Optimal Color Spaces for Balancing Digital Color Images

Geoffrey Woolfe, Kevin Spaulding, and Edward Giorgianni
Eastman Kodak Company
Rochester, New York

Abstract

Unlike the human visual system, image-capture sensors lack the ability to adapt to the ambient exposure conditions. Accordingly, post-capture adjustments are required to compensate for variations in the exposure level (brightness) and capture illuminant chromaticity (white balance). These adjustments are typically made by scaling linear RGB sensor signals or by applying additive adjustments to logarithmic encodings of these signals. Even if such adjustments result in perfect balance for a given reference color in the scene (such as a reference neutral gray card), some degree of color error will typically be introduced into other scene colors. Such errors can be particularly problematic if they result in significant hue errors or if the color name is changed. The magnitude and direction of color errors introduced by balance adjustments will depend on the color space in which the balance adjustments are applied. For additive RGB color spaces, the fundamental attribute affecting these color errors is the chromaticities of the color space primaries. This paper develops a method for characterizing such color errors, and evaluates the performance of a number of commonly used color spaces. A method to determine the optimal color space primaries for which the application of balance adjustments produces minimal color errors is also presented.

Introduction

Image-capture sensors including photographic film and solid-state sensors differ fundamentally from the human visual system in that they lack the ability to adapt to ambient illumination. While the human visual system adapts to both the luminance and chromatic content of the scene illuminant [1], image-capture systems require a number of compensation mechanisms to mimic such adaptation. These compensation mechanisms can be broadly divided into two classes—pre-exposure mechanisms and post-exposure mechanisms. Some degree of pre-exposure compensation to the average illumination level is accomplished using exposure controls—aperture and shutter speed—available on most cameras [2]. These controls provide a basic level of adjustment over the exposure given to the sensor(s), but they are seldom reliable enough to produce perfectly exposed pictures in every instance. Furthermore, camera exposure controls provide no mechanism for compensation for variations in the chromaticity of the ambient illumination. The only pre-exposure mechanism that can be used to correct for the illuminant chromaticity is the use of color-correction filters. Again, such corrections are generally approximate because color correction filters are designed for specific discrete illuminant spectral power distributions. Pre-exposure compensation mechanisms only provide a partial solution to the problem of illuminant variability; therefore, most images require further compensation for the scene illuminant variability using post-capture mechanisms. There are numerous techniques available to accomplish such compensation. In the traditional silver-halide-negative photography arena, post-exposure illuminant compensation has been accomplished by adjustment of the enlarger or printer lamphouse filtration and the printing exposure time. The magnitude of such adjustments can be determined by trial and error or by algorithmic corrections based on statistical measurements of the negative transmittance.

Digital imaging opens up many additional possibilities for post-exposure illuminant compensation. Such compensation is generally performed using mathematical transforms, either as a white-balance adjustment in a digital camera, an adjustment to the scanned densities of a color negative in a digital photofinishing system, or by color and lightness adjustments made with desktop photo-editing software. These transforms are generally applied to a numerical encoding of the captured colors. Many such encodings—both device-dependent and device-independent—are available to represent colors.

The intent of balancing transforms in digital imaging applications is usually to modify the image in a manner that is consistent with the adaptation mechanisms of the human visual system. Typically, simple transforms, consisting of either a multiplicative scaling factor in a linear color encoding or an additive shift in a logarithmic color encoding, are used to map a particular visually neutral reference in the scene to a set of aim coordinates in the color encoding. The same transform that is used to map the reference neutral is also applied to all other colors in the image. (Because most scenes do not contain a reference neutral patch that can be used for the balancing operation, a neutral reference is usually estimated by analysis of the image content.) Even if the balancing transform produces a perfectly balanced result for the reference neutral, the results obtained for other image colors may not be consistent with human visual system adaptation and would, therefore, be perceived as an error. The extent of these errors is a function of the magnitude of the balancing correction (and therefore a function of the scene illuminant) and of the color encoding in which the balancing transform is applied. This paper will investigate the characteristics and extent of such color errors for a range of typical scene illuminants and a selection of color spaces. It will also propose a method to determine an optimal set of primaries that define a color space in which color errors are minimized when balancing transforms are applied.

Results and Discussion

Scene Capture

In order to determine the color errors that result from a balancing operation, we need to define a set of color aims that the balancing operation is designed to achieve. To assist in developing an appropriate set of aims, it is instructive to consider the process of photographing a scene using a digital camera. One can consider the photographic process in very broad terms as a two-stage process involving capture of the scene colors, followed by rendering of the scene colors, to produce the most pleasing picture. It is reasonable to assume that the function of the scene-capture stage is to accurately record the color of objects in the scene. Unfortunately, capture sensors are generally not colorimetric devices. Consequently, it is usually not possible to capture perfectly accurate scene colorimetry, let alone accurate color appearance. Furthermore, as discussed above, capture sensors have a fixed response to light and do not adapt to the scene viewing environment. Therefore, even if perfectly accurate scene colorimetry could be obtained, the color appearance of the scene to a human observer would not be captured and could not be determined without a detailed knowledge of the state of visual adaptation of the observer. The former problem can only be addressed by optimal design of the capture-sensor spectral sensitivities, while the latter problem is the subject of this work. The analysis presented here will be performed assuming colorimetrically accurate capture sensors in order to enable the study of color balancing errors without the complication of accounting for capture-color errors. However, the analysis methods are easily generalized to incorporate the characteristics of non-colorimetric capture sensors.

Capture Sensors

Colorimetrically accurate sensors have spectral sensitivities that can be represented as exact linear combinations of the CIE standard color matching functions. For simplicity, let us assume that we have a set of RGB sensors having spectral sensitivities that match the CIE 1931 Standard 2° Colorimetric Observer color -matching functions (i.e., the red channel spectral sensitivity matches the $\bar{x}(\lambda)$ color-matching function, the green channel matches $\bar{y}(\lambda)$, and the blue channel matches $\bar{z}(\lambda)$). Let us further assume that the sensor is designed for a CIE standard D50 capture illuminant. That is, the sensor will return RGB values of [1, 1, 1] when a perfect diffuse reflector is captured under the CIE Standard D50 illuminant. (This effectively defines a scaling of the spectral sensitivity functions.)

This sensor represents a colorimetric device, therefore, there is an exact and simple transform to the CIE *XYZ* colorimetry of scene colors:

$$\begin{bmatrix} X \\ Y \\ Z \end{bmatrix} = \begin{bmatrix} 96.4210 & 0 & 0 \\ 0 & 100.0000 & 0 \\ 0 & 0 & 82.5319 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad (1)$$

In order to make use of this accurate scene colorimetry to produce correctly balanced, pleasing pictures, we need some knowledge of the adapting illuminant under which the scene is viewed. There are a number of ways this can be achieved, ranging from the inclusion of a standard reflector often called a “gray card” in the scene, to a number of sophisticated image analysis algorithms designed to determine the appropriate color balance of the scene. Regardless of how the scene balancing adjustments are determined, the goal of the balancing adjustment is to correct the RGB values of a perfect diffuse reflector in the scene to values of [1, 1, 1]. Because our objective is to investigate the effect of applying scene-balancing corrections in different color spaces, we will perform this study with a set of patch reflectances of real scene colors that includes a perfect, nonselective diffuse reflector. Therefore, the balance adjustments are simply defined by the transformation necessary to adjust the captured values for this patch to achieve the aim values.

Test Colors

The test colors used in this study comprised a set of 379 real object reflectance spectra, providing a reasonably uniform sampling of the real-world surface colors. The set included a representative number of human skin tones and other memory colors (e.g., blue sky and plant foliage). Eleven spectrally nonselective neutral test colors were also added to this color set. These neutrals included a perfect diffuse reflector for use in determining the necessary balance correction for each capture illuminant.

Capture Illuminants

In order to accurately assess the overall performance of a color space with respect to color balancing errors, a representative set of capture illuminants should be used. This set should include illuminants typically encountered in photographic applications. The capture illuminants used in this study included CIE standard daylight illuminants ranging in correlated color temperature from 4,000K to 18,000K, standard blackbody (Planckian) radiators with correlated color temperatures from 2,000K to 4,500K; and standard fluorescent illuminants F2–F12.

The RGB responses of the sensor with spectral sensitivity $S(\lambda)$, to any given scene object reflectance, $P(\lambda)$, under a given scene illuminant, $I(\lambda)$, can be calculated according to Eq. (2) where $D_{50}(\lambda)$ represents the spectral power distribution of the CIE standard D50 illuminant.

$$R = \int S_R(\lambda)P(\lambda)I(\lambda)d\lambda \Big/ \int S_R(\lambda)D_{50}(\lambda)d\lambda$$
$$G = \int S_G(\lambda)P(\lambda)I(\lambda)d\lambda \Big/ \int S_G(\lambda)D_{50}(\lambda)d\lambda \qquad (2)$$
$$R = \int S_B(\lambda)P(\lambda)I(\lambda)d\lambda \Big/ \int S_B(\lambda)D_{50}(\lambda)d\lambda$$

Balance Adjustment

Balancing Color Spaces

Because the captured scene colors are colorimetrically accurate, the captured sensor RGB values can be transformed into any other additive RGB color space using a simple linear transform in the form of a 3 x 3 matrix, without any loss of accuracy. This 3 x 3 matrix can be calculated from the chromaticities of the color space primaries and the white point. In this case, a D50 white point was used so that the *XYZ* tristimulus values of CIE standard D50 map to RGB values of [1, 1, 1]. The balancing correction can be applied with equal validity, but not necessarily equal performance, in any additive RGB color space. (Note that the balancing correction is applied in the linear RGB domain, therefore, any nonlinear encoding functions associated with the color space have no impact on the results presented here.) In this study, we have evaluated the color primaries associated with the *sRGB* [3] and *ROMM RGB* [4,5] color encodings. These results were compared to balancing directly using the *CIE XYZ* primaries. In addition, the color balancing error-analysis technique described here was used to determine a set of optimized color balancing primaries producing minimum color errors for a wide range of commonly encountered capture illuminants. The color balancing errors for these optimized primaries are also reported.

Balancing Correction

The balance correction applied to each color in the test set is the correction required to bring the RGB values of a spectrally nonselective perfect diffuse reflector to aim values of [1, 1, 1]. This correction will be applied in the linear RGB domain as a multiplicative scaling according to Eq. (3):

$$X_{bal} = \frac{X_{in}}{X_{PDR}} \qquad (3)$$

where X is either R, G, or B, X_{in} represents the input color in the balancing color space, X_{PDR} represents the captured color of the perfect diffuse reflector under the capture illuminant in the balancing color space, and X_{BAL} represents the balanced color value of the test color.

Calculation of Balancing Color Errors

Aims

In order to determine the relative performance of various color spaces when they are used for the application of scene balance corrections, we require a set of color balance aims for the color patches in our test set. The issue of aims for a color balancing operation requires some discussion. Perhaps the most obvious set of balancing aims for a set of captured scene colors is the set of color values that would have been obtained had the scene been captured under the design illuminant for the sensor. With this philosophy, the color of the balanced image would be independent of the illuminant under which it was captured. In the case of the example presented here, this would mean that regardless of the illuminant under which the scene was captured, the balancing operation should restore the captured colorimetry to the colorimetry that would have been obtained with a CIE standard D50 illuminant. This set of aims is referred to as "Aim Set 1" in the tables of results below.

Good arguments can also be made for alternate sets of aims. For example, another reasonable set of aims would be to match the color appearance of the scene as it was originally observed. Even assuming complete chromatic adaptation, the color appearance of the scene to a real human observer will depend somewhat on the actual scene illuminant. Therefore, according to this philosophy, different aims would be required for each capture illuminant. Accordingly, an appropriate set of balancing aims would be the corresponding colors under the sensor-design illuminant that would have the identical color appearance to the scene colors viewed under the capture illuminant. Such corresponding colors can be calculated using one of the many chromatic adaptation transforms that are available in the literature. In this study, we will also consider a second set of aims, where the CMCCAT2000 chromatic adaptation transform (adopted for the CIECAM02 color appearance model) will be used to generate the corresponding colors. This set of aims is referred to as "Aim Set 2" in the tables of results below.

Error Calculations

The comparison of the balanced test patches against the aim colors can be performed using one of several color difference metrics. For this study, the color differences were determined by converting both the aims and the balanced colors to the *CIELAB* color space, and calculating the *CIELAB* ΔE^*_{ab} between corresponding patches. Summary statistics, calculated from the full patch set are presented below. Slightly different results would be obtained using different color difference metrics, but the general analysis and optimization techniques would still be valid.

A small number of the colors in the test set lie outside the color gamut that can be represented using the *sRGB* primaries. For these out-of-gamut colors, color space conversion problems can result if the linear RGB values are less than zero. In these cases, negative RGB values were clipped to zero. This serves to slightly modify the error statistics for the *sRGB* primaries but, generally, the impact is relatively small. Linear RGB values greater than one do not result in conversion problems and are, therefore, not clipped.

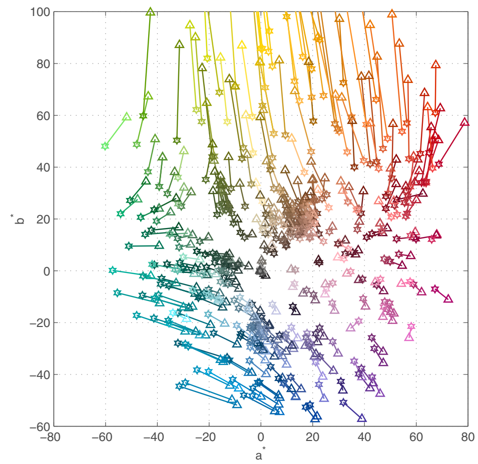

(a)

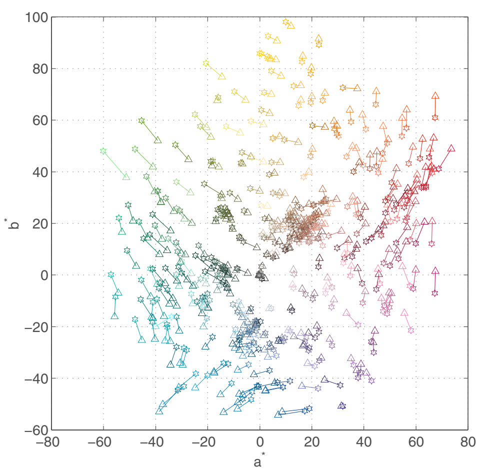

(b)

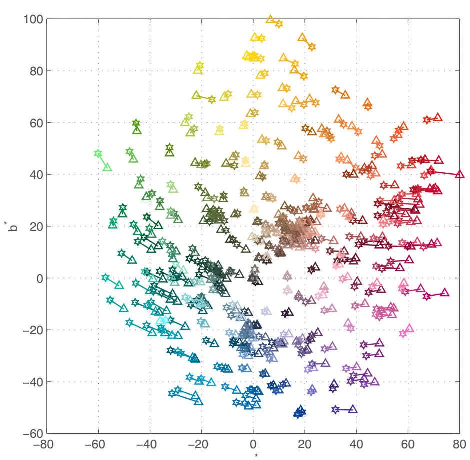

(c)

Figure 1. Balancing color errors for capture with 3,000K blackbody illuminant relative to Aim Set 1 for (a) sRGB primaries, (b) ROMM RGB and (c) optimized primaries.

Results

Capture Under a 3000K Blackbody Illuminant

Consider the case of our set of 379 scene test colors being captured under a 3000K blackbody (Planckian) radiator. This illuminant is representative of incandescent bulbs typically found in homes. The *xyz* chromaticity coordinates of this illuminant are [0.4369, 0.4041, 0.1590]. The color balance errors statistics resulting from balance corrections applied in the candidate color spaces are summarized in Table 1 below. The color errors resulting from balancing images captured under this illuminant are clearly quite dependent on the choice of balancing color space primaries. Both the *ROMM RGB* and the optimized primaries result in considerably smaller errors than the other sets of primaries, with the *sRGB* primaries producing the largest errors. This conclusion holds, regardless of which set of aims is used. Figure 1 shows a comparison of the color errors between the *sRGB* primaries, *ROMM RGB* and the optimal primaries. The aim colors, represented in the figure using star symbols, correspond to Aim Set 1. The balanced colors are shown as triangles. It can be seen that the optimal primaries produce much smaller errors, particularly for the saturated colors. With the *sRGB* primaries, there are some significant hue rotations, as well as chroma errors, in many regions of color space.

Table 1: Balancing color error statistics for capture with 3,000K blackbody illuminant

Primaries	Aim Set 1		Aim Set 2	
	RMS ΔE^*_{ab}	Max ΔE^*_{ab}	RMS ΔE^*_{ab}	Max ΔE^*_{ab}
sRGB	17.63	69.22	19.08	75.74
ROMM RGB	4.74	12.20	3.57	11.39
XYZ	6.40	17.87	5.38	19.27
Optimal	3.73	11.44	3.47	11.14

Capture Under 15,000K Daylight Illuminant

As a second example, we will consider the case of capture under a CIE standard daylight illuminant with a 15,000K color temperature. This illuminant was chosen as being representative of daylight from the sky with the sun occluded. (Although light from a clear north sky can reach correlated color temperatures as high as 40,000K.) The captured colors were color balanced using the same sets of additive RGB primaries, and the results are shown in Table 2. In this case, the optimal primaries clearly performed the best, with the *ROMM RGB* primaries being slightly better than the other two sets.

Table 2: Balancing color error statistics for capture with 15,000K daylight illuminant

Primaries	Aim Set 1		Aim Set 2	
	RMS ΔE^*_{ab}	Max ΔE^*_{ab}	RMS ΔE^*_{ab}	Max ΔE^*_{ab}
sRGB	10.69	53.97	11.71	50.68
ROMM RGB	8.66	19.88	6.34	20.34
XYZ	10.57	24.63	8.24	25.73
Optimal	3.80	13.23	3.35	8.78

Capture Under Fluorescent F11 Illuminant

The third example we will consider is for capture under a standard F11 fluorescent illuminant. This is a typical three-band fluorescent illuminant with *xyz* chromaticity coordinates of [0.3805, 0.3769, 0.2251] and a correlated color temperature of 4,000K. The color balance errors statistics resulting from balance corrections applied in the candidate color spaces are summarized in Table 3 below. These results show that applying balancing corrections using the *sRGB* primaries again results in larger color errors than in any of the other candidate color spaces. The optimized primaries deliver slightly worse performance than either *ROMM RGB* or *XYZ* primaries when evaluated against aim set 1, but perform slightly better than these other primaries when evaluated against aim set 2. (It is not surprising that the optimal primaries do not necessarily perform the best for all individual illuminants, since they represent a global optimum across a wide variety of illuminants.)

Table 3: Balancing color error statistics for capture with F11 fluorescent illuminant

Primaries	Aim Set 1		Aim Set 2	
	RMS ΔE^*_{ab}	Max ΔE^*_{ab}	RMS ΔE^*_{ab}	Max ΔE^*_{ab}
sRGB	8.97	38.65	5.28	32.54
ROMM RGB	5.39	19.03	1.66	4.60
XYZ	5.28	17.26	2.33	6.66
Optimal	6.16	23.95	1.33	4.25

Determination of Optimized Balancing Primaries

Given a metric for characterizing the color balance errors for a set of primaries, an optimal set of primaries producing the minimal color balance errors can be determined using conventional nonlinear optimization techniques. We performed such an optimization using the test colors and the RMS ΔE^*_{ab} error metric discussed above. The aims used in calculating the color balance errors were Aim Set 1. A Marquardt-Levenberg optimization technique was used to determine the optimal primaries for an ensemble of illuminants including 15 daylight illuminants (ranging from 4,000K–18,000K), 6 blackbody illuminants (ranging from 2,000K–4,500K), and 11 fluorescent illuminants (F2–F12). The resulting set of optimal balancing primaries is shown in Table 4.

Table 4. Optimal balancing primaries.

Primary	*x*	*y*
R	0.6625	0.3375
G	0.1806	0.7963
B	0.1391	0.0351

Conclusions

This work illustrates that the results obtained when color balancing a captured image are strongly dependent on the color space in which the balance correction is applied, as well as the characteristics of the capture illuminant. Evaluation of 3 representative capture illuminants using a sampling of 379 real scene reflectances showed that significantly lower errors could be obtained using the ROMM RGB primaries compared to the sRGB primaries. Balancing in the CIE XYZ color space gave only marginally larger errors than the ROMM RGB primaries. It was found that a color space based on an optimized set of primaries could provide significantly smaller color errors than any of these standard color spaces.

References

1. G. Wyszecki and W.S. Stiles, Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd Edition, John Wiley and Sons, 1982.
2. L. Stroebel, J. Compton, I. Current and R. Zakia, Photographic Materials and Processes, Focal Press, 1986.
3. "Multimedia Systems and Equipment – Part 2-1: Colour Measurement and Management – Default RGB Colour Space – sRGB," IEC 61966-2-1 (1999)
4. K.E. Spaulding, G.J. Woolfe and E.J. Giorgianni, Reference Input/Output Medium Metric RGB color encodings (RIMM/ROMM RGB), Prpc. PICS, pg. 155, (2000).
5. K.E. Spaulding, G.J. Woolfe and E.J. Giorgianni, Image States and Standard Color Encodings (RIMM/ROMM RGB), Proc. 8th Color Imaging Conference, IS&T, pg. 288, (2000).

Biography

Geoff Woolfe received his BS (with honors) and PhD in physical chemistry from the University of Melbourne (Australia) and his MS degree in imaging science from the Rochester Institute of Technology. He is currently a Senior Principal Research Scientist in the Imaging Science Division of the Eastman Kodak Research Laboratories.

His research interests include hardcopy/softcopy appearance matching of images, simulation and modeling of conventional and digital imaging systems, development of color imaging algorithms, preferred color image reproduction, color restoration of faded images, gamut mapping, computational color science and development of color control tools and color management systems.

Geoff Woolfe is the author of numerous scientific papers and patents in the fields of chemistry and color imaging.