

Tone Mapping For High Dynamic Range Displays

Laurence Meylan^a, Scott Daly^b, and Sabine Süsstrunk^a

^aEcole Polytechnique Fédérale de Lausanne (EPFL), Switzerland;

^bSharp Laboratories of America, Camas, WA

ABSTRACT

We address the problem of re-rendering images to high dynamic range (HDR) displays, which were originally tone-mapped to standard displays. As these new HDR displays have a much larger dynamic range than standard displays, an image rendered to standard monitors is likely to look too bright when displayed on a HDR monitor. Moreover, because of the operations performed during capture and rendering to standard displays, the specular highlights are likely to have been clipped or compressed, which causes a loss of realism. We propose a tone scale function to re-render images first tone-mapped to standard displays, that focuses on the representation of specular highlights. The shape of the tone scale function depends on the segmentation of the input image into its diffuse and specular components. In this article, we describe a method to perform this segmentation automatically. Our method detects specular highlights by using two low-pass filters of different sizes combined with morphological operators. The results show that our method successfully detects small and middle sized specular highlights. The locations of specular highlights define a mask used for the construction of the tone scale function. We then propose two ways of applying the tone scale, the global version that applies the same curve to each pixel in the image and the local version that uses spatial information given by the mask to apply the tone scale differently to diffuse and to specular pixels.

Keywords: Tone Mapping, Detection of Specular Highlights, High Dynamic Range Displays

1. INTRODUCTION

Recently, new high dynamic range (HDR) monitors have been developed that have a much larger dynamic range than standard CRTs or LCDs.¹ These monitors are composed of an array of ultra-bright LEDs that illuminate a standard LCD panel. Each LED can be controlled independently, which provides a gain in dynamic range given by the different LED intensities modulated by the LCD cell transmittivity. As these displays become more widely available, we are confronted with the need for new tone mapping algorithms adapted to these new display conditions.

In this article, we address the problem of re-rendering legacy images that have already been tone-mapped. These images have lost part of their dynamic range in the rendering and capturing process to standard displays. In particular, the specular highlights have usually been clipped or compressed. We propose to use a tone scale function that allocates more display range to the specular highlights than what was allocated in the image tone-mapped to standard displays. The goal of this tone scale function is to recover the “natural look” of the original scene on HDR displays to obtain a visually pleasing image. In a previous article,² we reported the results of a psychovisual experiment that validated the benefit of our tone scale function. We showed that allocating more display range to the specular highlights than what was allocated in the image rendered to standard displays provides a more natural visual sensation.

The shape of the proposed tone scale depends on a segmentation of the input image into its diffuse and specular components (illustrated in Figure 1). For the psychovisual experiment previously described,² the segmentation was performed manually. In this article, we provide a method to automatically segment the input image into its diffuse and specular components. We then use the result of the segmentation to construct the tone scale, which is applied to each pixel of the image. We also propose an alternative way of applying the tone scale that uses spatial information resulting from the segmentation.



Figure 1. Example of an image segmentation into its specular and diffuse components. **Left:** Original image. **Center:** Diffuse image. The specular part of the image is filled with black. **Right:** Specular image. The diffuse part of the image is filled with black.

This article is structured as follows: Section 2 first gives some background on the problem of detecting specular highlights in images. We then describe our algorithm to segment an image into its diffuse and specular component and show the results. Section 3 presents our proposed tone scale and how it is constructed using the segmentation presented in Section 2. Section 4 proposes a spatially variant method to apply the tone scale and illustrates the difference with the global technique. Section 5 concludes the article.

2. AUTOMATIC DETECTION OF SPECULAR HIGHLIGHTS

2.1. Background

Detecting specular highlights, which often means separating an image into diffuse and specular components, is not an easy task. It has been addressed by many researchers for various purposes. Klinker *et al.*^{3,4} separate diffuse from specular components using the distribution of color clusters in color spaces. They show that diffuse and specular surfaces form a “T” in the dichromatic plane of the RGB color space. The base of the “T” points to the color of the illuminant. However, this does not hold for saturated pixels since these pixels do not generally obey the characteristics of the dichromatic reflection model.

The specular highlights in an image are often clipped and this creates problems for most detection algorithms based on the dichromatic model. In [Tan et al. 2004],⁵ the goal is to estimate the illuminant using specular highlights. The first part of their method, based on earlier work,⁶ consists in finding the specular highlights by thresholding on brightness and saturation. They then estimate the illuminant using an inverse chromaticity space in which the correlation between illumination chromaticity and image chromaticity can be obtained. No special care is taken for clipped pixels. A method to discard clipped pixels is proposed in Park and Kak⁷ so that the distribution obeys the characteristics of the dichromatic model. Osadchy *et al.*⁸ use specularities for recognition of a known 3D object in an image. They detect specular highlight candidates using a threshold operation but also use 3D information of the known object to reject false candidates. They use additional assumptions that are worth mentioning as our technique is also based on them: specular highlights are usually monotonically decreasing and rotationally symmetric. Similar to these methods, our algorithm uses thresholding to discriminate between the diffuse and the specular components of an image. A key difference is that our threshold is determined by spatial filtering and morphological operators, which makes the detection of specular highlights more robust.

2.2. Automatic Segmentation and Definition of the Maximum Diffuse White

Our algorithm detects the specular highlights by applying two low-pass filters to the luminance channel of the input image, combined with morphological operators. The detection of specular highlights results in a binary map where 1 indicates a pixel belonging to a specular highlight. Figure 2 illustrates the process.

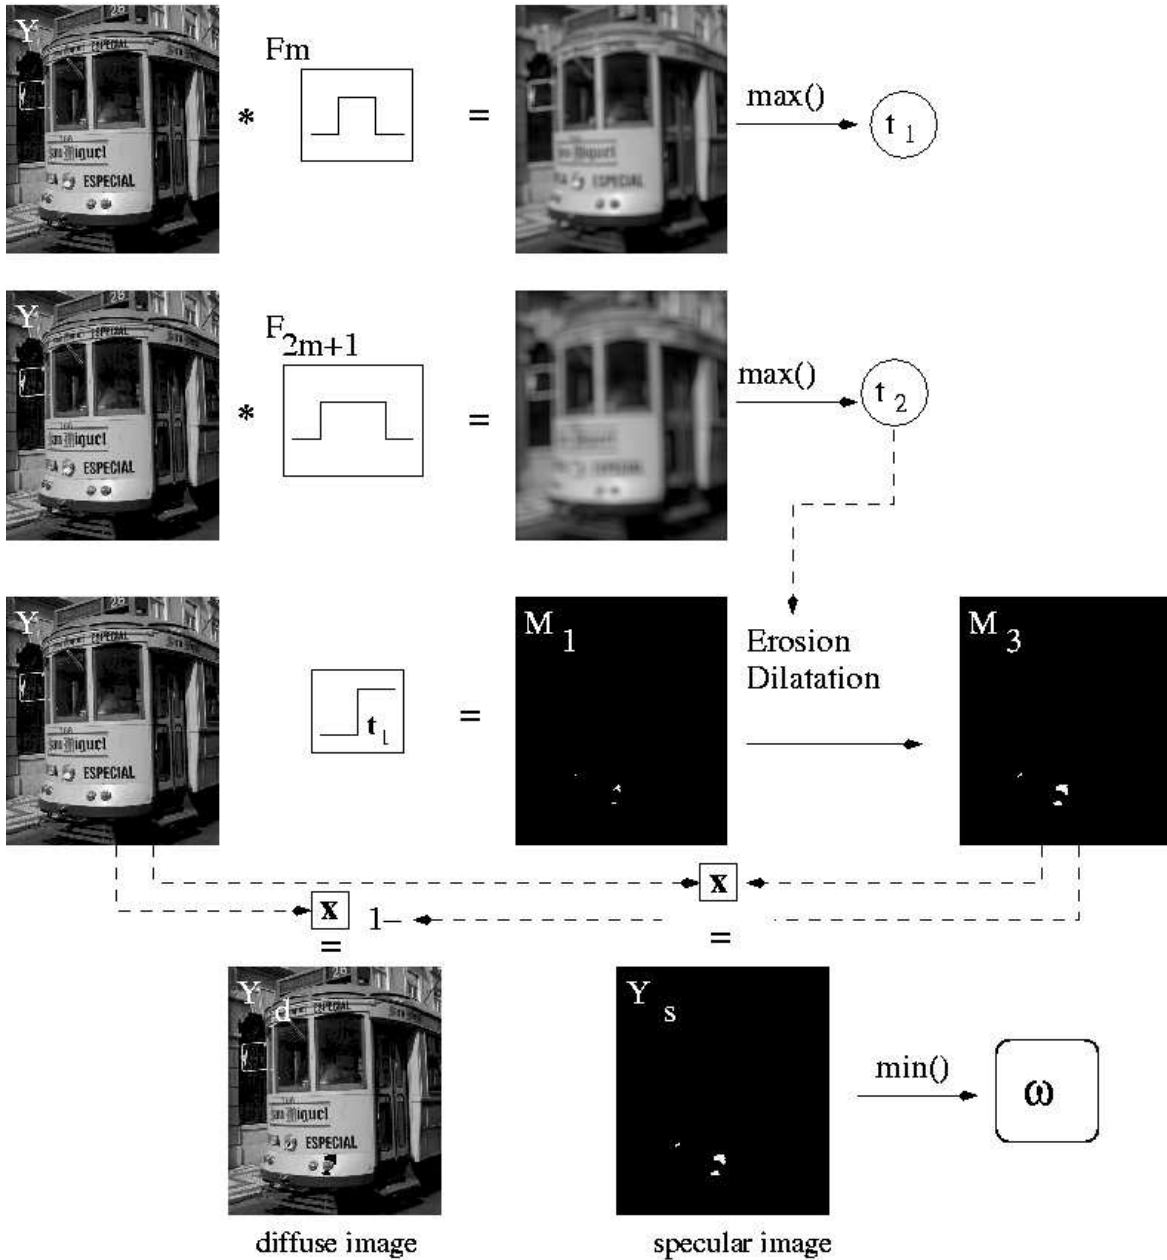


Figure 2. Definition of maximum diffuse white ω using two low-pass filters and morphological operators.

We use the assumption that specular highlights are small and bright. Under this assumption, the maximum pixel value of the low-pass filtered image gives the value of the maximum diffuse white and thus can serve as a threshold to segment between diffuse and specular pixels. However, it can occur that white diffuse parts of the image are brighter than the region bordering the specular highlights. We take that into account by using two low-pass filters of different sizes combined with morphological operators. The luminance channel Y of the input image I is low-passed by a first filter \mathcal{F}_m . Taking the maximum digital value of the low-passed image results in a first threshold t_1 (1) that is used to define the specular highlight candidates under the form of a binary map

$\mathcal{M}_1(2)$.

$$t_1 = \max_{p \in I} (Y(p) * \mathcal{F}_m), \quad (1)$$

where p is a pixel in the image I and \mathcal{F}_m is an $m \times m$ averaging filter. We use m equals $\frac{1}{50}$ of the image size in the vertical direction.

$$\mathcal{M}_1(p) = \begin{cases} 0 & \text{if } Y(p) \leq t_1 \\ 1 & \text{if } Y(p) > t_1 \end{cases} \quad (2)$$

Independently of this first step, the input image is low-passed again by a filter that is twice as large \mathcal{F}_{2m+1} . The maximum of the second low-pass filtered image defines a second threshold t_2 :

$$t_2 = \max_{p \in I} (Y(p) * \mathcal{F}_{2m+1}). \quad (3)$$

The specular highlight candidate map \mathcal{M}_1 obtained in (2) is now processed with morphological operators. At first, an erosion is applied in order to remove single pixels (4). After this first step, we consider that each group of white pixels corresponds to the center of a specular highlight. Dilation is performed iteratively on the specular highlight candidate map \mathcal{M}_2 constrained by the second threshold t_2 . That is, a pixel is turned from non-specular to specular if more than three of its neighbors are specular and if $Y(p)$ is greater than t_2 (specular regions are shown white in Figure 2). Iterations continue until a stable state is reached. Usually only a few iterations are necessary. The sequence of morphological operators is given by the following equations:

$$\mathcal{M}_2(p) = \begin{cases} 0 & \text{if } \mathcal{M}_1(p) = 0 \text{ or } \mathcal{M}_1(p) * \mathcal{H} < 1 \\ 1 & \text{otherwise} \end{cases}, \quad (4)$$

where

$$\mathcal{H} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}. \quad (5)$$

$$\mathcal{M}_3(p) = \begin{cases} 1 & \text{if } \mathcal{M}_2(p) = 1 \text{ or } (\mathcal{M}_2(p) * \mathcal{H} > 3 \text{ and } Y(p) > t_2) \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

The binary map thus obtained can be used as a mask to discriminate between diffuse and specular image. Y_d and Y_s are respectively described as:

$$Y_s(p) = Y(p) \cdot \mathcal{M}_3(p), \quad (7)$$

$$Y_d(p) = Y(p) \cdot (1 - \mathcal{M}_3(p)). \quad (8)$$

We can then define a threshold ω given by the minimum digital value of the specular image Y_s , which in turn determines the maximum diffuse white of the image. The advantage of using the minimum value of the specular image rather than the maximum value of the diffuse image is that it does not depend on the presence of a white surface in the image.

$$\omega = \min_{p \in P} (Y_s(p)), \quad (9)$$

where P is the set of pixels for which $Y_s(p) > 0$.

2.3. Application on Images

In Figure 3, we show the results of segmentation into diffuse and specular components (Y_s, Y_d) for images containing specular highlights. The detected specular highlights (given by the specular image) are shown in red. Small and middle sized specular highlights are detected successfully.



Figure 3. Result of the segmentation applied on images containing specular highlights. The detected specular highlights are shown in red. Small and middle sized specular highlights are detected successfully.

The detection of specular highlights is dependent on the size of the filter \mathcal{F}_m . Using two filters brings some flexibility but detecting large bright highlights remains a problem. In the presence of such a bright area, the maximum of both low-pass filtered images equals the maximum of the input image. An extension to our algorithm to handle larger specular highlights could be performed by using an iterative process, increasing the size of \mathcal{F}_m at each iteration.

3. THE TONE SCALE FUNCTION

Our method to detect specular highlights is used to define the shape of the re-rendering tone scale function. The proposed tone scale is a piecewise linear function composed of two slopes, one that scales the diffuse part of the input image and the other that scales the specular part (Figure 4). It transforms the normalized code values of the input image that is encoded linearly into normalized display values.

The tone scale function is uniquely defined by two parameters, ω and ρ . ω is the digital code value of the maximum diffuse white in the input image mapped to standard displays. It is computed by the method presented in Section 2.2. ρ is the percentage of display luminance allocated to ω after the application of the tone scale.

The tone scale function f is defined as follows:

$$f(Y(p)) = \begin{cases} s_1 \cdot Y(p) & \text{if } Y(p) \leq \omega \\ s_1 \cdot \omega + s_2 \cdot (Y(p) - \omega) & \text{if } Y(p) > \omega \end{cases}, \quad (10)$$

where




Figure 4. Piecewise linear tone scale function.

$$s_1 = \frac{\rho}{\omega}, \quad (11)$$

$$s_2 = \frac{1 - \rho}{Y_{max} - \omega}. \quad (12)$$

Y is the normalized luminance and p is a pixel in the image. The maximum digital value of the input image is noted Y_{max} . Y_{max} is normalized to 1 to make the method independent of the digital code value range.

A Smoothing Technique to Remove Unnatural Contours

A difficulty caused by our piece-wise linear function is that there is a sudden change of angle between the two slopes. This can cause unnatural contours (quantization effects) to appear around specular highlights.⁹ To remove these contours, we developed a smoothing algorithm that is applied in addition to the tone scale function. The smoothing algorithm uses the binary map given by the segmentation to slightly blur the tone-scaled image around the borders of specular highlights.

The principle of the smoothing algorithm is to smooth the contours around specular highlights by applying a low-pass filter only at the location of specular highlights (Figure 5). First, a thresholding operation is applied to the luminance channel of the original image Y to find the location of specular highlights:

$$\mathcal{B}_1(p) = \begin{cases} 0 & \text{if } Y(p) \leq \omega \\ 1 & \text{if } Y(p) > \omega. \end{cases} \quad (13)$$

$\mathcal{B}_1(p)$ is a binary map, where 1 indicates the presence of a specular highlight. p is a pixel in the image I .

Then, a dilation is applied to $\mathcal{B}_1(p)$. A pixel is turned from 0 to 1 if more than one of its neighbors are 1:

$$\mathcal{B}_2(p) = \begin{cases} 1 & \text{if } \mathcal{B}_1(p) = 1 \text{ or } (\mathcal{B}_1(p) * \mathcal{H} > 1) \\ 0 & \text{otherwise,} \end{cases} \quad (14)$$

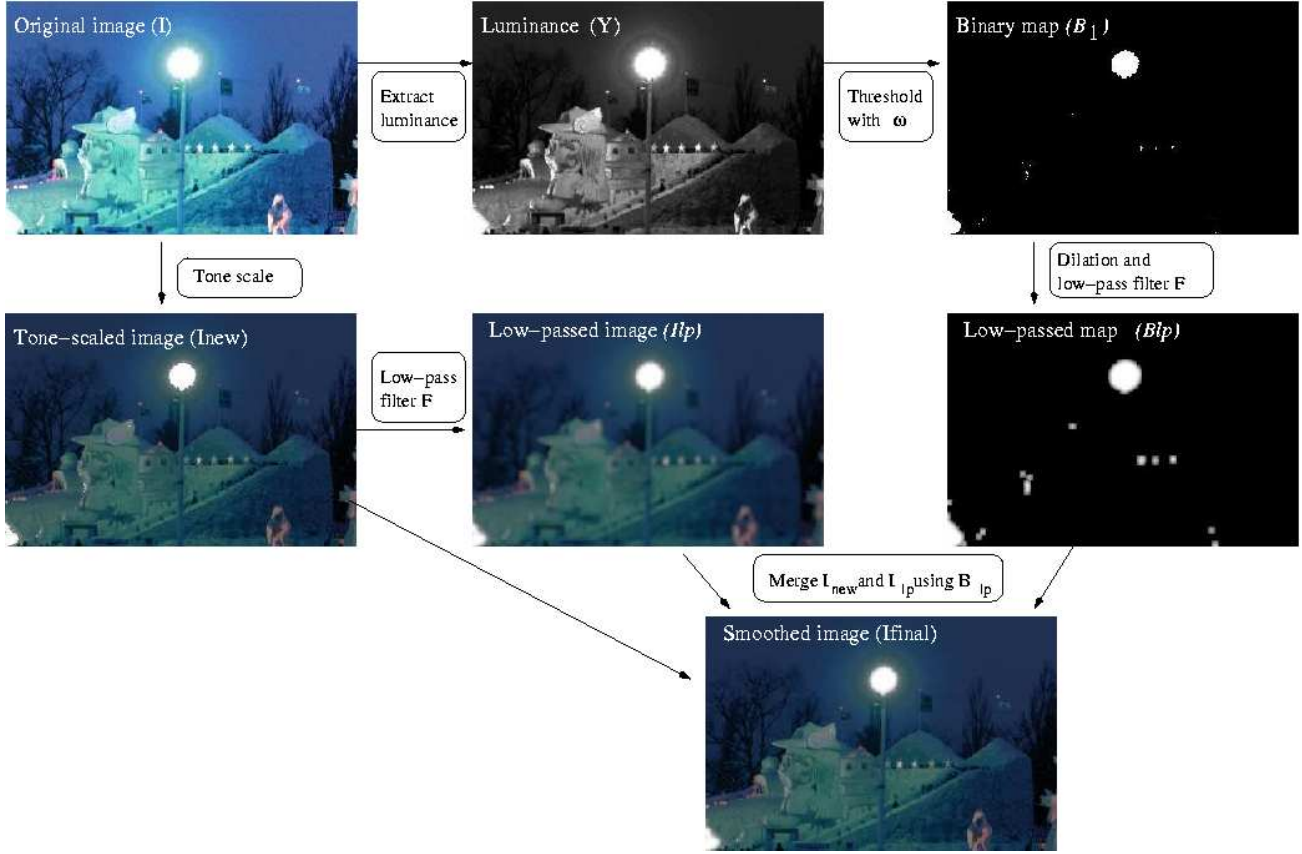


Figure 5. Illustration of the smoothing technique to remove unnatural contours (see text for explanation).

where

$$\mathcal{H} = \begin{pmatrix} 1 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 1 \end{pmatrix}. \quad (15)$$

Both the tone-scaled image I_{new} and the binary map $B_2(p)$ are then low-passed with a moving average filter of size 5 denoted by \mathcal{F}_5 .

$$\begin{aligned} I_{lp}(p) &= I_{new}(p) * \mathcal{F}_5 \\ B_{lp}(p) &= B_2(p) * \mathcal{F}_5 \end{aligned} \quad (16)$$

The final image I_{final} where contours are removed is simply computed by merging the tone-scaled image I_{new} and the low-passed tone-scaled image I_{lp} using the blurred mask B_{lp} as a matrix of weighting coefficients.

$$I_{final}(p) = I_{new}(p) \cdot (1 - B_{lp}(p)) + I_{lp}(p) \cdot B_{lp}(p) \quad (17)$$

The results of applying the smoothing algorithm to images tone-scaled with $\rho = 20\%$ are shown in Figure 6. We observe that the unnatural contours around the specular highlights are successfully removed without affecting the sharpness of the rest of the image.




Figure 6. Application of the smoothing algorithm on two images. **Left:** Image after tone scale ($\rho = 20\%$). **Right:** Tone-scaled image after smoothing. We observe that the unnatural contours around the specular highlights are successfully removed without affecting the sharpness of the rest of the image.

4. APPLYING THE TONE SCALE FUNCTION USING A MASK

The straightforward way to use our tone scale function is to apply it globally to each pixel of the image (18). This technique was used in² and leads to pleasing results. However, due to the way ω is computed using morphological operators, there could be diffuse parts of the image whose pixel values are greater than ω . In this case, it is advantageous to use the mask \mathcal{M}_3 (6) to scale the diffuse image and the specular image differently. The diffuse image is scaled by the first linear function even for pixels whose values are greater than ω and the specular image is scaled by the second linear function. This is illustrated in Figure 7. We call it local application of the tone scale. Global and local application of the tone scale are given by (18) and (19), respectively.

$$Y_{new}(p) = f(Y(p)) \quad (18)$$

$$Y_{new}(p) = \begin{cases} s_1 \cdot \omega + s_2 \cdot (Y(p) - \omega) & \text{if } \mathcal{M}_3(p) = 1 \\ s_1 \cdot Y(p) & \text{if } \mathcal{M}_3(p) = 0 \end{cases} \quad (19)$$

Figure 7 shows a comparison between global and local application of the tone scale function. For the local application, the diffuse image is scaled with the dotted slope (s_1) while the specular image is scaled by the solid slope (s_2). The difference can be seen in the zoomed images on the roof of the white truck. The truck is part of the diffuse image but has a luminance value greater than ω . It is scaled to a lower value when the tone scale is applied locally than when it is applied globally and appears darker in the image treated with the local method.

Another example is given in Figure 8. The top left image (a) shows the image scaled by our tone scale applied locally using the specular highlight candidate map (c). The top right image (b) shows the image scaled globally, not taking spatial information (\mathcal{M}_3) into account. Diffuse pixels whose values are greater than the computed ω are also boosted up. Differences between (a) and (b) are indicated by the white arrow. Another way of illustrating these differences is to compute the binary map $\mathcal{M}_d(p)$ (20). In $\mathcal{M}_d(p)$, all pixels whose value is greater than ω is set to 1, as illustrated by the bottom images (d) of Figure 8.

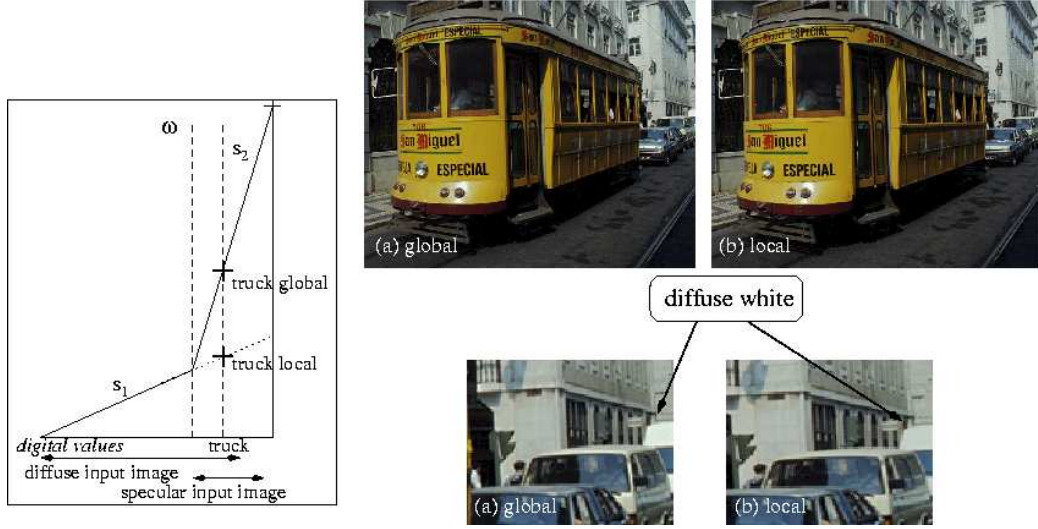


Figure 7. Comparison between global and local application of the tone scale. The truck is scaled to a lower value by the local application and appears darker than if it was scaled by the global method since it is not part of the specular image.

$$\mathcal{M}_d(p) = \begin{cases} 0 & \text{if } Y(p) \leq \omega \\ 1 & \text{if } Y(p) > \omega \end{cases} \quad (20)$$

We combined the automatic detection of ω with the local application of the tone scale and present the results. Figure 9 and 10 compare images originally tone-mapped to SDR display that are scaled linearly (left, a and c), with images treated by our piecewise linear tone scale function (right, b and d). According to the results obtained in a previous article,² we use $\rho = 67\%$, which allocates 33% of the output display range to the specular highlights. The first row (a,b) shows what happens if the images are shown on a SDR device or a print. We can see that the specular highlights of (b) have been boosted but at the expense of global brightness. The second row (c,d) simulates what would happen if these images were displayed on a HDR monitor. This is done by increasing the brightness of the images (a,b) similarly, but the only definitive way to assess the quality of the tone-scaled image is to look at it on a HDR display. What the simulation shows is that image (d) gives about the same impression of brightness than image (a) while (d) has brighter specular highlights (only visible when displayed on a HDR monitor). The simulation of the linearly scaled image displayed on a HDR device (c) shows that the image looks too bright due to the increase in brightness.

5. CONCLUSION

We propose the use of a tone scale function to enhance the appearance of specular highlights when the image is displayed on a HDR monitor. The shape of the tone scale is entirely defined by two parameters, the maximum diffuse white of the input image and the percentage of display luminance allocated to it. This article presents a method to automatically define the maximum diffuse white of an input image by segmenting the image into its diffuse and specular components. This segmentation uses two-low pass filters combined with morphological operators, which is more robust than using a single threshold. We then present a way to take advantage of this segmentation during the application of the tone scale function so that the diffuse and the specular components are treated differently.

ACKNOWLEDGMENTS

Part of this work was supported by the Swiss National Science Foundation under grant number 21-101681.




Figure 8. Comparison between local and global application of the tone scale function. A global application enhances the brightness of some diffuse pixels while a local application prevents that.

REFERENCES

1. H. Seetzen, W. Heidrich, W. Stuerzlinger, G. Ward, L. Whitehead, M. Trentacoste, A. Ghosh, and A. Vorozcov, "High dynamic range display systems," *ACM Transactions on Graphics (special issue SIGGRAPH 2004)* **23**, pp. 760–768, August 2004.
2. L. Meylan, S. Daly, and S. Süsstrunk, "The reproduction of specular highlights on high dynamic range displays," in *Proc. IS&T/SID Fourteenth Color Imaging Conference: Color Science, Systems, and Application*, pp. 333–338, (Scottsdale, AR), November 2006.
3. G. J. Klinker, S. A. Shafer, and T. Kanade, "Using a color reflection model to separate highlights from object color," in *Proc. First International Conference on Computer Vision (ICCV)*, pp. 145–150, (London, UK), 1987.
4. G. J. Klinker, S. A. Shafer, and T. Kanade, "The measurement of highlights in color images," *International Journal of Computer Vision* **2**, pp. 7–32, June 1988.
5. R. T. Tan, K. Nishino, and K. Ikeuchi, "Color constancy through inverse-intensity chromaticity space," *Journal of the Optical Society of America A* **21**, pp. 321–334, March 2004.
6. H.-C. Lee, "Method for computing the scene-illuminant chromaticity from specular highlights," *Journal of the Optical Society of America A* **3**, pp. 1694–1698, October 1986.
7. J. B. Park and A. C. Kak, "A truncated least squares approach to the detection of specular highlight in color images," in *Proc. International Conference on Robotics and Automation*, pp. 1397–1403, (Taipei, Taiwan), September 2003.
8. M. Osadchy, D. Jacobs, and R. Ramamoorthi, "Using specularities for recognition," in *Proc. Ninth International Conference on Computer Vision (ICCV)*, **2**, pp. 1512–1519, (Nice, France), 2003.
9. L. Meylan, *Tone Mapping for High Dynamic Range Images*. PhD thesis, Ecole Polytechnique Fédérale de Lausanne (EPFL), 2006.




Figure 9. Comparison between linear scaling (left) and our tone scale (right). First row: images displayed on an SDR device or print. Second row: Simulation of images (a,b) when displayed on a HDR device. Only the diffuse brightness of these images is simulated. The bright specular highlights of image (d) can not be reproduced.




Figure 10. Comparison between linear scaling (left) and our tone scale (right). First row: images displayed on an SDR device or print. Second row: Simulation of images (a,b) when displayed on a HDR device. Only the diffuse brightness of these images is simulated. The bright specular highlights of image (d) can not be reproduced.