

The Effect of Exposure on MaxRGB Color Constancy

Brian Funt and Lilong Shi
School of Computing Science
Simon Fraser University
Burnaby, British Columbia
Canada

Abstract

The performance of the MaxRGB illumination-estimation method for color constancy and automatic white balancing has been reported in the literature as being mediocre at best; however, MaxRGB has usually been tested on images of only 8-bits per channel. The question arises as to whether the method itself is inadequate, or rather whether it has simply been tested on data of inadequate dynamic range. To address this question, a database of sets of exposure-bracketed images was created. The image sets include exposures ranging from very underexposed to slightly overexposed. The color of the scene illumination was determined by taking an extra image of the scene containing 4 Gretag Macbeth mini Colorcheckers placed at an angle to one another. MaxRGB was then run on the images of increasing exposure. The results clearly show that its performance drops dramatically when the 14-bit exposure range of the Nikon D700 camera is exceeded, thereby resulting in clipping of high values. For those images exposed such that no clipping occurs, the median error in MaxRGB's estimate of the color of the scene illumination is found to be relatively small.

Introduction

MaxRGB has generally been reported [1][2][3] not to perform particularly well, but this may be the fault of inadequate dynamic range in the image data, not the method itself. This paper presents tests of MaxRGB on high dynamic range image data indicating that in fact MaxRGB works surprisingly well when provided with good image data.

MaxRGB is an extremely simple method of estimating the chromaticity of the scene illumination for color constancy and automatic white balancing. It is based on the assumption that the triple of maxima obtained independently from each of the three color channels represents the color of the illumination. In principle, the assumption will hold whenever there is a white surface in the scene, for example, and also when the scene contains three separate surfaces reflecting maximally in the R, G and B sensitivity ranges. However, in practice, most digital still

cameras are incapable of capturing the full dynamic range of a scene and choose exposures and tone reproduction curves that clip or compress high digital counts. As a result, the maximum R, G and B digital counts from an image generally do not faithfully represent the corresponding maximum scene luminances.

MaxRGB is a special and extremely limited case of Retinex [4]. In particular, it corresponds to McCann99 Retinex [5] when the number of iterations [6] is infinite, or path-based Retinex [7] without thresholding but with infinite paths. To the extent that MaxRGB is hampered by images of inadequate dynamic range, Retinex is likely to be similarly affected. In [8] some results are presented using artificial clipping of images that indicate this may be the case.

Although the clipping and compression resulting from the limitations of the dynamic range of standard image capture may be the cause of the historically poor performance of MaxRGB, perhaps the problem is instead that the assumption that 'white' is present in the scene is frequently violated. To determine which the more important factor is, we created a database of multiple-exposure images of 47 scenes (roughly 1/3 indoor, 2/3 outdoor) and measured the chromaticity of the illumination of each scene.

Experimental Setup

A Nikon D700 digital still camera was used to capture all the images in the test dataset. All images were recorded in Nikon's NEF raw data format [9]. They were decoded using dcraw [10] without demosaicing so that the original digital counts for each of the RGB channels were obtained. The camera outputs 14-bit data per channel so the range of possible digital counts is 0 to 16383. The raw data contains 4284x2844 14-bit values in an RGGB pattern. To create a color image the 2 G values were averaged, but no other demosaicing was done. The result is a 2142x1422 pixel image.

The camera's auto-bracketing was used to capture up to 9 images of exposures with a +1 EV (exposure value) difference between each in the sequence. The rate of capture was 5 frames per second. The exposure range was set to ensure that in each set there would be at least one image with maximum value less than 6000. During bracketing, the camera automatically adjusts the shutter speed and/or the aperture setting between frames in order to change the exposure by 1EV.

Two sets of bracketed images are taken for each scene. One set includes a frame in the foreground holding 4 Gretag Macbeth mini Colorcheckers at different angles. The second set of images is of the same scene, but without the Colorchecker frame. The Colorchecker frame has one Colorchecker in the center and holds the others at 45 degree angles to that central one.

Between taking the two image sets the camera is refocused and possibly moved slightly. For the first set, the focus is adjusted so the Colorchecker frame is in focus. For the second set, the focus is optimized for the scene overall. Figure 1 shows an example of a scene with and without the Colorchecker frame.

The Colorchecker frame is placed in the scene at a point where the illumination incident on it is expected to be typical of the scene illumination color in general. While all scenes contain some variation in the illumination color because of interreflections, scenes that clearly have strong variations in illumination color are avoided. For example, a room with interior tungsten lighting mixed with daylight entering through a window would be excluded.

The illumination chromaticity is determined by manually sampling the RGB digital counts from the 4 white patches of the Colorcheckers in the image from the bracketed set of maximum exposure that is not overexposed anywhere within the Colorchecker frame. The camera's NEF raw data is linear with respect to scene luminance. To be roughly similar to the requirements of an sRGB display, a gamma of 2.2 needs to be applied (i.e., $R^{1/2.2}$ etc. for R values normalized to a 0-1 range) to the linear data. The results below consider both the linear and gamma cases since the gamma value affects the error measures. Whether the RGB values are represented in linear (gamma = 1) or non-linear (gamma = 2.2) chromaticity in the corresponding space is computed as $r = R/(R+G+B)$, $g = G/(R+G+B)$, and $b = B/(R+G+B)$. Except when computing angles between chromaticities, only the r and g components are generally needed.




Figure 1. One frame each from the two bracketed sets of images of a single scene. The image on the left includes the frame holding the 4 Gretag Macbeth mini Colorcheckers, the one on the right excludes it. The Colorcheckers on the top and sides are at 45 degrees to the middle one.

Illumination Variation between Colorcheckers

Since the Colorchecker frame holds 4 Colorcheckers, we obtain measurements of the scene illumination from its grey patches from 4 different angles of incidence. Not surprisingly these measurements will not always agree. To evaluate the accuracy of a given illumination-estimation method, and MaxRGB in particular, requires a reliable ground-truth measurement of the illumination. For the tests described below, the average of the illumination chromaticities from the 4 Colorcheckers is used as the ground truth, but the average is a compromise. Taken over the 47 scenes, the median, mean, and maximum difference in angular error between the 6 possible pairings of the 4 Colorcheckers were 1.04, 1.79, 8.81 degrees respectively for gamma = 2.2, and 2.10, 3.48, 17.48 for corresponding linear case. The angle between two chromaticities r_a, g_a, b_a and r_b, g_b, b_b is computed as

$$\varepsilon_{i- angular} = \arccos \left[\frac{(r_a, g_a, b_a) \circ (r_e, g_e, b_e)}{\sqrt{r_a^2 + g_a^2 + b_a^2} \times \sqrt{r_e^2 + g_e^2 + b_e^2}} \right] \times \frac{2\pi}{360}$$

Since we cannot expect the performance of an illumination-estimation method to surpass direct measurement of the illumination, and all 4 Colorcheckers represent the chromaticity of the ‘true’ illumination, these numbers represent a lower bound on the mean, median and maximum illumination-estimation errors.

Results

Although each bracketed image set could be assembled into a single high dynamic range image, this is not necessary for testing MaxRGB. Instead, it was tested on the individual images from each bracketed image set (the ones without the Colorchecker frame in them) to measure its performance as a function of exposure. Its performance is measured in terms of the difference between the measured illumination chromaticity from the average of the 4 Colorchecker white patches and that estimated by MaxRGB. The chromaticity difference is evaluated in terms of both angular difference and Euclidean distance.

We evaluate the effect of clipping on MaxRGB performance by considering the sequence of bracketed images that have their maximum digital counts under a specified threshold. As

shown in Figure 2, for images without clipping, the illumination estimation error for MaxRGB remains low and relatively constant until there is a sharp rise at the point when the maximum intensity exceeds the 14-bit range of the camera. At that point, the high digital counts are clipped to the maximum value of 16383 and MaxRGB immediately fails. As more and more clipping occurs, we can expect MaxRGB eventually to approximate the Do-Nothing algorithm, which simply estimates the scene illumination as always being white, because the RGB maximum values in every 14-bit image will always be R=G=B=16383.

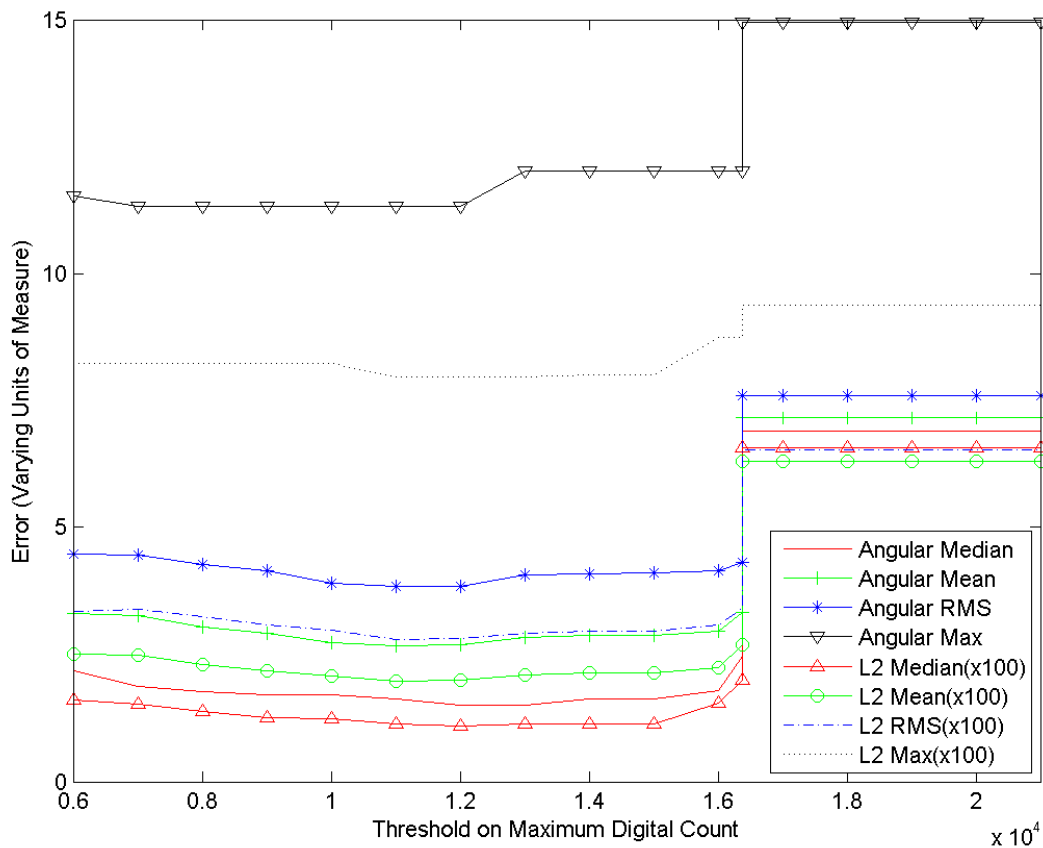


Figure 2. Median, mean, root mean square, and maximum of the angular difference and Euclidean L2 distances between the estimated illumination chromaticity and the measured illumination chromaticity as a function of the threshold on maximum digital count allowed within the image. Images have a gamma of 2.2 applied. The sharp increase in error occurs at the point at which the intensity exceeds the 14-bit range of the camera.

Although the plot in Figure 2 shows the error to be relatively constant under 16,000, there is a dip around 12,000. Using this as a cutoff, MaxRGB was tested on images with a maximum R, G, or B digital count of 12,000 or less. The results are tabulated in Table 1 for both the case of

linear and non-linear image data. The error for the non-linear case appears smaller, but this is mainly due to the fact that gamma compresses the RGB values, and hence the errors as well, rather than because the method actually works any better with non-linear image data.

	Angular Error in Degrees				L2 Distance (x 100)			
	Median	Mean	RMS	Max	Median	Mean	RMS	Max
Do-Nothing (gamma = 1)	14.73	15.87	16.24	29.49	14.70	14.44	14.57	20.98
MaxRGB with 12,000 cutoff (gamma = 1)	2.91	5.12	7.31	21.07	2.19	3.93	5.53	15.88
Do-Nothing (gamma = 2.2)	6.96	7.83	8.18	17.83	6.85	6.86	6.96	11.36
MaxRGB with 12,000 cutoff (gamma = 2.2)	1.50	2.68	3.84	11.33	1.09	1.99	2.81	7.96

Table 1. Performance of MaxRGB evaluated on both linear (gamma=1) and non-linear (gamma=2.2) image data in terms of angular error and Euclidean distance measures between the measured and estimated chromaticities of the illumination. The Do-Nothing error is the error in simply assuming the scene illumination is always white (i.e., estimating its chromaticity as $r=g=b=1/3$).

For the results in Table 1 the images were first spatially filtered using a 5-by-5 median filter. This was done to eliminate the possibility of a single noisy pixel creating a false reading for the maximum. The filter size was chosen by experimentation to find the optimal value.

Conclusion

MaxRGB was tested on images of varying exposure and found to work well when the exposure was such that high digital counts were preserved, not clipped. In fact the median angular error was not a lot greater than the variation in the measurements from the 4 Gretag Macbeth Colorcheckers. However, when the 14-bit exposure range of the camera was exceeded and clipping occurred, MaxRGB's performance immediately dropped significantly. These results show that the poor performance of MaxRGB previously reported in the literature may have more to do with the limited exposure range (8 to 14 bits per channel) of digital still cameras than the failure of MaxRGB's fundamental assumption that every scene must contain some region, or combination of regions, that reflects maximally in each of the R, G, and B sensitivity ranges.

References

- [1] Joost van de Weijer, Theo Gevers, and Arjan Gijsenij, "Edge-Based Color Constancy," *IEEE Transactions on Image Processing*, 16(9), Sept. 2007
- [2] G. Finlayson and E. Trezzi, "Shades of gray and colour constancy," in *Proc. IS&T/SID 12th Color Imaging Conf.*, pp. 37–41, Nov. 2004.
- [3] Kobus Barnard, Lindsay Martin, Adam Coath, and Brian Funt, "A Comparison of Computational Color Constancy Algorithms, Part 2; Experiments with Images," *IEEE Transactions on Image Processing*, 11(9), pp. 985-996, Sept. 2002.
- [4] Edwin Land and John McCann, "Lightness and Retinex Theory," *Journal of the Optical Society of America*, 61(1), Jan. 1971.
- [5] B. Funt, F. Ciurea, and J. J. McCann, "Retinex in Matlab," *Journal of Electronic Imaging* 13(1), 68–72, Jan. 2004.
- [6] Florian Ciurea and Brian Funt, "Tuning Retinex Parameters," *Journal of Electronic Imaging* 13(1), pp. 58–64, Jan. 2004.
- [7] Edwin H. Land, "The retinex theory of color vision," *Scientific American*, 237(6), pp. 108-128, Dec. 1977.
- [8] Brian Funt, Kobus Barnard and Lindsay Martin, "Is Machine Colour Constancy Good Enough?" *ECCV'98 5th European Conference on Computer Vision*, pp. 445-459, May 1998.
- [9] Nikon USA, <http://www.nikonusa.com/Learn-And-Explore/Nikon-Camera-Technology/ftlzi4ri/1/Nikon-Electronic-Forma-NEF.html>, accessed July 24, 2009
- [10] <http://en.wikipedia.org/wiki/Dcraw>, accessed July 24, 2009