The Influence of Luminance on Local Tone Mapping

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

Abstract

We study the influence of the choice of color space for local tone mapping methods. Many local tone mapping methods do not perform well when applied independently to the three color channels of an RGB image. A common solution is to only treat the luminance channel. However, the question of which color space provides the most suitable luminance definition has not been addressed. The correlation between luminance and chrominance is known to have an influence on the rendered image but the relation between a measure of correlation and the appearance of the image has not yet been found. We consider four color transforms and introduce a measure to evaluate how well they decorrelate luminance and chrominance information. We apply two local tone mapping algorithms to the luminance channel given by the four transforms and visually compare the results. As each transform leads to another luminance definition, the resulting color images will be different as well. Our results confirm that less correlation between luminance and chrominance results in better performance of the local tone mapping algorithms. Namely, they provide a better increase in local contrast in the luminance channel and less hue shifts. However, we show that a perfect decorrelation is not always necessary.

Introduction

Tone mapping methods are a critical step in the reproduction of images. These methods can be classified into two groups: global and local tone mapping. Global methods treat the image as a whole using a mapping function. One input value is mapped to one and only one output value, which depends on the mapping function that can be image dependent [1]. Local methods treat the image spatially using local operators. One input value can be mapped to different output values depending on the surrounding pixel values. Global methods provide satisfying results for most of the captured images. Nevertheless, when the dynamic range of the captured scene greatly exceeds the dynamic range of the display, a local tone mapping is necessary to render pleasing images.

Most local tone mapping algorithms are inspired by the Retinex theory of color vision [2]. Retinex aims to predict the sensation of color by making spatial comparisons of color surfaces across the image. In its first iteration, Retinex was applied independently to all three R,G,B channels. While treating R,G,B independently provides good results with global tone mapping methods, it becomes problematic for local tone mapping algorithms. Indeed, when applied locally such algorithms may create artifacts such as local graying out, hue shifts or color fringes, as illustrated in Figure 1. The left image of Figure 1 was obtained by applying the Multi-scale Retinex algorithm [3,4] to all three R,G,B channels independently, which tends to gray out the image. The right image was processed similarly with the Retinex-based adaptive filter algorithm [5,6]. Processing R,G,B independently causes a hue shift in this case.

A well-accepted solution to avoid these artifacts is to treat the luminance independently from the chrominance [5,6,7,8,9,10]. However, none of these methods investigate the influence of the chosen color transform on the appearance of the treated image.

Figure 1. Left: multi-scale Retinex applied to all R,G,B channels grays out the image. Right: The Retinex-based adaptive filter method applied to all R,G,B channels causes a hue shift.

In this article, we investigate the influence of the color space transformation in the case of luminance-based local tone mapping methods. In particular, we focus on surround-based Retinex methods, which compute new pixel values (Ψ_{new}) by taking the difference in the log domain between each pixel value and a weighted average of its surround:

$$\Psi_{new}(x,y) = \log(\Psi(x,y)) - \log(mask(x,y)), \qquad (1)$$

where Ψ is the luminance image. It is computed by a color transform applied to the input image, which is linear with respect to scene radiance. *mask* is the weighted average of pixel values in the surrounding of coordinate (x,y).

Our aim is to relate a measure of the correlation between luminance and chrominance with the color rendition of images treated by surround-based Retinex methods. We consider four color transforms and define a measure to evaluate how well they decorrelate luminance and chrominance. Then, we test two Retinex-based local tone mapping algorithms with the four different color transforms and relate the results with our measure. We show that there is a relation between the visual representation of the rendered image and the measure of correlation. Color artifacts become visible when luminance and chrominance are significantly correlated.

This article is structured as follows: Section 2 reviews background work about color rendition in the case of local tone mapping algorithms. Section 3 presents our measure and the four color transforms that we consider. Section 4 presents the two algorithms used for the test. Section 5 comments the images obtained with the two algorithms and the different color transforms. Conclusion and future work are given in section 6.

Background

Current methods provide solutions to overcome artifacts created by local tone mapping. Rahman et al. [3,4] discussed the graying out effect of surround-based Retinex algorithms and added a color restoration step to their Multi-Scale Retinex (MSR) algorithm. The Multi-Scale Retinex with Color Restoration (MSRCR) was studied by Funt and Barnard [7,11]. They concluded that the color restoration step compensates for the graying out effect by increasing the saturation, but has an unpredictable effect on the hue of the image. This graying out effect is due to the regional violations of the gray-world assumption intrinsic to the Retinex theory. Funt and Barnard thus suggest applying MSR to the luminance channel only. The treated luminance is then combined with the chrominance to obtain the final color image. They define the luminance as the average of the three color channels R,G,B. With this definition of luminance, some chromatic information remains in the luminance and vice-versa, which may lead to artifacts.

In a recent article [6], we presented a Retinex-based method that applies an adaptive filter to the luminance channel. The luminance is defined by a principal component analysis (PCA) computed over an RGB input image. The use of a PCA is motivated by the fact that it has properties that intrinsically lead to an opponent representation of colors, which makes it biologically plausible [12,13]. The first component is all positive and has the largest share of signal energy. It represents the achromatic channel, carrying luminance information. The second and third components are opponent and represent the chrominance information. Moreover, PCA provides an optimal decorrelation of the three color channels.

Fairchild and Johnson [8] developed a color appearance model (iCAM), which applies a local treatment to the luminance channel. The first stage of iCAM accounts for chromatic adaptation. Then, the image is transformed into an opponent representation. Only the luminance channel is processed to avoid the desaturation caused by the local treatment.

Sobol [10] also applies his Retinex-based algorithm to the luminance channel only. Unlike previously mentioned methods that define the luminance as weighted sum of R,G,B color channels, his luminance definition is given by the maximum between these three channels. The final color image is obtained by adding the new luminance to the log-encoded RGB image.

Thus, many local tone mapping methods first transform the input image into a luminance/chrominance representation and treat the luminance only:

$$\{\Psi, C_1, C_2\} = f_{cs}(\{R, G, B\}), \tag{2}$$

where Ψ, C_1 and C_2 are a function of R,G,B and f_{cs} is defined by the color transform considered.

In most cases, the luminance is defined by a weighted average of R,G,B color channels [3,4,5,7,8,9,11], except for Sobol's method [10]. Then, the final RGB image is obtained either by converting the luminance/chrominance image back to RGB (3) or by using a scaling technique where the ratio of the initial luminance and the treated luminance multiplies the three color channel (4). The particular case of Sobol's method adds the treated luminance to the log-encoded RGB image (5).

$$\{R, G, B\}_{new} = f_{cs}^{-1}(\{\Psi_{new}, C_1, C_2\}) \tag{3}$$

$$\{R, G, B\}_{new} = \frac{\Psi_{new}}{\Psi} \cdot \{R, G, B\} \tag{4}$$

$$\{R, G, B\}_{new} = \Psi_{new} + \{\log(R), \log(G), \log(B)\} \tag{5}$$

Here "·" and "+" are component per component operations.

In this article, we study the effect of different luminance definitions on the rendered image for the case of MSR [3,4] and Retinex-based adaptive filtering [5,6]. Our aim is to investigate the relationship between the decorrelation of luminance and chrominance information and the correct rendition of the color after a local tone mapping was applied to the luminance.

A measure of correlation

The four color transforms that we chose for our tests are described in Table 1. Each of them transforms the linear RGB input image in a luminance/chrominance encoding (2).

The first one, "f_{RGB}" transform, simply defines the green channel G as being the luminance and the red and blue channels R,B as being the chrominance. With this transform, the luminance is strongly correlated with the chrominance. The second transform is "f_{YUV}", which is a linear transform widely used for video processing [14]. The third one, "f_{Lab}", is used for perceptual experiment and is not a linear transformation. The last one "f_{PCA}" is an image-dependent, linear transform based on a PCA applied to the input image, which guarantees perfect decorrelation between components. The luminance is defined by the first principal component.

Table 1. The four color transforms tested.

	f_{cs}	Luminance	Chrominance	Transform
RGB	Linear	G	R,B	f_{RGB}
YUV	Linear	Y	u,v	f_{YUV}
Lab	Non-Linear	L	a,b	f_{Lab}
PCA	Linear	L: 1st principal component	C_1: 2nd principal component C_2: 3rd principal component	f_{pca} : Defined by the eigenvectors of the input's covariance matrix

We estimate a simple correlation measure MC obtained by computing the mean of the correlation coefficients between the luminance and the chrominance channels over a set of representative images S:

$$MC = \frac{1}{I}\sum_{i \in S}^{I} MC_i \ , \tag{6}$$

where I is number of representative images in the set S.

The correlation measure for one image i is given by the average correlation between luminance and chrominance channels:

$$MC_i = \frac{corr(L_i, C_{i,1}) + corr(L_i, C_{i,2})}{2}, \tag{7}$$

where the correlation between luminance and chrominance is defined by the normalized covariance (8),(9).

$$\mathrm{cov}(L_i, C_{i,x}) = \frac{1}{N}\sum_{p=1}^{N}(L_i(p) - \overline{L})(C_{i,x}(p) - \overline{C_{i,x}(p)}) \tag{8}$$

$$corr(L_i, C_{i,x}) = \frac{\mathrm{cov}(L_i, C_{i,x})}{\sqrt{\mathrm{cov}(L_i, L_i)\,\mathrm{cov}(C_{i,x}, C_{i,x})}} \tag{9}$$

In (8) and (9), $C_{i,x}$ represents one of the two color channel $C_{i,1}$ or $C_{i,2}$. N is the number of pixels in the image i.

Table 2 shows the measure of correlation of the four considered color transforms. Uncorrelated data results in a correlation coefficient of 0; equivalent data sets have a correlation coefficient of 1. A graphical representation is given in Figure 2. It shows how well the color transforms decorrelate the luminance from the chrominance. Our aim is to see how this measure is related to color rendering. In particular, we want to test if f_{PCA}, which is image dependent and guarantees perfect decorrelation between components, leads to the best reproduction.

Table 2. Measure of correlation.

MC_{RGB}	0.98
MC_{YUV}	0.28
MC_{Lab}	0.25
MC_{PCA}	0

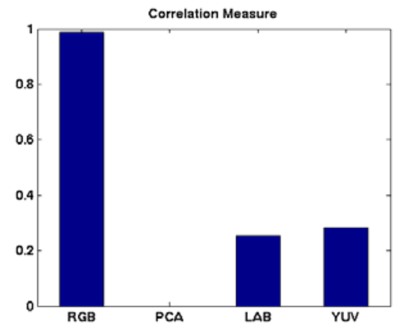

Figure 2. MC for f_{RGB}, f_{PCA}, f_{Lab}, f_{YUV}.

Two Luminance-based tone mapping methods

Two surround-based local tone mapping algorithms are used for our tests. We implemented the multi-scale Retinex (MSR) of Rahman et al. [3,4] for that purpose and used our Retinex-based adaptive filter method [5,6].

Multi-scale Retinex [3,4]:

MSR is a combination of three single-scale Retinex. It computes new pixel values by taking the difference in the log-domain between each pixel and a weighted average of its surround for three different surround sizes. The weighted average is obtained by convolving the input image with a two-dimensional filter, whose entries are given by a Gaussian function. The spatial constant of the Gaussian function varies according to the surround size. In their article, Rahman et al. apply MSR separately to R,G,B color channels and add a color restoration factor. We do not use the color restoration factor but apply MSR only to the luminance channel as suggested by [7,11].

Retinex-based adaptive filter [5,6]:

This algorithm is based on surround-based Retinex. It also computes the new image by taking the ratio between each pixel and a weighted average of its surround but differs in three ways:

1. As local processing tends to make pure white and pure black low contrast areas turn gray, the Retinex-based adaptive filter adds a $\beta(x,y)$ factor that weighs the mask in order to preserve white and black. The value of $\beta(x,y)$ depends on the input image value at position (x,y).

$$\Psi_{new}(x,y) = \log(\Psi(x,y)) - \beta(x,y)\log(mask(x,y)) \tag{10}$$

2. Instead of using a circular surround, the shape and the weights of the surround are adapted for each pixel depending on the position of high contrast edges in the image. In this way, halo artifacts common to other surround-based methods are prevented.

3. The Retinex-based adaptive filter is applied to the luminance channel only. Before doing the inverse color transformation to obtain the RGB image, the chrominance channels are multiplied by a factor to compensate for the decrease in saturation induced by the increase in image brightness.

We apply theses two local tone mapping algorithms to a set of test images. Each input image is first converted from RGB to a luminance/chrominance opponent representation using one of the color transforms of Table 1. Then, one of the local tone mapping methods is applied to the luminance channel. Finally, the chrominance channels and the new luminance are transformed back into RGB encoding. We visually compare the results and comment them in the next Section.

Discussion and results

Figure 3 shows the results obtained with MSR and the four color transformations defined in Table 1. It shows that different transforms, thus different luminance definitions, result in different output images[1]. The images obtained using f_{RGB} and MSR algorithm clearly shows a pink color shift in the center of the image. This is due to the fact the G channel is strongly correlated with the R and B channels. Then, there is little visible difference between the images obtained with f_{Lab}, f_{YUV} and f_{PCA}. The graying out of low contrast areas such as the sky is due to the local averaging induced by Retinex surround-based methods. As mentioned before the Retinex-based adaptive filter method prevents graying out by introducing a factor that weighs the mask depending on the input image values. This, in addition to the saturation compensation factor, results in visually more appealing images.

Figure 4 shows the same image treated by our Retinex-based adaptive filter method and the different color transforms. The image computed with f_{RGB} also presents a color shift in the center of the image. We can observe as well that f_{YUV} and f_{PCA} lead to images with a better increase in local contrast than f_{Lab}. In other words, the detail of the central part of the image is more visible. Moreover, the sky is slightly more saturated with f_{PCA} and f_{YUV} than with f_{Lab}. These differences may come from the fact that the f_{Lab} transform is non-linear with respect to the scene radiance. We see no difference between the image computed using f_{YUV} and the image computed using f_{PCA}.

Nevertheless, if we compare the results of f_{PCA} and f_{YUV} on another image (Figure 5), color differences appear. In Figure 5, the top image (f_{YUV}) appears greener than the bottom image (f_{PCA}). This hue shift becomes obvious if we look at the a, b chromaticity plane of these two images (Figure 6). The image treated using f_{YUV} is plotted in cyan, while the image treated using f_{PCA} is plotted in magenta. The cyan cloud is shifted to the left of the magenta cloud, i.e. the image treated using f_{YUV} tends to appear greener than the image treated using f_{PCA}. This causes the face of the person to look slightly green on the f_{YUV} image.

Figure 3, 4 and 5 show that the choice of the luminance definition on which a local tone mapping algorithm is applied plays a role for the image appearance. However, a small correlation between luminance and chrominance does not affect significantly the final result. f_{RGB} had the worst decorrelation measure and the images obtained using this transform clearly present color shift artifacts. f_{YUV} and f_{Lab} had good decorrelation measure but not as good as f_{PCA} that ensures perfect decorrelation between components. The fact that the luminance of Lab-encoded image is non-linear induces some differences with the f_{YUV} and f_{PCA} images. However, the images computed using f_{YUV} and f_{PCA} are very similar, which makes it difficult to judge which transforms leads to the best resulting image.

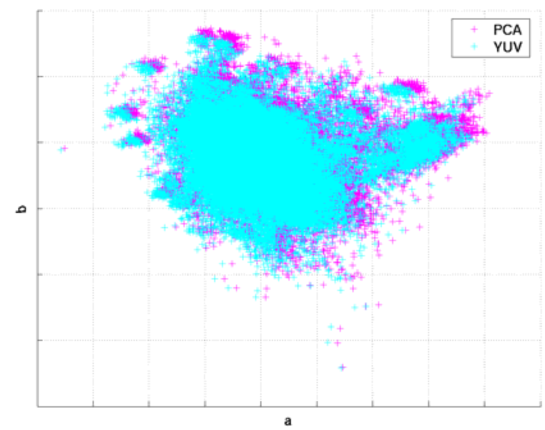

Figure 6. Plot of the a,b chromaticities of the two images of Figure 5. The image treated using f_{YUV} is plotted in cyan while the image treated using f_{PCA} is plotted in magenta.

It is important to note that f_{PCA} works well for natural images, which contain a reasonable diversity of colors. However, particular cases such as a singular color image would lead to an ill-conditioned transformation matrix and thus to the failure of the PCA algorithm. This does not happen when treating natural images even in the presence of color cast but is more likely to happen with synthetic images.

Conclusion

The goal of our study was to investigate the role of the luminance definition on the final image appearance in the case of luminance-based local tone mapping algorithms. For that purpose, we tested two algorithms and four color transforms (f_{RGB}, f_{YUV}, f_{Lab} and f_{PCA}).

[1] As printing the images might introduce other color shifts, we suggest the reader looks at the images on a display.

A measure of correlation was established for these four transforms. The f_{RGB} transform had highly correlated components. f_{YUV}, f_{Lab} were slightly correlated and f_{PCA} was designed to ensure perfect decorrelation for all images. Our observations were that there were little visible differences between the images treated using f_{YUV}, f_{Lab} and f_{PCA}. Local color shifts started to appear when using a transform where luminance and chrominance were highly correlated such as f_{RGB}. That suggests that there is a relationship between the amount of correlation between luminance and chrominance, and the quality of the image appearance. However, a perfect decorrelation does not seem to be necessary to obtain visually pleasing images.

References

1. J. Holm, Photographic tone and colour reproduction goals, in CIE Expert Symposium'96 on Colour Standrards for Image Technology, Vienna, Austria, March 1996, pp. 51-56.
2. E. Land, The Retinex, American Scientist, vol. 52, no. 2, pp. 247-264, June 1964.
3. 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, vol. 6, no. 7, pp. 965-976, July 1997.
4. Z. Rahman, D. D. Jobson, and G. A. Woodell, Retinex processing for automatic image enhancement, Journal of Electronic Imaging, vol. 13, no. 1, pp. 100-110, January 2004.
5. L. Meylan and S. Süsstrunk, Color image enhancement using a Retinex-based adaptive filter, in IS&T Second European Conference on Color in Graphics, Image, and Vision (CGIV 2004), Aachen, Germany, April 2004, pp. 359-261.
6. L. Meylan and S. Süsstrunk, High dynamic range image rendering with a Retinex-based adaptive filter, accepted to IEEE Transactions on Image Processing, May 2005.
7. B. Funt, K. Barnard, M. Brockington, and V. Cardei, Luminance-based multi-scale Retinex, in AIC Colour 97 Eigth Congree of the International Colour Association, vol. I, Kyoto, Japan, May 1997, pp. 330-333.
8. M. D. Fairchild and G. M. Johnson, iCAM framework for image appearance differences and quality, Journal of Electronic Imaging, vol. 13, no. 1, pp. 126-138, January 2004.
9. F. Durand and J. Dorsey, Fast bilateral filtering for the display of high-dynamic-range images, in SIGGRAPH 02 Annual Conference on Computer Graphics, Los Angeles, CA, July 2002, pp. 257-266.
10. R. Sobol, Improving the Retinex algorithm for rendering wide dynamic range photographs, Journal of Electronic Imaging, vol. 13, no. 1, pp. 65-74, January 2004.
11. K. Barnard and B. Funt, Investigations into multi-scale Retinex, in Colour Imaging in Multimedia '98, Derby, UK, March 1998, pp. 9-17.
12. D. Alleysson and S. Süsstrunk, Spatio-chromatic PCA of a mosaiced color image, in IS&T Second European Conference on Color in Graphics, Image, and Vision (CGIV 2004), Aachen, Germany, April 2004, pp. 311-314.
13. G. Buchsbaum and A. Gottschalk, Trichromacy, opponent colours coding and optimum colour information transmission in the retina, in Royal Society of London, Series B, biological Sciences, vol. 220, no. 1218, November 1983, pp. 89-113.
14. C. Poynton, Digital video and HDTV, algorithms and interfaces. San Francisco, California: Morgan Kaufmann publishers, 2003.

Author Biography

Laurence Meylan received her MS degree in computer sciences from the Ecole Polytechnique Fédérale de Lausanne (EPFL, Switzerland) in 2002. Since then, she has been a PhD candidate in the Laboratory of Audiovisual Communications in Lausanne (LCAV-EPFL), Switzerland. Her work has focused on the problem of high dynamic range imaging for both the creation and the rendering of high dynamic range images. In particular, she studied the color rendition of HDR images in the presence of local tone mapping algorithms.

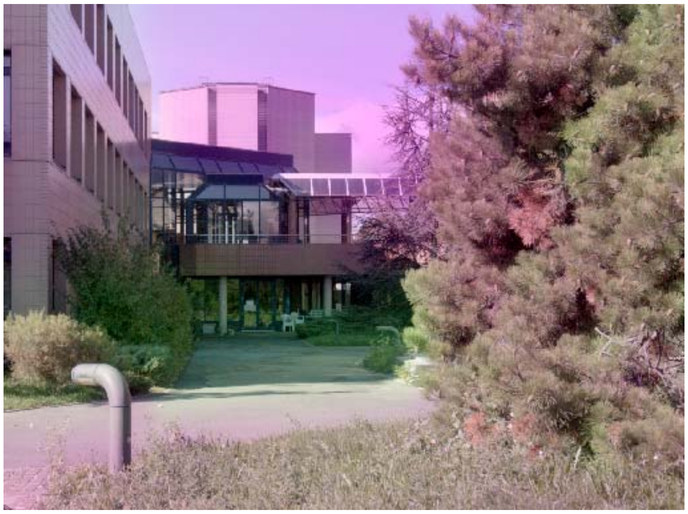

f_{RGB}

f_{Lab}

f_{YUV}

f_{PCA}

Figure 3. Results obtained with MSR.

Figure 4. Results obtained with the adaptive filter method.

f_{YUV}

f_{PCA}

Figure 5. Results obtained with the Retinex-based adaptive filter method. The top image (f_{YUV}) appears greener than the bottom image (f_{PCA}).