LED Multispectral Imaging: Reconstruction of Reflectance Spectra

March 21, 2008
Stanford University
Elza John, David Lau, Stephanie Leung

INTRODUCTION & MOTIVATION

When we take photos with digital cameras and view them on display devices (such as LCD monitors), we want it to be an accurate representation of the original scene. However, the camera only gives us a limited set of data that does not allow for a very accurate reconstruction of the image. We can see this by comparing the data that the camera captures with the true information that is available in a scene. A digital camera records only three values for each point (pixel) in the image: a red channel value, a blue channel value, and a green channel value. However, at every point there is actually a smooth curve of spectral reflectance spanning across wavelengths. Additionally, the information that a camera captures from a scene depends on the illuminant and the particular characteristics of a camera's sensors, among other factors. With only three numbers for each pixel, a digital camera does not accurately capture the original properties of the scene.

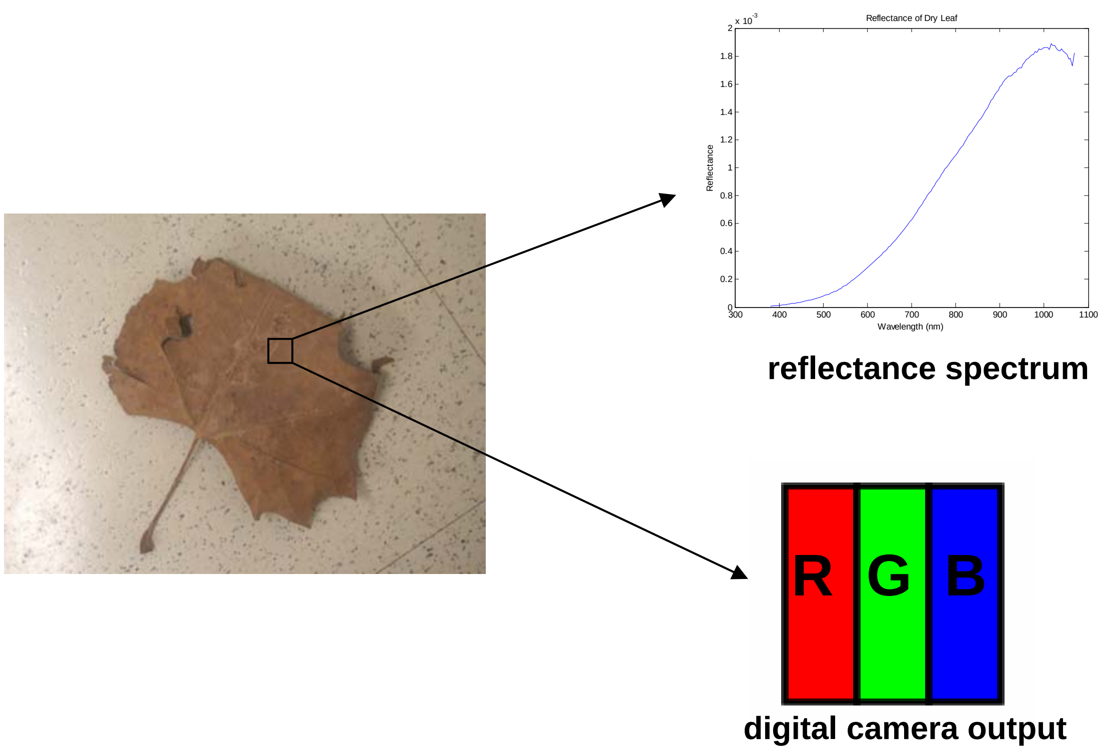

Figure 1: Comparison of digital camera output and reflectance spectrum.

Reconstructing the original reflectance is important because it allows us to accurately represent the colors and intensity of a scene on different displays, and also to accurately render the scene under different luminants, such as daylight and tungsten and halogen lights. One method for more accurately estimating the reflectance using a consumer end digital camera is through multispectral imaging, an approach that involves using illuminants of different wavelengths to capture a single scene, and then combining that information together to use for reconstructing the spectral reflectance. Now instead of having just three pieces of data for each pixel, we have 3 x the number of illuminants for each pixel.

Previously in this class, a group worked on a similar project on multispectral imaging. Their project, "Accurate Prediction of Scene Radiance Spectra," used two color filters, blue and yellow, and had no IR data. In contrast, our project uses an LED lighting system, which Max Klein built for his project in the class, which illuminates the scene with eight different LEDs, including the IR region. The benefit of using this LED system is that the switching of colors is very quick compared to the procedure when using color filters. This is most important for taking pictures of people, who must stay as still as possible during a sequence of pictures.

For our project, we used eight LEDs: UV, blue, cyan, green, amber, red, infrared centered at 880nm, and infrared centered at 940 nm. Below is a plot of the LED spectra:

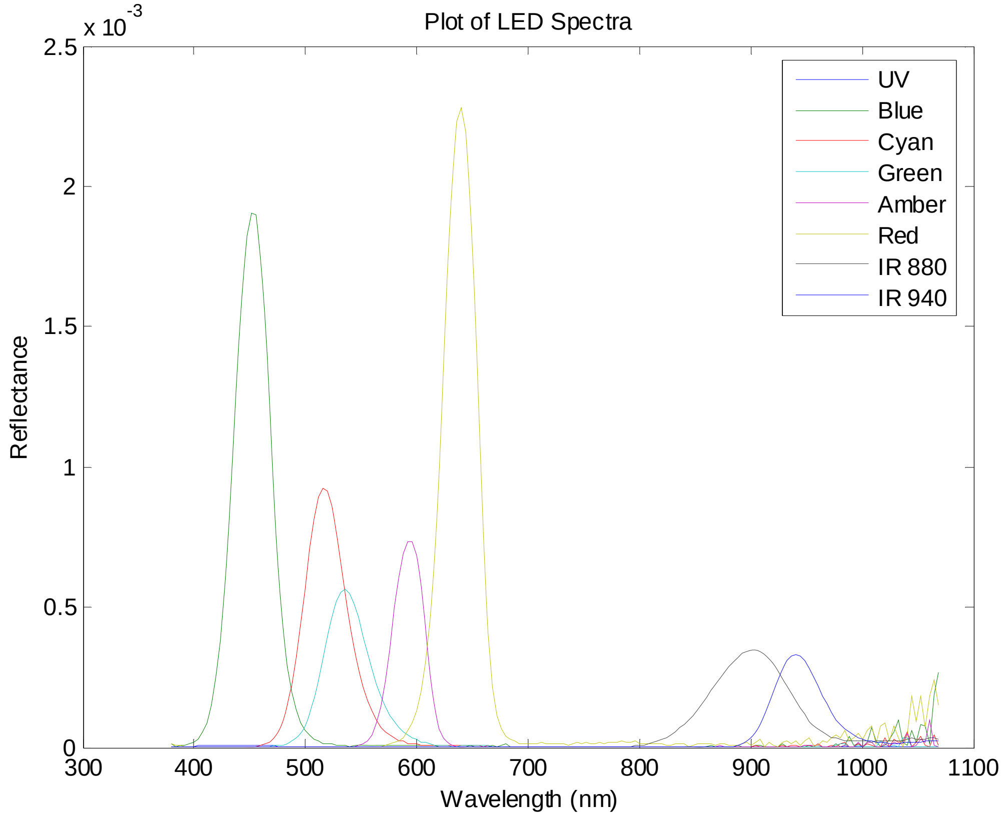

Figure 2: Plot of LED spectra.

Note that the reflectance of the UV LEDs were much more dim (by orders of magnitude) than the other colors, so it barely shows up on this plot. The LED spectra are reproduced below, individually:

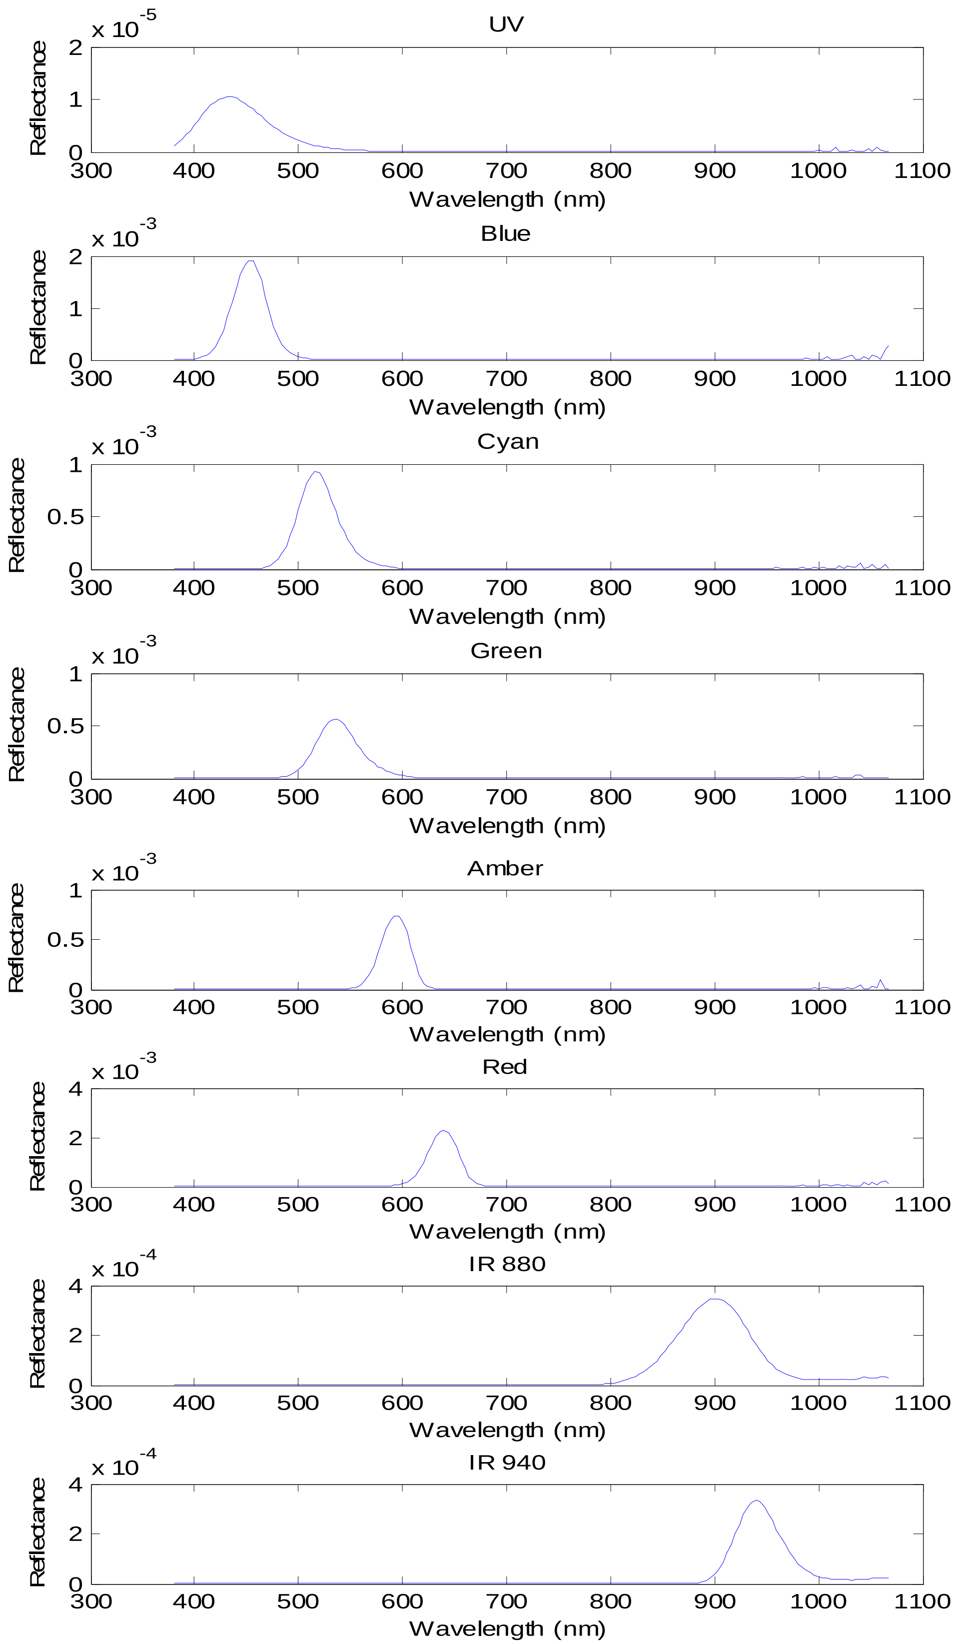

Figure 3: Plot of LED spectra.

METHOD OVERVIEW

Measurement of Reflectances

As a first step in our project we collected data of reflectances of various surfaces. We started out by measuring the reflectances of the skin of different people and then plotting them using the ISET simulation kit. We also measured reflectances of objects like leaves, wood, hair, fabric and fruits. The procedure followed for measuring the reflectances was by focusing the spectrophotometer on the subject .The spectrophotometer actually measures the radiance from the surfaces .When light falls on a surface some of it gets radiated and the spectral power distribution of this radiated light is what is measured by the spectrophotometer and is called the radiance of the surface. We divide out the illuminant from this radiance value to get the reflectance since

$$\text{Radiance}=\text{Reflectance} * \text{Illuminant}$$

The illuminant is initially measured by placing a standard white surface at the same spot as the subjects.

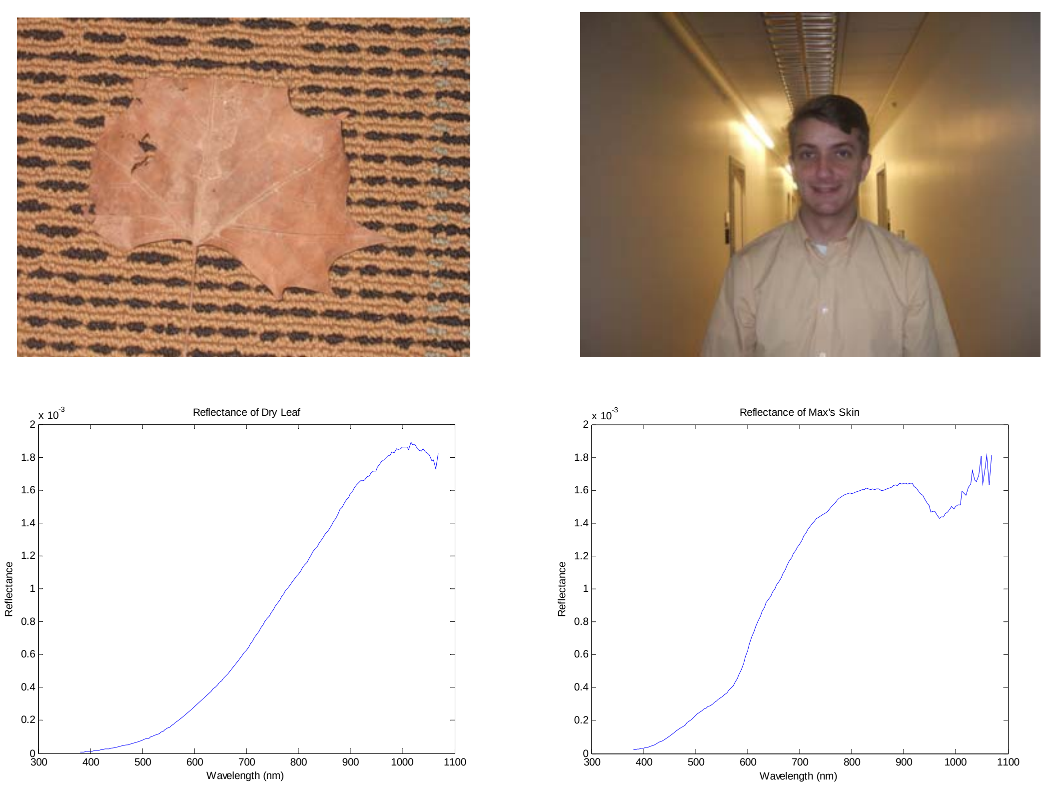

Figure 3: Measurement subjects and their corresponding reflectance spectra.

We can observe the differences in the reflectance plots shown above. The reflectance of the skin shows a dip in the infra-red region while this dip is absent in the case of the dry leaf. This is because the skin has a lot of water content and this absorbs some of the IR light incident on it. However in the case of the dry leaf there is almost no water content and hence no dip is observed.

Spectrophotometer Setup

The spectrophotometer setup is shown in the figure below. The spectrometer is turned on and the proper settings made. Once the spectrophotometer is ready it is focused to a particular point on the surface to be measured. The radiance data is then measured by the spectrophotometer and this is stored on to a connected computer. The previously measured illuminant values are also stored on the computer. From these values the reflectance is calculated as mentioned above. The reflectance spectrum of the surface is plotted using the ISET simulation tool. Before the measurements are made only the required light sources are turned on. The subject has to stay stationary.

Figure 4: Radiance of a piece of wood, under tungsten light, being measured using a spectrophotometer.

After the data had been collected we used the principles of linear algebra especially singular value decomposition on our data to simulate the reflectances of the different surfaces. We then compared these simulated results with the actual measurements through various plots created using MATLAB. These procedures are described in detail in the following section. Our ultimate goal through this project has been multispectral imaging where we have tried to represent an image using more than just one set of R, G, and B values. We went about this by using an LED setup made of 24 LEDs each of colors at eight different wavelengths, namely blue, cyan, green, amber, red, IR (at two different wavelengths of 880nm and 940nm) and UV. The LED setup was used to illuminate the scene with each color successively and the images of a particular object were captured after each illumination using the Nikon D200 whose IR filters have been removed.

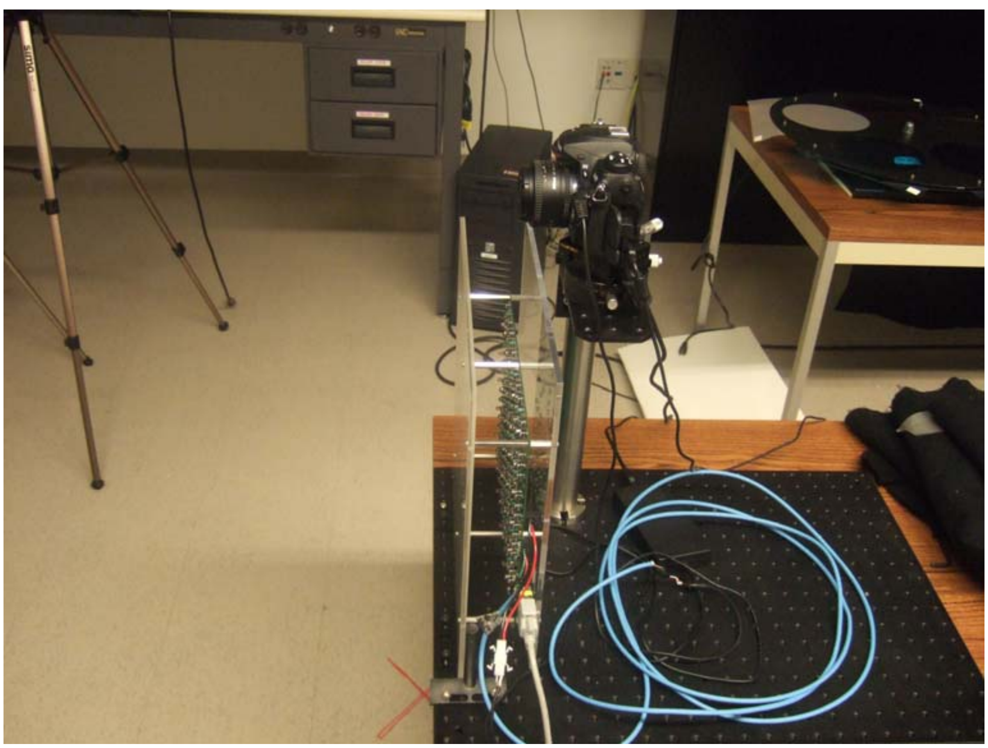

Figure 5: The LED circuit and Nikon D200 setup used to capture multispectral images.

We used this to capture images of people as well as of objects. The advantage of multispectral imaging is that we are now able to represent each image with 24 R, G, B values three each of the eight different wavelengths. Practically though we can get only get up to 20 values since some of the spectra overlap. This still produces a better result than just using one set of R, G, and B values.

Spatial Map and Scaling

The observation that the intensity at different points of the image captured is non-uniform led us to normalize the images.

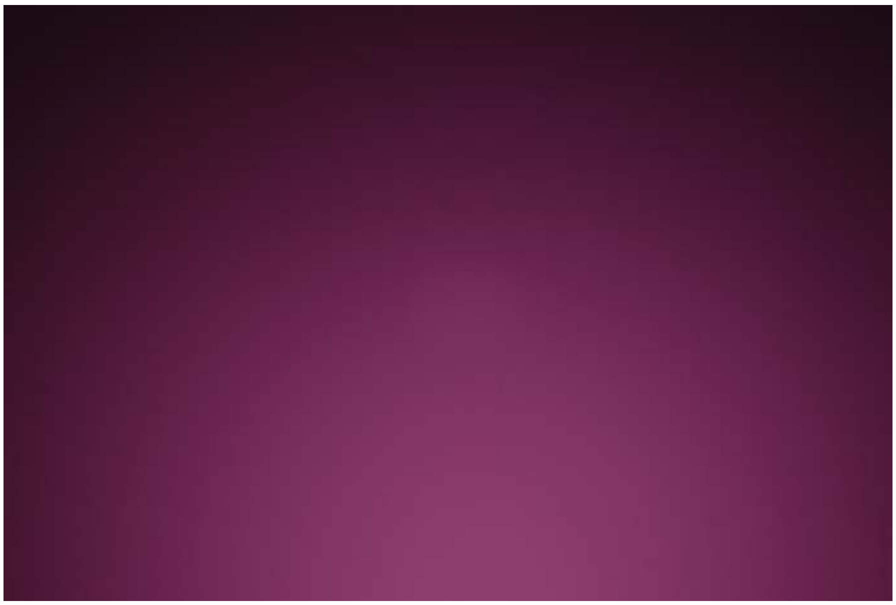

Figure 5: Image of a uniform gray surface illuminated using IR LED.

We went about this by measuring the radiance of a uniform gray surface under the eight different lights .This was repeated six times and the results were averaged using MATLAB to get a better approximation. Then we created a spatial map for each of the R,G and B channels of the eight different colors and applied these to the images of the gray surface in order to make the intensity of the images uniform. The spatial map was created so that it could also be applied to the multispectral images of people, that we had captured using the Nikon D200.

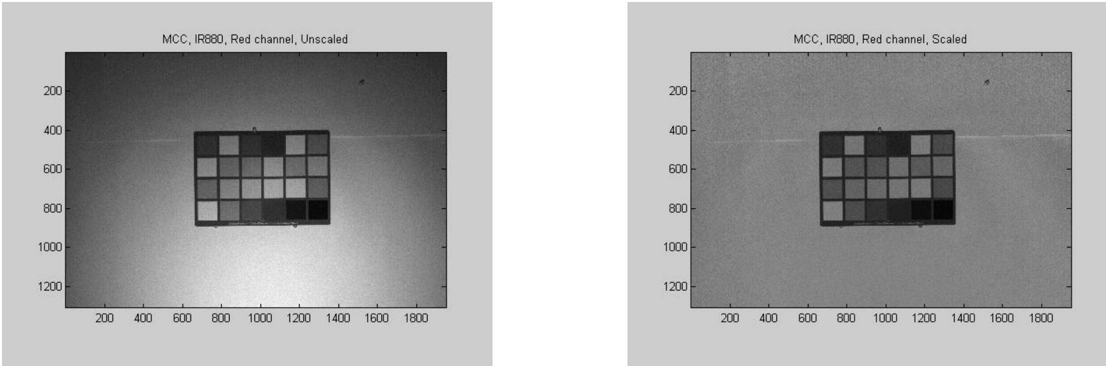

Unscaled Image: IR LED, R channel   *Scaled Image*: IR LED, R channel

Figure 6: Comparison of scaled and unscaled images.

METHOD DETAILS AND RESULTS

Our system that we are modeling is c = F*x, where x is a column vector of our reflectances, c is the column vector of intensities coming from the camera at each combination of camera's RGB color channel and LED. Therefore, F is our system, which is generated from each of the LED spectrums and the camera's RGB spectrum. In our case, we are using 8 LEDs with 3 camera color channels (RGB) and take samples from 173 different wavelengths (from 380nm to 1068nm with a spacing of 4nm), so the size of c is 24×1, F is 24×173, and x is 173×1. If we have more than one measurement, we can also state this as C = F*X, where F is the same but C is $24\times p$, x is $173\times p$, where p is the number of measurements.

Reconstruction using simulated data

Prior to using the Nikon camera and the LED system to capture scenes, we used simulations to predict what the outcome would be. We first simulate c by generating F (from known or measured spectra) and x (measured reflectance). Then, we can try to approximate x from our simulated c, which is essentially reconstructing the original reflectances using the simulated camera output. Since F is an underdetermined system, we cannot perfectly reconstruct x. We start out with 173 variables (x), but the system outputs only 24 variables (c), so there is no way to determine our initial variables (if the initial variables are independent). To find an approximation of x, which we'll call $\hat{x}$, we can solve the following constrained minimization problem:

minimize $\| L^*\hat{x} \|$
subject to: $\| c - F^*\hat{x} \| <=$ maxerror
$0 \le \hat{x} \le 1$

where L is the matrix which is the "discrete derivative" of $\hat{x}$, and maxerror (a user-determined constant) is the maximum allowed error between the measurement data and our results from F* $\hat{x}$. Usually, maxerror is set based on the noise characteristics of our system, if known. Roughly speaking, this problem is constraining the difference between the measured value c and the calculated F* $\hat{x}$ to a certain error, but since there are many solutions to this, we want to choose the value of $\hat{x}$ that minimizes the "roughness" of $\hat{x}$.

Another method to obtain $\hat{x}$ is by generating the pseudoinverse of F, denoted $F^{\dagger}$, and letting $\hat{x} = F^{\dagger} * c$. This is not as good an estimate since we know that reflectance data are not independent: they are generally smooth and do not vary quickly. However, we will demonstrate a robust version of this method later due to issues with scaling during our reconstruction.

Here are some results of our $\hat{x}$ from constrained minimization problem from *simulated* c:

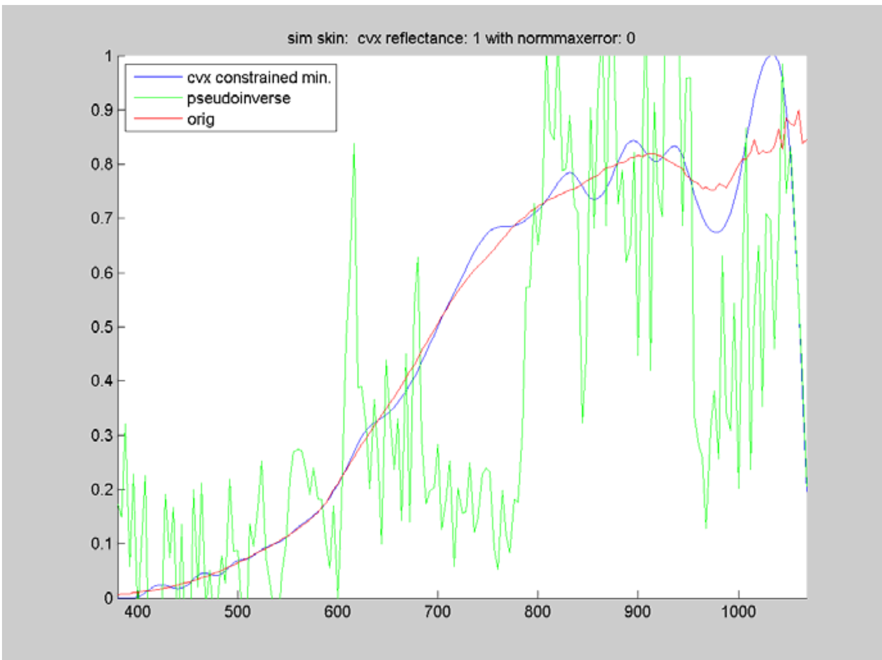

Figure 7: Plot of our original measured skin reflectance overlayed with estimated reflectance from the constrained minimization method and pseudoinverse method.

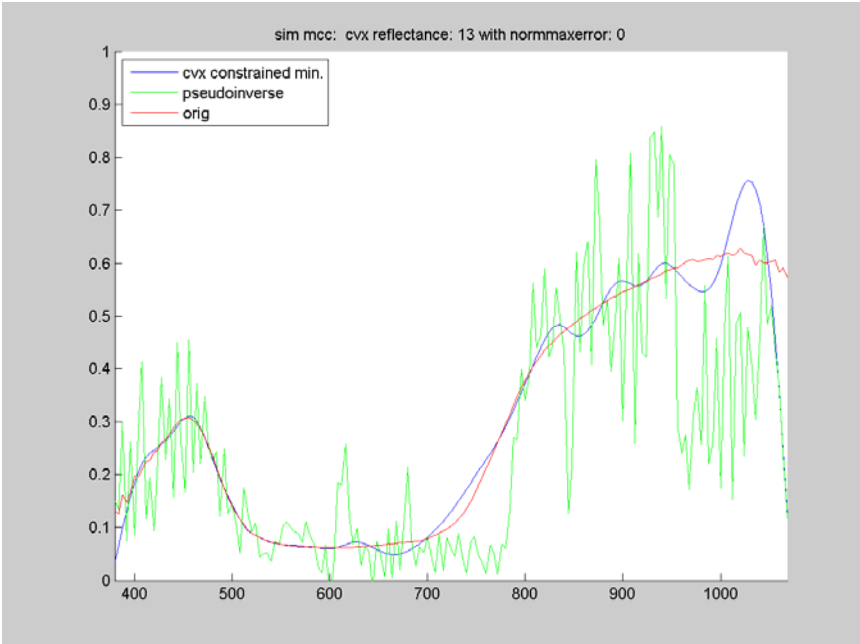

Figure 8: Plot of the reflectance of the blue square of the MCC overlayed with estimated reflectance from the constrained minimization method and pseudoinverse method.

The results of our simulations showed that the constrained minimum method provides a relatively close match to the original reflectance. We will see in the next section how close our actual results are to the original reflectances.

Reconstruction using actual data

After we obtained actual data from the camera, we noticed that we had illumination scaling issues. Our constrained minimization problem (used for the previous method) did not solve it well, because each point in the scene has a different illumination and thus different F matrix. This is due to scaling changes compounded in each system, surface geometry, and angles of illumination, etc.). To solve this, our system F needs to be adapted for each scene point. Since we did not have enough lead time on our data, we could not iron out these illumination scaling problems in time.

However, we used a faster but less robust way to generate these data which gave good results: the regularized pseudoinverse method. Recall $\hat{x} = F^{\dagger} * c$, where $F^{\dagger}$ is the pseudoinverse of F. Once $F^{\dagger}$ is determined, we can simply multiply each of our camera measurements c, by $F^{\dagger}$ to get our estimated x. This single matrix multiply is much faster than solving the optimization problem, but generally it is not as optimal. We calculate $F^{\dagger}$ in the following way. We take many camera measurements (through our real system) of known reflectances. Since we have many measurements, our $F^{\dagger}$ estimate will be relatively accurate. This solves our previous issue of illumination scaling, because we take real data from objects with known reflectances, so our illumination scaling is automatically embedded in our $F^{\dagger}$ estimate.

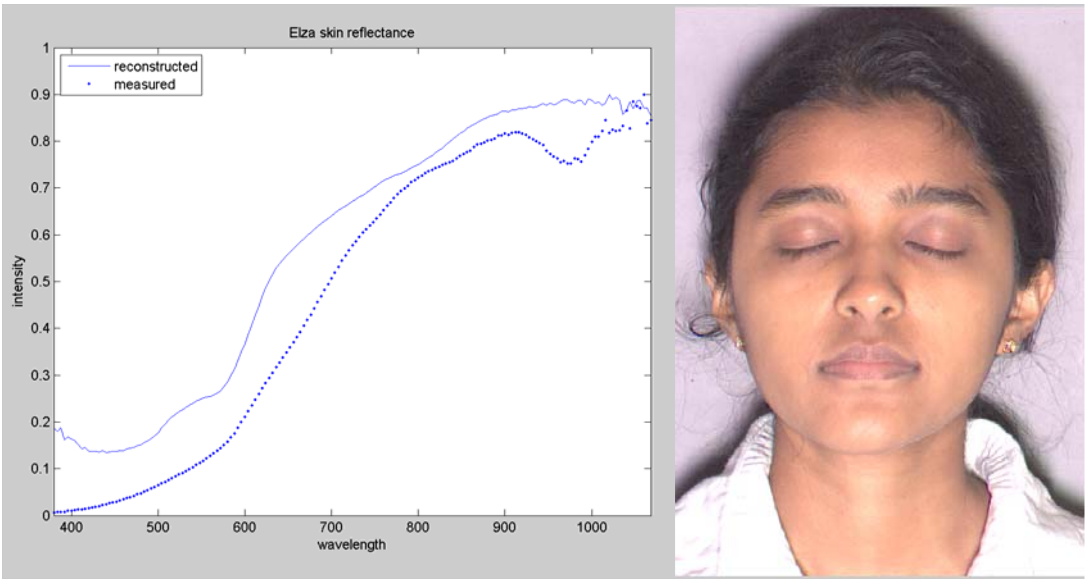

Figure 9: Left: Plot of reconstructed reflectance and measured reflectance of Elza's image. Right: Image generated from pixel-wise reconstruction of reflectances (mapped to RGB using D65).

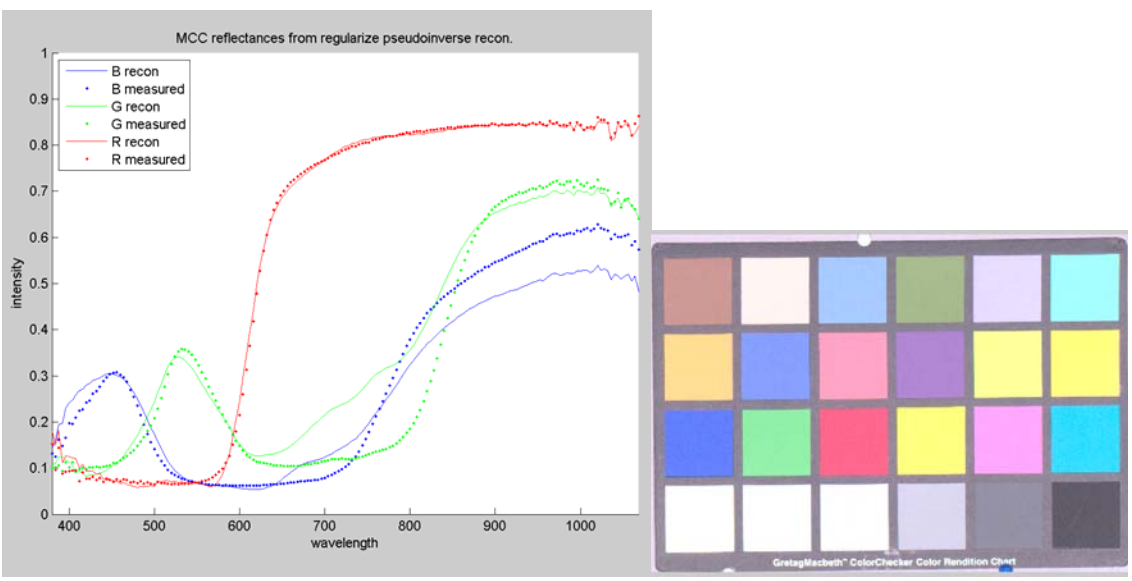

Figure 10: Left: Plot of reconstructed reflectance and measured reflectance for blue, green, and red patches of MCC. Right: Image of MCC generated from pixel-wise reconstruction of reflectances (mapped to RGB using D65).

CONCLUSION:

With accurate models, these techniques can estimate reflectance quite well. Even with the speedy, non-optimal regularized pseudoinverse technique, we can get good results. Our current LEDs cover the range decently, but if more accurate reflectances are needed, one should add another LED to cover the gaps in the spectrum where the radiance is not measured. However, it is difficult to find LEDs that cover this range because there is no market for it. Still, even with this gap in LED spectra, we were able to achieve good results.

REFERENCES

[1] Clinton, Kevin, Faruque, Jessica, Wati, Cecylia. “Accurate Prediction of Scene Radiance Spectra.” PSYCH 221 Final Projects. 2006. Stanford Center for Image Systems Engineering. <http://scien.stanford.edu/class/psych221/projects/06/clintonk/index.htm>

[2]Shimano, Noriyuki. Recovery of spectral reflectances of objects being imaged by multispectral cameras. 19 September 2007. Journal of the Optical Society of America A.

APPENDIX I

See Psych 221 Website.

APPENDIX II
Division of Labor

Below are how we divided specialized tasks, and all other work was contributed to equally by all three members of the team.

Elza John
- Spectrophotometer reflectance measurements
- Nikon D200 reflectance measurements
- Post processing of Nikon D200 images

David Lau
- Spectrophotometer reflectance measurements
- Constrained minimization code
- Reconstruction, image formation, and plotting

Stephanie Leung

- Spectrophotometer reflectance measurements
- Nikon D200 reflectance measurements
- Post processing of Nikon D200 images

ACKNOWLEDGEMENTS

We would like to thank Manu Parmar and Joyce Farrell for their help and guidance.