Volume xx (200y), Number z, pp. 1–12

Beyond Tone Mapping: Enhanced Depiction of Tone Mapped HDR Images

Kaleigh Smith, Grzegorz Krawczyk, Karol Myszkowski, and Hans-Peter Seidel

MPI Informatik, Saarbrücken, Germany

Abstract

High Dynamic Range (HDR) images capture the full range of luminance present in real world scenes, and unlike Low Dynamic Range (LDR) images, can simultaneously contain detailed information in the deepest of shadows and the brightest of light sources. For display or aesthetic purposes, it is often necessary to perform tone mapping, which creates LDR depictions of HDR images at the cost of contrast information loss. The purpose of this work is two-fold: to analyze a displayed LDR image against its original HDR counterpart in terms of perceived contrast distortion, and to enhance the LDR depiction with perceptually driven colour adjustments to restore the original HDR contrast information. For analysis, we present a novel algorithm for the characterization of tone mapping distortion in terms of observed loss of global contrast, and loss of contour and texture details. We classify existing tone mapping operators accordingly. We measure both distortions with perceptual metrics that enable the automatic and meaningful enhancement of LDR depictions. For image enhancement, we identify artistic and photographic colour techniques from which we derive adjustments that create contrast with colour. The enhanced LDR image is an improved depiction of the original HDR image with restored contrast information.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]: Picture/Image Generation I.4.0 [Image Processing and Computer Vision]: GeneralImage processing software

1. Introduction

High Dynamic Range (HDR) images accurately describe the wide range of luminance visible in the real world. Because their dynamic range is broad enough to represent the true range of luminosity in a scene (between 3 to 12 orders of magnitude), HDR images capture details that are perceived by the human visual system (HVS) but missed by standard photographic techniques.

HDR images are well known to the computer graphics research community, and the recent introduction of HDR image creation and editing capabilities into most common image editing software ensures that HDR images will become an increasingly common form for storing and manipulating visual information. Additionally, the high quantity of information in HDR images can prevent editing artifacts and should improve the performance of image processing algorithms. As such, there is new interest in HDR processing techniques and methods for exploiting the expanded information contained in HDR images.

Tone mapping is the first and most developed research area in HDR image processing [RWPD05]. Tone mapping compresses the wide dynamic range to a narrower range for display and aesthetic purposes thus creating an LDR depiction of an HDR image. For a majority of existing tone mapping operators this is achieved through the reduction of physical contrast in LDR images. However, perceived image contrast is not only a function of the dynamic range of the tone mapped image, but also depends significantly on other image attributes such as lightness, hue, chroma, and sharpness [CF03, Hun95]. This means that by skillfully tuning these attributes, the losses in physical contrast due to tone mapping can be restored as perceived contrast.

In order to restore physical contrast, we must first determine the contrast distortion between the HDR image and the tone mapping: how much perceived contrast has been lost and where it should be restored. While much work has been done in the subjective evaluation of different tone mapping operators [LCTS05, YBMS05], to our knowledge, we present the first feature-based characterization and objective perceptual measure of tone mapping distortion. Since distortion results from balancing the trade-off between preserving global contrast and preserving details, we create perceptual metrics for *Global Contrast Change* and *Detail Visibility Change* between an HDR image and its tone mapped LDR counterpart.

The change in ratio between brightest and darkest points of an image is a traditional definition of global contrast change that is necessarily adjusted by tone mapping, and so would not be considered a distortion. Contrary to this definition and others, such as one using the multi-resolution definition given by [MNN*05], we consider global contrast change to be a characteristic defined by the shape of the tone mapping function, thus removing the emphasis on extreme brights and darks which have less impact on the impression of global contrast. Our definition of global contrast change is more closely related to image comprehension, which according to Gestalt theorists, involves the cognitive task of separating the image into recognizable objects, most importantly, the separation of foreground objects from the background [Liv02]. As such, a decrease in global contrast may make comprehension of the LDR image more difficult, indicating a loss in visual communication efficacy. We define Detail Visibility Change as the reduction, disappearance or exaggeration of high frequency contrasts in the LDR image compared to the HDR original. To obtain results that accurately represent these two distortions, we develop novel methods to measure their perceptual aspects and we analyze several tone mappings accordingly.

Once we localize contrast distortions resulting from tone mapping and estimate the magnitude of those distortions using a perceptually linear and meaningful scale, we can try to restore the perceived contrast of the original HDR image. While various contrast correction techniques could be envisioned [CF03], we adjust perceived contrast by operating directly on chroma according to our distortion measuments, thus avoiding substantial changes to the tone mapped luminance. Technically, chroma has a clear mathematical definitions (for example in the perceptually uniform colour space CIELUV) and can be easily set on the pixel level. The spatial aspect of controlling chroma for neighbouring pixels or more distant image regions has interesting uses in perceived contrast manipulation and is widely exploited by artists and photographers [Liv02].

Taking inspiration from painting and photography, we suggest chroma operations that enhance the appearance of global contrast and detail contrast. We enhance perceived global contrast by applying countershading to encourage image efficacy and to create the impression of greater global contrast. Countershading is the juxtaposition of opposing gradients to create an exaggerated difference at a feature boundary, often the boundary between foreground objects and the background. It is a techniques used by renowned photographer Pete Turner, who characteristically creates photographs with saturation gradients applied to

Figure 1: *Left: countershading with a saturation gradient of the sky in Orange Wall and Sea, by Pete Turner (www.peteturner.com). Right: countershading in Breakfast Still-Life, Willem Claesz Heda. Both from the Web Gallery of Art (www.wga.hu).*

backgrounds, and by Dutch still life painter Willem Claesz Heda, who creates a diagonal countershading of the background, to make the bright foreground regions seem brighter and the dark regions darker, thus creating an impression of greater dynamic range and strengthening object silhouettes (Figure 1). We adjust the colour of high frequency details to introduce variation in saturation, a technique Michelangelo employed to emphasize highlights, contours and texture details and Monet used to distinguish nearly shapeless details (Figure 2) [Liv02].

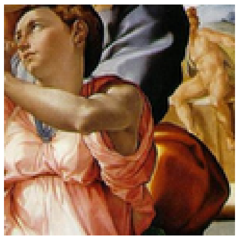

Figure 2: *Left: saturated details in The Holy Family with the infant St. John the Baptist (the Doni tondo), Michelangelo Buonarroti. Right: Monet flower details almost entirely distinguished by colours, from The Artist's House.*

The paper is structured as follows. We first refer to the related work on image enhancements, artistic colour techniques and visibility metrics in Section 2. In Section 3, we present our perceptual metrics for contrast distortion between HDR image and its displayed LDR counterpart, and perform the analysis of existing tone mapping algorithms according to these distortions. Next, in Section 4 we propose enhancements of tone mapped images by colour adjustment. Finally, we discuss our results in Section 5 and conclude the paper in Section 6.

2. Related Work

The image enhancement aspect of this work relies extensively on the use of colour in imagery and is related to image recolourization. Colour is a prominent attribute for effective visual communication and its use is addressed in a variety of fields including colour appearance modelling, scientific visualization and image processing. The work most closely related to our treatment of colour is image recolouring, which transfers colours between images, introduces colours into a greyscale image or quantizes the number of colours in an image [RAGS01, GH03, RGW05].

Given that colour is an inherent attribute of image quality, one would assume that tone mapping operators perform some colour enhancement. However, the majority of tone mappings compress only luminance values, and are not concerned with color issues. Two notable exceptions are the iCAM model [FJ03] and the multi-scale adaptation model [PFFG98], both of which are advanced image appearance models that incorporate colour appearance modeling [Hun95]. In this work, we add colour enhancement atop of an arbitrary tone mapping operator, which compensates for perceived contrast losses due to the physical contrast compression, the major task of the traditional tone mapping operator (refer to [RWPD05] for a detailed survey). Such a traditional operator is often designed to produce images that "look good" or to obtain a perceptual match between the image and the corresponding real world scenes. The success of meeting these goals depends heavily on particular HDR image characteristics and as such, it is difficult to single out one existing operator that consistently performs best [RWPD05].

Each tone mapping operator takes form as a collection of certain image processing operations, whose impact on the perceived image quality or fidelity to the real world appearance is not well understood. Recent psychophysical studies attempt to evaluate tone mapping operators in terms of subject preference or fidelity of the real world scene depiction [KYJF04, LCTS05, YBMS05]. In such studies each operator is treated as a "black box" and its performance is compared on the whole with respect to other operators, without an attempt at understanding the reasons for subjects' judgments. While some studies of tone mapping operators go further and take into account the reproduction of overall brightness, global contrast or details (local contrast) in dark and bright image regions [LCTS05, YBMS05], they remain focused on comparing the operator performance for each of these tasks. These studies provide no deeper analysis of how the pixels of an HDR image have been transformed by tone mapping and in what way the outcome of such a transformation depends on image content. Such analysis could help in understanding how particular image characteristics, such as contrast or brightness, are locally distorted by tone mapping and determining the impact of such distortions on perceived image quality.

In this work, instead of subjective analysis, we focus on global and local contrast distortions between HDR image and its tone mapped counterpart. For this purpose we must evaluate the magnitude of these distortions along a perceptually meaningful scale. A number of perception-based visible difference (fidelity) metrics for image pairs have been developed, mostly for image compression and color reproduction applications (refer to [Win05] for a recent survey of such metrics). State of the art fidelity metrics such as the Visible Differences Predictor (VDP) [Dal93] or the Sarnoff Visual Discrimination Model (VDM) [Lub95] include many important characteristics of the HVS, such as eye optic imperfections, luminance masking, the contrast sensitivity function (CSF), and pattern masking, making them very general metrics. However, such complex metrics may perform worse than simpler metrics specialized for the task of detecting well-defined distortion types, such as blocking artifacts that arise in image compression [Win05]. The majority of existing fidelity metrics are based on HVS models developed through threshold psychophysical experiments whose goal is to determine the magnitude of a simple stimulus so that it becomes just noticeable. Such metrics successfully detect the presence of perceivable image distortions, but perform poorly in estimating the magnitude of suprathreshold distortions and predicting their distraction to the human observer [CH03]. With its spatial features for estimating imperceptible texture details, the iCAM model [FJ03] is an exception, however, since the magnitude of perceptual responses to local contrast is not available, it can not be used to determine the change in detail visibility.

In this work we are mostly concerned with one well defined suprathreshold distortion: contrast compression due to tone mapping. Since fidelity metrics dealing with image pairs of drastically different dynamic ranges have not so far been proposed, and since we have found existing models to be ill-suited for our purposes, we develop custom fidelity metrics for comparing perceived contrast differences between an original HDR image and its tone mapped LDR counterpart.

3. Tone Mapping Distortions

All successful tone mapping operators balance the trade-off between loyal reproduction of the luminance range and preservation of details. One can argue the that *photographic tone reproduction* operator [RSSF02] best reproduces global contrast, while the *gradient domain compression* [FLW02] operator best preserves details. However, the accuracy of such statements may depend on the particular HDR image, and as concluded by evaluations of tone mapping operators [YBMS05, LCTS05], it is difficult for one tone mapping operator to be well-suited to all types of images. Regardless of technique, each tone mapping operator introduces a degree of distortion into the resulting LDR tone mapped image. Drawing conclusions from previous evaluations and our own observations, we identify two major contrast distortions resulting from tone mapping:

Global Contrast Change the ratio between lightest and darkest areas of the HDR is reduced in the LDR,

Detail Visibility Change (textures and contours) the high frequency contrasts of the HDR image become less prominent, disappear, or become exaggerated in the LDR.

A significant Global Contrast Change is undesirable not only for esthetic reasons, but also because of changes in image understandability, despite good detail visibility. Certain specialized tone mapping operators assign a wider dynamic range to detailed regions to preserve textures and contours, which results in a narrower dynamic range available for global luminance changes, decreasing the ratio between lightest and darkest areas. Detail Visibility Change occurs either because a region becomes entirely saturated or because an area is mapped to very few or very low brightness levels. The second case is especially interesting from the perceptual point of view, because the physical contrasts still exist in the LDR image, however the details are invisible to the human observer.

3.1. Distortion Metrics

Our goal is to determine the apparent distortion in detail visibility and global contrast change which were introduced during the tone mapping of HDR image. We focus on the luminance compression aspect of the operators. Instead of analyzing particular algorithms one by one, we consider tone mapping as an unknown transformation applied to the luminance of an HDR image, resulting in an LDR image. To do so, we use knowledge of human perception to compare a real world or synthetic scene, captured as an HDR image, to its LDR tone mapping as depicted on display device. The output of our metric consists of a single value representing the global contrast change factor and a map representing the magnitude of change in detail visibility. The units of the detail visibility map are Just Noticeable Differences (JND), which allows for an informed use of this information for potential perceptually based corrections.

To compare images of significantly different dynamic ranges we compare the luminance of an HDR image, denoted as Y, to the luminance shown on a display device, denoted as L. To accurately predict the displayed luminance, we assume that sufficient characteristics of the display device are known so that we can calculate the luminance value in cd/m^2 of each LDR image pixel. For an sRGB monitor, this requires black and white levels increased by an ambient illumination level. Similarly, a photometrically calibrated HDR image is desirable.

We transform the gamma corrected luminance values[†] y of the LDR image to display luminance values L. Given the display black L_{black} and white L_{white} levels in cd/m^2 and assuming sRGB response, the transformation is the following:

$$L = L_{black} + \mathrm{sRGB}^{-1}(y) \cdot (L_{white} - L_{black}). \tag{1}$$

If the absolute luminance values of an HDR image are unknown, we align the relative HDR values Y to the LDR values L according to the average logarithmic luminance, a method often used as an adaptation estimate in tone mapping [DMAC03, RSSF02].

Global Contrast Change

Global contrast can be measured as a ratio of maximum to minimum displayable luminance. Tone mapping algorithms, however, generally use the whole display dynamic range which, according to above definition, always results in maximum global contrast. Yet images resulting from different tone mapping operators with identical ratios can create starkly different impressions of global contrast, meaning that such a naïve measure is not appropriate. The variety in global contrast impression comes from the different shapes of tone mapping functions, and therefore it is sensible to analyze these functions to obtain a global contrast estimate. Unfortunately these functions are either unknown or not well-defined, as in the case of *gradient domain compression*. However, we argue that a general approximation of the tone mapping function is sufficient for estimating global contrast. In our metric, we approximate the tone mapping function using linear regression in the brightness domain:

$$L_B \approx TM(Y_B) = C \cdot Y_B + B \tag{2}$$

where C and B are estimated coefficients, and Y_B and L_B approximate brightness following the Weber-Fechner Law ($Y_B = log_{10}Y, L_B = log_{10}L$).

Given the tone mapping function approximation, we calculate the display luminance values corresponding to the minimum and maximum luminance of the HDR image. In our opinion, these values are more reliable for global contrast estimation in the LDR image than actual minimum and maximum values. The calculated values reflect the general tendency of brightness mapping rather than being a product of a detail enhancing procedure which is independent of global contrast relations. We calculate the global brightness contrast ΔL_B using the tone mapping function estimation from Equation (2):

$$\Delta L_B = TM(\max(Y_B)) - TM(\min(Y_B)), \tag{3}$$

where the result of tone mapping function is clamped to the minimum and maximum displayable values.

Finally, to calculate the Global Contrast Change $\mathscr{C}$ we relate the global contrast in LDR image to its original HDR:

$$\mathscr{C} = \frac{\Delta L_B}{\Delta Y_B}, \tag{4}$$

where ΔY_B is a difference between the maximum and minimum brightness in the HDR image. $\mathscr{C} < 1$ indicates lower

[†] image luminance is calculated from the RGB channels according to the [ITU90] standard.

global contrast in the LDR image, whereas $\mathscr{C} > 1$ indicates higher global contrast. By deduction, the global contrast change $\mathscr{C}$ is equivalent to the C coefficient from the brightness mapping estimation (2).

The result of applying our measure of Global Contrast Change to two tone mappings (one global and one local) is shown in Figure 3. While both methods make use of the entire available dynamic range, the shapes of their mapping functions differ: the global mapping function is well-defined, as opposed to the non-uniform and scattered local mapping function. Higher global contrast is obtained with the global tone mapping method, whereas the detail preserving local method exhibits a smaller ratio between bright and dark areas (the function approximation is nearly flat).

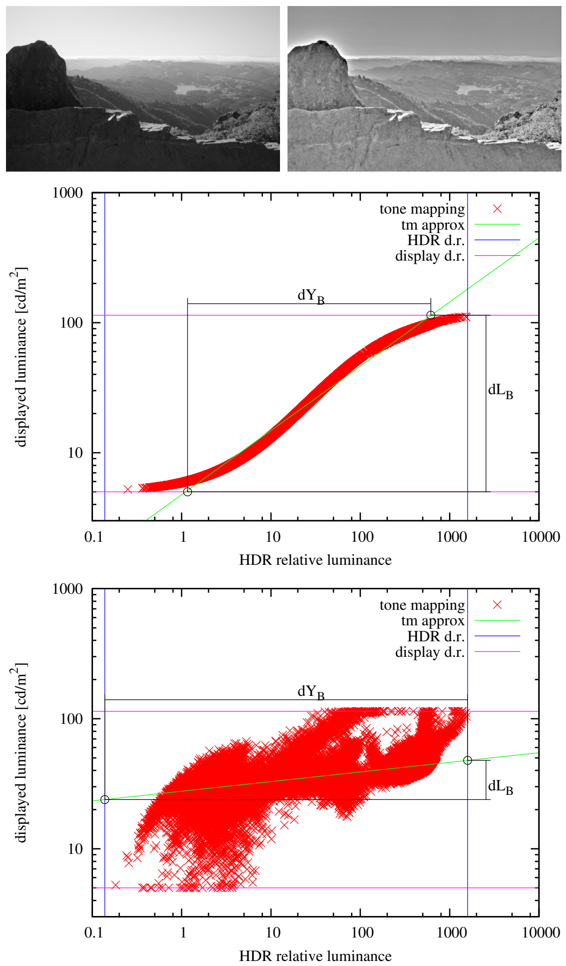

Figure 3: *Global Contrast estimation for global (left image, top plot) and local (right image, bottom plot) tone mapping. Each plot shows pixel-by-pixel mapping between HDR and LDR, linear brightness mapping estimation, dynamic ranges (d.r.) of LDR and HDR, and contrast measures [DW00].*

Details Visibility Change

Details of textures and contours can be described as the high frequency contrasts between a pixel and its adapting field. Visibility, the response of the HVS to the magnitude of such contrasts, is not linear and depends on the adaptation level. Contrast visibility can be analyzed in terms of contrast detection and contrast discrimination. We use contrast detection for identifying visible details in both the HDR and LDR images, and we use contrast discrimination for identifying the magnitude of visible difference in detail contrast between the HDR and LDR images.

We start by identifying high frequency contrasts that presumably create texture and contour details in the image. For each pixel Y_i we estimate the adapting luminance Y_i^{sp} in its neighbouring area and calculate the contrast expressed as a logarithmic ratio of luminance values:

$$G(Y_i, Y_i^{sp}) = \log_{10} \frac{\max(Y_i, Y_i^{sp})}{\min(Y_i, Y_i^{sp})}. \quad (5)$$

We simulate the adaptation to low spatial frequencies in an image and we take special care to prevent the influence of significantly different luminance values on an adaptation level. We obtain the adaptation map Y^{sp} by processing the HDR image with a low pass bilateral filter in the logarithmic domain. Such a filter removes high frequencies while preserving high contrast edges. The adaptation map is refined by eliminating frequencies above 20 cycles per pixel and preserving edges of logarithmic contrast ratio higher than 0.25. We calculate the high frequency contrasts of the LDR image in the same way. It is important to note that the particular choice of the bilateral filter for estimating the adaptation map is not mandatory. Other algorithms known from tone mapping can be used as well, as long as they do not introduce artifacts at the high contrast edges.

To estimate the Details Visibility Change between two images of significantly different dynamic range, knowledge of the hypothetical HVS response to given physical contrasts under given adaptation conditions is required. A reasonable prediction for a full range of contrast values is given by the following transducer function that is derived and approximated by Mantiuk et al. [MMS06]:

$$T(G) = 54.09288 \cdot G^{0.41850}, \quad (6)$$

with the following properties:

$$T(0) = 0 \quad \text{and} \quad T(G_{threshold}) = 1. \quad (7)$$

The transducer function estimates the HVS response to physical contrast in Just Noticeable Difference (JND) units. Thus for a given contrast threshold, $G_{threshold}$, a transducer value equals 1 JND. It is important to note that this measure holds for suprathreshold measurements, since it not only estimates the detection, but also the magnitude of change.

The approximation given by Equation (6) has been derived with the assumption of 1% contrast detection threshold, i.e.

$G_{threshold} = \log_{10}(1.01)$. Although such an assumption is often made in image processing for LDR, the detection threshold depends on an adapting luminance level and is described by the Threshold Versus Intensity (TVI) function [CIE81]. The TVI function shows that this threshold varies in the luminance range of displays and the dynamic range in HDR is often high enough to make this 1% assumption for the detection threshold inaccurate. We therefore derive a scaling factor $t(Y^{sp})$ for the transducer function (6) which adjusts its properties (7) to match the TVI function for given an adapting luminance:

$$\mathrm{t}(Y^{sp}) = \frac{\log_{10} 1.01}{\log_{10} \frac{Y^{sp} + \mathrm{tvi}(Y^{sp})}{Y^{sp}}}. \quad (8)$$

Such a scaling factor is appropriate because the approximation of the transducer function (6) was derived with starting conditions from (7), and since the influence of the threshold is multiplicative [MMS06]. Figure 4 illustrates the magnitude of change in the HVS response depending on the adapting luminance. The response changes by a factor of almost 1 order of magnitude within the visible range of luminance on a display. In practice, the scaling factor reduces the response to contrast in the dark areas of an image.

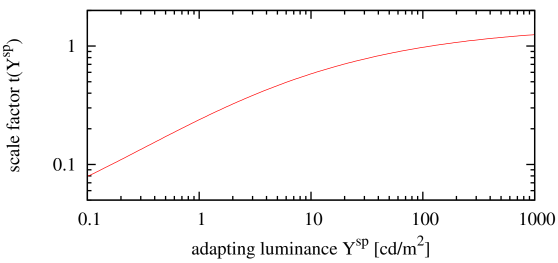

Figure 4: *Plot of a scale factor from Equation (8). Luminance range of a typical LCD display is* 2 *to* 200 cd/m^2.

Given the scaled transducer function, we can estimate the hypothetical response of the HVS to the high frequency contrasts measured with equation (5):

$$T^*(Y_i, Y^{sp}) = T(G(Y_i, Y_i^{sp})) \cdot t(Y_i^{sp}). \quad (9)$$

The response T^* is expressed in JND units, which means that a detail Y_i is visible under given luminance conditions only if $T^* > 1$. Given this relation, we are able to estimate the details of a displayed LDR image and the details of an HDR image which would be visible to a human observer. Furthermore, since the transducer function is a suprathreshold measure, we are able to estimate change by comparing the magnitude of detail visibility in a displayed LDR image to its HDR version (spatial arguments are omitted for brevity):

$$\Delta\overline{T^*}(Y_i, L_i) = \begin{cases} 1 & \text{for } \overline{T^*}(Y_i) > 1 > \overline{T^*}(L_i), \\ 0 & \text{for } \|\overline{T^*}(Y_i) - \overline{T^*}(L_i)\| < 1, \\ \overline{T^*}(Y_i) - \overline{T^*}(L_i) & \text{otherwise.} \end{cases} \quad (10)$$

For practical reasons, we consider the average detail visibility measure over its neighbouring pixels, denoted as $\overline{T^*}$, because we are interested in general detail visibility in a certain arbitrary small area. As shown in Equation 10, we consider three cases of detail visibility change. When a response to high frequency contrast in the HDR image is attenuated from above 1 JND to below 1 JND in the tone mapped image, the change is 1 JND. When the difference in response is below 1 JND, the change is deemed invisible and is set to 0. In all other cases, the magnitude of Detail Visibility Change is set to the difference in responses T^*. We illustrate the performance of this measure in Figure 5.

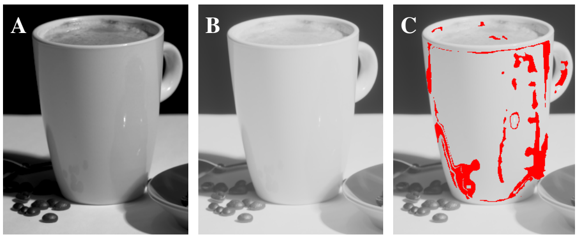

Figure 5: *Detail Visibility. HDR image (A) contains subtle reflection on a surface of the cup. A global tone mapping (B) reveals the coffee beans in the shadow but the reflection details become indiscernible. The areas of image with lost details are predicted by our metric (C), where red colour marks* $\Delta\overline{T^*} > 1$.

3.2. Analysis of Distortions

We analyzed the performance of 8 tone mapping methods in terms of Global Contrast Change and Detail Visibility Change using the presented metrics. The analysis was performed on a set of 18 HDR images with an average dynamic range of approximately 4 orders of magnitude and a resolution between 0.5 and 4 megapixels. The set contained a variety of scenes with differing lighting conditions and included panoramic images. We tested the following global (spatially uniform) tone mapping algorithms: *gamma correction* ($\gamma = 2.2$), *adaptive logarithmic mapping* [DMAC03], *photographic tone reproduction (global)* [RSSF02], *photoreceptor* [RD05]; and the following local (detail preserving algorithms): *gradient domain compression* [FLW02], *bilateral filtering* [DD02], *lightness perception* [KMS05], *photographic tone reproduction (local)* [RSSF02]. The tone mapped LDR images were obtained either from the authors of these methods or by using publicly available implementations [pfs]. Tone mapping parameters were fine tuned whenever default values did not produce satisfactory images.

In practice, the contrast detection component of our Detail Visibility Change metric required calibration to correctly estimate the visibility of subtle details in extreme dark and light regions. We introduced a scaling factor to Equation 8 to increase the predicted response of the HVS to contrasts, and found that a value of 1.89 led to satisfactory predictions in our set of test images. The display characteristics corresponded to a typical consumer LCD with an sRGB response, black level at $2.5cd/m^2$, and white level at $210cd/m^2$ measured in office illumination conditions.

In our analysis, we measure the Global Contrast Change according to Equation (4). However, in the case of Detail Visibility Change we limit possible analysis to the case when visible details in the HDR image become invisible in the tone mapped image. This is a significantly more important case than contrast magnitude change. Following [YBMS05], we perform Detail Visibility analysis separately on dark and light areas. To distinguish these areas, we compare the HDR pixel luminance to the logarithmic average luminance of the HDR image.

The results of analysis are summarized on plots shown in Figure 6. In terms of Global Contrast Change, the advantage of the *photographic tone reproduction (local & global)* methods is clearly visible – global contrast impression is conveyed almost without any change. These methods were also among the top rated in other studies [LCTS05, YBMS05]. However, they result in the loss of detail information in dark areas. This is particularly interesting for the local version of the operator. Although physical high frequency contrast has been preserved in dark areas, the luminance has been mapped to very low values, making the detection of these details impossible. The advantage of the local version is obvious in the light areas, where tone mapping led to a verbatim detail preservation. There is a visible tendency of local methods to be less efficient in limiting the Global Contrast Change, while being better for detail visibility. A notable example is the *gradient domain compression*, which preserves details in dark areas at the cost of a significant reduction to the ratio between light and dark areas. Clearly, the trade-off between detail visibility and global contrast seems unavoidable. One interesting exception is *adaptive logarithmic mapping*, which is able to limit the change in global contrast while performing exceptionally well at preserving detail visibility in dark areas.

From the standard deviation bars of Figure 6, it can be seen that the performance of each operator varied over the image set, meaning that a universal tone mapping operator has not yet been invented. Since the discovery of a universal operator seems unlikely, instead of developing a new algorithm, we decide to counter the distortions with enhancements to the tone mapped LDR images using the distortion information obtained from our Global Contrast Change metric and Detail Visibility Change map.

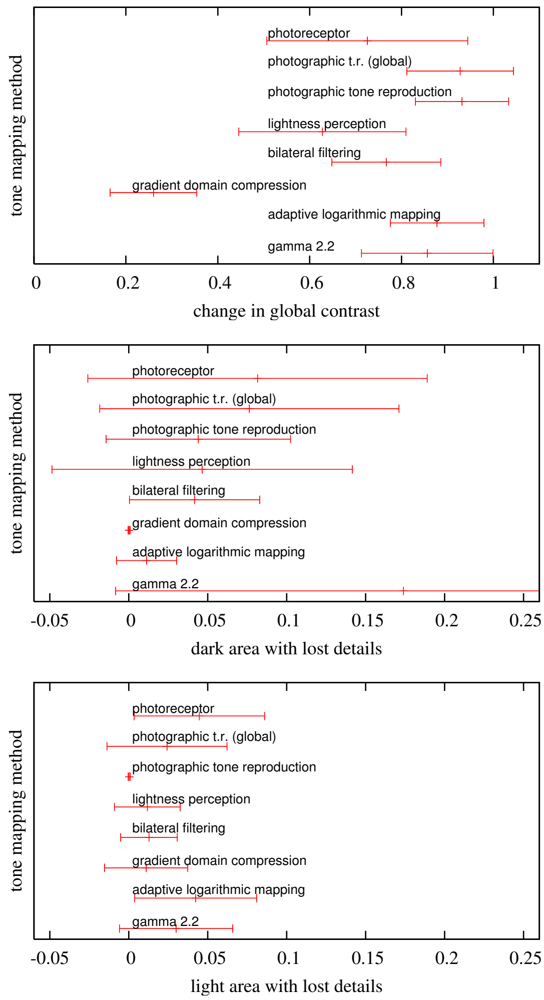

Figure 6: *Analysis of distortions. Top: Global Contrast Change, value* 1 *represents no contrast change,* < 1 *denotes contrast attenuation. Middle: Detail Visibility in dark areas expressed as a relative area where details became invisible after tone mapping. Bottom: Detail Visibility in light areas. Each bar represents mean value and standard deviation.*

4. Contrast Restoration by Colour Adjustment

Following the measurement of tone mapping distortion in terms of change in *Global Contrast* and *Details Visibility*, we approach the problem of compensating for these distortions. For each type of distortion, we identify a *restoration technique* that uses the contrast change information to introduce new contrasts that restore the original contrast information. Our technique for restoring global contrast is countershading, and our technique to restore detail visibility is a per-pixel contrast increase between detail pixels and their immediate neighbourhood. Since tone mapping involves a trade-off between detail preservation and global contrast compression, one restoration technique will be prominent. When both distortions are present, they will be slight, allowing the combination of both restorations. In these cases, it is best to begin with detail restoration so as not to disrupt the global effect of countershading. Although in this paper we introduce only colour contrasts, the restoration techniques are general tools and can be used to add any kind of visual contrast.

Visual contrast is created in many ways [CF03], colour being one very important contributor. Colour contrast is the perceived difference that arises from the juxtaposition of two different colours. Such contrast is ideal for restoring lost luminance contrast resulting from HDR compression because colour is a flexible and aesthetic property, and because colour is often treated too casually by tone mapping operators. An additional reason for working with colour is to limit changes to the high quality luminance compression resulting from the tone mapping. Of the seven general types of colour contrast [Itt61], we focus on *contrast of saturation* and *contrast of complements*, illustrated in Figure 7. We choose to create contrast of saturation and complements by scaling chroma, a perceptual measure of colourfulness, with negative scale values moving the colour towards a saturated opponent colour. We choose to adjust colours by chroma scaling because its independence from the lightness channel provides a straightforward way to limit luminance modification.

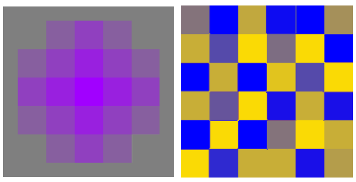

Figure 7: Contrast of saturation *and* contrast of complements *with varying luminance (Weber State University).*

Chroma scaling strengthens image colourfulness, a common trend in photography, and althrough increased colourfulness leads to somewhat unnatural images, up to a maximum colourfulness point they are still preferred by humans [FdB97]. This phenomena exists partly because images are usually judged without direct reference to the original scene and memory for coloured objects can be unreliable [Bar60], so manipulated chroma increase often remains unnoticed while the perceived image quality is consistently ranked higher.

We work in the approximately uniform perceptual CIE $L^*u^*v^*$ colour space where axis L^* represents perceived lightness and u^* and v^* are roughly decorrelated chromatic axes coinciding with red/green and yellow/blue opponent hue pairs. This space is ideal for our image enhancement because luminance is related to a perceptual scale of lightness and because the space provides a correlate of chroma, C^*_{uv}, defined as

$$C^*_{uv} = (u^{*2} + v^{*2})^{1/2}, \tag{11}$$

which can be interpreted as a perceptual measure of colourfulness with respect to a white of similar brightness [Hun95]. Scaling chroma to mC^* by scaling both u^* and v^* by m increases or decreases the perceived colourfulness without changing hue angle h_{uv} or lightness L^*, where h_{uv} is a correlate of hue defined by $\arctan(v^*/u^*)$. When $m \geq 1$, chroma increases (colour becomes saturated with respect L^*), and when $m < 1$, colours become desaturated until they are achromatic and then become saturated in the opponent hue. The colour difference in $L^*u^*v^*$ between two colours differing only by scaled chroma is then

$$\Delta C^*_{uv} = |C^*_{uv}(m-1)|. \tag{12}$$

4.1. Global Contrast Technique

Global Contrast Restoration introduces countershading to the image to enhance the perceived dynamic range, thus making the LDR global contrast impression approach that of the original global contrast. We define countershading as the juxtaposition of gradients on either side of large feature boundaries (often foreground/background) creating an higher contrast border that gives the illusion of greater global contrast. Artists employ countershading, recall Figure 1, often creating a controlled halo at large feature boundaries to increase the perceived brightness difference between the feature and its surround, helping the HVS perform the cognitive task of segmenting features from the background [Tum99]. Visual perception has attempted to explain this phenomemon, with the effect being known as the Craik-Cornsweet-O'Brien illusion: a local attribute (the Cornsweet edge) has a global effect, or as the perceived brightness of two adjacent regions is affected by the contrast at their border, Figure 8 [KM88].

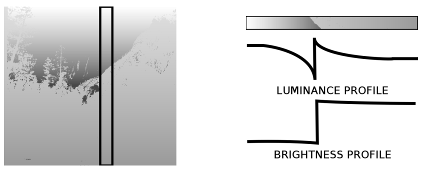

Figure 8: *Example of countershading with luminance values (left), the luminance profile results in a different brightness profile.*

Because we work with chroma, we do not make exact use of the specific Cornsweet illusion. Instead, we create a Cornsweet-style edge of chroma contrast along the border between the foreground and background of the image by applying chroma scale values m defined as:

$$m_{i,j} = \begin{cases} a \cdot exp(-d^2/\sigma^2) + 1, & I_{i,j} \in \text{Segment A} \\ a \cdot (1 - exp(-d^2/\sigma^2)) + 1, & I_{i,j} \in \text{Segment B} \end{cases} \tag{13}$$

where d is the shortest distance to the border ($0 \le d \le 1$), a is the amplitude of the scaling, and σ specifies the width of the slope of scale values. The amplitude of the scaling is $a = 2(1 - \mathscr{C})$, for $\mathscr{C} \le 1$, otherwise no enhancement is required. The width of the slope, σ, is set according to the image size (we use $\sigma = 0.5$ for all our images as they have similar size). We obtain the border by partitioning the HDR luminance Y into two segments using *K-means* image segmentation, which was satisfactory for our test images. A more elaborate image segmentation technique can be used for in challenging cases, for instance, when a simple segmentation returns unadjoining regions or when regions do not meet the image boundaries [KMS05].

4.2. Details Visibility Technique

Detail Visibility Restoration is a per-pixel operation that directly increases the visual contrast between detail pixels and their surrounding neighbourhood, thereby improving their salience in the image. This operation works much like modulating a base signal with a detail signal. To create contrast by chroma scaling, we locally increase pixel chroma proportionally to the detail visibility loss mask $\Delta\overline{T}^*$ from Equation 10. We operate only on HDR high frequency pixels (determined with bilateral filter as in Section 3.1), making our approach different from a global chroma increase, which is commonly performed to improve the overall perceived quality of images [FdB97].

Given a tone mapped LDR image I in $L^*u^*v^*$, we first set pixels with undefined C^*_{uv} or C^*_{uv} drastically different from the HDR image to properly scaled HDR u^* and v^* values, thus reintroducing chromatic information that has been lost due to clamping. We define the enhanced image I' as

$$I' = (L^*_{uv}(I),\ mC^*_{uv}(I),\ h_{uv}(I)) \tag{14}$$

for scale values m that are determined by relating the JND values of $\Delta\overline{T}^*$ to $\Delta C^*_{uv}(I', I) = |C^*_{uv}(I)(m-1)|$. We have found through experiments that one JND is approximately $6.89\ \Delta C^*_{uv}$ units.

$$m = \pm\left(\frac{6.89\ \overline{T}^*(I_Y, I_L)}{C^*_{uv}(I)} + 1\right) \tag{15}$$

Since humans favour increased saturation, we favour the positive scale values, and selectively use negative chroma scale values if they do not break the impression of naturalness [deR96]. The amount of affected pixels is controlled by parameters on the bilateral filter. We set a minimum reliable chroma value so as not to scale any unreliable pixels, thus avoiding the enhancement of noise. Additionally, the white point of the image drastically impacts the overall impression of the image, so we do not modify pixels within the 99*th* luminance percentile.

5. Results and Discussion

The resulting enhanced images maintain a natural quality, and as such, some enhancements can be subtle. For this reason, and because of colour infidelity in print, the results are best visualized on screen. A tone mapping from *bilateral filtering* [DD02] exhibits light global contrast loss of $\mathscr{C} = 0.8747$. While the effect of countershading to restore global contrast is subtle in print, on screen the increased chroma contrast at the horizon serves to emphasize the separation between the sky and the mountains, as shown in Figure 9. The Strasbourg image, resulting from a gradient method tone mapping [MMS06], exhibits greater global contrast loss ($\mathscr{C} = 0.47626$) and so the effect of countershading is correspondingly stronger, as shown in Figure 12.

Figure 9: *Global contrast enhancement of* bilateral filter *tone mapping (top) results in an enhanced image (bottom).*

We begin by illustrating detail visibility restoration on a poor quality LDR image resulting from simple gamma correction in Figure 10. This didactic result exemplifies how our enhancement technique reintroduces details and chromatic information into areas where they have been lost, in this case, enhancing the bleached sky surrounding the sun. Figure 11 shows the details visibility loss mask for the café image tone mapped by *photoreceptor* [RD05], and two enlargements depicting the flower details and the more detailed distant landscape. The original and enhanced LDR images are shown in Figure 13, where details in the outdoor areas have closer appearance to the original HDR. In Figure 14, we show the subtle effect of restoring details to the tree image, tone mapped by *adaptive logarithmic mapping* [DMAC03], which results in brighter, more detailed background trees and textured sand, as in the original HDR (best viewed on screen).

Figure 10: *Gamma corrected LDR $\gamma = 2.2$ (top), and enhanced LDR (bottom) with HDR chroma and details.*

Figure 11: *Café image Details Visibility Change mask (leftmost); Enlargements of original (left) and enhanced (right).*

6. Conclusions

In this work we take a non-standard approach to the problem of depicting HDR images for LDR display. Instead of developing yet another algorithm, we provide the means to enhance the depiction of an HDR image produced by an arbitrary tone mapping algorithm, thus restoring original contrast information. Based on experience and conclusions from previous work we identified two major distortions introduced to luminance while tone mapping: Global Contrast Change and Detail Visibility Change. To our knowledge, we present the first objective perceptual metrics for the measure of contrast distortions between an HDR image and its LDR depiction. To construct these metrics, we extended the transducer function to handle HDR luminance levels. We analyzed selected tone mapping operators using our metrics and we provided an indicative characterization of these operators in terms of global contrast and detail preservation in dark and light regions.

Driven by these metrics, we present techniques for creating enhanced images that restore the original HDR contrast information with colour contrast achieved by chroma scaling. The increased colour contrast augments detail visibility, and countershading encourages the prominence of foreground objects, thus reclaiming the loss in perceived global contrast. Luminance values are not affected by our adjustments so that we do not distort the users choice of desired tone mapping. Instead, using the perceptually meaningful distortion measures, we introduce corrections to the LDR image to compensate for distortions while preserving the preferred tone mapped luminance.

Since only luminance values are evaluated by our distortion metrics, their application is most suitable for the luminance-based subset of tone mapping operators. Consideration to colour contrast and an additional metric for analyzing HVS colour reproduction could further improve the existing metrics. Our techniques for distortion detection and magnitude evaluation can be used with other methods of perceived contrast enhancement [CF03], including luminance manipulation, an enhancement method recently exploited in [LCD06]. It would also be interesting to design an algorithm for the effective combination of perceived contrast restoration methods. We recognize that subjective experiments comparing the original HDR to its enhanced LDR are required to fully evaluate the success of our approach, and we consider this an important part of our future work.

Acknowledgments

We would like to acknowledge OpenEXR, Greg Ward, and SpheronVR for making their images available, and we thank Rafał Mantiuk and the anonymous reviewers for their valuable comments concerning this work.

References

[Bar60] BARTLESON C.: Memory colors of familiar objects. *J. Opt. Soc. Am. 50*, 1 (1960), 73–77. 8

[CF03] CALABRIA A., FAIRCHILD M.: Perceived image contrast and observer preference: The effects of lightness, chroma, and sharpness manipulations on contrast perception. *J. Imag. Sci.and Tech. 47* (2003), 479–493. 1, 2, 8, 10

[CH03] CHANDLER D., HEMAMI S.: Suprathreshold image compression based on contrast allocation and global precedence. In *Human Vision and Electronic Imaging VIII* (2003), SPIE, volume 5007, pp. 73–86. 3

[CIE81] CIE: *An Analytical Model for Describing the Influence of Lighting Parameters Upon Visual Performance*, vol. 1. Technical Foundations. Intl. Org. for Standardization, 1981. 6

[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. 3

[DD02] DURAND F., DORSEY J.: Fast bilateral filtering for the display of high-dynamic-range images. *ACM Transactions on Graphics 21*, 3 (July 2002), 257–266. 6, 9

[deR96] DERIDDER H.: Naturalness and image quality: saturation and lightness variation in color images of natural scenes. *J. Imag. Sci.and Tech. 40*, 6 (1996), 487–493. 9

[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. 4, 6, 10

[DW00] DICARLO J. M., WANDELL B. A.: Rendering high dynamic range images. In *Sensors and Camera Systems for Scientific, Industrial, and Digital Photography Applications* (May 2000), vol. 3965 of *Processings of SPIE*, pp. 392–401. 5

[FdB97] FEDOROVSKAYA E., DERIDDER H., BLOMMAERT F.: Chroma variations and perceived quality of color images of natural scenes. *Color Research and Application 22*, 2 (1997), 96–110. 8, 9

[FJ03] FAIRCHILD M., JOHNSON G.: Image appearance modeling. In *Human Vision and Electronic Imaging VIII* (2003), SPIE, volume 5007, pp. 149–160. 3

[FLW02] FATTAL R., LISCHINSKI D., WERMAN M.: Gradient domain high dynamic range compression. *ACM Transactions on Graphics 21*, 3 (July 2002), 249–256. 3, 6

[GH03] GREENFIELD G. R., HOUSE D. H.: Image recoloring induced by palette color associations. In *Proceedings of WSCG* (2003), pp. 189–196. 3

[Hun95] HUNT R.: *The Reproduction of Color. 5th edition*. Fountain Press, Kingston-upon-Thames, 1995. 1, 3, 8

[Itt61] ITTEN J.: *The Art of Color*. Van Nostrand Reinhold, New York, 1961. 8

[ITU90] ITU .: *Basic Parameter Values for the HDTV Standard for the Studio and for International Programme Exchange*. ITU-R Recommendation BT.709, Geneva, 1990. 4

[KM88] KINGDOM F., MOULDEN B.: Border effects on brightness: A review of findings, models and issues. *Spatial Vision 3* (1988). 8

[KMS05] KRAWCZYK G., MYSZKOWSKI K., SEIDEL H.-P.: Lightness perception in tone reproduction for high dynamic range images. *Computer Graphics Forum 24*, 3 (2005). 6, 9

[KYJF04] KUANG J., YAMAGUCHI H., JOHNSON G., FAIRCHILD M.: Testing hdr rendering algorithms. In *IS&T/SID Color Imaging Conference* (2004), pp. 315–320. 3

[LCD06] LUFT T., COLDITZ C., DEUSSEN O.: Image enhancement by unsharp masking the depth buffer. *ACM Transactions on Graphics* (2006). To Appear. 10

[LCTS05] LEDDA P., CHALMERS A., TROSCIANKO T., SEETZEN H.: Evaluation of tone mapping operators using a high dynamic range display. *ACM Transactions on Graphics 24*, 3 (2005), 640–648. 1, 3, 7

[Liv02] LIVINGSTONE M.: *Vision and Art: The Biology of Seeing*. Harry N. Abrams, 2002. 2

[Lub95] LUBIN J.: A visual discrimination model for imaging system design and development. In *Vision models for target detection and recognition* (1995), E. P., (Ed.), World Scientific, pp. 245–283. 3

[MMS06] MANTIUK R., MYSZKOWSKI K., SEIDEL H.-P.: A perceptual framework for contrast processing of high dynamic range images. *ACM Trans. Applied Perception 3*, 3 (2006). 5, 6, 9

[MNN*05] MATKOVIC K., NEUMANN L., NEUMANN A., PSIK T., PURGATHOFER W.: Global contrast factor-a new approach to image contrast. In *Computational Aesthetics in Graphics, Visualization and Imaging 2005* (May 2005), L. Neumann M. Sbert B. G. W. P., (Ed.), Eurographics Association, pp. 159–168. 2

[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. *Computer Graphics 32*, Annual Conference Series (1998), 287–298. 3

[pfs] PFSTMO: *Open Source Implementation Of Tone Mapping Operators*. http://www.mpii.mpg.de/resources/tmo/. 6

[RAGS01] REINHARD E., ASHIKHMIN M., GOOCH B., SHIRLEY P.: Color transfer between images. *IEEE Comput. Graph. Appl. 21*, 5 (2001), 34–41. 3

[RD05] REINHARD E., DEVLIN K.: Dynamic range reduction inspired by photoreceptor physiology. *IEEE Transactions on Visualization and Computer Graphics 11*, 1 (2005), 13–24. 6, 9

[RGW05] RASCHE K., GEIST R., WESTALL J.: Re-coloring images for gamuts of lower dimension. *Computer Graphics Forum, proceedings of Eurographics 2005 24*, 3 (2005). 3

[RSSF02] REINHARD E., STARK M., SHIRLEY P., FERWERDA J.: Photographic tone reproduction for digital images. *ACM Transactions on Graphics 21*, 3 (2002), 267–276. 3, 4, 6

[RWPD05] REINHARD E., WARD G., PATTANAIK S., DEBEVEC P.: *High Dynamic Range Imaging: Acquisition, Display, and Image-Based Lighting*. Morgan Kauffman, 2005. 1, 3

[Tum99] TUMBLIN J.: *Three Methods For Detail-Preserving Contrast Reduction For Displayed Images*. PhD thesis, Georgia Institute of Technology, 1999. 8

[Win05] WINKLER S.: *Digital Video Quality: Vision Models and Metrics*. John Wiley & Sons, Ltd, West Sussex, England, 2005. 3

[YBMS05] YOSHIDA A., BLANZ V., MYSZKOWSKI K., SEIDEL H.-P.: Perceptual evaluation of tone mapping operators with real-world scenes. In *Human Vision and Electronic Imaging X* (2005), SPIE, volume 5666, pp. 192–203. 1, 3, 7

Figure 12: *Enhance global contrast: Strasbourg image resulting from a gradient method tone mapping (left), and with countershading (right). In this example, the higher contrast created with chroma helps to evoke a greater sense of scene depth.*

Figure 13: *Detail restoration: Café image resulting from photoreceptor tone mapping (left), and with enhanced details (right). Notice that the flowers, chairs and umbrellas are more visible and the distant landscape contains more details and depth.*

Figure 14: *Detail restoration: Tree image resulting from adaptive logarithmic mapping (left), and with enhanced details (right). Notice that the background trees are more pronounced and there is additional texture in the sandy areas.*