Contents:

Introduction

For our project, we will replicate the Stanford method for correcting the quantum efficiency (QE) of a sensor by using a LiveScript in MatLab. The foundation of our script will draw heavily upon the Figure_03.mlx file found in ISETCornellBox which loads and processes both the measured and predicted RGB data of a Google Pixel4a camera when images of the MacBeth ColorChecker (MCC) are captured under three different illuminants: tungsten (A), daylight (Day), and cool white fluorescent (CWF). Using this data, we will find a 3x3 matrix that transforms the predicted RGB values so that they match closely with the measured RGB values, correcting the sensor QE.

Once we successfully demonstrate how to correct the sensor QE, we will test different methods, such as linear, diagonal, and nonnegative linear regression, for finding the 3x3 transformation matrix and see how well each one works. Furthermore, we will find the confidence interval of the Stanford Method for correcting the QE of a sensor.

Finally, we will explore how using the data from only one type of light impacts the effectiveness of the transformation matrix in correcting the sensor QE. For each illuminant, we will find the transform using only the data gathered under that illuminant and then apply the transform to the full set of data. We will compare the three results to determine which illuminant produces the most effective transform.

Background

There are various camera properties that are used to evaluate image quality, and one of those properties is color. Color calibration creates a model of the spectral QE that relates spectral radiance to the digital camera values. By predicting the RGB values of each patch of the MCC under the three different illuminants and comparing them to measured RGB values, we can assess the accuracy of models used and test different methods to correct the predicted RGB values so that they better align with the measured data.

Methods

Overview

The paper *Validation of Physics-Based Image Systems Simulation With 3-D Scenes* by Zheng Lyu, Thomas Goossens, Brian A. Wandell, and Joyce Farrell, discusses how computer software can be used to predict elements of images captured by different cameras and how to correct simulated data to better align with measurement results. Section III, Part C1 discusses how the Sony IMX363 sensor found in the Google Pixel 4a camera was calibrated. Measurement data was collected by capturing images of an MCC under three different illuminants. The MCC consists of 24 patches, so when we record the RGB values for each patch, we obtain a 24x3 matrix for each illuminant. Accumulating all of the measurement data for all three illuminants together results in a 72x3 matrix. The three columns of this matrix are defined as (*r*', *g*', *b*').

Predicting RGB Values and Finding *M* to Correct the Sensor QE

The predicted RGB values are found by multiplying the sensor QE data with the spectral radiance. As seen in the original LiveScript, radiance is obtained by multiplying the data from the scene by each of the different light sources. The Pixel4a sensor QE data is already stored, so by multiplying it by the newly calculated radiance, we can obtain predicted RGB values. These predicted RGB values also make up a 72x3 matrix, and the three columns of this matrix are defined as (*r*, *g*, *b*).

In order to account for discrepancies between predicted values and measurement results, a 3x3 matrix *M* that transforms the spectral QE was found by taking the least square error between (*r*', *g*', *b*') and (*r*, *g*, *b*)

$$min\, M \;=\; ||[r'\;\; g'\;\; b'] - [r\;\; g\;\; b]M||^2$$

The diagonal entries of *M* indicate channel gain, and the off-diagonal values indicate crosstalk between channels. Multiplying the predicted RGB values with *M* results in the corrected sensor QE that better aligns with the measurement data.

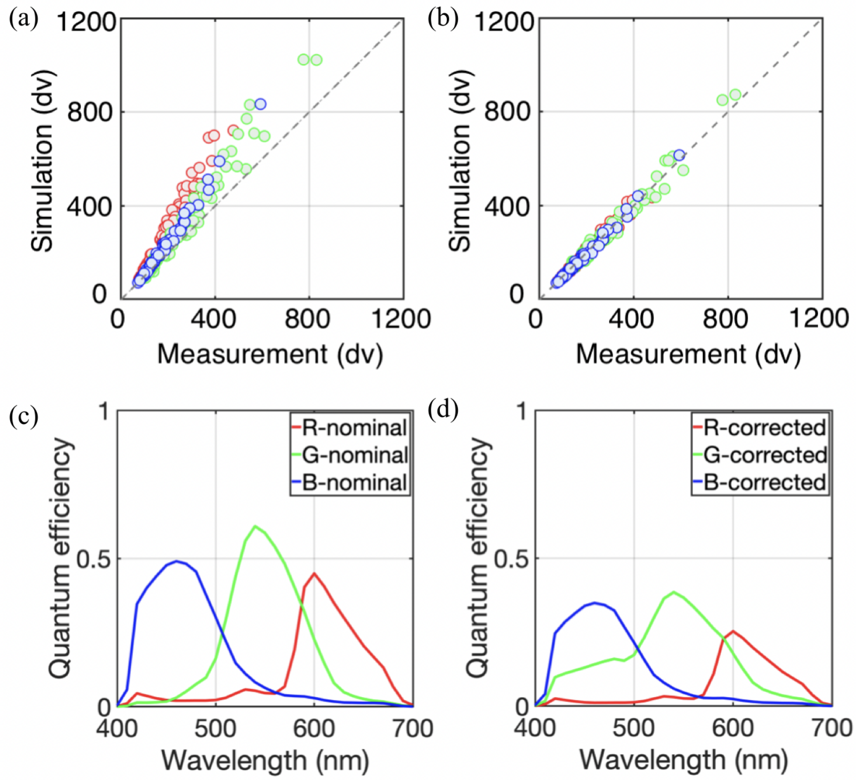

Figure 1. Sensor spectral QE calibration results from Lyu et al. (a) Scatter plot comparing measured and predicted RGB values of the MCC. (b) Scatter plot comparing measured and predicted RGB values after applying the transformation matrix *M* to correct the QE of the sensor. (c) Uncorrected color filter QE. (d) Corrected color filter QE.

The transformation matrix *M* found for correction is

$$M = \begin{bmatrix} 0.5636 & 0.0807 & 0.0069 \\ 0 & 0.5917 & 0 \\ 0 & 0.2470 & 0.7098 \end{bmatrix}$$

From this matrix, we see that channel gains are not the only elements contributing to the differences between predicted and measured RGB values, but optical crosstalk between channels is also a factor that needs to be taken into account.

Results

Replicating the Stanford Method for Correcting the QE of a Sensor

To replicate the Stanford Method for correcting the QE of a sensor, we imported the data collecting and processing sections of the Figure_03.mlx file into our own LiveScript. The first section reads in the Pixel4a data and obtains the measured RGB values of each of the 24 MCC patches under the three different illuminants by taking the mean values in the mid center of each patch. The next section reads in the data of the three illuminants and multiplies each of them by the scene being observed to find the radiance. As mentioned previously, the Pixel4a sensor QE data is already stored, so by multiplying it by the radiances found, we can obtain the predicted RGB values.

We store the measured RGB values in a 72x3 matrix called measRGB, and we store the predicted RGB values in another 72x3 matrix called predRGB. As seen in Figure 2, we then plot Predicted vs Measured data, and we obtain the same results as shown in Figure 1a.

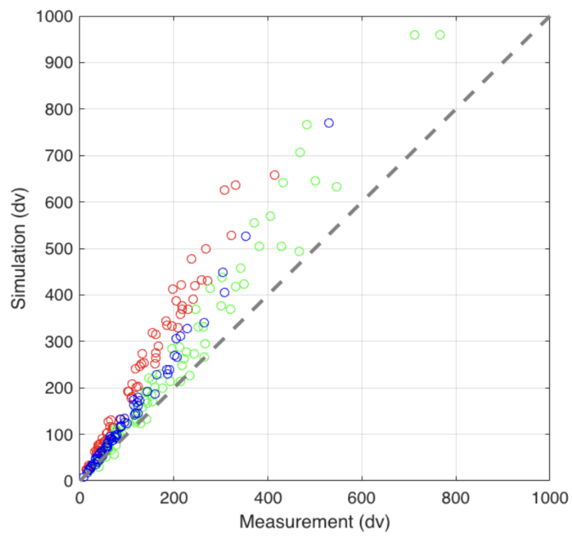

Figure 2. Scatter plot comparing measured and predicted RGB values of the MCC using our LiveScript.

From Figure 2, we see that there are discrepancies between predicted and measured results because the data points deviate significantly from the identity line which is represented by the dashed gray line. Therefore, we must find a transformation matrix *M* that puts the data on the identity line.

There are three methods we considered for finding *M*: The linear method, the diagonal method, and the nonnegative method. The linear method maps data in one space to data in another space, the diagonal method finds a diagonal matrix that only changes the gain of each channel, and the nonnegative matrix finds a matrix in which each element is a nonnegative number. The Lyu et al. paper uses the nonnegative approach.

The Linear Method

In our LiveScript, we drew upon the cbMccFit function found in ISETCornellBox to find *M* using these different methods. First, we applied the linear method and found *M* to be

$$M = \begin{bmatrix} 0.6002 & 0.0314 & 0.0104 \\ -0.0090 & 0.6359 & -0.0082 \\ -0.0543 & 0.2231 & 0.7219 \end{bmatrix}$$

Multiplying predicted RGB values matrix by *M*, we obtain a new 72x3 matrix consisting of corrected predicted RGB values. When we plot the Corrected Predicted vs Measured data, we obtain the results shown in Figure 3.

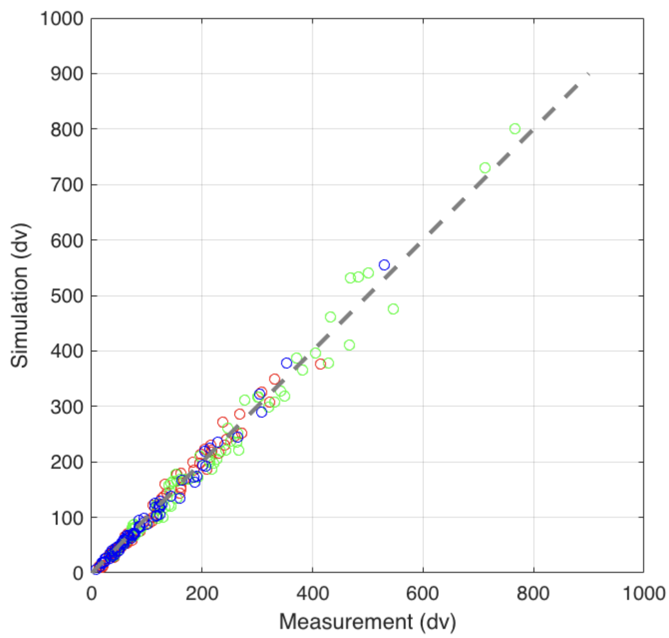

Figure 3. Scatter plot comparing measured and predicted RGB values after applying the transformation matrix *M* found using the linear method to correct the QE of the sensor using our LiveScript.

From Figure 3, we see the discrepancies between predicted and measured results seen in Figure 2 have been corrected, and the data points more closely align with the identity line now.

The last section of the LiveScript plots and compares both the uncorrected and corrected sensor QEs. We used the same section found in the Figure_03.mlx file to obtain the original filter data and plotted transmissibility vs wavelength for RGB as seen in Figure 4.

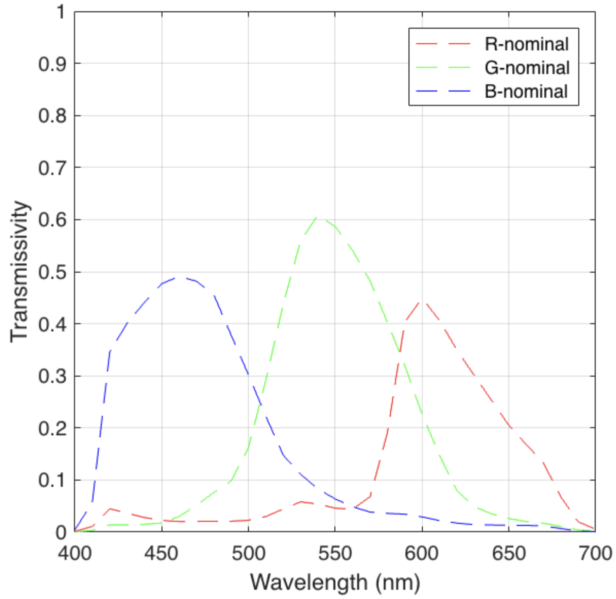

Figure 4. Uncorrected color filter QE using our LiveScript

Figure 4 produces the same plot we see in Figure 1c. By applying the transformation matrix *M* to the uncorrected color filter QE, we obtain the corrected sensor QE that can be seen in Figure 5.

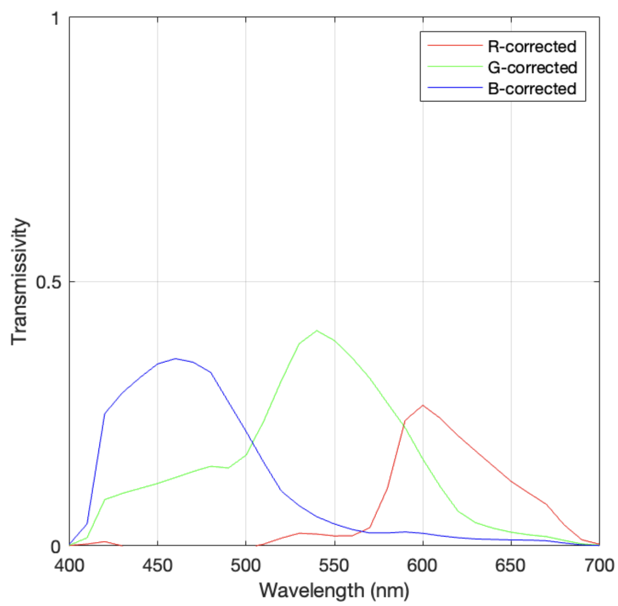

Figure 5. Corrected color filter QE using the linear method to find the transformation matrix *M* and our LiveScript.

The Diagonal Method

Next, we applied the diagonal method to obtain the transformation matrix and found *M* to be

$$M = \begin{bmatrix} 0.5613 & 0 & 0 \\ 0 & 0.7653 & 0 \\ 0 & 0 & 0.7205 \end{bmatrix}$$

Multiplying predicted RGB values matrix by *M*, we obtain a new 72x3 matrix consisting of corrected predicted RGB values. When we plot the Corrected Predicted vs Measured data, we obtain the results shown in Figure 6.

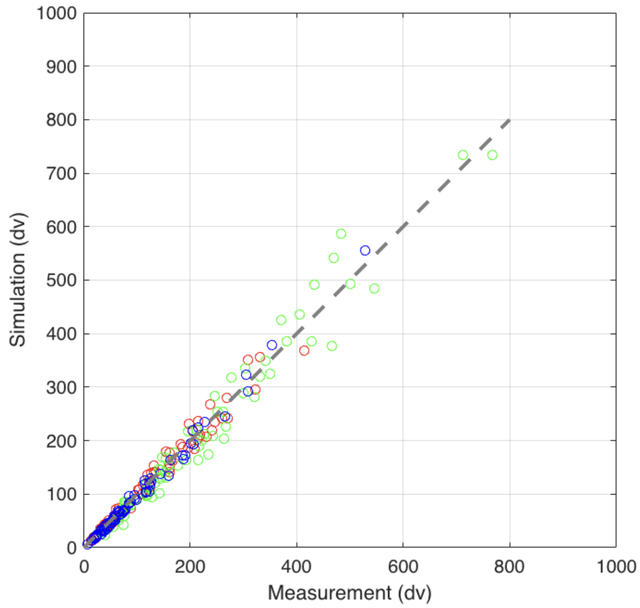

Figure 6. Scatter plot comparing measured and predicted RGB values after applying the transformation matrix *M* found using the linear method to correct the QE of the sensor using our LiveScript.

From Figure 6, we see the discrepancies between predicted and measured results seen in Figure 2 have been corrected, and the data points more closely align with the identity line now; however, for G data points around the 400-600 dv range, there are still some noticeable deviations from the identity line, especially when compared to the results from Figure 3. where the linear method was used to find *M*. This indicates that the linear method may be more effective than the diagonal method when it comes to finding an *M* that corrects the predicted RGB values.

By applying the transformation matrix *M* to the uncorrected color filter QE, we obtain the corrected sensor QE that can be seen in Figure 7.

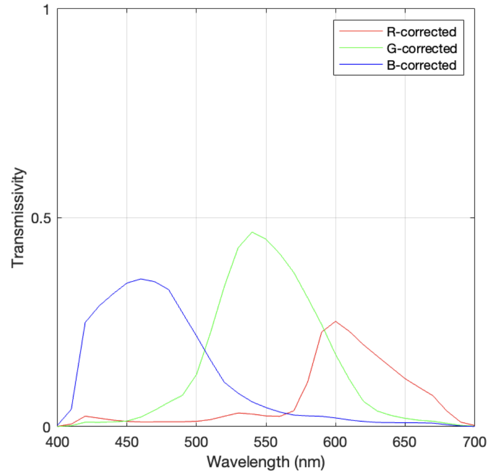

Figure 7. Corrected color filter QE using the diagonal method to find the transformation matrix *M* and our LiveScript.

The Nonnegative Method

Lastly, we applied the nonnegative method to obtain the transformation matrix and found *M* to be

$$M = \begin{bmatrix} 0.5613 & 0.0314 & 0.0038 \\ 0 & 0.6359 & 0 \\ 0 & 0.2231 & 0.7163 \end{bmatrix}$$

This transformation matrix is comparable to the one found in the Lyu et al. paper where the nonnegative method for finding *M* was also used. Multiplying predicted RGB values matrix by *M*, we obtain a new 72x3 matrix consisting of corrected predicted RGB values. When we plot the Corrected Predicted vs Measured data, we obtain the results shown in Figure 8.

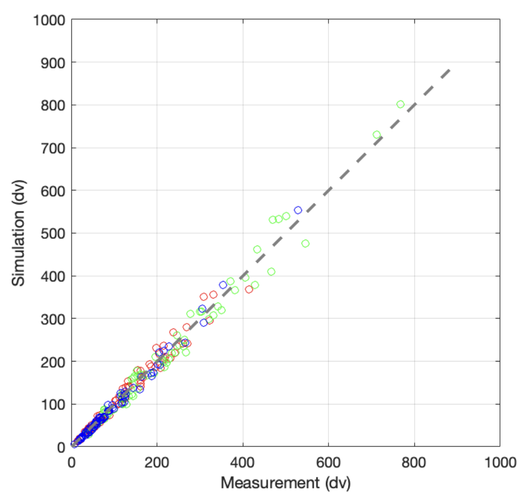

Figure 8. Scatter plot comparing measured and predicted RGB values after applying the transformation matrix *M* found using the nonnegative method to correct the QE of the sensor using our LiveScript.

From Figure 8, we see the discrepancies between predicted and measured results seen in Figure 2 have been corrected, and the data points more closely align with the identity line now. Additionally, the results are similar to those seen in Figure 3 where the linear method was used to find *M*. This indicates that the nonnegative method may be more effective than the diagonal method but equally as effective than the linear method when it comes to finding an *M* that corrects the predicted RGB values. It is, however, easier to interpret the corrections being made by *M* when the nonnegative method is used because differences in channel gain are clearly defined by the diagonal elements and the amount of crosstalk between channels can also be interpreted from any nonzero off-diagonal elements. As expected, the results seen in Figure 8 closely resemble the results found in the Lyu et al. paper as seen in Figure 1b where the nonnegative method for finding *M* was also used.

By applying the transformation matrix *M* to the uncorrected color filter QE, we obtain the corrected sensor QE that can be seen in Figure 9.

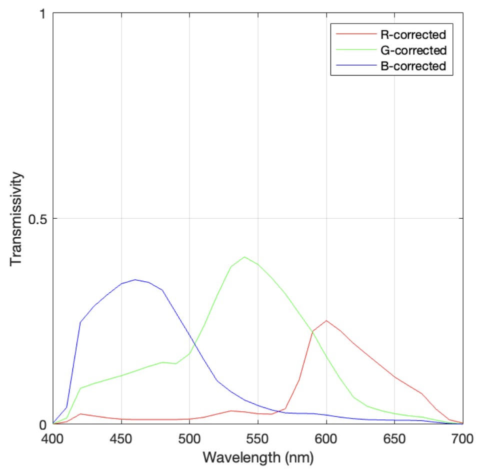

Figure 9. Corrected color filter QE using the nonnegative method to find the transformation matrix *M* and our LiveScript.

As expected, the results seen in Figure 9 closely resemble the results found in the Lyu et al. paper as seen in Figure 1d where the nonnegative method for finding *M* was also used.

Finding the Confidence Interval of the Stanford Method for Correcting the QE of a Sensor

To assess the reliability of the Stanford method for correcting the QE of a sensor, we repeated the method 100 times and plotted all outcomes of the corrected sensor QE on one graph. To obtain random predicted RGB values for each iteration, we generated three 24x3 matrices of random coefficients between 0.75 and 1.25 to randomly change the elements of the original predRGB by at most ±25%. The elements of these matrices are defined as rand1, rand2, and rand3 and are multiplied by the respective elements of the three 24x3 matrices of predicted RGB values obtained by using the three different illuminants. The resulting matrices are combined to create the 72x3 matrix defined as predRGB.

Any of the three methods for finding the transformation matrix *M* can be used, however, we decided to use the nonnegative method for all iterations. Once *M* is found for the given set of predicted RGB values, we apply *M* to the uncorrected color filter QE to obtain the corrected color filter QE before moving onto the next set of random predicted RGB values.

When we plot the corrected color filter QE for all 100 iterations, we obtain the results shown in Figure 10.

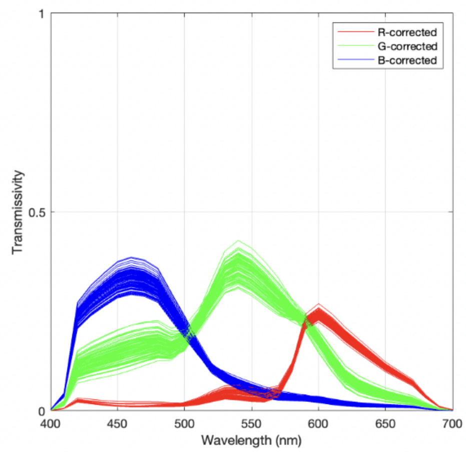

Figure 10. Corrected color filter QEs for 100 iterations using the nonnegative method to find the transformation matrix *M* and our LiveScript.

To determine the confidence interval of the Stanford Method, we looked where there appeared to be the largest spread of results. We concluded that this occurs at around 450nm for the green curve. Figure 11 shows approximately the highest and lowest levels of transmissivity for green light at 450nm. As seen in the figure, these values are about 0.21 and 0.09 respectively, suggesting a range of at most 0.12. Therefore, we can expect the Stanford Method to provide a result that falls within ±0.06 transmissivity of the corrected color filter QE for a given sensor.

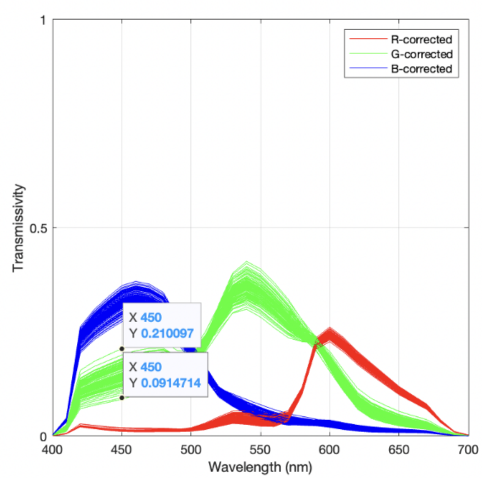

Figure 11. The approximate highest and approximate lowest levels of transmissivity for green light at 450nm.

Testing the Effectiveness of Using Only One Illuminant

To investigate the effectiveness of a transform found using data from only one illuminant, we created a second LiveScript. The first two sections of this LiveScript are the same as the LiveScript for part one of this project. The first section reads in the Pixel4a data and obtains the measured RGB values of each of the 24 MCC patches under the three different illuminants. The next section reads in the data of the three illuminants and multiplies each of them by the scene being observed to find the radiance. Multiplying the Pixel4a sensor QE data by the radiances found results in the predicted RGB values.

The next three sections of the LiveScript calculate the transformation matrix, M, using only data from illuminant A, Day and CWF, respectively. First, the data for the selected illuminant is isolated using indexing. Then, M is calculated using this data and the Stanford method previously described. The LiveScript allows M to be calculated using the linear, diagonal or nonnegative method. For simplicity, however, we will only include results obtained using the nonnegative method since this produced the most effective and easily understandable results in the first part of this project. M is multiplied by the full set of predRGB data and the corrected predicted values are plotted versus the measured values to assess the effectiveness of the transform. Finally, M is used to correct the sensor QE.

The final section plots the spectral power distribution (SPD) of each of the three illuminants. This is done to aid in the analysis of the results.

Illuminant A

First, we applied the nonnegative method to only the illuminant A data to obtain the transformation matrix and found *M* to be

$$M = \begin{bmatrix} 0.5853 & 0 & 0 \\ 0.0349 & 0.6641 & 0 \\ 0 & 0.4951 & 0.7909 \end{bmatrix}$$

Multiplying the full set of predicted RGB values by *M*, we obtained a new 72x3 matrix consisting of corrected predicted RGB values. We plotted the Corrected Predicted vs Measured data and obtained the results shown in Figure 12.

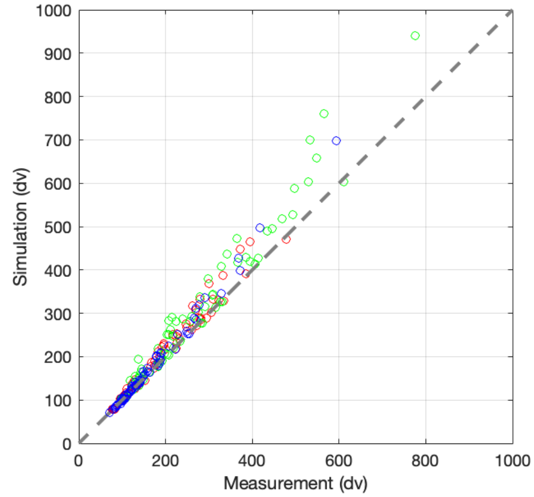

Figure 12. Scatter plot comparing measured and predicted RGB values after applying the transformation matrix M found using the nonnegative method and data from only the A illuminant

Figure 12 shows that the transformation matrix found using only illuminant A data corrects some of the discrepancies, but does not correct all of them. It performs noticeably worse than the transform incorporating all the data, especially in the 300-800 dv range. Based on the SPD of illuminant A, shown in Figure 13, we hypothesize that this transform is less effective at correcting the blue and green values because it has very low radiance in the blue and green parts of the spectrum.

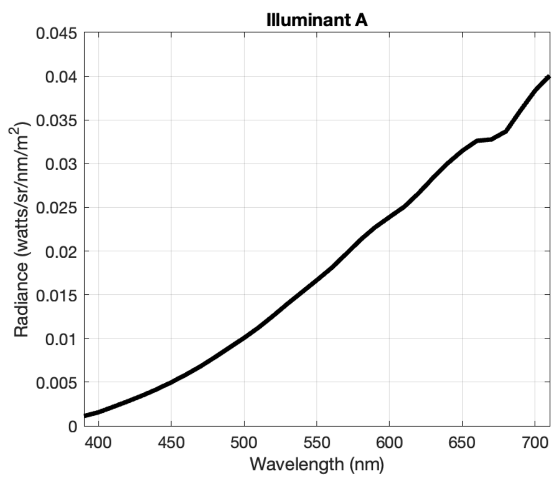

Figure 13. SPD of illuminant A

Illuminant Day

Next, we applied the nonnegative method to only the illuminant Day data to obtain the transformation matrix and found *M* to be

$$M = \begin{bmatrix} 0.5063 & 0 & 0 \\ 0 & 0.6384 & 0 \\ 0 & 0.2319 & 0.7115 \end{bmatrix}$$

Then, we multiplied the full set of predicted RGB values by *M* to obtain the corrected predicted RGB values. The plot of the Corrected Predicted vs Measured data is shown in Figure 14.

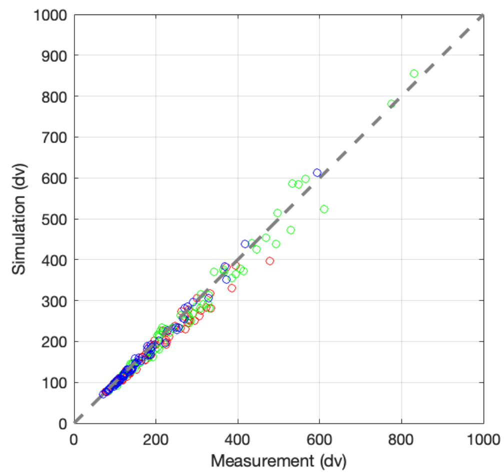

Figure 14. Scatter plot comparing measured and predicted RGB values after applying the transformation matrix M found using the nonnegative method and data from only the Day illuminant

Figure 14 shows that illuminant Day performs better than illuminant A. The SPD of illuminant Day, as shown in Figure 15, has a more even distribution and significantly high levels of radiance at lower wavelengths, when compared to illuminant A.

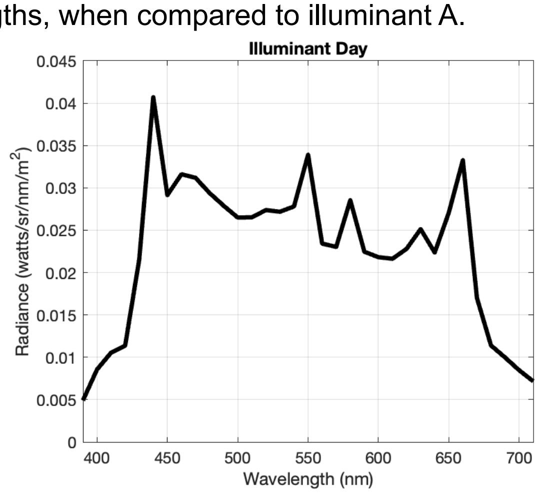

Figure 15. SPD of illuminant Day

Illuminant CWF

Finally, we applied the nonnegative method to only the illuminant CWF data to obtain the transformation matrix and found *M* to be

$$M = \begin{bmatrix} 0.5544 & 0 & 0 \\ 0 & 0.6384 & 0 \\ 0 & 0.2954 & 0.7267 \end{bmatrix}$$

This transform was multiplied by the full set of predicted RGB values. This was then plotted versus the measured values, as shown in Figure 16.

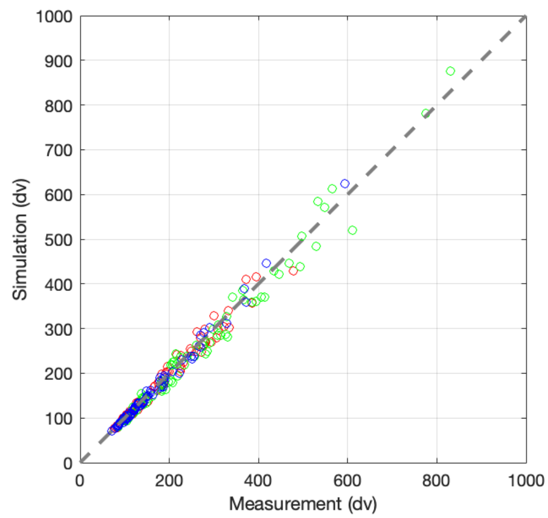

Figure 16. Scatter plot comparing measured and predicted RGB values after applying the transformation matrix M found using the nonnegative method and data from only the CWF illuminant

Figure 16 shows that the transform using just the illuminant CWF data produces very similar results to the illuminant Day transform, as well as the transform found using the full set of data. There are still some noticeable discrepancies in the 400-600 dv range, however, similar discrepancies are also present in Figures 8 and 14.

These results were somewhat surprising to us, given the SPD of illuminant CWF, as shown in Figure 17. The distribution of illuminant CWF is less even than that of illuminant Day and has relatively low radiance at higher wavelengths. After examining Figure 2, the plot of the uncorrected data, and all the transformation matrices we found throughout this project, we noticed that the gain of red values always needed to be decreased more than that of blue and green. We hypothesize that this could contribute to why illuminant CWF performed well, despite having low radiance in the part of the spectrum corresponding to red. It might be interesting for future projects to explore this more in depth.

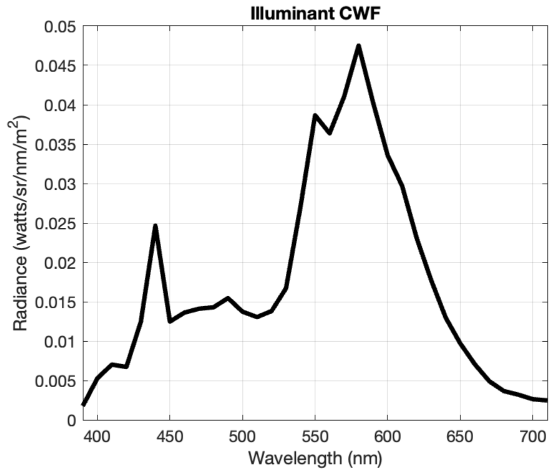

Figure 17. SPD of illuminant CWF

Conclusions

In this project, we replicated the Stanford Method for correcting the QE of a sensor by using predicted RGB, measured RGB, and filter data obtained from the Lyu et al. paper to find a transformation matrix *M* that corrects the sensor QE. We used three different methods to find *M*: The linear method, the diagonal method, and the nonnegative method. Based on our results, we found that the linear and nonnegative methods were the most effective in correcting the sensor QE because the results from both methods were not only the most similar to the results from the Lyu et al. paper, but also because the corrected data aligned closest with the identity line. Of the two methods, we concluded that the nonnegative method, which is also used in the Lyu et al. paper, is favorable because not only are all matrix elements positive, but any nonzero off-diagonal elements give us insight about crosstalk happening between channels. By repeating the Stanford Method for 100 different iterations of predRGB, we concluded that we can expect the method to provide a result that falls within ±0.06 transmissivity of the corrected color filter QE for a given sensor.

For the second part of this project, we used data from only one illuminant and the nonnegative method to find M. We then applied M to the full set of data and plotted the results. We repeated this for all three illuminants, A, Day and CWF. We found that illuminants Day and CWF produced the most effective transform. Additionally, we found that both were fairly similar to the M found using the nonnegative method and the full set of data. The performance of the CWF transform was somewhat surprising given its uneven SPD distribution.

References

"Color Calibration Project." *Image Systems Engineering @ Stanford University*, Confluence & Carta.

Lyu, Zheng, et al. "Validation of Physics-Based Image Systems Simulation with 3-D Scenes." *IEEE Sensors Journal*, vol. 22, no. 20, 15 Oct. 2022, pp. 19400–19410., https://doi.org/10.1109/jsen.2022.3199699.

Lyu and Wandell, ISETCornellBox, 2022, GitHub repository, https://github.com/iset/isetcornellbox

Notable Files:

Figure_03.zip (Foundation of our LiveScript for this project)

cbMccFit.zip (Script of different methods used to find transformation matrix *M*)

Tominaga, Shoji, et al. "Measurement and Estimation of Spectral Sensitivity Functions for Mobile Phone Cameras." *Sensors*, vol. 21, no. 15, 22 June 2021, p. 4985., https://doi.org/10.3390/s21154985.

Wandell, et al. ISETCAM, 2022, GitHub repository, https://github.com/iset/isetcam

Appendix

Presentation Slides: PSYCH 221 Final Presentation

FinalProjectPart1.mlx.zip

This zip file contains the script and data for the section under Results titled "Replicating the Stanford Method for Correcting the QE of a Sensor"

FinalProjectConfidenceInt.mlx.zip

This zip file contains the script and data for the section under Results titled "Finding the Confidence Interval of the Stanford Method for Correcting the QE of a Sensor"

FinalProjectPart2.mlx.zip

This zip file contains the script and data for the section under Results titled "Testing the Effectiveness of Using Only One Illuminant"