

Color Image Enhancement Using A Retinex-Based Adaptive Filter

Laurence Meylan and Sabine Süsstrunk
LCAV, EPFL, Ecole Polytechnique Fédérale de Lausanne
Switzerland

Abstract

We present a new adaptation of Retinex to enhance the rendering of high dynamic range digital color images. The image is processed using an adaptive Gaussian filter. The shape of the filter basis is adapted to follow the high contrasted edges of the image. In this way, the artifacts introduced by a circularly symmetric filter at the border of high contrasted areas are reduced. This method provides a way of rendering natural images that is inspired by human local adaptation. It is included into a framework that takes raw linear images or radiance maps and outputs 24-bit images rendered for display.

Introduction

Rendering an image that looks “nice” according to human preferences is not a trivial task. The image captured by a digital device often differs from the human perception of the original scene. This is due to the fact that the captured image contains the information given by the physical values of light while humans tend to perceive the relations between objects. Indeed, the human visual system (HVS) is a complex non-linear mechanism that determines the perceived color by spatial comparisons of color signals across a scene and not with absolute values.^{9,10} Applying a similar processing on a captured natural image would bring the reproduced image closer to what the observer remembers.

In this paper, we address the problem of rendering images representing high dynamic range scenes. A scene is said to have a high dynamic range when the ratio between the highest luminance to the lowest luminance by far exceeds the one of the capture or output device. The dynamic range of such a scene has to be compressed to fit the one of the device, which often causes a loss of details in areas of low or high illumination. Recent developments made possible the capture of high dynamic range scenes.⁴ The principle is to capture multiple pictures of the same scene with different exposure times. A so-called radiance-map is built from the acquired pictures. Nevertheless, the problem of mapping the radiance map values to the output device’s dynamic range remains.

Our aim is to mimic the processing of the human visual system on a scene to render high dynamic range images. We take inspiration from an existing model of color vision called Retinex and adapt it to a new

algorithm. We test our method on a set of high dynamic range images of natural scenes.

This article is structured as follows: In section 2, we give an overview of Retinex and its adaptations to computational models. Section 3 has two parts. The first one presents a global framework for rendering high dynamic range images. The second one explains the Retinex-based method that represents the core of the framework. Section 4 presents the results obtained with natural images and section 5 contains our conclusion.

Retinex and Its Adaptations to Computational Models

Retinex theory intends to explain how the visual system extracts reliable information from the world despite changes of illumination. It was developed by Land and is based on a series of psychophysical experiments.^{9,10} Retinex determines the perceived color by spatial comparisons of color surfaces across the whole image. This processing takes place independently in each waveband.

Retinex has been used as a theoretical basis for computational models that have been adapted for color image rendering. Nevertheless, all of these models have drawbacks in their implementation. Three trends can be identified among the computational Retinex adaptations.

A first set of algorithms computes the new value of a pixel using subsequent additions of pixel differences along a set of random one-dimensional paths.^{1,11} These methods have been adapted toward more efficient computations using matrices. The second set of methods computes the pixel values using a recurrent iterative formula.^{3,6,15} They provide good quality images but have a major drawback: the number of iterations is not clearly defined and can strongly influence the final result. If the chosen number of iterations is too large, the resulting image converges toward the original. The best image is obtained with an optimal number of comparisons. Although an on-line stopping method exists,⁷ it is still a difficult task to determine automatically the number of iterations.

A third set of methods are “surround-based”. The algorithm of Rizzi et al. computes the value of a pixel given the sum of differences between the treated pixel and the other pixels in a defined surround.¹⁴ The influence of surrounding pixels is determined by a weighting function. A different surround-based approach is the one of Rahman

et al.^{8,13} Their method computes the enhanced image using a weighted sum of single-scale Retinex and a color restoration factor. A single scale Retinex is computed by subtracting the log-encoded image to a weighted average of the linear image. The weighted average is a convolution of the image with a low-pass filter, also called a surround function. The filter coefficients are determined by a Gaussian low-pass filter.

Common problems of the “surround-based” Retinex methods are that the surround has a fixed circular basis and that the filter is circularly symmetric; the weights vary only radially. Such a fixed shape leads to artifacts around high contrasted edges. In a previous article, we proposed a “surround-based” method that rendered the image using a single filter.¹² The filter weights were defined by an addition of Gaussian functions. The method presented here is an evolution of the previous one. It reduces the artifacts by using a filter, whose spatial constant is decreased in presence of edges. The aim is to reduce the influence of high contrasted areas.

A New Retinex Adaptation Using An Adaptive Filter

In this section, we propose a framework for enhancing color images. This framework can be used to enhance traditional 24-bit images as well as to compress high dynamic range images that are linear RGB image derived from raw format or from multiple exposure technique. The first part describes the global framework illustrated in Figure 1 and the second part explains the Retinex-based adaptive filtering.

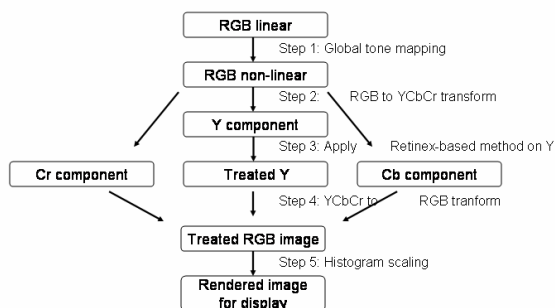


Figure 1. A global tone mapping is first applied on the input linear image. The RGB non-linear image is transformed to a YCbCr image and a Retinex-based algorithm is applied on the luminance channel Y. The resulting image is transformed back to RGB. It is then scaled and displayed on the output device.

The Global Framework

The first step of this algorithm is to apply a global tone mapping on the linear image. This step can be assimilated in the first adaptation stage of the HVS where an adaptation to global illumination takes place. The non-linear RGB image is then transformed into YCbCr color space. Only the luminance component Y is treated. The two chrominance components, Cb and Cr are left unchanged. Then, we apply the adaptive filter on the Y component. This step is described in the next section. At

step 4, the treated luminance component and the two chrominance components are transformed back to RGB color space. Finally, the RGB image is scaled to the output device dynamic range using histogram scaling.

The Adaptive Filter Algorithm Based on Retinex

In this section, we describe the Retinex-based algorithm that is applied on the luminance channel Y. Our algorithm takes inspiration from the Retinex theory in the way it determines the new pixel value by computing the ratio of the treated pixel to a weighted average of other pixels in the image. Let the treated luminance component be defined as:

$$R_Y = \log_{10}(I)'_Y - \log_{10}(mask) \tag{1}$$

where $\log_{10}(I)'_Y$ is the Y component of the non-linear image computed at step 2 and transformed into YCbCr color space. The last term of equation 1, called the mask, is a matrix that represents for each pixel the weighted average of its surround. An important point is how this surround and its corresponding weights are defined.

A traditional approach is to define the mask using a convolution of the image with a filter.^{8,12,13}

$$mask = I_Y * F \tag{2}$$

where F is a low-pass filter that is circularly symmetric. F is entirely defined by a 1-dimensional function that is rotated around the z axis. The 1-dimensional curve is usually defined by a simple Gaussian or a composition of Gaussian functions.

Our method uses a filter that is not circularly symmetric. The 1-dimensional curve that is rotated around the z axis varies with the rotation angle. An illustration of a circularly symmetric and a circularly non-symmetric filter is given in Figure 2.

The adaptive filter’s coefficients are defined by exploring the surround radially for each rotation angle. The radial 1-dimensional function is a Gaussian curve whose spatial constant varies as a function of the image’s local contrast. The spatial constant has an initial value σ given by equation 3. If a high contrasted edge is crossed along the radius, σ is divided by 2.

$$\sigma = \frac{\max(size(I))}{8} \tag{3}$$




Figure 2. Left: Circularly symmetric filter. Right: Circularly non-symmetric filter.

Our motivation to introduce a non-circular filter is that it reduces the influence of neighboring areas whose luminance produces a high contrast, which would lead to artifacts. For example, if a dim area is near a light source,

the influence of the light source should be reduced to avoid black halos around its contour. Furthermore, it seems natural that human local adaptation follows image contours to adapt to a local surround that depends on the image content.

Since the filter's weights and support are adapted for each pixel, the mask is computed sequentially pixel after pixel. $mask(x,y)$ is the weighted sum of elements in the surround of the pixel of coordinate (x,y) .

$$mask(x,y) = \int_{\theta=0}^{360} \int_{r=0}^{r_{max}} I_Y(x + \cos(\theta), y + \sin(\theta)) \cdot e^{-\frac{r^2}{\sigma^2}} \quad (4)$$

$$r_{max} = \max(\text{size}(I)) \quad (5)$$

where σ is the Gaussian spatial constant that varies along the radial direction. In this way, the filter's support approximately follows the image's high contrast edges. These edges are detected using the Canny algorithm.² The Canny method finds edges by looking for global maxima of the image's gradient. It detects strong and weak edges. Weak edges appear in the output only if they are connected to strong edges. Figure 3 shows the results of a "Canny" edge detection on a log-encoded image.



Figure 3. Left: Log-encoded image. Right: "Canny" edge detection.

Results

We applied our method to a set of high dynamic range images. This set includes radiance maps that were downloaded from the Internet^{5*} and raw linear images taken with a Canon Power-Shot G2.

Results are shown in Figure 5. The left images represent the log-encoded originals and the right images are the output of our algorithm.[†] It is visible that applying our adaptive filter method on the log-encoded image retrieves details in dark areas and increases the contrast in bright areas. Nevertheless, our method tends to desaturate the image as we apply our algorithm only on the

luminance channel. A processing of the chrominance is still to be introduced.



Figure 4. Left: Treated image without varying σ . Right: Image treated with our method. The varying σ prevents the black T-shirt from becoming gray. It also keeps the areas surrounding the window from becoming too dark.

Figure 4 illustrates the effect of varying σ along the radial direction. The image on the left is computed with our method but using a fixed sigma. The image on the right is the image obtained with a varying σ .

In the future, we want to use a more general segmentation than edge detection. This will allow a finer definition of the variation of the Gaussian spatial constant. We are currently working on improving the definition of the 1-dimensional function that determines the filter coefficients.

Conclusion

We propose a method for rendering natural color images that takes inspiration from some aspects of the biological vision system. We base our work on the Retinex theory of color vision and previously developed Retinex-based computational models. Our method computes the new intensity value for each pixel, given by the ratio of the treated pixel to a weighted average of a surrounding area. The novelty of this method is that the surround function is not circularly symmetric but its shape follows the image's high contrasted edges. This method allows reducing artifacts such as black halos around light sources but increases significantly the computational time.

The Retinex-based method is applied on the luminance channel only. It is embedded in a framework that includes pre- and post-processing. The framework takes raw linear images and compresses the dynamic range so that the rendered image reproduces better what the observer remembers of the original scene.



Figure 5. Left: Log-encoded image. Right: Treated Image.

Acknowledgment

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

References

1. D. Brainard and B. Wandell, Analysis of the Retinex theory of color vision, *Journal of the Optical Society of America*, 3, 1651-1661, 1986.
 2. J. Canny, A Computational Approach to Edge Detection, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8 (6), pp. 679-698, 1986.
 3. T. Cooper, Modifications to retinex to relax reset nonlinearity and implement segmentation constraints, in *IS&T/SPIE Electronic Imaging 2002, The Human Vision and Electronic Imaging VII Conference*, Vol 4662, pp. 349-357, 2002.
 4. P. E. Debevec and J. Malik, Recovering high dynamic range radiance maps from photographs," in *SIGGRAPH'97*, pp. 369-378, 1997.
 5. F. Xiao, J. DiCarlo, P. Catysse, and B. Wandell, High dynamic range imaging of natural scenes," in *IS&T/SID Tenth Color Imaging Conference*, pp. 337-342, 2002.
 6. B. Funt, F. Ciuera, and J. McCann, Retinex in matlab, in *IS&T/SID Eighth Color Imaging Conference*, pp. 112-121, 2000.
 7. B. Funt, F. Ciurea, and J. McCann, Tuning Retinex parameters, in *IS&T/SPIE Electronic Imaging 2002, The Human Vision and Electronic Imaging VII Conference*, Vol 4662, pp. 358-366, 2002.
 8. D. Jobson, Z. Rahman, and G. Woodell, A multiscale retinex for bridging the gap between color images and the human observation of scenes, *IEEE Transactions on Image Processing*, 6(7), pp. 965-976, July 1997.
 9. E. Land, The Retinex, *American Scientist*, 52, pp. 247-264, 1964.
 10. E. Land, The Retinex theory of colour vision, *Proceedings of The Royal Institution of Great Britain*, 47, pp. 23-58, 1974.
 11. E. Land, Recent advances in retinex theory, *Vision Research*, 26(1), pp. 7-21, 1986.
 12. L. Meylan and S. Süsstrunk, Bio-inspired color image enhancement, in *IS&T/SPIE Electronic Imaging 2004. The Human Vision and Electronic Imaging IX Conference*, Vol 5292, 2004.
 13. Z. Rahman, D. Jobson, and G. Woodell, Retinex processing for automatic image enhancement, in *IS&T/SPIE Electronic Imaging 2002. The Human Vision and Electronic Imaging VII Conference*, Vol 4662, pp. 390-401, 2002.
 14. A. Rizzi, C. Gatta, and D. Marini, A new algorithm for unsupervised global and local color correction, *Pattern Recognition Letters*, 24(11), pp. 1663-1677, July 2003.
 15. R. Sobol. Improving the retinex algorithm for rendering wide dynamic range photographs, in *IS&T/SPIE Electronic Imaging 2002. The Human Vision and Electronic Imaging VII Conference*, Vol 4662, pp. 341-348, 2002.
- * <http://white.stanford.edu/hdri/>
 † Images are available at <http://ivrg.epfl.ch/index.php?name=CGIV04>

Biography

Laurence Meylan received her Master degree in computer science at EPFL (Ecole Polytechnique Fédérale de Lausanne) in 2002. Since 2002, she has worked at the LCAV (Laboratory of audiovisual communication, EPFL) as a PhD student. Her work has focused on image rendering and color processing. In particular, she has investigated the Retinex theory and its adaptations for image rendering.