This is a preprint of
Proc. IS&T Color Imaging Conference, Albuquerque, November, 2013

Chromaticity limits in color constancy calculations

John McCann, McCann Imaging, Belmont, USA;
Vassilios Vonikakis, Advanced Digital Sciences Center, Singapore;
Cristian Bonanomi, and Alessandro Rizzi, University of Milan, Italy

Chromaticity limits in color constancy calculations

John McCann, McCann Imaging, Belmont, USA; Vassilios Vonikakis, Advanced Digital Sciences Center, Singapore; Cristian Bonanomi, and Alessandro Rizzi, University of Milan, Italy

Abstract

There are two very different kinds of color constancy. One kind studies the ability of humans to be insensitive to the spectral composition of scene illumination. The second studies computer vision techniques for calculating the surface reflectances of objects in variable illumination. Camera-measured chromaticity has been used as a tool in computer vision scene analysis. This paper measures the ColorChecker test target in uniform illumination to verify the accuracy of scene capture. We identify the limitations of sRGB camera standards, the dynamic range limits of RAW scene captures, and the presence of camera veiling glare in areas darker than middle gray. Measurements of scene radiances and chromaticities with spot meters are much more accurate than camera capture due to scene-dependent veiling glare. Camera capture data must be verified by calibration.

Introduction

Studies of color constancy in complex scenes were originally done with simple images that allowed radiometric measurements of all image segments in the field of view. (Land, 1964: Land & McCann, 1971) McCann , McKee and Taylor (1976) introduced the field of Computational Color Constancy, using digital input arrays of 20 by 24 pixels. This size seems ridiculously small now, but it was very large for the time. They showed that:

- 1. Observer color constancy matches correlated with Scaled Integrated Reflectance of Mondrian areas calculated using an algorithm that made spatial comparisons.
- 2. The subtle changes in departures from perfect constancy were modeled well by cone crosstalk in the spatial comparisons.

As digital imaging advanced it became possible to automatically capture arrays of millions of digital values from complex scenes. As well, Computational Color Constancy has split into two distinct domains:

- Human Color Constancy (HCC) studies the ability of humans to be insensitive to the spectral composition of scene illumination. Its goal is to calculate the appearance of scene segments given only accurate radiances from each segment. No additional information, such as the radiance of the illumination is required, as in CIECAM models. The ground truth of HCC is the pyschophysical measurement of the appearance of each image segment.
- Computer Vision Color Constancy (CVCC) studies techniques for estimating the surface reflectance of objects in variable illumination. Its goal is to separate the reflectance and illumination components from the input array of scene radiances. If successful, these algorithms use the information from the entire scene to find an object's surface reflectance. The ground truth of CVCC is the physical measurement of the surface reflectance of each image segment.

Experiments measuring the human appearance of constant surface reflectances show considerable variation depending on scene content. Innumerable examples include: simultaneous contrast, color assimilation, and 3-D Mondrians (Albers, 1962; Parraman, et al., 2009, 2010). Computer vision's goal is to identify the surface, regardless of its appearance to humans. Thus, the two distinct kinds of Color Constancy do not share the same ground truth. They either have different fundamental mechanisms, or they have very different implementations. If they use the same underlying mechanism, then that mechanism would have to compute very different results. A single reflectance surface is seen in HCC to vary considerably with scene content, while the challenge to CVCC is to estimate the same constant reflectance in all scene contents.

Image capture

A common problem in both HCC and CVCC is the need for accurate data of scene radiances as the input for the models. The early spotmeter technique to measure simple targets was replaced by digital scans of high-dynamic-range film images (McCann, 1988); and more recently by multiple exposures using electronic imaging. Papers by Debevec and Malik (1997), Mitsunaga and Nayar (1999), Robertson et al. (2003), and Grossberg and Nayar (2004) propose calibration methods for standard digital images. Funt & Shi (2010) describe the advantages of using DCRAW software to extract RAW camera data that is linear and closer to the camera sensor's response. Xiong et al. (2012) and Kim et al. (2012) describe techniques for converting standard images to RAW for further processing. The common thread is that these papers attempt to remove the camera response functions from its digital data to measure accurate scene radiances.

Surface reflectance by first finding illumination

Helmholtz (1924) introduced the idea that constancy could be explained by finding the illumination first. If that were accomplished by some means, the quanta catch of the receptors divided by the quanta catch from the illumination equals a measure of surface reflectance. For Human Color Constancy (HCC) that approach could provide an alternative partial explanation of McCann et al.(1976), but not subsequent vision measurements. (McCann, 2012, chapter 27.5) For CVCC, that approach works within strict bounds imposed on the illumination. Obviously, it can work perfectly in illumination that is both spatially and spectrally uniform. Under these conditions there is a singular description of illumination falling on all objects in the scene.

Real scenes do not have uniform illumination. One CVCC approach assumes that the illumination is spectrally uniform, namely the illuminant has only a single spectrum falling on all objects, but variable in intensity. In such uniform spectral illuminants we can use chromaticity - a measure of spectral composition - to describe any intensity of that spectra. However, if the scene contains more than one spectral illuminant, such as sunlight and skylight, or colored reflections from a colored surface, then a single chromaticity value does not describe the illuminant on all areas.

Chromaticity as a Constancy Tool

Chromaticity is a tool used frequently in Computer Vision Color Constancy (Funt et al.1998; Finlayson et al. 2001; Ebner, 2007; Yao, 2008; Funt & Shi 2010; Yang et al. 2011; Gevers et al. 2012; Jiang et al. 2012; Ratnasingam et al. 2012)

Chromaticity is the projection of the three-dimensional color solid onto a plane defined by the RGB components. Position in the plane, defined by *r,g* are calculated

$$r = R / (R + G + B)$$
$$g = G / (R + G + B)$$

where R, G, B are the digital color values from the camera image. These chromaticity values are specific to the camera system and file format. Chromaticity is the ratio involving a sum. That requires strict linearity of input information. Camera chromaticity should not be confused with colorimetric chromaticities (x, y) that represent camera independent transforms of color-matching functions X,Y,Z. Further, we will not assume that all cameras have strict compliance to sRGB standards.

Color Constancy in Color Photography

Human Color Constancy (HCC) is a critical part of the success of color photography. HCC is extremely insensitive to the chromaticity values produced by photographic reproductions. Successful reproductions show a very large range of shifts in chromaticity from a single surface reflectance. This is seen in the standard photographs of the high-chroma BGRYMC squares in the ColorChecker target using multiple exposures.(Figure 1, left) A major limitation of camera chromaticity constancy in good color photography is the use of color masking. It has been in use since patented in 1889. (Albert, 1889; Friedman, 1944; Spencer, 1966; Yule, 1977)

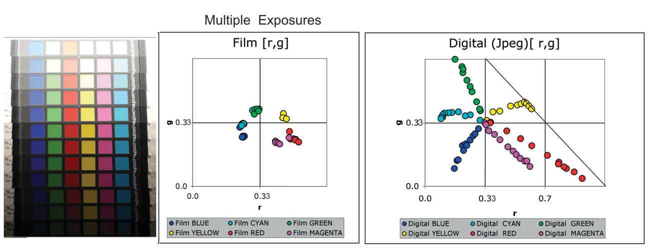

Figure 1 (left) Composite of segments of 12 different jpeg images of BGRYMC ColorChecker squares made with different exposures; (center) camera chromaticity estimates using RGB film separations without color masking; (right) digital camera chromaticities measurements from left jpeg images with color masking.

Reproducing color using Maxwell's (1964) color photographic technique shown at the Royal Institution in 1861 (3 B&W spectral separations) has no color masking. Three black and white photographs scanned and converted to linear (slope 1.0) RGB digital records have constant chromaticities with variations in exposure. Figure 1(center) plots the chromaticities from such films for BGRYMC ColorChecker squares. These chromaticities are constant with variable film exposures. The reason for constant chromaticities is that this system recreates the captured scene information, with output equal to input.

The problem is that color pictures without color masking do not look good. Color masking is the principle tool for "improving" scene capture information. Increasing chroma improves observer preference evaluations, and sells more cameras. It expands the color separation information to fill the viewing medias's color gamut using nonlinear processing. Color negative films and standard digital cameras have built in color masking.

The chroma *r,g* values plotted in Figure 1 (right) were calculated from the RGB separation values read by ImageJ from standard images (left). The photograph segments on the left are all appropriate renditions for human vision. Other than being a little too light, or a little too dark, they work well to do the task of reproducing the scene for humans. The fact that these chromaticities from multiple exposures show such extreme variations is of little concern to humans. That also means that chromaticity alone is of little use in predicting color appearance. However if accurate, it is highly useful in Computer Vision Color Constancy.

RAW and RAW* images

RAW format images were introduced to allow photographers more control of their image rendering. They provide the photographer the ability to control in software the many automatic image processing operations usually performed in camera firmware. The RAW format stores images that are much closer to linear records of the light falling on the camera sensor. There is no international standard for RAW. Each company provides a different software package that gives the photographer more control, but that does not mean that all RAW images are linear.

A better digital camera approach to measure scene radiances is to use RAW image data that is verified to be linear. Funt and Shi (2010) used DCRAW freeware library to read camera RAW and convert it to a black and white image that is as close as possible to the sensor's response. It is a single image that contains the different sensor responses to the Bayer pattern mosaic.

We used the (LibRaw 2013) Image Decoder Library, which is built on DCRAW. More specifically, we used the "unprocessed" function of LibRaw which outputs the unprocessed data of the RAW file, without applying any processing such as demosaicing, denoising, white balance, gamma modification, enhancement, compression, or min/max normalization. This results in an image with visible Bayer pattern and not normalized R, G and B responses (G>R>B). For this reason, we used separate calibration RGB LUTs to scale and linearize the sensor responses in order to remove any tone scale nonlinearities, and attain the same output response for achromatic patches. We refer to the output of this process as RAW*, described in McCann and Vonikakis (2012). All images referred to as RAW* have no outlier pixels, because we experimentally verified linearity.

Limitations of Camera Chromaticity

We saw above in Figure 1 that standard digital photographs, representing preferred reproductions, are highly nonlinear and cannot be used for accurate illuminant estimation in CVCC algorithms. Nevertheless, camera color-balance algorithms can use this data successfully. They process images expecting further processing by human observers' color constancy.

Figure 2 compares the use of color spaces of standard image formats with a linear RAW format. Figure 2 (left) shows a 3-D plot of RGB digits from the best standard image of the ColorChecker target with color masking. The color squares are plotted in a color cube (256 levels per axis). The digits fill much of the color space because of color balance (placement of achromatics along the cube diagonal) and color masking (amplification of chroma information).

Figure 2 (right) shows RGB digits in the best exposed RAW photograph. It has more digits (14 bit, scaled to match jpeg). We measured the range of accurate scene data using a Canon 60D digital camera with an EF-S17-85mm zoom lens. The RAW data uses a small fraction of the color space volume. It fills much of the achromatic range, but is very limited in chroma.

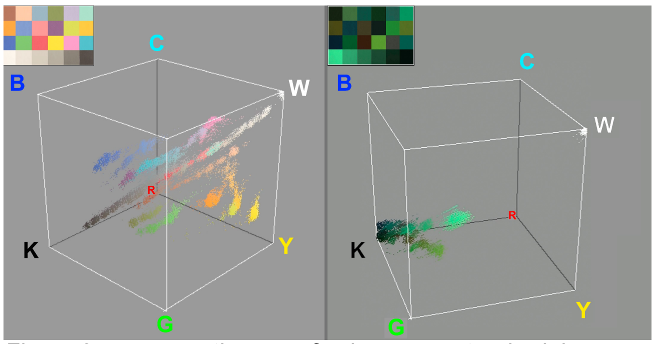

Figure 2 compares the use of color space standard image formats (jpeg) and linear RAW formats

There are many algorithms that calculate statistics of spectral values of illumination using camera chromaticity values from the scene. Because the range of possible linear camera chromaticities is so small compared to luminance, the requirement for accuracy of these measurements is critical. A lot depends on the intended discrimination of the CVCC algorithm. If the goal is to differentiate hues (red vs. cyan) accuracy is less important. However, if the goal is to differentiate individual chips in the Munsell Book of Color, then great precision is required. Although desirable, full spectral correlations are not usually attempted in CVCC because most digital cameras have broad-band Bayer pattern RGB sensors. As well, many papers use short cuts of studying scene averages, or database correlations, rather than measuring the accuracy of individual areas to determine the precision of color constancy algorithms. Sometimes papers look for correlations between average calculated chromaticities and average scene chromaticities.

The CVCC goal is to calculate measured surface values and requires a calculated physical number, independent of the human observer. The analysis of CVCC algorithms needs to carefully avoid visual inspection of pictures of scenes processed by an algorithm. Visual inspection incorporates human color-constancy in the analysis. We cannot judge the contribution of the CVCC algorithm separate from the HCC contribution.

Any CVCC prediction that involves an average, or statistical analysis of chromaticity values should restrict the sample set to accurate chromaticity values (outlier free data). Our current study shows that the range of accurate chromaticity values is much smaller than we had first thought.

Calibrated RAW digits

We made a series of 48 exposures of a Munsell ColorChecker using a Canon 60D camera (IS0 100) with an EF -S17- 85mm lens. We turned off all automatic firmware. The exposures covered a range of 16 stops with 1/3 stop increments. We used ImageJ to measure the RAW RGB values in the different areas in the Bayer pattern black and white image from DCRAW.

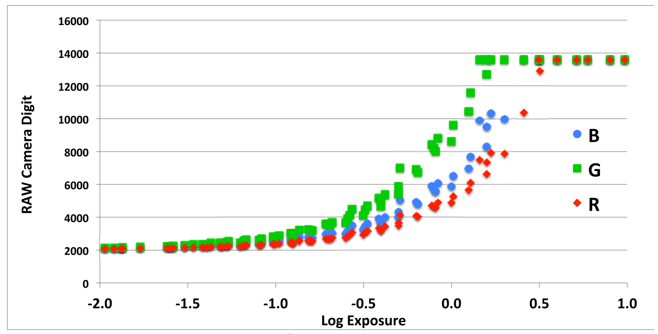

Figure 3 plots the RAW RGB digits vs log exposure.

The maximum RAW digital value was 13,584 out of a possible 16,384 addresses in a 14-bit digital value. We found the minimum RAW value to be 2049. This pedestal, or black floor limit, is the lowest value found in the camera response. This value is also found when one takes a digital image in a light-free dark room. Thus, this camera used a range 11,034 (13,584-2,049) digits for all three color separations. This is 70% of the 14-bit range of digits.

We calculated the relative exposure for each achromatic ColorChecker square using the relative Y reflectance values (Pascal, 2002) and the relative exposure time with constant aperture. The plots of log relative exposure vs. RGB camera digit are shown in Figure 3. These 11,034 camera digits represent the exposure G range of 2,951:1 (3.5 log units); the B range of 5,012:1(3.7 log units); and R range of 6,310:1 (3.8 log units).

We found good fit of the RGB data using the following linear functions:

$$RAWG = 3.908 * e + 2049$$

$$RAWB = 2.301 * e + 2049$$

$$RAWR = 1.828 * e + 2049$$

when e is the % maximum exposure at the camera's maximum saturation digital value, in this case 13,584.

In figure 4, the green line plots the entire usable range between 13,584 and 2049. The red dashed line plots the G separation range used by the achromatic squares in the ColorChecker target, with icons for white, middle gray (G3) and black.

This analysis uses the achromatic squares in uniform illumination. This has a range of luminance of 29:1. In the case of this *Best RAW* exposure, there are 896 digits available between the ColorChecker White and max saturation value. As well, there are 274 digits available between the ColorChecker black and the minimum black level pedestal. Together they are 8% of the 14-bit data range. The unused digits above CCwhite are needed to render nonuniform illumination; e.g. highlights and light sources in real scenes. The unused digits below CCblack are needed for shadows. Both are needed to calculate the color balance of the image, since

RAW digit values are recorded before the color balancing camera firmware. RAW digits for a particular scene are independent of the color balance settings of the camera. Only 10% of the usable range of digits is used to meet the many different challenges of real-scene dynamic range image capture. With the exception of photographs of a foggy day, the ColorChecker target in uniform spatial and spectral illumination is one of the smallest-range scenes encountered in camera capture.

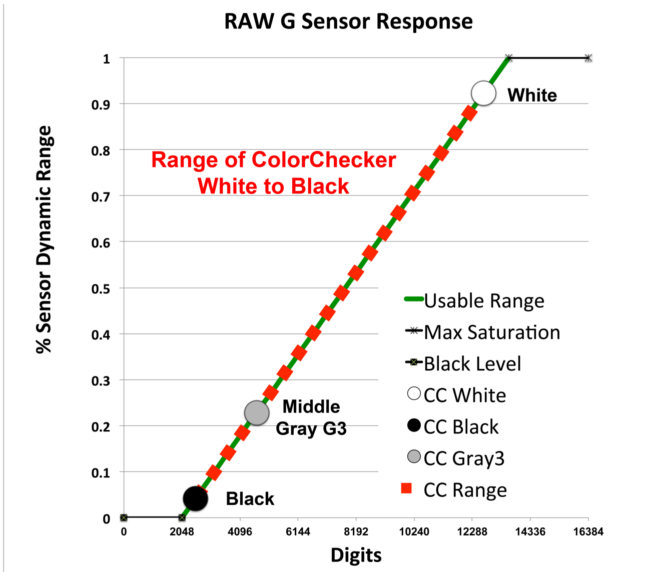

Fig 4 plots the % usable RAW sensor data vs. digit.

ColorChecker Chromaticities minus Pedestal

We can use this ideal range shown in Figures 4 to calculate chromaticity *r,g* from RAW* camera digits. Here we subtract the dark minimum pedestal value from the stored digital camera response If the additive constant black level is not subtracted, then chromaticities vary with exposure, meaning that these *r,g* values are inappropriate for CVCC algorithms. This pedestal subtraction is standard practice among camera designers, but it is not always discussed in CVCC papers using camera chromaticity.

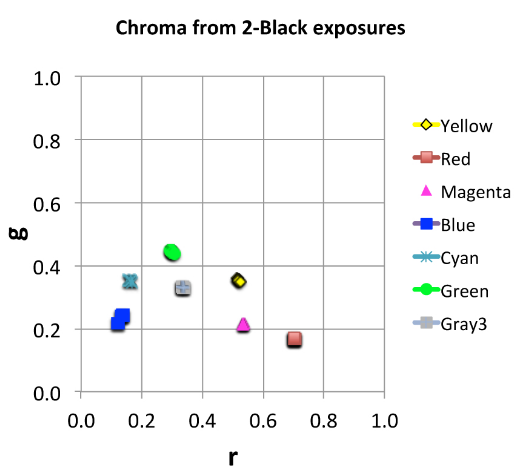

Figure 5 shows that camera chromaticities can correlate with surface reflectance using linear RAW camera digits.*

LibRaw Image Decoder Library has a switch (called "black level subtraction") that, when set, does something similar. Under some conditions with some cameras, RAW digits can have the pedestal included.

The consequence of subtracting 2049 from RAW camera digits is to amplify the range of the digital values. The G, B, R digital ranges increase from 7:1(13,584/2,049) to 11,384:1.

Using the ColorChecker test target we found excellent linearity over a range of exposures. In this ideal range, the camera

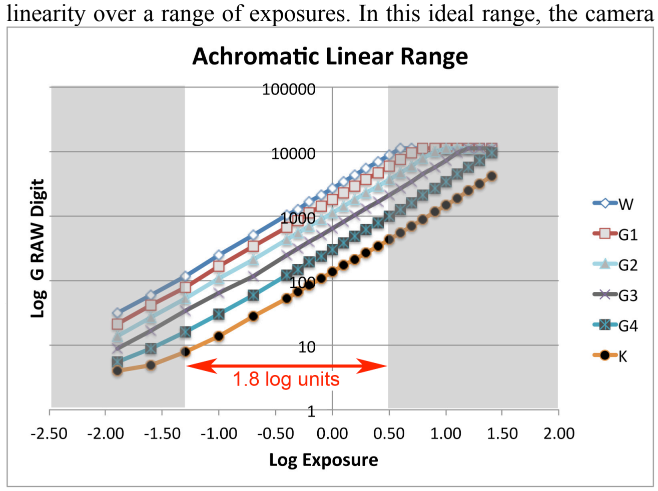

Figure 6. RAW G camera digits with variable exposure for the six achromatic squares in the ColorChecker. The graph's gray backgrounds identify exposures in which part of the scene has a nonlinear response. The graph's white background identifies the 1.8 log unit range of sensor exposure in which the entire sensor response is linear.

We see the black reflectance square (K) vary between max and min digits over a range of exposures of 3.5 log units. Each of the gray squares has a similar range that is displaced by the range of reflectance values. The window for all 6 gray squares is only 1.8 log units. Above and below that range the errors introduced by both max saturation, and minimum black level disrupts linearity. As soon as we reach these min-max limits, we begin to alter the average values of the scene, and their average chromaticities in CVCC analysis.

We saw above that the camera dynamic ranges were greater for B and for R response functions, but for any chromaticity measurements all three components GRAW, BRAW and RRAW must be strictly linear.

Review

If we review the calibration experiments so far we have measured that:

- Ordinary camera images that look good have large amounts of color masking that increase the chroma of all non-achromatic stimuli.
- Color masking is a chromatic amplification that distorts a color sample's reflectance chromaticity.
- A manufacturer's RAW software performs similar chromatic amplifications of scene data and cannot be used to measure scene chromaticities.

- RAW extracting libraries such as DCRAW or LibRaw provide the closest access to the achromatic sensor response as filtered by Bayer pattern color filters.
- Between saturation at very high exposures, and the noise pedestal for low exposures, the camera's response is linear.
- The RAW R, G, B responses have different slopes depending on the spectral content of the light.(Figure 3)
- The linear RAW R, G, B responses for chromaticity measurements are necessary, but not sufficient. The RGB response functions need to be color balanced, so that achromatic scene objects have equal RAW R, G, B digits. (RAW* calibration procedure - (McCann and Vonikakis 2012))
- Any camera digits in each of the RAW R,G,B images that reach min or max linear limits are outliers and they reduce the number of pixels that are usable camera chromaticity responses.

Under the above conditions with DCRAW/LibRaw data camera chromaticities are proportional to scene chromaticities and can be used for calculating the average illumination and performing image analysis. For the low-range ColorChecker scene, the range of exposures that is free of outliers is limited to total 1.8 log units in the Canon 60D.

Camera Glare

The remaining task is to test for veiling glare in the camera. McCann and Rizzi (2007) measured that scene content limited the dynamic range of cameras. They used an achromatic transparent target with a dynamic range of 18,619: 1. They found that glare from the entire scene added to the desired scene radiances. The effect of glare cannot be removed because it is the sum of all the very small amounts of scattered light from every pixel. The amount of glare is dependent on that pixel's intensity and distance from each scattering pixel. The camera green digit Gc is the sum of scene radiance Gs and veiling glare v. It is the sum of glare from all other pixels and light from outside the camera sensors field of view.

$$G_c = G_s + v$$

$$v = \sum_{1,1}^{x,y} f(G_s, d) + L$$

where x is the maximum horizontal pixel location address, and y is the maximum vertical pixel location address. Typically the product of x*y equals millions of pixels. Veiling glare is the scene radiance G_s convolved with the glare spread function (GSF) for every other scene pixel in the camera's field of view. Rizzi & McCann, 2009). Glare is a function of the distance (d) between the location (x_i,y_i) of G_s. L represents the glare light that falls on the camera lens from light sources outside the camera's field of view. The forward calculation using all the accurate scene radiances and the camera's glare spread function can calculate the resulting image with glare as long as we know all the scene radiances at each pixel without glare. That is the easy part.

The reverse process is not possible. If we need to calculate the accurate radiance Gs for a scene pixel from the camera radiance digits, we have to accurately calculate the values of millions of separate $v_{x,y}$ contributions from every pixel, and the glare L from light falling on the lens from outside the camera's field of view. No one has shown a real solution to this problem. Optical experts agree that this inverse calculation is not possible, as stated explicitly in the ISO standard for measuring veiling glare. (ISO 9358, 1994) As yet, claims of approximate solutions do not stand up to optical imaging standards.

Testing for Glare in Ordinary Photos

The following experiments measure the influence of veiling glare in three different sets of ordinary photographs. Our experiments used the ColorChecker achromatic papers that have a radiance range of only 29:1 in uniform illumination.

Beach Scene

To illustrate the effect of scene content we made RAW* G images of four scenes with different scene content. In all cases the lens was shielded from the sun, and the camera pointed away from the sun. The pictures were made with a different camera (Canon D60) with subtracted dark pedestal. We used a Canon EF 50 mm F/1.8 II primary lens having only five optical elements, so as to minimize glare. The scenes photographed, shown in Figure 7 bottom were:

- ColorChecker inside an automobile shaded by trees
- Closeup of the Color Checker on a beach
- Same scene - greater distance (1/9 area)
- Same scene taken much further back (1/144 area).

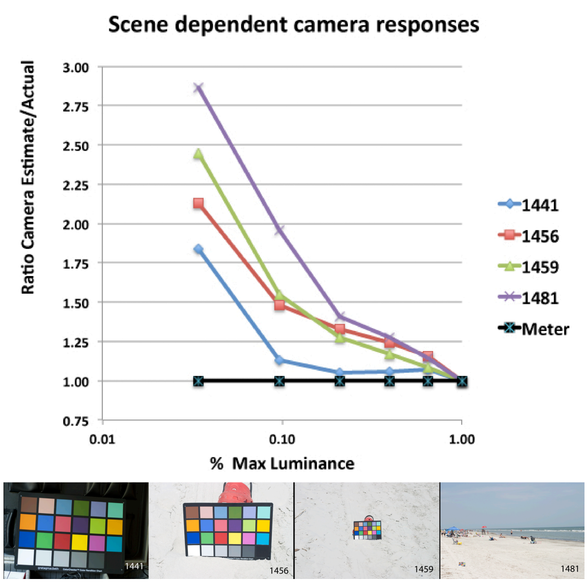

Figure 7 (top) Ratio of camera estimate/actual vs. % max luminance;(bottom) The scenes: CC on black, and on the a beach at different distances.

The black line in Figure 7 plots the meter readings of the reflectance of the six achromatic square along the bottom of the ColorChecker. The horizontal axis is % maximum luminance read by a Konica Minolta C100 meter for those six squares.Our meter readings match the values reported by Pascal (2003). The vertical axis is the ratio of camera-estimated reflectance to actual reflectance. Theses estimates were made from RAW*G data known to be in the linear camera response region.

Photo 1441 (Figure7, left) taken inside an automobile on a dark background shows good correlation for five achromatic squares. The camera reflectance estimate for black is 1.75 times actual. The beach scenes (1456, 1459, & 1481) showed larger errors for all squares, reaching 2.85 overestimate for black reflectance in Photo 1481. The beach scenes show overestimates of middle-gray square greater than 1.25 times. Scene content introduces substantial errors due to veiling glare even though the scene is low-contrast and captured in uniform illumination. A beach scene has a distribution of image luminance values that are nearly all at the maximum luminance value, thus providing more scattered light in the camera's image plane. The ColorChecker in a dark environment has less veiling glare and smaller errors from scene content.

Uniform illumination in the laboratory

We also made four sets of exposures with a Canon 60D camera varying scene content. They are:

- Wall [Center]: a ColorChecker (CC) target in uniform daylight on a white wall on all sides. [Scene range= 29:1]
- Wall [Upper left]: We moved the ColorChecker so that it was in the upper left of the camera's field of view. The Black square was in the center of the field of view.
- Wall [Lower right]: We moved the ColorChecker so that it was in the lower right of the Camera's field of view. The Black square was in the corner of the field of view.
- Window: a CC target in uniform reflected light on a window diffuser that has 3 stops (0.8 log units) more luminance than the white in the CC. The bright window background is out of the camera's field of view. [Scene range= 249:1; Camera field of view range 29:1]

Figure 8 plots the data for these four scenes from G RAW* camera digits (with verified linearity) after the dark-noise pedestal was removed. The horizontal axis is the % maximum luminance in the camera's field of view. We analyzed only the 6 achromatic grayscale squares because we have measurements of reflectance independent of camera's spectral sensitivity.(Pascal, 2003) Two additional sets of spotmeter measurements confirmed the Pascal data. Accurate camera estimates of scene reflectance fall on the black line with a ratio of 1.0.

The camera estimate of scene reflectance was the G RAW* camera digit divided by the camera digit value of the white square. The graph's vertical axis plots the ratio of the camera-estimated to actual reflectance. Figure 8 shows that scene content, *in and out of the camera's field of view*, influences camera estimates of scene reflectances.

The three wall pictures show that camera estimates of reflectance vary with position in the camera's field of view. When the ColorChecker is in the top-left corner of the field of view, the black square falls in the center of the picture, surrounded by the white wall on three sides. It has the largest overestimate of reflectance. The lower right picture places the black square in the corner and has the lowest overestimate. The center values fall between the others.

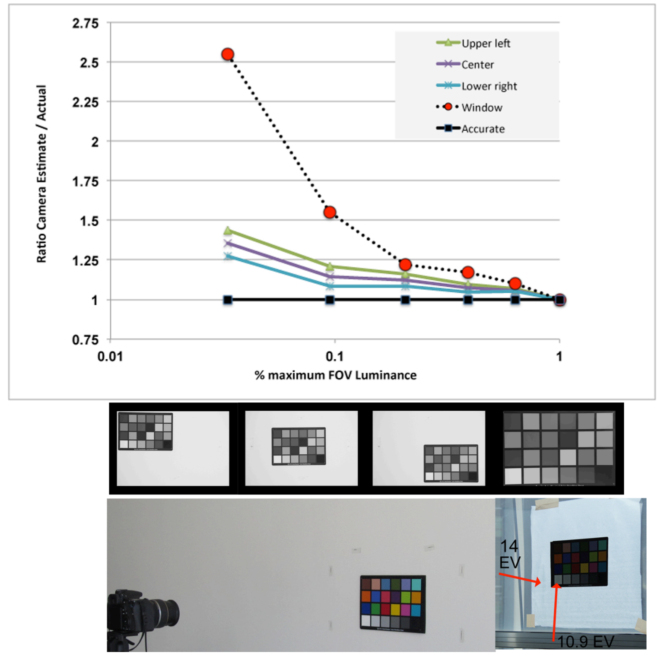

Figure 8 (top) Ratio of camera estimate/actual vs. %max luminance in the field of view; (middle) The field of view (FOV) images of G separations captured by the camera; (bottom) The scenes: CC on a large white wall and window.

The fourth image in Figure 8 studies the effect of light outside the camera's field of view. The color checker was placed on a window in front of a translucent paper diffuser. The opaque ColorChecker was illuminated by light reflected off a wall behind the camera. The white paper was illuminated by that light plus transmitted light from outside the window. The white surround had 8 times the luminance as the white in the ColorChecker.

Additional light outside the field of view (Window) showed an even larger departure from accurate surface reflectance measurements. The area around the ColorChecker was only 8 times higher luminance than the white square. That results in a scene range of only 249:1. That is not a large value in HDR imaging. Nevertheless, the camera estimate of black reflectance was 250% the actual value.

Glare affects chromaticity

A third pair of pictures using a Panasonic DMC-G2K camera showed a similar change in camera responses to scene content.

Figure 9 plots the same comparison photographs of an outdoor scene with a ColorChecker on white stairs; and a second picture in a dark room with a black background. The Window data from Figure 8 is repeated here for comparison.

In Figure 9 the camera-estimated black paper is only 1.5 times actual, when photographed in a dark room on a black background. When photographed outside on bright white marble steps, the black estimate is 2 times actual. This is another example of variable camera-reflectance estimates from different scene content.

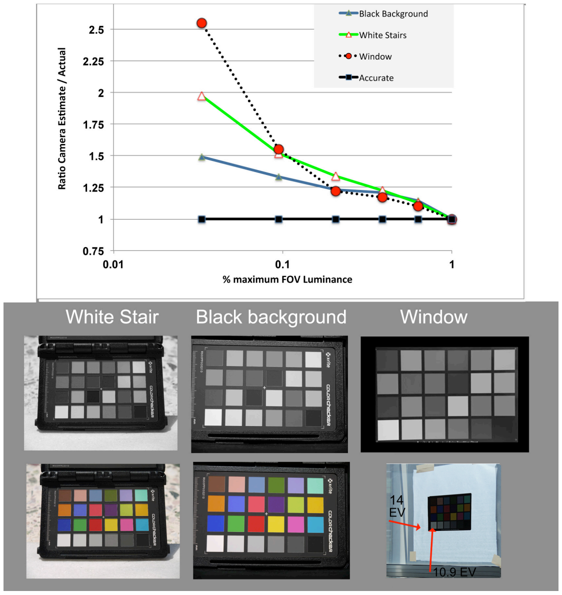

Figure 9 (top) Ratio of camera-estimated/actual reflectance vs. % max luminance in the field of view; (middle) The field of view (FOV) images of G separation captured by the camera; (bottom) the scenes: daylight on white stairs, darkroom on black, Window from Figure 8.

Figure 10 plots the Panasonic DMC-G2K camera chromaticities of Figure 9 scenes for the RYGCBM ColorChecker squares. The linear RAW* data came from exposures using different spectral illuminants, so the Gray3(middle gray) digits were used to color balance the images. The data show that camera calculated chromaticities from linear RAW* camera digits vary with scene content. The results show that photographs of a ColorChecker taken: a) on white marble stairs, and b) in a darkroom on black, render different camera chromaticities. Camera chromaticities of constant reflectances vary with scene content. These departures from accurate reflectance values correspond to the effects of veiling glare.

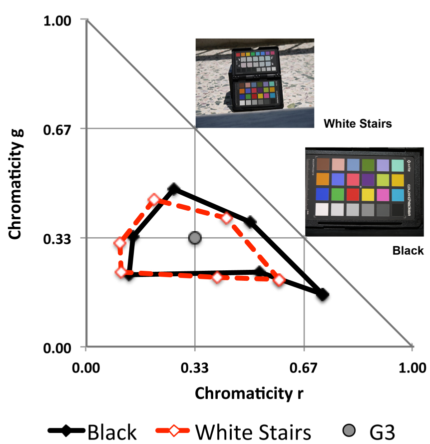

Figure 10 shows variable camera chromaticities with scene content: White Stairs; and Black. G3 is the chromaticity of middle gray. The white stairs reduces the range of camera chromaticities.

In human vision departures from accurate scene values in the retinal image are of no consequence. Human spatial image processing is insensitive to such glare distortions. Human neural contrast mechanism tends to counteract glare. Images with the lowest retinal contrast have the highest apparent contrast. (McCann & Rizzi, 2012)

Discussion

CVCC has the task of identifying an object's reflectance in an unknown illuminant. The accuracy of a CVCC algorithm is its ability to discriminate between an object and other similar objects. A frequently used approach is to calculate the chromaticity of the illuminant from the array of all scene radiances. Finding the illuminant provides the path to finding the object's surface reflectance. Chromaticity has the significant advantage that it is indifferent to the amount of illumination, as long as there is no variation in its spectra. Chromaticity CVCC algorithms convert local changes in spectral illumination directly into object's spectral reflectance changes. Scenes with multiple spectral illuminants, such as sunlight and sky light, or scenes with light reflected from colored objects, introduce large errors in CVCC results.

The data reported here shows that camera chromaticity correlates with surface reflectance in Figures 1(left) and 5. These experiments used uniform spatial and spectral illumination. These graphs restricted their data to chromaticities that were calculated from RGB RAW* digits that had been verified as linear responses to scene radiances. Our data also shows that many camera characteristics can introduce nonlinear responses that are critical to CVCC research. They include:

- Color masking in standard images
- RAW image processing (in camera software)
- Dark value minimal pedestal
- Camera veiling glare (caused by scene content)

Accurate calculations of CVCC require camera digits with verified linearity. Otherwise, any or all of these nonlinearities will introduce error in these calculations.

Camera Glare

Sensors count photons. In theory, when we vary the exposure time with constant aperture, we alter the photon count falling on the sensor, and the sensor responds in a proportional manner. The measurements in Figures 7, 8 & 9 show a different behavior. These plots show that the ratio of camera-estimated reflectance to actual reflectance (telephotomer readings) varies depending on the scene content. We observed this behavior for all of the different exposures, showing that this effect is consistent with camera veiling glare. Glare adds the same scene-dependent parasite images (Sowerby, 1956; Kingslake, 1992; McCann & Rizzi, 2012) to every exposure. We have seen that Canon camera's RAW* digits (minus pedestal) are extremely linear near its minimum value. The image falling on the sensor is the sum of the scene radiance image plus the veiling glare image. For each individual pixel the veiling glare value is a constant fraction of the scene radiance value. It does not matter that this fraction is different for every pixel. Our measurements showed that the proportion of glare to scene radiance does not change with exposure. That is the behavior we observe in there measurements. That is evidence that veiling glare is large enough to alter scene data needed to calculate accurate chromaticity values. Camera reflectance estimates show between

10% and 20% error below middle gray (scene radiance is only 0.7 log units below maximum scene radiance) for these camera configurations. Recall that the ColorChecker in uniform illumination is an abnormally low-contrast scene. Nevertheless, it is severely affected by glare introduced in the captured image by the scene content.

The glare problem becomes acute when one considers chromaticity accuracy in real scenes with real illumination. The most difficult part of the problem is that accuracy is a function of scene content, with no compensation mechanism for the variable size of veiling glare affecting camera estimates of scene radiances. The experiments in Figures 1, 3, 6, 7, 8, and 9 used uniform illumination. However, real scenes are characterized by non-uniformities in both spatial and spectral light distributions. This means that the effects of veiling glare, on the estimated radiances, can be considerably greater in real life scenes compared to the ones reported in this study.

Veiling glare is responsive to scene radiances, not objects' reflectances. That means that any cause of nonuniform illumination will affect the veiling glare contribution to a camera chromaticity value. For example, a white reflectance square in a very deep shadow will have a significant camera veiling glare component added to its scene radiance component. This will limit the dynamic range and consequently, the accurate calculation of camera chromaticity values. Glare can come from any light falling on the lens, whether or not the source is included in the camera image's field of view (ISO 9358, 1994). Unfortunately, there are no meaningful assumptions we can make about the statistics of scenes we could capture with our cameras, in order to predict the contribution of glare. If we make any assumption about the distribution of light in all scenes, we can always find a scene that is an exception to that rule. The camera can be aimed at scenes with any, and all possible statistics.

There are viable and frequently used techniques to subsample the camera image to remove undesirable *outlier* data. These techniques work well in color photography for human viewing, but not in CVCC. The rationale for subsampling needs no more justification that most people will not complain about the picture's image quality. However, to strictly verify a scientific principle one needs a precise rationale for subsampling to eliminate the effects of *outliers*. When we demonstrate that an algorithm works in a subsample of the scene, we must be careful not to assert that it works for the entire scene. Again, a lot depends on the goal for differentiating colored surfaces. If the goal is to identify chips in the Munsell Book, the accuracy of camera chromaticity must be very high. Our experiments show that the best ground truth for evaluating CVCC algorithms is spotmeter reflectance measurements.

Multiple Exposures

Over the last 25 years many papers have used multiple exposures in electronic imaging to increase the dynamic range of captured scene information. (Ochi and Yamanaka, 1985; Alston et al., 1987; Mann, 1993; Debevec and Malik, 1997; Reinhard et al., 2006) As seen in our measurements above, accurate scene information from this technique is limited by veiling glare. The captured range strictly depends on scene content. Some scenes, such as beaches, show that the limit of accurate information is reached with light-gray reflectances in uniform illumination.

The very valuable information found in multiple exposures is the linearity measurements described in Figure 6. By testing the difference of digital values between different exposures, we can easily measure the range of linearity of a camera response to an individual scene. Again, verified linearity is essential if one uses chromaticity in scene analysis.

Conclusions

Computer Vision Color Constancy sometimes employs a pixel-based analysis of the image statistics of camera chromaticities. We measured the range of conditions in which camera digits generated accurate chromaticities. First, we showed the need for using captured data before the color masking processes found in standard images. Second, RAW image data that bypassed standard camera processing, using DCRAW, or LibRAW, software, gives us access to linear camera data. The largest linear camera response we measured here was an exposure range of flux of 1.8 log units. This is the range of G sensor digit between black (digit = 7) and white (digit = 11,384) using a Canon 60D camera. Third, comparisons of camera measurements vs. spotmeter radiance readings show that scene-dependent veiling glare reduces the range of accuracy of chromaticity measurements. Consequently, camera estimates of constant scene reflectances vary with scene content, since veiling glare affects radiance estimates, that in turn affect chromaticity. Our study shows that reflectance chromaticity has very small limits of accuracy when measured by a camera. More importantly, these limits are much smaller than the 3.0 log unit range of luminances in many typical real-world scenes.

Computer Vision Color Constancy is best evaluated in real scenes by comparing camera captured reflectance estimates with actual reflectance measurements as ground truth. That way we can get a true evaluation of the algorithm's accuracy, taking into account glare and linearity issues found in cameras.

Acknowledgments

This study is partially supported by the research grant for ADSC's Human Sixth Sense Programme from Singapore's Agency for Science, Technology and Research (A*STAR). We want to thank Brian Funt and Mary McCann for advice and helpful discussions.

References

Albers, J. "The Interaction of Colors", New Haven, CT: Yale Un. Press (1962).

Albert E, German Patent 101379 (1889).

Alston L, Levinstone D, Plummer W, "Exposure Control System for an Electronic Imaging Camera Having Increased Dynamic Range", US Patent 4,647,975, filing date 10/30/1985, issue date 3/3/1987.

Debevec PE, Malik J., "Recovering high dynamic range radiance maps from photographs", SIGGRAPH 97 Conf. Proc., 369-78 (1997).

Ebner M, Color Constancy, (Wiley, Chichester, 2007).

Finlayson GD, Hordley SD, Hubel PM, "Color by correlation:" IEEE Trans.Pattern Analysis and Machine Intel., 23, 1209-1221 (2007).

Friedman JS,"Chapter 19, Masking", in History of Color Photography, (The American Photo Publishing Co, Boston, 1944) 273-295.

Funt B, Barnard K & Martin L, "Is Machine Colour Constancy Good Enough?", Proc. 5th European Conference Computer Vision (1998).

Funt B & Shi L, "The rehabilitation of MAXRGB", Proc. IS&T/SID Color Imaging Conference, 18, 256-259 (2010).

Gevers T, Arjan Gijsenij A, Joost van de Weijer J, & Geusebroek JM, Color in Computer Vision, (Wiley, Chichester 2012).

Grossberg M, Nayar S, "Modeling the space of camera response functions". Pat,Anal, &Machine Intel, IEEE Trans; 26(10):1272-82 (2004).

Jiang T, Nguyen D and Kuhnert KD, "Auto White Balance Using the Coincidence of Chromaticity Histograms", Signal Image Tech & Internet Based Sys, 201-208 (2012).

Kim SJ, Lin HT, Lu Z, Susstrunk S, Lin S & Brown MS, "A New In-Camera Imaging Model for ColorComputer Vision and its Application", IEEE Trans. Pat.. Anal & Mach. Intel, 34,2289-2302 (2012).

Land EH, "The Retinex", Am. Scientist, 52, 247-264 (1964).

Land E & McCann JJ Lightness and Retinex Theory, J. opt. Soc. Am. 61, 1 (1971).

“LibRaw” Image Decoder Library <www.libraw.org/> (2013)

ISO 9358: "1994 Standard Optics and Optical Instrument: Veiling Glare of Image Forming Systems. Definitions and Methods of Measurement" IISO, Geneva (1994).

Kingslake R, Optics in Photography, (SPIE Press, Bellingham,1992).

Mann S, "Compositing Multiple Pictures of the Same Scene", Proc. IS&T Annual Meeting, 46, 50,-52 (1993).

Maxwell JC, The Scientific Papers of James Clerk Maxwell, ed. Niven, (Dover, New York, 1965) 410–44.

McCann JJ, McKee S & Taylor T. "Quantitative Studies in Retinex Theory, A Comparison BetweenTheoretical Predictions and Observer Responses to Color Mondrian Experiments", Vision Res., 16, 445-58 (1976).

McCann J & Rizzi A, "Camera and Visual Veiling Glare in HDR Images", J. Soc. Info. Display 15(9), 721–30 (2007).

McCann J & Rizzi A The Art and Science of HDR Imaging, (Wiley, Chichester, 2012).

McCann J & Vonikakis V, "Accurate Information vs. Looks Good: Scientific vs. Preferred Rendering", Proc. CGIV, 6. 231-238 (2012).

Mitsunaga T, Nayar S., Radiometric self calibration. In: Proc. of CVPR; vol. 2. 1999, p. 374-80 (1999).

Ochi S & Yamanaka S, US Patent, 4,541,016, filed 12/29/1982, issued 9/10/(1985).

Parraman C, Rizzi A, & McCann J, "Colour Appearance and Colour Rendering of HDR Scenes: An Experiment", Proc. IS&T/SPIE Electronic Imaging, San Jose, 7241-26 (2009).

Parraman C, McCann J, & Rizzi A, "Artist's Colour Rendering of HDR Scenes in 3-D Mondrian Colour-constancy Experiments", Proc IS&T/SPIE Electronic Imaging, San Jose, 7528-31 (2010).

Pascale D, "RGB Coordinates of the Macbeth Color Checker. http://www.babelcolor.com/mainlevel/download.htm (2002)

Ratnasingam, S & McGinnity TM, "Chromaticity Space for Illuminant Invariant Recognition", IEEE Trans. Im. Proc, 21, 3612-23 (2012).

Reinhard E, Ward G, Pattanaik S, Debevec P, High Dynamic Range Imaging Acquisition, Display& Image-Based Lighting (Amsterdam: Elsevier, Morgan Kaufmann, 2006).

Rizzi A & McCann J, "Glare-limited appearances in HDR images, J. Soc. Info. Display 17, 3-12 (2009).

Robertson M, Borman S, Stevenson R, "Estimation-theoretic approach to dynamic range enhancement using multiple exposures", Journal of Electronic Imaging 12(2), 219 -28 (2003).

Sowerby A, Dictionary of Photography, 18th ed. (Philosophical Library, New York, 1956) p 568.

Spencer D.A., Chapter XIV Masking and Colour Correction, in Colour Photography in Practice, The Focal Press, London, 297-328, (1966).

von Helmholtz, H, Physiological Optics, Vol. II, J. P. Southall, ed., Optical Society of America, 264-65 (1924).

Xiong Y, Saenko K, Darrell T, & Zickler T, "From Pixels to Physics: Probabilistic Color De-rendering", Proc. IEEE Comp. Vis. & Pattern Recog. (CVPR), (2012).

Yang Q , Wang S, Ahuja N and Yang R, "A Uniform Framework for Estimating Illumination Chromaticity, Correspondence, and Specular Reflection", IEEE Trans. Im. Proc., 20, 53-63 (2011).

Yao L, "Estimation Illumination Chromaticity", IEEE Int. Symp. Intel. Inform.Tech. Appl., 2, 756-760 (2008).

Yule, J A C, Colour Reproduction in the Graphic Arts, in Color Photography, in Neblette's Handbook, 7th ed., Van Nostrand Reinhold, New York, 466-480, (1977).

Author Biographies

John McCann (B.A. in Biology Harvard,1964) worked in, and later managed, the Vision Research Lab at Polaroid from 1961 to 1996. He is a Fellow of IS&T and OSA; past President IS&T and Artists Foundation, Boston; IS&T/OSA Land Medalist, and IS&T Honorary Member. He continues to consult and study vision.

Dr. Vassilios Vonikakis (Diploma degree &PhD in Electrical and Computer Engineering, Democritus University of Thrace (DUTH),2002 and 2008). Currently he is a Post-Doctoral fellow in the Advanced Digital Sciences Center (ADSC), Singapore, working in the field of Multimedia and Image Retrieval. His research interests include Biological Vision, Computational Photography, HDR Imaging, Multimedia, and Content Based Image Retrieval.

Cristian Bonanomi is research fellow at the Dept. of Computer Science of Università degli Studi di Milano. He received a Ph D. in Computer Science in 2011. His research interests include computer graphics, human perception, color appearance models and digital photography.

Alessandro Rizzi is Associate Professor,Dept. of Computer Science, Università degli Studi di Milano teaching Fundamentals of Digital Imaging, Multimedia Video and Human-Computer Interaction. Since 1990 he has worked on digital imaging and vision. He is particularly interested in color, and in the appearance issues of digital imaging and lighting. He is one of the founders of the Italian Color Group and member of several program committees of conferences related to color and digital imaging.

This is a preprint of
Proc. IS&T Color Imaging Conference, Albuquerque, November, 2013

Chromaticity limits in color constancy calculations

John McCann, McCann Imaging, Belmont, USA;
Vassilios Vonikakis, Advanced Digital Sciences Center, Singapore;
Cristian Bonanomi, and Alessandro Rizzi, University of Milan, Italy