

Color Constancy for Multi-Illuminants High-Dynamic-Range Scenes

Jiangtao Kuang, Weihua Xiong, OmniVision Technology Inc., Santa Clara, California, USA

Abstract

Multiple illuminants with different color temperatures within a scene provide a complicated situation for color constancy and automatic white balance (AWB) algorithms in digital photography. This problem gets even worse in high-dynamic-range (HDR) imaging since a large scale of luminance information is able to be captured and thus it is more likely to be influenced by different illuminants in the scene. Under mixing lighting, a single global adjustment of colors may not yield a good result, since this approach tends to exaggerate the color difference for each illuminant as compared to what observed with the human eye, or only partially remove color cast in the image, making one lighting area look better while others look worse. A local auto white balance algorithm that adjusts colors pixel-by-pixel based on its local area was proposed to solve this problem. For a specific pixel, illumination is estimated from the color information from its neighboring pixels that is weighted by the spatial distance, luminance intensity difference and chromaticity. Experiments on synthetic and real images show that this algorithm performs significantly better than other global and local AWB algorithms when evaluated in terms of the accuracy with which correct surface object colors are estimated.

Introduction

In digital photography, white balance algorithms are designed to remove illuminants' color casts, so that objects that appear white in person are rendered white in the photos. Scenes lit by multiple illuminants with different color temperature, for example, mixed lighting of indoor and natural light, often provide more complicate situations for auto white balance (AWB) in imaging systems. Human eyes are very good at judging what is white in different color lighting, and even locally adapting to different illuminants in the same scene simultaneously by saccadic eye motion. However, digital cameras often have great difficulties with auto white balance in this situation. The majority of the illumination-estimation methods^[1-5] usually calculates the average color temperature for the entire scene and uses it for white balance. While this approach is often acceptable in most cases, it tends to exaggerate the color difference for each illuminant as compared with the real-world counterpart, or only partially correct white balance for one illuminant, making one lighting area looking better while others look worse. Figure 1 shows a church lit by a warm incandescent light under a moonlit sky. White balance based on the moonlit sky brings out the warm color temperature of the artificial lighting. However, white balance based on the church generates a sky with an unrealistic blue appearance. It suggests that a single color balance setting is insufficient. The colors must, to some extent at least, be adjusted locally to account for the local variation in scene illumination.



Figure 1. (a) white balance for the moonlit sky; (b) white balance for the church

Retinex^[6,7] estimates the illumination color locally for each pixel by making comparisons between the pixel and other neighboring image pixels and adjusts colors on a pixel-by-pixel basis. The nearby pixels are given more weight than distant pixels in the estimation. The “reset” operation in Retinex assumes that locally the most reflective surface is “white” and other pixels’ colors are adjusted relative to this white. This assumption produces the effect of strongly influencing illumination estimate by the colors in each pixel’s neighborhood. To improve Retinex’s illuminant estimation, Xiong et al.^[8] integrated Retinex with 3 dimensional spatial-edge information reconstructed from stereo images, in order to prohibit its local neighbouring pixels’ comparison across the spatial edge. While the effectiveness of using spatial information was demonstrated for some scenes, it is still difficult to identify where the illuminations’ change occur, since they are not always separated by spatial edges or changes in surface orientation. Besides, it requires stereo images derived from two or more images of the same scene captured at different viewpoints simultaneously, which is also impractical for real-life applications. They^[9] then proposed a hybrid solution to combine the benefits of local and global constancy solutions, Retinex and SVR. Under mixed lighting, Retinex first mitigates the illumination difference to be more uniform by its local illumination estimation, followed by SVR to cancel out the illumination’s effect globally.

Multiple illuminants color constancy becomes more critical in HDR imaging as a larger scale of luminance is able to be captured and thus more likely to be influenced by different lighting. For instance, a common scenario occurs when one is indoors and looks across the room and through a window to the outdoors. Traditional dynamic range image devices cannot capture all luminance information, leaving either the indoors portion of the scene too dark to discern or the outdoors portion of the scene saturated. This might explain why it has been a less serious problem for multi-

illuminants white balance over the years. HDR imaging technologies have advanced so as to capture and store both indoors and outdoors luminance in a HDR image with limited dynamic range image devices using multi-exposure techniques,^[10,11] or directly with HDR image sensors.^[12,13] Therefore, different color temperature lighting has to be considered in HDR image rendering. Unfortunately, little research has been conducted to solve this problem. As a part of HDR color image processing and rendering flowchart, a local illuminant estimation and correction AWB algorithm was thus proposed and tested in this paper.

Algorithms

The assumption in human vision has been widely accepted in imaging research that an image is regarded as a product of the reflectance and the illuminances, shown in Eq. 1.

$$I(x, y) = R(x, y)L(x, y) \quad (1)$$

where $I(x, y)$ is the intensity, $R(x, y)$ is the reflectance and $L(x, y)$ is the illumination at each point (x, y) . A common assumption for the simplicity of calculating R and L is that L varies slowly while R can change abruptly. A low-pass Gaussian filtering of the image is often regarded as the illuminances of the scene.

However, in a HDR scene, multiple illuminants lighting often comes with dramatic luminance intensity change, which is less likely caused by objects' reflectance change. A Gaussian filter removes details well, but it also smoothes across the sharp edges where lighting change. To improve local illuminant color estimation, we want to preserve sharp intensity edges in the low-passed filtering to avoid incorrect color adjustments along abrupt illuminant changes.

The "white" image is thus obtained using an edge-stopping filter called the bilateral filter, which is first proposed by Durand and Dorsey^[14] in their HDR tone-mapping operator. The Bilateral filter is an anisotropic filter, where each pixel is weighted by the product of a Gaussian filtering in the spatial domain and another Gaussian filtering in the intensity domain that decreases the weight of pixels with large intensity differences. Therefore, the bilateral filter effectively blurs an image while keeps sharp edges intact, and thus avoids color bleeding along sharp lighting changes caused by different color temperature illuminants. The pixel intensity calculations are performed in the log domain because the pixel differences directly correspond to the perceptual contrast.

Furthermore, for a robust estimation of the illumination light, we can restrict our search to a range of likely illuminant pixels and excluded those with high saturation pixels in the calculation. Indeed, this way of constraining the possible estimates produces the best color constancy algorithms.^[15] Figure 2 shows the Planckian locus of black-body radiators on the xy chromaticity diagram.^[16] Finlayson^[15] demonstrated that the illuminant chromaticities of typical light sources fall on a long thin "band" around the Planckian locus in chromaticity space. This means that we can put weights to each pixel which contributes to the illumination estimate according to the distance to the Planckian locus in xy chromaticity diagram. The nearby pixels are given more weight than distant pixels in the estimation.

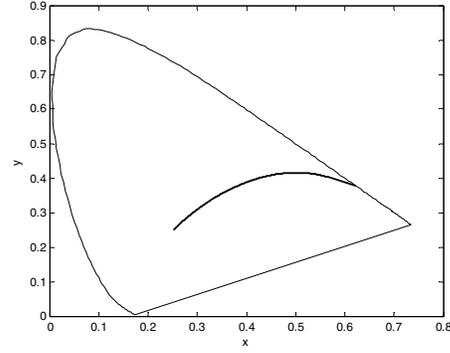


Figure 2. Planckian locus plotted in xy chromaticity space

A trilateral-like filtering was developed to incorporate the functions above, given in Eq. 2 and Eq. 3.

$$J'_s = \frac{1}{k(s)} \sum_{p \in \Omega} f(p-s)g(I_p - I_s)w(p)J_p \quad (2)$$

where $k(s)$ is a normalization term:

$$k(s) = \sum_{p \in \Omega} f(p-s)g(I_p - I_s)w(p) \quad (3)$$

I_s is the luminance intensity value for pixel s , which is influenced mainly by pixels that are close spatially and that have a similar intensity and that locate around Planckian locus in the chromaticity diagram. J_p is the pixel value for a specific color channel from the local area of pixel s and J_s' is its corresponding output values. The filtering is carried out on each color channel independently. $f()$ is a Gaussian function in the spatial domain with the kernel size σ_f . Ideally σ_f is specified with device independent coordinates such as cycle-per-degree according to the viewing conditions. Since they are often not available in the application, σ_f is set to an empirical value of 1/3 of the image size. $w()$ is another Gaussian function in chromaticity diagram with its scale σ_w empirically set to a constant value of 0.04. $g()$ is a Tukey's biweight function^[17] (Eq. 4) in the intensity domain, which completely stops diffusion across edges. The trilateral filtering is speeded up using nearest neighbor down-sampling and up-sampling in the implementation.

$$g(x) = \begin{cases} x \left(1 - \frac{x^2}{c^2} \right) & \text{for } |x| < c \\ 0 & \text{for } |x| \geq c \end{cases} \quad (4)$$

After obtaining the "white" image, color adjustment is conducted on a pixel-by-pixel basis.

Testing the Local AWB Algorithm

In this section, we apply the local AWB algorithm to multiple illuminants HDR images. Results are compared to Gray-World method and Retinex algorithm to demonstrate its performance. For the Retinex, the iteration time is set to 1, as Xiong etc.^[8,9] has shown that one iteration always give out the best performance for most scenes.

Synthetic Image Experiments

Our first experiments are based on synthetic images that model a HDR scene with two quite distinct illuminants lighting different parts of the scene. We generate synthetic scenes composed of patches of different reflectance by randomly selecting reflectances from a publicly accessible database^[18]. The patches are divided into two sections by an irregular boundary representing where the illumination changes. RGB values for the patches are calculated by using two illumination spectra, CIE A on the left, CIE D65 daylight on the right, and sensor sensitivity functions of a commercial CMOS camera color balanced equal-energy white. The ground-truth image is generated by correcting each illuminated part individually to an equal-energy white illuminant. Luminance of the right part is increased 1000 times to simulate a HDR scene.

The top left image of Figure 3 shows the linear rendering of the synthetic HDR image. To indicate different color temperature lighting areas, the left part is increased by 1000 times in the top right image, with a white line demarcating the boundary between two illuminants. Note that this line is not included in the original input image. The same process is applied to all results to make it easier to see their performance. Comparing to the ground-truth image, the local AWB algorithm successfully removed color cast in both illumination areas. Gray-world method, as a popular global AWB algorithm, removes some of the bluish cast from the right side of the image, but introduces more warm color to the left side. Retinex removes some of color cast in both sides, but tends to over-correct and de-saturate the color patches.



Figure 3. Synthetic image results. From left to right and top to bottom: input HDR image, image with balanced luminance and white line superimposed to indicate the illumination boundary, ground-truth image, Local AWB algorithm result, gray-world result, Retinex result.

To quantitatively evaluate the results, geometric distance between the white balanced image and the ground-truth image is calculated at each pixel in rg-chromaticity ($r=R/(R+G+B)$, $g=G/(R+G+B)$) space. The angle in degree between two color vectors in rgb space ($b=1-r-g$) is also investigated.^[9] These errors are defined by the following formulas, where subscript ‘p’ indicates the result from a specific AWB method and ‘g’ indicates the ground-truth image.

$$Ed = \sqrt{(r_p - r_g)^2 + (g_p - g_g)^2} \quad (5)$$

$$Ea = \cos^{-1} \left[\frac{(r_p, g_p, b_p) \bullet (r_g, g_g, b_g)}{\sqrt{r_p^2 + g_p^2 + b_p^2} * \sqrt{r_g^2 + g_g^2 + b_g^2}} \right] \quad (6)$$

Three statistics are computed on the distribution of errors across all the pixels in an image: the median, the RMS (root mean square) and the mean of the 99.9 percentile of the largest errors, denoted MMax. In contrast to a single maximum error, MMax is more robust to exclude outliers.

The error results are shown in Table 1. The local AWB algorithm has significantly outperformed Retinex and Gray-world method in color accuracy. The MMax distance and angular errors of local AWB are less than 50% of those from Retinex and Gray-world. Other statistics also illustrate the same trend. It is interesting to find out that Gray-world performs acceptably for the median statistics but much worse in MMax and RMS. The reason becomes obvious, since the bright part is dominant in the Gray-world calculation for HDR images, the gray-world algorithm correctly color adjusts for the bright illumination while making the dim illumination worse.

Table 1: Comparison for the synthetic image of MMax, RMS and median errors measured on a pixel-by-pixel basis between the ground-truth image and the processed images by three AWB algorithms

	Distance (* 10 ²)			Angular		
	MMax	Med	RMS	MMax	Med	RMS
Local AWB	7.20	2.32	2.73	7.86	2.49	3.14
Retinex	16.37	3.94	5.33	20.83	4.94	6.62
Gray-world	13.72	2.68	7.65	19.10	2.57	9.42

Real Image Experiments

The local AWB algorithm was then evaluated using real-world HDR images. First we built up a multiple illuminants HDR scene in the laboratory. The advantage of the laboratory scenes is that it is possible to obtain a ground-truth image with which to evaluate the white balance color accuracy errors. Outside the laboratory, it is difficult to make enough measurements of the illumination distribution to obtain the ground-truth image. We have to evaluate the result images by a direct comparison against the corresponding real-world counterparts; otherwise, we have to count on the prior experience of the scenes.

The laboratory scene (Figure 4), designated Lab Scene, has two distinct illuminations similar to those found indoors and outdoors, where the left part was illuminated by bright incandescent lights and the right part was under dim daylight lighting in a light booth. They are carefully separated so that the

illuminations' color temperature was approximately uniform in each individual part. Two Macbeth Color Checkers were included under each illumination. They provided "white" information to generate the ground-truth image, which was obtained by manually white balancing for each illumination. The luminance range of this scene is from 1 cd/m² to 2000 cd/m². This scene was photographed using Canon 400D digital camera with different shutter speeds ranging from 1/250 to 1 second. All captured images were stored with 12-bit raw data for the construction of HDR images with camera response curve using the multiple exposure combination method proposed by Robertson et al.^[11]

The input HDR image, ground-truth image and white balance result images by Local AWB algorithm, Retinex and Gray-world method are shown in Figure 4. To present these HDR images in the paper, all images except the Retinex one were rendered with HDR tone-mapping algorithm iCAM06,^[19] and since Retinex is claimed to be an tone-mapping operator as well, no further rendering method was applied to its result here. All algorithms have successfully removed the yellowish cast in the incandescent lighting, while the local AWB algorithm is most effective for removing the bluish cast in the daylight lighting. The result of the local AWB algorithm is more close to the real experience when observed with the human eye.



Figure 4. Lab Scene image results. From left to right and top to bottom: input HDR image with incandescent lights and daylight illumination, ground-truth image with manually white balance for individual illumination, Local AWB algorithm result, Retinex result and Gray-world method result.

The numerical results presented in Table 2 show that Retinex and Gray-world method perform with relative similar accuracy for this image, while the local AWB algorithm outperforms each of the others taken individually.

Table 2: Comparison for the Lab Scene of MMax, RMS and median errors measured on a pixel-by-pixel basis between the ground-truth image and the processed images by three AWB algorithms

	Distance (* 10 ²)			Angular		
	MMax	Med	RMS	MMax	Med	RMS
Local AWB	15.19	1.89	6.17	19.36	2.36	8.07
Retinex	21.07	2.20	6.54	23.95	3.27	8.37
Gray-world	18.39	1.65	7.19	23.29	2.31	9.38

We designed a second HDR scene (Figure 5) to incorporate a typical multi-illuminant scenario of being indoors in a room with a window to the outdoors. The indoor objects are lit with a warm incandescent table lamp, while the outdoor ones are lit by sky blue light. Note that natural scenes often have mixed lighting conditions

instead of two distinct illuminations. For example, the wall under the window is illuminated by the mixed light from indoors and outdoors simultaneously. Even for the outdoor objects, the buildings are lit with higher color temperature lighting than the concrete road since they are in the shadows. Comparing to the Lab Scene, it is more challenging for white balance algorithms. The HDR image was generated with the same multiple-exposure technique stating above.

Figure 5 shows the input HDR image and white balance results from three algorithms. Again, to present HDR images on the paper, we applied iCAM06 HDR tone-mapping algorithm to the images except the Retinex one. The result image of the local AWB algorithm has a very natural appearance in both indoors and outdoors. Retinex has a strong local effect by de-saturating color surfaces, for instance, turning the blue sky completely white. The Gray-world method tends to have a balance between indoors and outdoors; however, the indoor wall under the window becomes green and the outdoor buildings have a blue cast. Intrinsically the Gray-world method cannot solve the white balance problem under mixed lighting conditions.



Figure 5. Window scene and results. From left to right and top to bottom: input HDR image with indoor incandescent lights and outdoor daylight illumination, Local AWB algorithm result, Retinex result and Gray-world method result.

Conclusions and Discussions

HDR images are often of multiple illuminations, which is a challenge for automatic white balance in HDR image processing. Global AWB algorithms have an intrinsic limitation to have unsatisfied white balance for each individual lighting part, leaving uncorrected color cast in one part or several parts. Retinex is one method that makes local adjustments for the illumination, but it tends to over correct the color cast, generating unrealistic de-saturation and halo effects. We have proposed a local AWB algorithm in this paper to solve this problem. The Illuminant is locally estimated from the color information from its neighbor pixels that is weighted by their spatial distance, intensity difference and chromaticity distance from the Plankian locus in chromaticity diagram. Experimental results from synthetic images and real images demonstrated that the local AWB algorithm is effective to remove the color cast under different illuminations by generating natural appearance. It outperformed Retinex and Gray-world method in our test.

We have not addressed the problem of incomplete chromatic adaptation under different color temperature illuminations; instead, we assumed that we have a complete white balance for each individual lighting area. However, this assumption is probably not true from the perceptual experience. For some situations, we may prefer to leave the color temperatures as is; on the other hand, some situations may not even have a truly "correct" white balance, and will depend upon where we focus on or where color accuracy is most important. Psychophysical experimental data will be very helpful in the research of multi-illuminant color constancy.

References

- [1] V. Cardei, B. Funt, K. Barnard, "Estimating the Scene Illumination Chromaticity Using a Neural Network," *Journal of the Optical Society of America A*, Vol. 19, No.12, pp. 2374-2386, (2002)
- [2] B. Funt, W.H. Xiong, "Estimating Illumination Chromaticity via Support Vector Regression," *Proc. 12th Color Imaging Conference*, pp. 47-52, (2004)
- [3] G.D. Finlayson, S.D. Hordley, P.M. Hubel, "Color by Correlation: A Simple, Unifying Framework for Color Constancy," *IEEE Trans. Pattern Anal. Mach. Intel.* Vol. 23, No.11, pp.1209-1221, (2001)
- [4] D. Forsyth, "A novel algorithm for color constancy", *Int. J. of Comp. Vis.*, Vol. 5, pp. 5-36, (1990)
- [5] C. Rosenberg, M. Hebert, and S. Thrun, "Color constancy using KL-divergence", *Proc. Eighth IEEE International Conference on Computer Vision (ICCV '01)*, Vol. 1, July, pp. 239 – 246, (2001)
- [6] E. H. Land and J. McCann, Lightness and Retinex Theory, *Journal of the Optical Society of America*, Vol. 61, No. 1, pg. 1-11 (1971).
- [7] B. Funt, F. Ciurea., J. McCann., "Retinex in Matlab," *Journal of the Electronic Imaging*, pp.48- 57, (2004)
- [8] Xiong, W. and Funt. B., "Stereo Retinex," CRV'2006 Third Canadian Conference on Computer and Robot Vision. (2006)
- [9] Xiong, W., and Funt, B., "Color Constancy for Multiple-Illuminant Scenes using Retinex and SVR," *Proc. Imaging Science and Technology Fourteenth Color Imaging Conference* pp. 304-308. (2006)
- [10] P. E. Debevec and J. Malik, Recovering High Dynamic Range Radiance Maps from Photographs, *Proc. SIGGRAPH '97*, pg. 369-378 (1997).
- [11] M.A. Robertson, S. Borman, R.L. Stevenson, Dynamic range improvement through multiple exposures, *IEEE International Conference on Image Processing*, (1999).
- [12] F. Xiao, J. M. DiCarlo, P. B. Catrysse and B. A. Wandell, High Dynamic Range Imaging of Natural Scenes, *IS&T/SID 10th Color Imaging Conference*, pg. 337-342 (2002).
- [13] S. K. Nayar and T. Mitsunaga, High Dynamic Range Imaging: Spatially Varying Pixel Exposures, *Proc. IEEE CVPR*, Vol. 1, pg. 472-479 (2000)
- [14] F. Durand, J. Dorsey, Fast bilateral filtering for the display of high-dynamic-range image. In *Proceedings of ACM SIGGRAPH 2002, Computer Graphics Proceedings, Annual Conference Proceedings*, pg. 257-266 (2002).
- [15] G. D. Finlayson and G. Schaefer, "Solving for color constancy using a constrained dichromatic reflection model," *Int. J. Comput. Vis.* 42, 127-144 (2001).
- [16] CIE. Colorimetry. CIE Publications 15.2. Commission International de L'Eclairage, 2nd Edition. (1986)
- [17] M. Black, G. Sapiro, D. Marimont and D. Heeger. Robust anisotropic diffusion. *IEEE Trans. Image Processing* 7, 3 (Mar.), 421-432. (1998)
- [18] K. Barnard, L. Martin, B. Funt, and A. Coath., "A Data Set for Color Research," *Color Research and Application*, Vol. 27, No. 3, pp.148-152, (2002)
- [19] J. Kuang, G.M. Johnson, M.D. Fairchild, iCAM06: A refined image appearance model for HDR image rendering, *Journal of Visualization and Communication*. (2007)

Author Biography

Jiangtao Kuang received his B.S. degree in optical engineering from Zhejiang University, China (2000) and his PhD in imaging science from Rochester Institute of Technology (2006). Since then he has worked in OmniVision Tech. Inc. at Santa Clara, CA. His work has focused on the research and development of high-dynamic-range CMOS image sensor, digital color image processing algorithms and image quality.