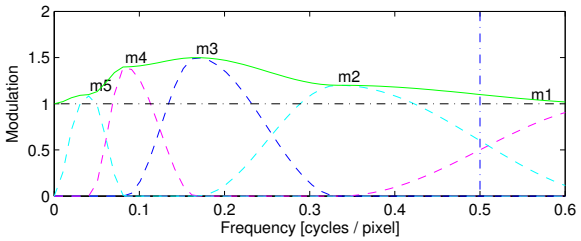

Figure 5: *Modulation transfer function (solid green) as a linear combination of five band-pass filters (dashed lines). m1–m5 are the modulations of each filter.*

The tone curve and color saturation correction is sufficient to model most of the global tone mapping operators. To model spatially-varying (local) TMOs, we need to introduce an additional spatial operator. We choose the modulation transfer function (MTF), which is commonly used to model optical systems. The MTF is usually depicted as a 1D function of spatial frequency, such as the one shown in Figure 5 (solid green line). It specifies which spatial frequencies to amplify or compress and can be thought of as a selection of low- and high-pass filters that can sharpen or blur an image as needed.

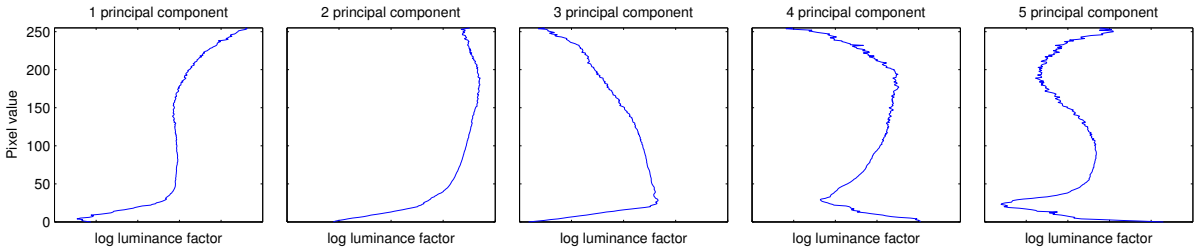

Figure 3: *Principal component for the set of typical tone curves. The first principal components resembles the S-shaped curve, commonly used in tone operators.*

#	Label	Name	Reference
1.	LogMap	Adaptive Logarithmic Mapping For Displaying High Contrast Scenes	[DMAC03]
2.	Photo	Photographic Tone Reproduction for Digital Images	[RSSF02]
3.	VisAdapt	A Model of Visual Adaptation for Realistic Image Synthesis	[FPSG96]
4.	GradDom	Gradient Domain High Dynamic Range Compression	[FLW02]
5.	BriAda	Tone Reproduction for Realistic Images	[TR93]
6.	HistEq	A Visibility Matching Tone Reproduction Operator for High Dynamic Range Scenes	[WLRP97]
7.	Retinex	Retinex adapted to tone-mapping	[DMMS02, p.8]
8.	ContMap	A perceptual framework for contrast processing of high dynamic range images	[MMS06, sec.4]
9.	Bilat	Fast Bilateral Filtering for the Display of High-Dynamic-Range Images	[DD02]
10.	Ashik	A Tone Mapping Algorithm for High Contrast Images	[Ash02]
11.	RQuant	Quantization Techniques for the Visualization of High Dynamic Range Pictures	[Sch94]
12.	ContEq	as in 8.	[MMS06, sec.5]
13.	PSLA	Photoshop TMO: Local adaptation	n/a
14.	PSHC	Photoshop TMO: Highlight compression	n/a
15.	TimeAdapt	Time-dependent Visual Adaptation For Fast Realistic Image Display	[PTYG00]

Table 1: *Tone mapping operators used in this study and their labels.*

To reduce the number of model parameters, we do not use a continuous MTF, but a linear combination of five parameters and five basis functions. We choose the filters used in the Cortex transform [Wat87] with the modifications from [Dal93] as the basis functions. The filters become wider for higher frequencies, which is consistent with both the findings on spatial selective pathways in the human visual system and with the distribution of energy in natural scenes (higher frequencies contain lower energy, which is compensated by wider filters). We use only five band-pass filters and do not modulate the base band, since the base band is adjusted by the tone curve. The band-pass filters used in our generic tone-mapping model are visualized as dashed lines in Figure 5 and as dotted lines on each plot in the right panes of Figure 6. The choice of particular basis functions is not essential in our application, thus the Cortex transform could be also replaced with differences of Gaussians (DoG) or wavelets.

Although the MTF cannot fully substitute all local image operations used in local TMOs, it gives a surprisingly good approximation of them, as we show in Figure and discuss in Section 4.

3.3. Fitting procedure

To find the parameters of the generic TMO efficiently, we split the fitting procedure into two parts: first we fit the tone curve and find the saturation correction parameter s using the Levenberg-Marquardt method:

$$\arg\max_{b,c,d_l,d_h,s} \sum_{k=1,2,3} \left| C_{LDR} - TC(C_{HDR}; b, c, d_l, d_h) \cdot R^s \right|^2 \quad (6)$$

where C_{LDR} is the input LDR image and k is the index of a color channel (red, green and blue). For simplicity we skip the summation over pixels. Then, we find the five parameters of the MTF by solving a linear least-squares problem:

$$\arg\max_{m1,..,m5} \left| HP\left[mean\left(C_{LDR}/R^s\right)\right] - MTF(L_{HDR}; m1,..,m5) \right|^2 \quad (7)$$

where $m1,..,m5$ are MTF coefficients and the *mean* function averages values across color channels. $HP[]$ is the high-pass filter that removes the base-band, which is also not included in the MTF function (because the base band is modified by the tone curve). We exclude the pixels from the summation that are clipped in the tone-mapped image C_{LDR}.

We compared our results with the complete solution, where we minimized for all parameters at once, and did not notice any significant improvement. The two part procedure

can find the best fit in less than 15 seconds for a half-megapixel image. This timing can be further improved when only a subset of pixels is used for fitting the tone curve.

4. Analysis of tone-mapping operators

We fit our generic tone-mapping operator to the results of 12 popular TMOs and several images. Wherever possible, we use the images generated by the authors of the operators that are available on Internet [tmo], thus avoiding problems with unfaithful implementations and reducing bias from parameter selection. For unavailable tone-mapped images, we run the tone-mapping algorithms using the default parameters and the `pfstmo` software [pfs] or the tone-mappers bundled with the book [RWPD05]. Table 1 lists all TMOs that have been used in this study, together with the labels used as references.

Based on the parameters of the generic TMO, we can visualize the characteristics of a dozen TMOs in a unified manner in Figure 6. The visualization for each image pair consists of three parts: an original tone-mapped image and the result of the best fit; tone curve and scatter plot of pixel luminance values; and the modulation transfer function. We also indicate how close the fit is to the tone-mapped image in terms of the peak-signal-to-noise ratio (PSNR) and the structural similarity index (SSIM) [WBSS04]. A more extensive set of results together with full-resolution images can be found at `http:http://www.mpi-inf.mpg.de/resources/hdr/generic_tmo/`.

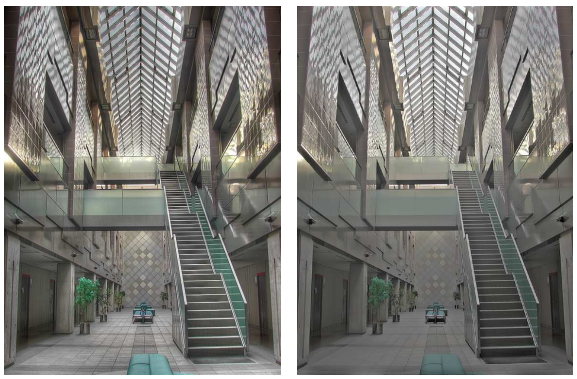

Figure 7: *Example of a TMO result that does not give a good fit to the generic TMO. The original result of* GradDom *(left) is visibly sharper and has boosted shadows and highlights.*

For most tone-mapping operators and images, the fit of the generic TMO is surprisingly good, with the SSIM index [WBSS04] close to one (SSIM=1 denotes highest quality), and results that cannot be distinguished from the original TMO. Some examples of nearly perfect fits are shown in Figure . The fit is, however, worse for the TMOs that employ strongly non-linear image processing, such as gradient compression in the case of the *GradDom* and *ContEq* operators (refer to Table 1). An example of a poor fit is shown in Figure 7. Even though a plain MTF cannot produce as strong a local contrast boost as these highly non-linear operators, the resulting approximations are still useful in such applications as backward-compatible HDR video compression, discussed in Section 5. Another limitation of the model are the operators that employ user-assisted or automatic segmentation and apply radically different TMOs to each region of an image, such as [KMS05, LFUS06]. We excluded such operators from our study.

Figure 6 offers many insights into tone-mapping operators, regardless of their complexity and without any need to analyze the underlying image processing operations. Local and global TMOs can easily be distinguished since local operators have a scattered LDR/HDR luminance plot (the middle panels in Figure 6) and amplified medium and high frequencies on the MTF plot (the right panel). The stronger the sharpening effect, the higher the boost of the medium and high frequencies. Most of the tone-mapping operators compress the low frequency (base band) contrast, with the compression ratio from $c' = 0.30$, resulting in relatively flat-looking images, up to the ratio $c' = 0.88$, resulting in high-contrast images but also clipping of brightest and darkest pixels. The exception is the *HistEq* operator ($c' = 1.28$), which is based on histogram equalization, which causes contrast amplification in the well-represented medium gray-levels. The majority of the TMOs have a tendency to allocate a higher dynamic range for the lower midtones ($d_l > d_h$), resulting in lower contrast for shades than for bright regions (refer also to Figure 4). If this proportion is reversed, for example as in the *Ashik* operator, the shades are brightened with a characteristic flash-like fill-in effect. The color saturation compensation is similar for most TMOs ($s \approx 0.45$).

We can use our fitting procedure to investigate the TMOs on which no information is available. To demonstrate this, we tone-map several images using Adobe® Photoshop's® CS2 *local adaptation* and *highlight compression* methods. The fit results shown in Figure 8 indicate that the highlight compression method is a typical global TMO based on a sigmoidal tone curve. The local adaptation method is a local TMO with a moderate sharpening effect, as shown on the MTF plot.

Tone-mapping operators offer two more dimensions to explore: image content and tone-mapping parameters. To analyze them, we compute fits for the *Bilat* TMO. We chose the *Bilat* because it is a local TMO with interesting characteristic and is well approximated with the generic TMO. The results of fitting the *Bilat* TMO to 80 images at four different settings of the base-layer contrast compression parameter are shown in Figure 9. The most pronounced effect of image content is the horizontal shift of the tone curve (the image brightness parameter), but also the noticeable change of contrast. The value of the base-layer contrast compres-

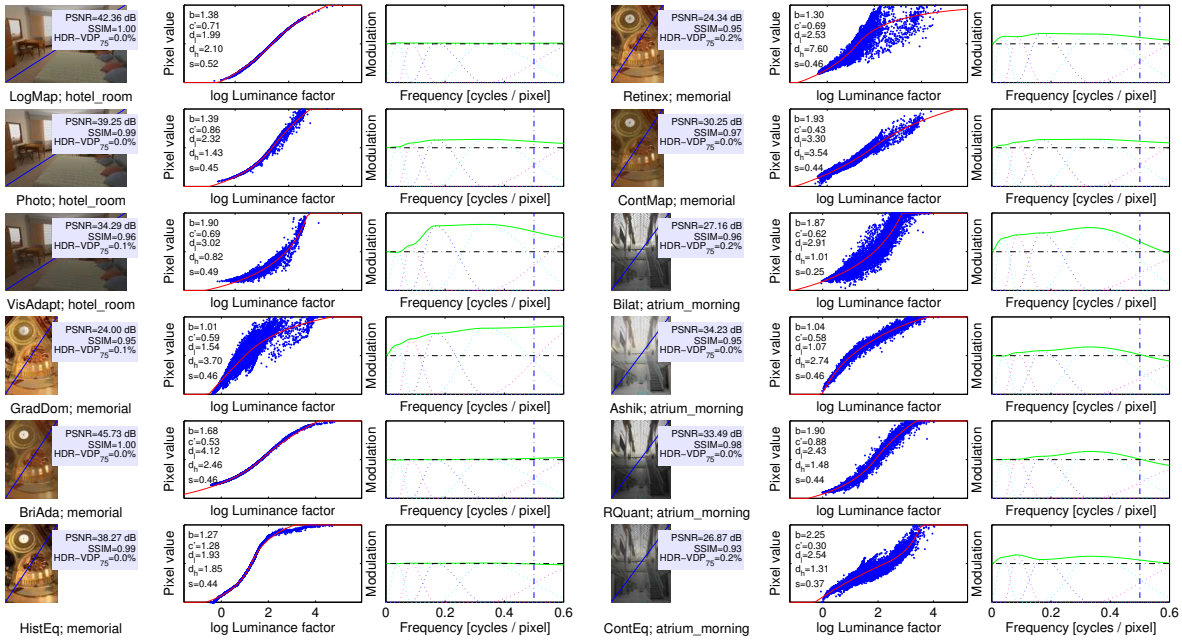

Figure 6: *Results of fitting several tone-mapped images to the generic operator. For each TMO: left – the original tone-mapped image (left-top) and its fit with the generic TMO (right-bottom); middle – fitted tone curve (red) and the scatter plot showing the relation between LDR/HDR luminance values; right – MTF of the fitted generic TMO. The blue vertical line indicates the Nyquist frequency, the black horizontal – the value 1 (no change), and the dotted curves – the base functions of the MTF.*

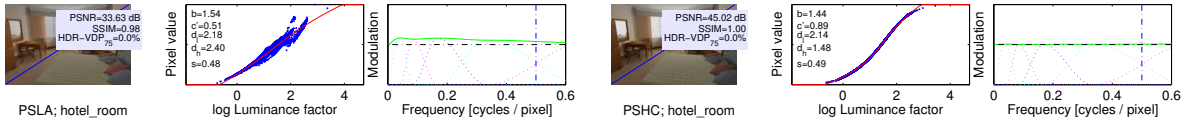

Figure 8: *Results of the fitting procedure for Adobe Photoshop's CS2 tone mapping algorithms.*

sion parameter elevates MTF coefficients. Since the overall MTF coefficients also change from image to image, we can infer that the *Bilat* TMO sharpens an image adaptively depending on its content. The effect of the base-layer contrast compression parameter can be better observed in Figure 10. The parameter twists the tone curve by affecting the contrast parameter c and elevates the MTF. The magnitude of the MTF elevation depends on image content.

The presented examples of the analysis may seem trivial for well-known tone-mapping operators, but they may be very useful for complex TMOs that cannot be easily understood by analyzing their formulas or the resulting images. Such analysis may also be applied to unknown optical and digital systems, such as digital cameras, where faithful modeling of the entire system is not always practical.

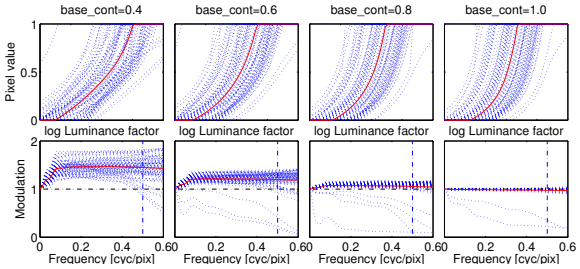

Figure 9: *Results of fitting* Bilat *at 4 TMO parameter settings (columns) to 80 images. Top: tone curves; bottom: MTFs. Red-continuous lines: averaged over images; Blue-dotted lines – separate results for each image.*

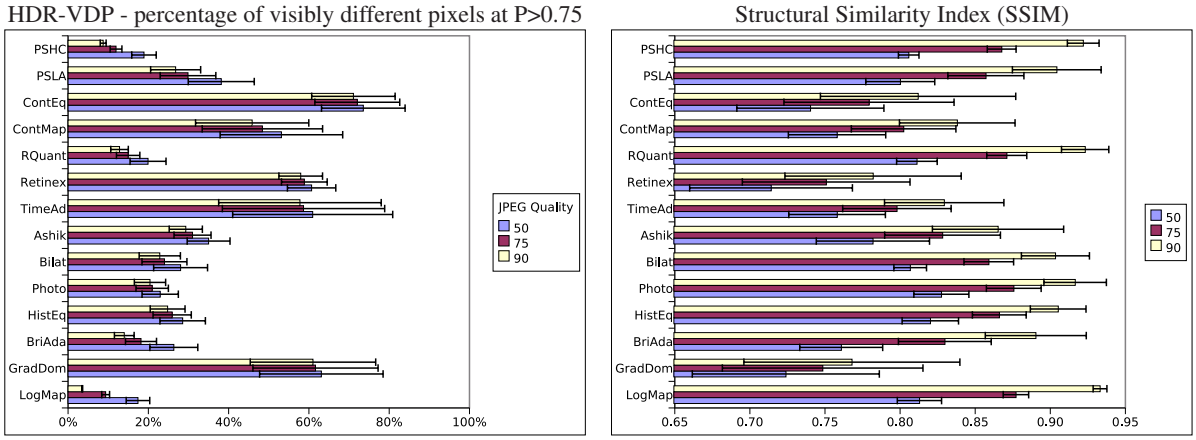

Figure 11: *Distortion measures for the HDR images reconstructed from their tone-mapped counterparts. Color bars denote JPEG compression quality: 50%, 75% or 90%, where 100% is the best quality. For the HDR-VDP lower values denote better quality. For the SSIM higher values denote higher quality.*

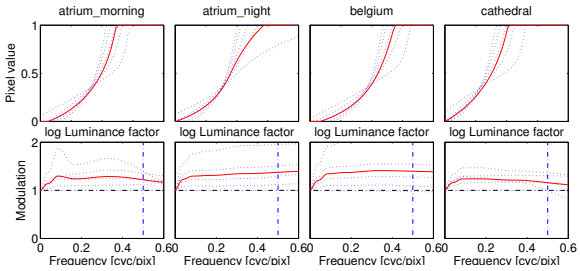

Figure 10: *Results of fitting* Bilat *at 5 TMO parameter settings to 4 images (columns). The notation is the same as in Figure 9.*

5. HDR Image Compression

One potential application of the generic TMO is backward-compatible compression of HDR images. A tone-mapped image is encoded using an ordinary JPEG compression, but also supplied with the 10 parameters of the generic TMO. Since the generic TMO is fully reversible, we can use these parameters to reverse tone-map and reconstruct the original HDR image. Such encoding results in very small overhead (only ten coefficients for reconstructing the HDR image) and can adapt to almost any TMO. This is however also a lossy compression method with three sources of possible distortions: lossy compression of the tone-mapped image, inaccurate fitting of the generic TMO, and quantization and clipping errors due to tone mapping. Most of these distortions can be reduced if we additionally encode a small residual image containing the differences between the original and the reconstructed HDR image, as proposed in [WS04] and [MEMS06].

Assuming that no residual image is used, we investigate which tone-mapping algorithms are the most suitable for this kind of HDR image compression. We reconstruct HDR images from their tone-mapped versions after they have been distorted by JPEG at three compression quality levels. This reconstruction is conducted for 14 tone-mapping operators, and the results are averaged over a set of 5 images. The results in Figure 11 indicate that the global TMOs result in the least distorted reconstruction, which, however, quickly deteriorates if the JPEG compression quality is low. Such sensitivity to JPEG compression can be explained by the contrast stretching performed when reversing the TMO, which amplifies JPEG compression artifacts. The best performance can be observed for those global TMOs that do not clip many pixels (*LogMap*, *Photo*), and much worse for those TMOs that distort the pixels in highlights and shades (*HistEq*). The local TMOs that involve sharpening produce the worst reconstruction. This may be due to a bad fit of the generic TMO, or to the sharpening, which amplifies high spatial frequencies, which are heavily distorted by the JPEG quantization matrix. High frequency boosting also causes worse JPEG compression performance, as shown in Figure 12. Therefore, the TMOs that sharpen an image should be avoided in backward-compatible HDR image compression.

A similar approach to compressing HDR images and video has been previously proposed by Ward and Simmons [WS04], Li et al. [LSA05] and Mantiuk et al. [MEMS06]. The difference is that both [WS04] and [LSA05] improve compression by distorting a tone-mapped image, which is not desirable in many applications. We extend and generalize the approach of Mantiuk et al. by adding spatial modulation to compensate for local TMOs and reduce the number of parameters that describe the tone curve. Our study also states which operators are the most suitable for backward-compatible HDR compression, and explains why.

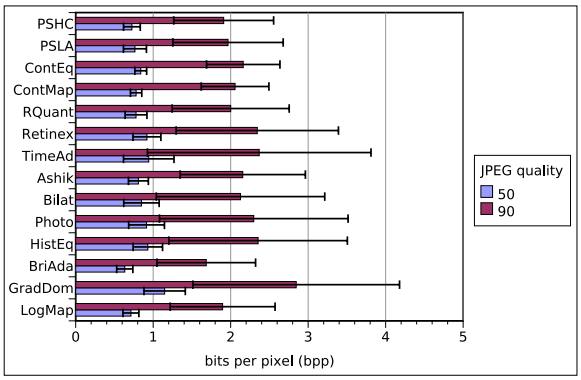

Figure 12: *Compression performance for a range of tone-mapping operators. The 75% quality results have been omitted for better clarity.*

6. Synthesis of tone-mapping operators

Each tone-mapping operator has been carefully designed to reach a certain aesthetic or application-specific goal, but because of a large number of the available operators, it is often difficult to choose the best one to apply. This section shows that we need not to decide on a single TMO, and can use several operators to tone-map a single image.

Since the generic TMO contains 10 parameters that are difficult to control, we replace them with a smaller number of new parameters that alter specific TMO characteristics. In particular, we want to mimic the behavior of selected TMOs by maximizing the chance that a set of new parameters gives a result that is close to one of these TMOs. As in [BV99], we assume that each original parameter is a linear combination of new parameters and employ PCA to find these combinations. Formally, we want to find a matrix A that lets us compute the original generic TMO parameters, P, based on the new parameters, N:

$$P = A \cdot N \tag{8}$$

where the vector N possibly has fewer elements than the 10-element vector $P = [b\ c\ d_l\ d_h\ s\ m_1\ \ldots\ m_5]^T$. For the principal component analysis we create a database of generic TMO parameters (P), based on 40 HDR images tone-mapped with 6 TMOs (*HistEq*, *Bilat*, *Ashik*, *Photo*, *LogMap*, *TimeAdapt*). The TMOs were selected to result in a good fit to the generic TMO.

As shown in Figure 9, the parameters of the generic TMO are determined not only by a particular TMO and its parameters, but also by image content. Since we want to be able to tone-map any image, we need to remove the effect of image content. To do this, we check if a linear combination of selected percentiles of luminance values has a significant effect on any of the TMO parameters. We choose percentiles, as they are the simplest description of image content, together with an image histogram, from which the percentiles can be derived. The ANOVA test indicates that only the image brightness parameter (b) is affected by the percentiles ($F = 5.68$, $\alpha < 0.01$) and therefore we fit a linear model for image brightness only. The image dependent component of brightness is predicted by:

$$b_{id} = \begin{matrix}[P_1\ P_{25}\ P_{50}\ P_{75}\ P_{99}\ 1]\cdot \\ [-0.23\ \ 0.6\ \ 0.84\ \ -0.03\ \ 0.15\ \ -0.32]\end{matrix} \tag{9}$$

where P_n is the n-th percentile of logarithmic image luminance and "$\cdot$" is the dot product. The value of b_{id} is subtracted from the brightness parameters of the vector P before it is passed to the PCA. When tone-mapping an image, b_{id} is added to the brightness parameter. Image content obviously also affects other TMO parameters, but a these effects cannot be predicted by a simple linear combination of percentiles.

To perform the PCA, all variables should be given in the same units, which is not the case for the parameters of the generic TMO. We normalize the generic TMO parameters before PCA by subtracting the mean value and dividing by the standard deviation. For the parameters that share the same units ($\{b, d_l, d_h\}$ and the MTF coefficients) we compute the standard deviation for the entire group.

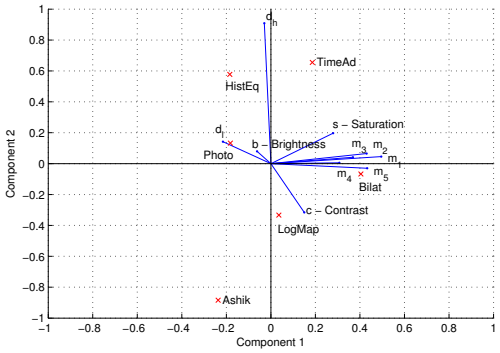

Figure 13: *First two principal components for our database of generic TMO parameters.*

The two first principal components for our database of TMOs and images ($A_{:,1}$, $A_{:,2}$) are illustrated in Figure 13. The Figure also shows the mean parameter values for each TMO. The plot indicates that the first principal component is responsible for sharpness (MTF parameters m_1–m_5) and color saturation (s), while the second component stands for the compression of highlights (d_h) and contrast (c). Brightness (b) and contrast (c) are inversely correlated (higher brightness is compensated by lower contrast), a correlation also found in the study by [YMMS06].

We can use the first two principal components (PC) to explore the space of possible TMOs. Figure 14 shows a collection of tone-mapped images, for which parameters were selected based on the first two PC. These images were not

used for computing the statistical model. The top-center image is a mixture of *HistEq* and *TimeAdapt* (refer to Figure 13 for the PC coordinates of TMOs), while the middle-right image is closer to *Bilat*. Similarly we can combine the result of any pre-learned TMO for an arbitrary image.

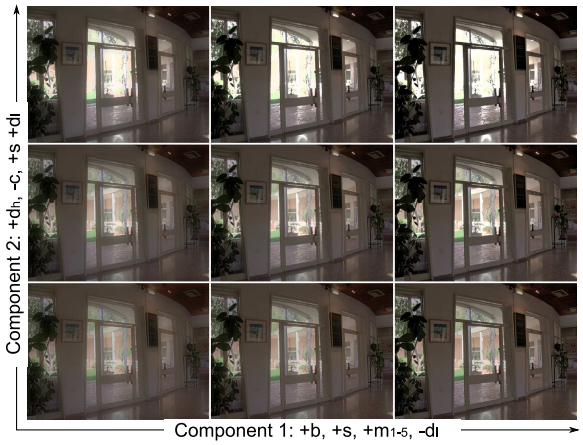

Figure 14: *Images tone-mapped using a linear combination of selected operators. The image in the center is the mean result of all TMOs, and the axis are the same as in Figure 13. The axis labels indicate positive or negative correlation with the genetic TMO parameters.*

This section shows how a new statistical TMO can learn solely based on the results of other operators. This technique can be used to build computationally less expensive operators that mimic the behavior of more complex ones, to study the parameter space of existing TMOs, or to combine features of several TMOs into a single operator.

7. Discussion

An interesting observation in this study is that non-linear operators, such as the bilateral filter, can be approximated by linear filters. It does not mean that linear filters can produce exactly same results as those computationally expensive non-linear operators, but rather that with a quality factor $SSIM \geq 0.95$, only expert observers can notice and appreciate the difference between the linear approximation and the result of the non-linear operator. This suggests that for many practical applications, well-designed linear filters may be used instead of expensive non-linear operators.

The major difficulty of our approach is that TMO processing is influenced by image content and therefore the generic TMO parameters will vary from image to image. Since a simple brightness estimator b_{id} does not account for all image effects, our TMO synthesis approach cannot accurately replicate original TMO results unless more advanced models predicting the generic TMO parameters are employed. This limitation, however, does not affect other applications of the generic TMO, such as backward-compatible HDR image compression or quantitative TMO analysis, where both the tone-mapped and HDR images are known.

We tried to improve the accuracy of the approximation by introducing additional non-linearity after the MTF step. The non-linearity was a power function applied to each band with a separate exponent. Although we achieved about 5% improvement in the SSIM and HDR-VDP scores, the fitting was considerably slower and in many cases unstable due to local minima. Such tone-mapping was also non-invertible and therefore not suitable for the backward-compatible HDR compression. We expect that better approximation would require an even more complex tone mapping operator that could account for local changes in the mapping function. However, designing an invertible operator that guarantees stable fitting is not a trivial task.

8. Conclusions

This paper demonstrates that many TMOs can be approximated by a single model consisting of a tone curve followed by a spatial modulation function. This indicates that these TMOs employ very similar image processing, and the major difference comes from the strategy used for choosing the set of parameters that gives desirable results. The method outlined in this paper is useful for the analysis and quantitative comparison of tone-mapping operators. The proposed generic model can be used in backward-compatible HDR image encoding to reverse tone-mapping operators and thus to recover HDR content. We also demonstrate how the behavior of existing operators can be learned from examples and reproduced on other images.

In future work we would like to improve the statistical model described in Section 6 to make it as independent of image content as possible. Faster implementation of the optimization procedure and the generic operator based on Gaussian pyramids would open possibilities for new applications, such as HDR video compression and video tone-mapping.

Acknowledgments

We thank Karol Myszkowski, Grzegorz Krawczyk, Gernot Ziegler, Kaleigh Smith, Akiko Yoshida and anonymous reviewers for their valuable comments. The images used for this study are the courtesy of Greg Ward, Paul Debevec, Dani Lischinski and Frédéric Drago.

References

[Ash02] ASHIKHMIN M.: A tone mapping algorithm for high contrast images. In *Rendering Techniques 2002: 13th Eurographics Workshop on Rendering* (2002), pp. 145–156.

[BV99] BLANZ V., VETTER T.: A morphable model for the synthesis of 3D faces. In *Proc. of SIGGRAPH '99* (1999), pp. 187–194.

[CWNA06] CADIK M., WIMMER M., NEUMANN L., ARTUSI A.: Image attributes and quality for evaluation of tone mapping operators. In *Proc. of the 14th Pacific Conf. on Comp. Graph. and Applications* (2006), pp. 35–44.

[Dal93] DALY S.: The Visible Differences Predictor: An algorithm for the assessment of image fidelity. In *Digital Image and Human Vision* (1993), Watson A., (Ed.), Cambridge, MA: MIT Press, pp. 179–206.

[DD02] DURAND F., DORSEY J.: Fast bilateral filtering for the display of high-dynamic-range images. *ACM Trans. on Graph. 21*, 3 (2002), 257–266.

[DMAC03] DRAGO F., MYSZKOWSKI K., ANNEN T., CHIBA N.: Adaptive logarithmic mapping for displaying high contrast scenes. *Computer Graphics Forum, Proceedings of Eurographics 2003 22*, 3 (2003), 419–426.

[DMMS02] DRAGO F., MARTENS W., MYSZKOWSKI K., SEIDEL H.-P.: *Perceptual Evaluation of Tone Mapping Operators with Regard to Similarity and Preference*. Technical Report MPI-I-2002-4-002, Max-Planck-Institut fuer Informatik, Oct. 2002.

[FLW02] FATTAL R., LISCHINSKI D., WERMAN M.: Gradient domain high dynamic range compression. *ACM Trans. on Graph. 21*, 3 (2002), 249–256.

[FPSG96] FERWERDA J., PATTANAIK S., SHIRLEY P., GREENBERG D.: A model of visual adaptation for realistic image synthesis. In *Proc. of SIGGRAPH 96* (1996), pp. 249–258.

[KMS05] KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Lightness perception in tone reproduction for high dynamic range images. *Comp. Graph. Forum 24*, 3 (2005), 635–645.

[KMS07] KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Contrast restoration by adaptive countershading. *Computer Graphics Forum 26*, 3 (2007), 581–590.

[KYJF04] KUANG J., YAMAGISHI H., JOHNSON G., FAIRCHILD M.: Testing HDR rendering algorithms. In *IS&T/SID Color Imaging Conference* (2004), pp. 315–320.

[LCTS05] LEDDA P., CHALMERS A., TROSCIANKO T., SEETZEN H.: Evaluation of tone mapping operators using a high dynamic range display. *ACM Trans. on Graph. 24*, 3 (2005), 640–648.

[LFUS06] LISCHINSKI D., FARBMAN Z., UYTTENDAELE M., SZELISKI R.: Interactive local adjustment of tonal values. *ACM Trans. Graph. 25*, 3 (2006), 646–653.

[LSA05] LI Y., SHARAN L., ADELSON E. H.: Compressing and companding high dynamic range images with subband architectures. *ACM Trans. on Graph. 24*, 3 (2005), 836–844.

[McC07] MCCANN J.: Art, science, and appearance in HDR. *Journal of the Soc. for Inf. Display 15*, 9 (2007), 709–719.

[MEMS06] MANTIUK R., EFREMOV A., MYSZKOWSKI K., SEIDEL H.-P.: Backward compatible high dynamic range MPEG video compression. *ACM Trans. on Graphics 25*, 3 (2006), 713–723.

[MMS06] MANTIUK R., MYSZKOWSKI K., SEIDEL H.-P.: A perceptual framework for contrast processing of high dynamic range images. *ACM Trans. Appl. Percept. 3*, 3 (2006), 286–308.

[PFFG98] PATTANAIK S. N., FERWERDA J. A., FAIRCHILD M. D., GREENBERG D. P.: A multiscale model of adaptation and spatial vision for realistic image display. In *Siggraph 1998, Computer Graphics Proceedings* (1998), pp. 287–298.

[pfs] PFSTMO: Library of tone-mapping operators. http://www.mpi-inf.mpg.de/resources/tmo/.

[PTYG00] PATTANAIK S., TUMBLIN J., YEE H., GREENBERG D.: Time-dependent visual adaptation for realistic image display. In *Proc. of SIGGRAPH 2000* (2000), pp. 47–54.

[RSSF02] REINHARD E., STARK M., SHIRLEY P., FERWERDA J.: Photographic tone reproduction for digital images. *ACM Trans. on Graph. 21*, 3 (2002), 267–276.

[RTS*07] REMPEL A. G., TRENTACOSTE M., SEETZEN H., YOUNG H. D., HEIDRICH W., WHITEHEAD L., WARD G.: LDR2HDR: On-the-fly reverse tone mapping of legacy video and photographs. *ACM Transactions on Graphics (Proc. SIGGRAPH) 26*, 3 (2007). Article 39.

[RWPD05] REINHARD E., WARD G., PATTANAIK S., DEBEVEC P.: *High Dynamic Range Imaging. Data Acquisition, Manipulation, and Display*. Morgan Kaufmann, 2005.

[Sch94] SCHLICK C.: Quantization techniques for the visualization of high dynamic range pictures. In *Photorealistic Rendering Techniques* (1994), Eurographics, Springer-Verlag Berlin Heidelberg New York, pp. 7–20.

[SKMS06] SMITH K., KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Beyond tone mapping: Enhanced depiction of tone mapped HDR images. *Computer Graphics Forum 25*, 3 (2006), 427–438.

[tmo] TMOIMG:: Gallery of reference tone-mapped images generated by TMO authors. http://www.mpi-inf.mpg.de/resources/tmo/NewExperiment/TmoOverview.html.

[TR93] TUMBLIN J., RUSHMEIER H. E.: Tone reproduction for realistic images. *IEEE Computer Graphics and Applications 13*, 6 (Nov. 1993), 42–48.

[TT99] TUMBLIN J., TURK G.: LCIS: A boundary hierarchy for detail-preserving contrast reduction. In *Siggraph 1999, Computer Graphics Proceedings* (1999), pp. 83–90.

[Wat87] WATSON A.: The cortex transform: Rapid computation of simulated neural images. *Comp. Vis. Graph. and Image Proc. 39* (1987), 311–327.

[WBSS04] WANG Z., BOVIK A., SHEIKH H., SIMONCELLI E.: Image quality assessment: from error visibility to structural similarity. *Image Processing, IEEE Transactions on 13*, 4 (2004), 600–612.

[WLRP97] WARD LARSON G., RUSHMEIER H., PIATKO C.: A visibility matching tone reproduction operator for high dynamic range scenes. *IEEE Transactions on Visualization and Computer Graphics 3*, 4 (1997), 291–306.

[WS04] WARD G., SIMMONS M.: Subband encoding of high dynamic range imagery. In *APGV '04: 1st Symposium on Applied Perception in Graphics and Visualization* (2004), pp. 83–90.

[YMMS06] YOSHIDA A., MANTIUK R., MYSZKOWSKI K., SEIDEL H.-P.: Analysis of reproducing real-world appearance on displays of varying dynamic range. *Computer Graphics Forum 25*, 3 (2006), 415–426.