

Optimizing Color Matching in a Lighting Reproduction System for Complex Subject and Illuminant Spectra

A. Wenger, T. Hawkins and P. Debevec

University of Southern California - Institute for Creative Technologies

Abstract

This paper presents a technique for improving color matching results in an LED-based lighting reproduction system for complex light source spectra. In our technique, we use measurements of the spectral response curve of the camera system as well as one or more spectral reflectance measurements of the illuminated object to optimize the color matching. We demonstrate our technique using two LED-based light sources: an off-the-shelf 3-channel RGB LED light source and a custom-built 9-channel multi-spectral LED light source. We use our technique to reproduce complex lighting spectra including both fluorescent and tungsten illumination using a Macbeth color checker chart and a human face as test subjects. We show that by using knowledge of the camera spectral response and/or the spectral reflectance of the subjects that we can significantly improve the accuracy of the color matching using either the 3-channel or the 9-channel light, achieving acceptable matches for the 3-channel source and very close matches for the multi-spectral 9-channel source.

Categories and Subject Descriptors (according to ACM CCS): I.3.3 [Computer Graphics]; I.3.7 [Computer Graphics]: Color, shading, shadowing and texture

1. Introduction

The process of lighting reproduction described in Debevec et al⁶ involves using computer-controlled light sources to illuminate a real-world subject as it would appear within a particular real-world environment. The light sources, aimed toward the subject from many directions, are driven to various intensities and colors to best approximate the illumination within the environment. One application for this technique is to realistically composite the subject into a scene, for example to composite an actor in a studio into a faraway location such as a cathedral. When the actor is illuminated by a close approximation of the lighting originally present in the cathedral, then such a composite believably shows the actor standing within the cathedral.

A noted challenge in lighting reproduction is that real-world illumination and subjects have complex spectral properties - lighting and reflectance are functions of wavelength across the visible spectrum, often with significant variation. In contrast, the light sources in lighting reproduction systems have so far used only three channels of illumination color - red, green, and blue - produced by appropriately

colored LEDs. Although RGB colors are commonly used throughout computer graphics, computing the color of light reflecting from a surface can only be performed accurately if the spectrum of the illuminant and the spectral reflectance of the surface are known, even when the end result is considered with respect to its RGB color. In lighting reproduction, the result is that it is not possible to accurately reproduce a subject's appearance under complex real-world illumination spectra such as tungsten and fluorescent lighting using just RGB lights.

In this paper, we introduce three techniques for producing better color matches for complex illumination spectra in a lighting reproduction system. First, we build and demonstrate an LED-based light source using 9 spectral channels rather than three. We show that by driving the 9 channels to optimally match the spectrum of the original illuminants (Spectral Illuminant Matching or SIM), we can produce improved color matches.

Our second technique, Metameric Illuminant Matching (MIM), leverages the fact that the usual goal of a lighting reproduction system is to reproduce the subject's appearance

under the illumination as seen by a particular RGB imaging system. We show that by measuring the camera system's spectral response curves, we can adjust the intensities of an RGB light's color channels to more faithfully reproduce the effects of illumination from complex spectra. Specifically, we drive the color channels of each light source to produce a metameric match with the target illuminant as seen by the camera system. We show that setting the light colors in this way produces better color matches for both the 3- and 9-channel lights than are obtained by matching the lights to the illuminant spectra directly.

Our third and final technique, Metameric Reflectance Matching (MRM), leverages the fact that some subjects will exhibit a limited set of different spectral reflectances. For example, the spectral reflectances across a person's face are largely similar. Knowing the dominant spectral reflectances of the subject in addition to the spectral response of the camera, we can drive the color channels of the light source so that the subject's reflection of the light is metameric to the subject's reflection of the original illuminant. We show that this technique produces the most accurate lighting reproduction results, and that for suitable subjects this technique allows the 3-channel RGB light to produce results which are nearly as accurate as the 9-channel light.

2. Background and Related Work

As mentioned in the introduction, the lighting reproduction approach described in Debevec et al ⁶ used a sphere of inward pointing RGB LED light sources to reproduce captured lighting environments. One of the problems identified in this work was the difficulty in achieving an accurate color match for complex input spectra.

Considering full spectral illumination beyond RGB has been addressed several times in computer graphics; Fairchild et al ¹ give an overview of spectral color imaging versus metameric color imaging. Borges ⁴ and later Hall ⁸ showed that for natural scenes with spectrally smooth illuminants the trichromatic approximation works well for the first reflection of light but degrades with multiple reflections. Peercy ¹⁵ efficiently performed full spectral rendering using a spectral basis for the scene based on principal component analysis of the illuminants and reflectances in the scene. Ward et al ¹⁶ used spectral prefiltering to efficiently approximate full spectral rendering using the Sharp RGB space. Drew and Finlayson ¹⁷ showed how to perform spectral sharpening on the spectral response curves of the camera sensor with positivity constraints. The benefit of such a transformation lies in improved performances of many computer vision and color image processing algorithms. Drew and Finlayson further develop this idea in ¹⁸ to perform multi-spectral calculations using principal component vectors that were "sharpened". The sharpening allows for a simple multiplication of the basis coefficients instead of a full linear transform to model illuminant changes.

Berns et al ³ propose a multi-spectral color reproduction paradigm using multi-spectral image capture and a spectral-based printing system. The authors state that the only way to assure a color match for all observers and across changes in illumination is to achieve a spectral match. This observation also applies to lighting reproduction.

Based on the observations made in this previous work, we investigate the color matching capabilities for lighting reproduction of an RGB light and of a custom-built 9-channel multi-spectral light, made from LEDs with different spectral outputs.

Matching colors for a given camera requires the knowledge of the camera's intensity response and spectral sensitivity. The recovery of the camera's intensity response curve is based on Debevec and Malik ⁷. To recover the spectral sensitivities Hardeberg et al ⁹ use a set of filters and Imai ¹¹ uses a monochromator or a set of interference filters to capture a series of photographs from which they reconstruct the spectral response curve of the camera. Similarly, we employ a series of color glass longpass filters to determine the camera response curve. As in ⁹ we also use the principal eigenvector method to invert the spectral system.

3. The Spectral Illumination and Camera Model

In order to generate a color match for a specific observer with reproduced light we need to define a model describing how the observer senses light and how the light is reproduced. The spectral camera model we use is a general spectral model for image acquisition systems, similar to the model used by Hardeberg et al ⁹. The components of the spectral camera and light model are shown in Figure 1.

The parameters of the spectral camera model include the spectral power distribution of the light source denoted by $I(\lambda)$, the spectral reflectance of the observed object $r(\lambda)$, the spectral properties of the optics $o(\lambda)$, the spectral transmittance of the k th filter $\phi_k(\lambda)$ and the spectral sensitivity of the imaging sensor $s(\lambda)$. Furthermore we model the nonlinear response of the imaging sensor of the k th channel with Γ_k^c . Linear pixel values c_k can be obtained by applying the inverse function of Γ_k^c to the nonlinear pixel values \check{c}_k . Observed pixel values are then determined by the following equation:

$$\check{c}_k = \Gamma_k^c \left(t_{img} \cdot \int_{\lambda_{min}}^{\lambda_{max}} I(\lambda) \cdot r(\lambda) \cdot o(\lambda) \cdot \phi_k(\lambda) \cdot s(\lambda) d\lambda + \epsilon_k \right)$$

$$c_k = \Gamma_k^{c-1}(\check{c}_k)$$

The above spectral model further models camera noise ϵ_k of the k th channel as additive noise. Smaller pixel values c_k are relatively much more affected by noise than larger pixel values. For an in-depth discussion on noise in multi-spectral color image acquisition we refer the interested reader to Burns ¹⁰. The integration time (i.e. shutter speed) t_{img} is the

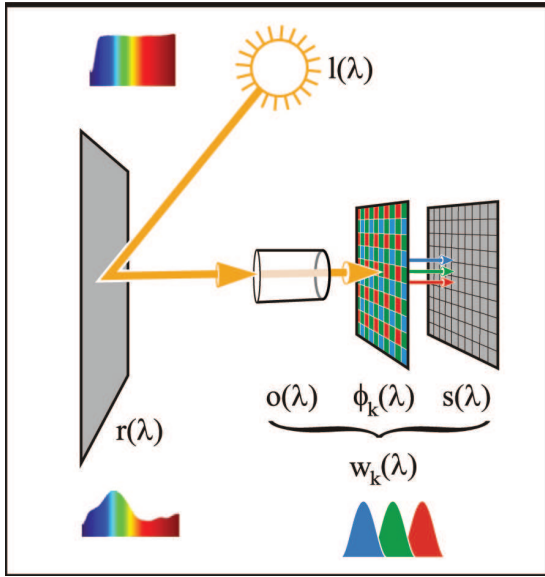


Figure 1: Spectral Camera Model. The spectral camera model describes how light is recorded by the image acquisition system. The main parameters in the model are the spectral power distribution of a light source $I(\lambda)$, the spectral reflectance of an observed material $r(\lambda)$, the spectral properties of the optics $o(\lambda)$, the spectral transmittance of the filters $\phi_k(\lambda)$ and the spectral sensitivity of the sensor $s(\lambda)$. The camera system is described by $w_k(\lambda) = o(\lambda) \cdot \phi_k(\lambda) \cdot s(\lambda)$ which includes the optics, the filters and the sensor.

only addition to Hardeberg’s ⁹ spectral camera model. The shutter speed becomes relevant in Section 5 to compensate for brightness disparities.

The spectral model for the reproduction light source consists of a small number of positive valued functions $b_j(\lambda)$. These are the spectral power distributions of the individual channels of our physical lights, shown in Figures 6 and 8. The final light output is a weighted sum $\sum_j \Gamma_j^{l-1}(p_j) \cdot \mathbf{b}_j$ of those functions, where the weights can only take on non-negative values ($p_j \geq 0$). The nonnegativity constraint on the weights represents the inability of the lights to produce negative light. The function Γ_j^l models the nonlinearity in the light output of the j th channel of the reproduction light sources. Figures 7 and 9 show the nonlinear behavior of the RGB and 9-channel light sources respectively.

If we know the observer’s spectral response curve, we only need to generate light that produces a metamer for the observer; there is no need to actually reproduce the spectrum of either the incident or reflected light. In Section 5.3, we show how we can use the dominant reflectances of the subject to derive the intensity for the individual channels of the light source to produce such a metameric color match.

4. Equipment and Calibration

In our work we use two main pieces of equipment, the camera system and the light sources. In this section we describe this equipment and the procedures we employed for calibrating the equipment’s illumination and sensing characteristics.

4.1. Camera System

The camera system used in our experiments is a Canon EOS D60 digital SLR camera with an 85mm Canon EF lens. The images are shot in RAW format in manual mode at ISO 100, with an aperture of f/4. Shutter speed is varied to produce properly exposed images. The 12 bit per channel RAW files are converted to floating point images using a raw image converter modified from ¹². The conversion process takes the exposure time t_{intg} , the nonlinearity of the sensor Γ_k^{c-1} , and the thermal noise ϵ_k into account to produce radiometric images from the camera data.

Our lighting reproduction system requires the camera to be radiometrically calibrated for its intensity response function Γ_k^c and aggregate spectral response function w_k as in Figure 1. The next two sections describe the intensity and spectral response curve recovery.

4.1.1. Intensity Response Curve

The intensity response curve of the k th channel Γ_k^c shows how the camera responds to different light intensity levels. We recover Γ_k^{c-1} using a series of differently exposed photographs at $\frac{1}{3}$ stop increments of a constant target. Graphing the resulting pixel values against exposure time produced Γ_k^{c-1} which specifies how to map pixel values to linear light levels. Figure 2 shows the recovered inverse intensity response curves for the red, green and blue channels of the Canon EOS D60. The curves are close to linear up to about 80% of the maximum pixel value, at which point nonlinearities due to sensor saturation become evident.




Figure 2: Inverse intensity response curves Γ_k^{c-1} of a Canon EOS D60. The inverse intensity response curves for the red, green and blue channels are close to linear until the upper 20 percent of the pixel value range.

4.1.2. Spectral Response Curve

The spectral response curves $w_k(\lambda) = o(\lambda) \cdot \phi_k(\lambda) \cdot s(\lambda)$ describe the sensitivity of the camera channels to light of different wavelengths. We assume $w_k(\lambda)$ is constant across the image sensor. We recover the spectral response curve by taking a series of photographs with 20 different glass filters in front of the lens. Figure 3 shows the spectral transmissivity of the different filters. To recover the spectral response curve




Figure 3: Color glass filter spectra. The spectra of the 19 Schott color glass longpass filters and the IR cutoff filter from Edmund Optics used in the spectral response curve reconstruction.

we invert the following system:

$$\check{c}_{i,k} = \Gamma_k^c \left(t_{intg} \cdot \int_{\lambda_{min}}^{\lambda_{max}} l(\lambda) \cdot r(\lambda) \cdot f_i(\lambda) \cdot w_k(\lambda) d\lambda + \epsilon_k \right)$$

For discretely sampled spectra this can be written in matrix notation:

$$\Gamma_k^{c-1}(\check{c}_k) - \epsilon_k = A \cdot w_k$$

The Matrix A holds in its rows the transmittance spectra of the filters f_i modulated by the light source spectrum l and the reflectance spectrum r . Due to the linear dependence in the filter transmittance spectra and the presence of noise in the acquired photographs the inversion of the above system is nontrivial. As Hardeberg⁹, we use the principal eigenvector method to invert the system. The spectral response curves w_k recovered for the Canon EOS D60 in Figure 4 were obtained by using 6 principle eigenvectors.

4.2. Light Sources

The two light sources we use in our lighting reproduction experiments are based on computer controllable LEDs. The first light source we use is an off the shelf 3-channel RGB LED light source, whereas the second light is a custom-built 9-channel light source.

4.2.1. 3- channel RGB Light

The 3-channel RGB light source is a Color Kinetics ColorBlast 6 (see Figure 5 on the left) driven by a Color Kinetics PDS-150e power/data supply. As for the camera system,




Figure 4: Estimated spectral response curves of the Canon EOS D60 camera. The spectral response curves were recovered with a principal eigenvector method using 6 principle eigenvectors.

we calibrated the spectral power distribution and light output curve of the light source. We measured the spectrum of each of the LEDs with a Photo Research PR-650 spectroradiometer. The spectral power distribution of the individual channels is shown in Figure 6. Note the wide gap between




Figure 5: Reproduction light sources. On the left is the 3-channel RGB light source. The spectral power distributions of the individual channels of the RGB light can be seen in Figure 6. On the right is the 9-channel light source. The spectral power distribution of its individual channels can be seen in Figure 8.

the red and the green channel—the light source generates very little light in the yellow part of the spectrum between 560nm and 600nm. The gap between the green and blue channels is far less pronounced.

We recovered the light output curve Γ_j^l of the light's j th channel by sending increasing values to the light source controller and measuring the light's intensity output with the PR-650. Figure 7 shows the measured light output curves for each of the channels, which allow us to compensate for the light's nonlinear intensity output behavior.

4.2.2. 9-Channel Multi-Spectral Light

The 9-channel multi-spectral light (Figure 5 on the right) is a custom built source based on three ColorBlast 6 light




Figure 6: Spectral power distribution of the 3-channel RGB light source. The red, green and blue spectral power distributions leave a significant gap between green and red, where there is no light output.




Figure 7: Light output curves of the 3-channel light source. All the light output curves exhibit a very similar non-linear behavior.

sources. We replaced the original ColorBlast 6 LEDs with a wider range of LED colors to obtain finer control over the spectral output of the light. We used white, royal blue, blue, cyan, green, amber, red-orange and red Luxeon Star/O emitters from Lumileds. The three ColorBlast 6 light sources provide 9 channels for the 8 differently colored LEDs. As there are only eight differently colored LEDs available we put white LEDs in two of the channels and placed gel filters in front of the LED's optics. One white channel is covered with Lee filter #101 and the other channel is covered with Lee filter #104. The two slightly distinct yellow filters help fill the gap near 560nm for which no super bright LEDs are readily available. Figure 8 shows the spectral power distributions of the 9 color channels.

As for the 3-channel light source, we also measured the light output curves Γ_j^L for the 9-channel light, the results of which are shown in Figure 9. Again all the 9 channels exhibit a nonlinear behavior. The one curve that deviates significantly from the other curves belongs to the amber color channel; we are not clear why this channel is so different from the others.

5. Color Matching Methods

With our equipment calibrated, we were able to design techniques for optimally driving the LED light sources so they most closely achieve the desired lighting reproduction effects. In the next three sections we present three different methods for determining light channel intensities to match the effect that the target illuminants would have on the subject.

Each of our three color matching methods determine light source channel intensities which optimally meet particular criteria, such as that the spectrum of the reproduced light optimally matches the spectrum of the target illuminant in a least squares sense. Because we cannot drive the light sources with negative values, we cannot determine the light source channel intensities using linear system techniques. Consequently, we use a conjugate gradient optimization method¹⁴ that was modified to enforce positivity in the light source color channels to determine which channel intensities optimize the color matching criteria. Furthermore,




Figure 8: Spectral power distribution of the 9-channel multi-spectral light source. All the spectral power distributions are relatively narrow except for the two yellowish channels that are based on filtered white LEDs which have a broader peak.




Figure 9: Light output curves for the channels of the 9-channel multi-spectral light source. All of the channels exhibit a similar nonlinear behavior except for the amber channel (dashed line).

when the target illuminant (such as a halogen bulb) is much brighter than what the LED light sources can generate, the techniques match the spectral shape of the curve up to a scale factor instead of in an absolute sense. In this paper, we compensate for this scale factor when needed by exposing the image using a proportionally longer shutter speed t_{intg} . The output values of the following three color matching methods are linear light channel intensities, which we map to the appropriate light control values using the measured light output curves Γ_k^l .

5.1. Spectral Illuminant Matching (SIM)

Our first approach is based on the fact that if we can spectrally match the target illuminant we are guaranteed that any possible reflectance will look correct for any observer. This approach is attractive since it is not dependent on the spectral response characteristics of the camera system or the subject. The only information we need is the spectral power distribution of the target illuminant and the properties of the reproduction light source. The problem of finding the optimal reproduction parameters \mathbf{p} given a specific target illuminant spectrum \mathbf{l} can be set up as a minimization of the sum of the square residuals of the reproduction light spectra \mathbf{b}_j to the target illuminant spectrum \mathbf{l} :

$$\min \sum_i \left(\sum_j p_j b_{j,i} - l_i \right)^2 \quad \left| \quad p_j \geq 0 \quad \forall j \right.$$

Where j is the index over the color channels of the reproduction light and i is the index over the spectral samples. Determining optimal parameters \mathbf{p} for the above system with the constrained optimization solver yields results such as the dotted curves in Figure 12 for matching tungsten and fluorescent illuminant spectra.

As the curves show, with a limited number of channels we cannot generate a very close spectral match to the target lighting. This mismatch can lead to errors in the lighting reproduction process, and motivates our second technique.

5.2. Metameric Illuminant Matching (MIM)

Our second method leverages knowledge of the spectral response curves of the camera system to improve lighting reproduction given that it is not possible to closely match the spectrum of the target illuminant. The idea is to match the output of the reproduction light $\sum_j p_j \mathbf{b}_j$ metamERICALLY to the target illuminant \mathbf{l} with respect to the camera's spectral response curves \mathbf{w}_k . With the above considerations the problem can be set up as a minimization of the sum of the square residuals of the reproduction light color channels observed by the camera system $\sum_i w_{k,i} \sum_j p_j b_{j,i}$ to the target illuminant observed by the camera system $\sum_i w_{k,i} l_i$:

$$\min \sum_k \left(\sum_i w_{k,i} \sum_j p_j b_{j,i} - \sum_i l_i w_{k,i} \right)^2 \quad \left| \quad p_j \geq 0 \quad \forall j \right.$$

Where i and j are indices over the same domain as before and k denotes the index over the color channels of the camera system. The dashed curves in Figure 12 show results employing this method to achieve a color match for the tungsten and fluorescent light sources with respect to the Canon EOS D60's response curves. While the spectra do not match any more closely, the appearance of the original and reproduced illuminants to the camera system are matched as closely as possible.

5.3. Metameric Reflectance Matching (MRM)

Our final method improves on the lighting reproduction quality of the previous method by additionally taking into account the spectral reflectances of the subject. By measuring key spectral reflectances \mathbf{r}_n and using those as a part of the optimization we can specifically aim to match the appearance of the subject under the reproduced illumination to its appearance under the target illumination, again with respect to the spectral responses of the given camera system \mathbf{w}_k . The minimization is set up as the sum of square relative differences between the key spectra \mathbf{r}_n illuminated with the target light spectrum \mathbf{l} observed by the camera system described by \mathbf{w}_k , and the key spectra \mathbf{r}_n illuminated with the reproduction light spectra $\sum_j p_j \mathbf{b}_j$ observed by \mathbf{w}_k :

$$\min \sum_n \sum_k \left(\frac{\sum_i r_{n,i} w_{k,i} \sum_j p_j b_{j,i} - \sum_i r_{n,i} l_i w_{k,i}}{\sum_i r_{n,i} l_i w_{k,i}} \right)^2 \quad \left| \quad p_j \geq 0 \quad \forall j \right.$$

Where i , j and k are indices over the same domain as in the above methods and n denotes the index over the number of measured key reflectances. The solid curves in Figure 12 are the result of optimizing the illuminated appearance of all the color swatches on the Macbeth Color Checker chart for tungsten and fluorescent illumination.

6. Results and Discussion

We demonstrate the color matching capabilities of the three color matching methods with a Macbeth Color Checker chart and the reproduction of a mixed lighting environment illuminating a person's face. For these tests, we use two different light sources: a fluorescent and a tungsten light. Figure 10 shows the spectral power distributions for these two light sources.

6.1. The Macbeth Color Checker Chart

With the Macbeth Color Checker chart we evaluated the color matching capabilities of both light sources using all three color matching methods. Figure 11 shows the spectral reflectance of the color swatches on a Macbeth Color Checker chart. The experiment is set up in a dark room with the camera system perpendicular to the Macbeth Color Checker chart at a distance such that the Macbeth chart fills the camera frame. The light source is positioned near the




Figure 10: Target illuminant spectra. Spectra of the tungsten (smooth) and fluorescent (spiky) light sources used as our target illuminants.




Figure 11: Macbeth Color Checker chart. Macbeth Color Checker chart with the spectral reflectance of each color swatch.

camera. First the Macbeth chart is lit by each of the two target illuminants. Then we switch to the reproduction light sources which we drive with the calculated light control values. In total we take 14 photographs, one with each of the target illuminants, six with the 3-channel light and six with the 9-channel light. The six photographs for the two reproduction lights consist of three images with the different color matching methods each reproducing the tungsten or fluorescent target illuminant.

Figure 12 shows the spectral power distributions calculated with the three color matching methods for each of the target illuminants.

For each experimental condition, we analyze the corresponding photograph to extract representative pixel values for each patch of the Macbeth chart. In addition, we can use our estimates of the illuminant, reflectance, and camera characteristics to compute theoretical predictions for these pixel values.

Figure 14 shows the theoretical (computed) and actual (photographed) colors for the Macbeth chart under each

source illuminant and for each of the color matching methods with both the RGB and 9-channel light sources. For each test condition, the left side of the color swatch is the computed or photographed color with the target illuminant, and the right side of the swatch is the computed or photographed color with the specified reproduction light source and color matching method.

The theoretical results can be used to judge the error inherent in using a particular reproduction light source and a particular color matching method, while the experimental results include these errors as well as additional experimental errors. We discuss some of the potential sources of experimental error in Section 6.3. Our theoretical and experimental results are largely similar, with the theoretical error typically representing about half the total experimental error. The qualitative conclusions drawn from both sets of results are similar.

A first observation is that in general the 9-channel light performs better than the RGB light, which is expected since it provides more degrees of freedom to achieve the color matching. The second observation is that the SIM method yields the poorest results, particularly for the RGB light. An explanation for this result can be found in Figure 12. Looking at the spectral power distribution produced by the spectral match method with respect to the spectral power distribution of the target illuminant, we can see a significant difference in light output particularly in the red region of the spectrum for the RGB light which leads to a blueish-green tint. We find that the MIM and MRM methods provide improved matches for both reproduction lights.

To numerically evaluate the color matching performance of the different methods and reproduction light sources, we calculated the relative error $\epsilon_{i,k}$ for each of the color swatches i in color channel k of the photograph lit by the target illuminant l and lit by the reproduction light r . The error is calculated as follows:

$$\epsilon_{i,k} = (\alpha \cdot c_{i,k}^r - c_{i,k}^l) / c_{i,k}^l$$

Where $c_{i,k}^r$ is the pixel value of the k th color channel in the i th color swatch for the photograph taken with the reproduction light r and $c_{i,k}^l$ with the target illuminant l . The normalization factor α is determined for each test condition by averaging the pixel values over all color swatches with both the target illuminant and the reproduction light. The scale factor α is then simply the ratio between the average calculated from the target illuminant photograph and the average calculated from the reproduction light photograph. It should be noted that a single α is computed combining information from all three color channels, so this is strictly a brightness scaling and not a color correction. Table 1 holds the root mean square error (RMS) for each target illuminant reproduced by each reproduction light with each color matching method over all the color swatches and over the three color channels.

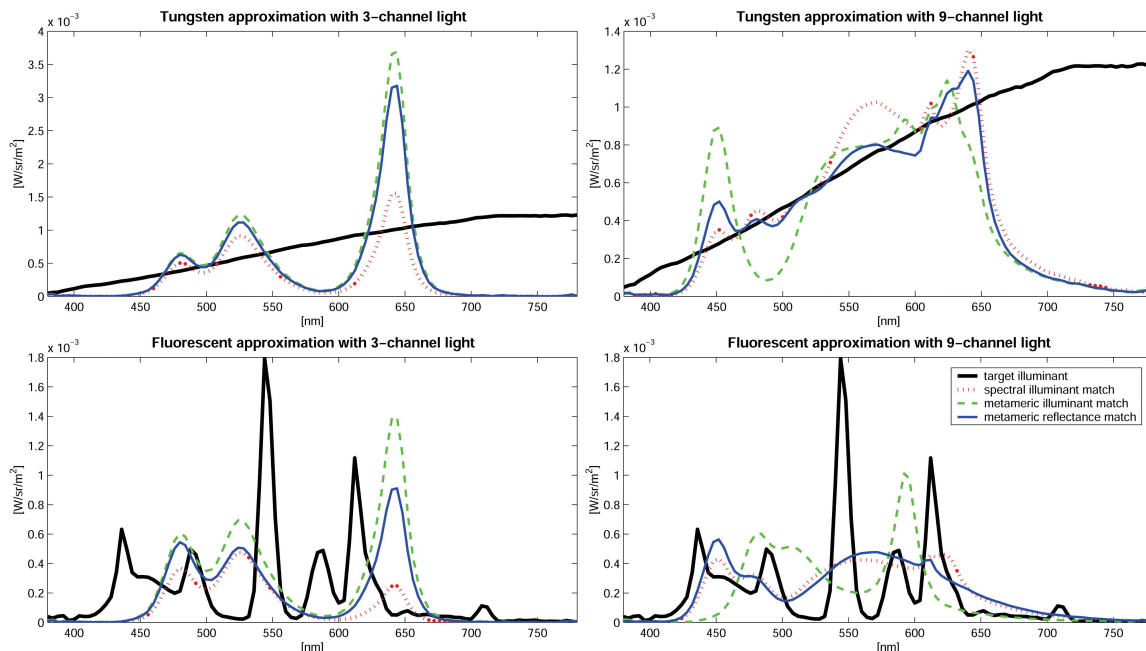


Figure 12: Results from the three different optimization methods. The black heavy curve is the spectral power distribution of the target illuminant. The dotted curve is the approximation using the spectral illuminant matching method, the dashed curve shows the result of the metamerism illuminant matching method, and the solid curve is the result of the metamerism reflectance matching method.

The data in Table 1 confirm that having more control over the spectral output of the reproduction light (as with the 9-channel light) produces closer color matches. In the theoretical case, the errors for the 9-channel light are in the neighborhood of 1-2%, making them almost imperceptible.

Furthermore, the data shows that using additional knowledge of the camera spectral response improves the color matching, particularly for the 3-channel light. The error for the 3-channel light drops from 7.5% and 9.5% for tungsten and fluorescent with the SIM method to about 5% with the MIM and MRM methods.

For the RGB light source the MRM method performs no better than the MIM method. For the Macbeth chart this is expected, since the three channels do not provide enough degrees of freedom to match to the wide variety of reflectance samples of the chart. For the 9-channel light the metamerism reflectance matching outperforms the metamerism illumination matching in the theoretical results, but this advantage is not apparent in the experimental results. This is unsurprising given that the apparent experimental error is considerably larger than the predicted reduction in error.

6.2. Mixed Lighting Environment for a Face

Our second experiment matches to a person's face lit from the left with tungsten light and from the right with fluores-

cent light as in Figure 15(d). This experiment tests how well we can color match skin for two very different light sources in a single photograph. As we only had a single RGB and 9-channel light source available, we acquired two separate images with the light source in each position and added them together to produce the effect of having two light sources on at once⁵ (this required the test subject to sit still while we moved the light). For the metamerism reflectance method we measured 4 different spots on our test subject's face from the forehead, cheek, lips and chin. These four spectral skin reflectances are shown in Figure 13. It is noteworthy that the four facial spectral reflectances are very similar.

Figure 15(d) shows the test subject lit by the original lighting setup with tungsten light to the left and fluorescent light to the right. In the top row (a)-(c) the illumination is generated by the 3-channel RGB light and in the bottom row (e)-(g) the light is generated by the 9-channel light.

The 9-channel light clearly performs better than the 3-channel light over all the color matching methods. For the two SIM methods (a) and (e) we can see a large discrepancy in performance. This observation does not come as a surprise because the 9-channel light provides much more variability for spectrally matching illuminants.

The color matching results improve for the 3-channel light as more information is taken into account. There is a sig-

Tungsten illuminant approximation with 3-channel light								
	red		green		blue		sum	
method	theory	experiment	theory	experiment	theory	experiment	theory	experiment
SIM	0.0689	0.1178	0.0262	0.0356	0.0400	0.0426	0.0484	0.0752
MIM	0.0334	0.0691	0.0236	0.0351	0.0291	0.0400	0.0290	0.0503
MRM	0.0326	0.0760	0.0233	0.0348	0.0292	0.0402	0.0286	0.0535

Tungsten illuminant approximation with 9-channel light								
	red		green		blue		sum	
method	theory	experiment	theory	experiment	theory	experiment	theory	experiment
SIM	0.0068	0.0366	0.0158	0.0384	0.0244	0.0428	0.0173	0.0394
MIM	0.0056	0.0359	0.0065	0.0351	0.0176	0.0402	0.0113	0.0371
MRM	0.0024	0.0367	0.0023	0.0352	0.0044	0.0370	0.0032	0.0363

Fluorescent illuminant approximation with 3-channel light								
	red		green		blue		sum	
method	theory	experiment	theory	experiment	theory	experiment	theory	experiment
SIM	0.1171	0.1526	0.0291	0.0320	0.0527	0.0544	0.0760	0.0953
MIM	0.0407	0.0628	0.0272	0.0285	0.0413	0.0504	0.0370	0.0493
MRM	0.0425	0.0790	0.0283	0.0282	0.0433	0.0503	0.0387	0.0565

Fluorescent illuminant approximation with 9-channel light								
	red		green		blue		sum	
method	theory	experiment	theory	experiment	theory	experiment	theory	experiment
SIM	0.0215	0.0376	0.0144	0.0320	0.0143	0.0371	0.0171	0.0357
MIM	0.0204	0.0451	0.0150	0.0366	0.0349	0.0520	0.0249	0.0450
MRM	0.0133	0.0364	0.0142	0.0328	0.0058	0.0323	0.0117	0.0339

Table 1: Theoretical and experimental RMS errors for the Macbeth Color Checker chart experiment. The four tables hold the theoretical and experimental performances for each of the color matching methods.

nificant improvement from the very poorly performing SIM method (a) to the MIM method (b) and a noticeable improvement from the MIM method to the MRM method (c). For the 9-channel light it is harder to make out a clear order of the performances. The SIM method (e) seems to perform slightly worse than the MIM and MRM method (f) and (g); the image has a noticeable greenish tint. The MIM and MRM method (f) and (g) perform very similarly.

It is at first surprising that the 3-channel light performs almost as well for the face as the 9-channel light in the MIM and particularly in the MRM method. This can be explained by the fact that the four measured spectra of the test subject are very similar and we essentially are metamERICALLY matching a single spectrum. As the camera's three channels spectrally correspond well to the 3-channel RGB light's three channels (see Figures 7 and 4) the three RGB light channels provide the necessary degrees of freedom to metamERICALLY

match the reflectance of the skin of the test subject. However, the 9-channel light better matches other reflectance spectra present in the images—for example, the subject's blue shirt and jade necklace are more closely matched by the 9-channel light than by the 3-channel light regardless of the color matching method.

6.3. Potential Sources of Experimental Error

As our methods of color matching are predictive methods the accuracy of the color match is highly dependent on the quality of the recovered characteristics of the equipment used. We have seen in Figure 14 and Table 1 that our experimental results do not agree perfectly with the theoretical predictions. The experimental results reflect roughly an additional 2-3% error over the theoretical predictions. This error is likely due to the accumulation of small errors in the




Figure 13: Skin reflectance spectra considered for the metameric reflectance method. The four spectra are spectral reflectance measurements of skin at the forehead, the cheek, the lips and the chin of the test subject.

estimation of the camera and LED intensity responses and spectral responses.

Using our approach of intensity response curve recovery we rely on the shutter speed information stored in the raw photograph. We do not know how accurately this information reflects the actual shutter speed of the camera which could lead to a miscalibration of the intensity response of the camera system. Incorrectly calibrating the intensity response also would impact the spectral response recovery, as the pixel values for this process are assumed to be linear.

The recovery method for the camera spectral response uses 19 longpass filters and an IR cutoff filter. The actual information for the recovery process lies in the differences between the photographs taken with the different filters and the differences in the spectral transmittances of the filters. This can lead to very small pixel value differences which are significantly affected by camera noise making the recovery of the spectral response difficult.

Some of the error in the experimental results is likely due to the experimental setup itself, and in particular to the assumption that the LED light sources produce a completely even field of illumination over the subject.

7. Future Work

The experiments performed in this paper show that taking spectral information into account for color matching in lighting reproduction is a promising step towards producing well-matched color composites using either three-channel or more complex light sources. Based on the results presented in this paper, there are several improvements that could be made to our lighting reproduction system.

Because of the dependence of our calculations on the accuracy of our measured intensity response and spectral response curves, it would be desirable to either find an extremely accurate method for measuring these curves or to

devise a method of performing an overall system calibration that does not depend on every component of the system being perfectly calibrated. The metameric color matching methods in particular would benefit from a more accurate spectral response curve for the camera.

We suspect that spatial variation in the intensity of the various LED's for the reproduction lights may be responsible for much of the difference between our theoretical and experimental results. Characterizing this variation or removing it by adding additional diffusers to the LED's might improve our experimental results.

Avenues for the future include incorporating the color matching methods and multi-spectral light source into a lighting reproduction systems such as the Light Stage device proposed in ⁶. This will require the acquisition of multi-spectral lighting environments to drive the individual light sources. Designing a multi-spectral light probe acquisition device would be part of this process.

Another interesting future task involves investigating post-processing of the image color channels to improve the color match. This would reduce the dependency on the calibration accuracy of the equipment, which could potentially yield more accurate color matches.

It would further be interesting to try different light sources in the reproduction process, such as using filtered incandescents instead of LEDs, using a designated light source such as an HMI light to model very bright sources like sunlight, or using video projectors to reproduce spatially varying illumination. Any of these would involve applying our color matching methods to new target illuminants and new reproduction lights. We note that the additional light output from a designated sun reproduction light would be very helpful for a larger scale Light Stage device as it would greatly reduce the dynamic range the LED or filtered incandescent light sources have to reproduce.

8. Conclusion

In this paper we have presented three techniques for improving the color matching of spectrally complex illuminants incident upon spectrally complex surfaces using LED light sources with limited numbers of color channels. Our results have shown that by taking both the camera spectral response and the subject's spectral reflectance into account, we are able to achieve color matches that are reasonably accurate even for three-channel RGB light sources. Our results also clearly show that having more control over the reproduction light spectrum using the 9-channel light yields much better results regardless of the employed color matching method. The results of this work bode well for applying lighting reproduction systems in domains where precise color matching is an important design component.

Acknowledgements

We gratefully acknowledge Marc Brownlow and Brian Emerson for their graphic design work on the figures, and Andrew Gardner, Gina Justice, and Andrea Wenger (seen in Fig. 15) for sitting as illumination subjects. We thank Kevin Dowling at Color Kinetics, Inc. for helping facilitate this research and Greg Ward for helpful spectral rendering discussions, as well as Lora Chen for production support and Laurie Swanson for coordination help. We further thank Richard Lindheim, Neil Sullivan, James Blake, and Mike Macedonia for their support of this project. This work has been sponsored by the University of Southern California Office of the Provost and U.S. Army contract number DAAD19-99-D-0046; the content of this information does not necessarily reflect the position or policy of the sponsors and no official endorsement should be inferred.

References

1. Mark D. Fairchild, Mitchell R. Rosen, and Garrett M. Johnson, "Spectral and Metameric Color Imaging", *Technical Report, Munsell Color Science Laboratory*, August 2001.
2. G. Wyszecki and W. S. Stiles, "Color Science: Concepts and Methods, Quantitative Data and Formulas", 2nd ed. Wiley, New York, 1982.
3. Roy S. Berns, Francisco H. Imai, Peter D. Burns, and Di-Y. Tzeng, "Multi-spectral-based Color Reproduction Research at the Munsell Color Science Laboratory", *Electronic Imaging*, pp. 14–25, 1998.
4. Carlos F. Borges, "Trichromatic approximation for computer graphics illumination models", *Computer Graphics (Proceedings of SIGGRAPH 91)* **25**(4), pp. 101–104, July 1991.
5. Paul Debevec and Tim Hawkins and Chris Tchou and Haarm-Pieter Duiker and Westley Sarokin and Mark Sagar, "Acquiring the Reflectance Field of a Human Face", *Proceedings of SIGGRAPH 2000*, pp. 145–156, July 2000.
6. Paul Debevec, Andreas Wenger, Chris Tchou, Andrew Gardner, Jamie Waese, and Tim Hawkins, "A Lighting Reproduction Approach to Live-Action Compositing", *ACM Transactions on Graphics (Proc. of SIGGRAPH '02)* **21**(3), pp. 547–556, 2002.
7. Paul E. Debevec and Jitendra Malik, "Recovering High Dynamic Range Radiance Maps from Photographs", *SIGGRAPH 97*, pp. 369–378, August 1997.
8. Roy Hall, "Comparing Spectral Color Computation Methods", *IEEE Computer Graphics and Applications* **19**(4), pp. 36–45, July/August 1999.
9. Jon Yngve Hardeberg, Hans Brettel, and Francis Schmitt, "Spectral characterisation of electronic cameras", *Electronic Imaging*, pp. 100–109, 1998.
10. P. D. Burns, "Analysis of image noise in multispectral color acquisition", PhD thesis, Center for Imaging Science, Rochester Institute of Technology, 1997.
11. Francisco H. Imai, "Multi-spectral Image Acquisition and Spectral Reconstruction using a Trichromatic Digital Camera System associated with absorption filters", *Munsell Color Science Laboratory, Rochester Institute of Technology*, August 1998.
12. Dave Coffin, "Raw Digital Photo Decoding in Linux", <http://www2.primushost.com/%7Edcoffin/powershot/>, 2003.
13. Chris Tchou and Paul Debevec, HDR Shop, Available at <http://www.debevec.org/HDRShop>, 2001.
14. William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery, "Numerical Recipes in C++: The Art of Scientific Computing", second edition, Cambridge University Press, 2002.
15. Mark S. Peercy, "Linear Color Representations for Full Spectral Rendering", *Computer Graphics*, **27**(3), pp. 191–198, 1993.
16. Greg Ward and Elena Eydelberg-Vilshin, "Picture Perfect RGB Rendering Using Spectral Prefiltering and Sharp Color Primaries", *Thirteenth Eurographics Workshop on Rendering*, 2002.
17. Mark S. Drew and Graham D. Finlayson, "Spectral sharpening with positivity", *Journal of the Optical Society of America A* **17**(8), pp.1361–1370, 2000.
18. Mark S. Drew and Graham D. Finlayson, "Multispectral Processing Without Spectra", *Journal of the Optical Society of America A*, 2003.




Figure 14: Theoretical and experimental reproduction comparison for the Macbeth Color Checker chart.




Figure 15: Mixed lighting environment reproduction comparison for a person's face. A person's face illuminated with tungsten light from the left and fluorescent light from the right. The center photograph is the reference photograph with the original lighting. The top row of photographs is shot with illumination reproduced with the RGB light source. The bottom row of photographs is shot with illumination from the 9-channel light source.