Estimating the Spectral Reflectance of Natural Imagery Using Color Image Features

Josh Hyman* josh@cs.ucla.edu
Mark Hansen* cocteau@stat.ucla.edu
Deborah Estrin* destrin@cs.udla.edu

Abstract

Relative spectral reflectance is an illumination invarient image feature that is related to many ecological phenomena that are difficult to measure, such as plant CO_2 uptake. We describe a procedure to estimate the relative spectral reflectance of known subject using color image features. Through application, we show that this procedure produces accurate estimates in the presence of changing field conditions. Using this procedure, we can use imagers as sensors to measure natural phenomena that cannot be easily measured using any other available sensing modality.

1 Introduction

There are many important natural phenomena that traditional sensors cannot measure directly. For example, accurately measuring a plant's rate of photosynthesis (release or absorption of CO_2) requires encasing part or all of the plant in a chamber, regulating the air flow, and measuring the composition of the air leaving the chamber. When direct measurement is difficult, imagers are the missing input required to accurately model natural phenomena.

Imagers are typically avoided in traditional sensing applications because they produce large quantities of uncalibrated data. The form of calibration required for an imager-based ecological sensor is dissimilar to that of typical sensors; there is no conveniently accessible reference that can be used to calibrate an imager used as a CO_2 sensor, for example. We aim to take the first step in this calibration process: estimate the spectral reflectance of a known subject using an imager.

We choose spectral reflectance because it had been shown that CO_2 uptake is related to the plant's spectral reflectance [5]. Other application have used spectral reflectance to successfully distinguish soil from vegitation [11] and clouds from land and ice sheets [14]. Legleiter et. al. [9] even used spectral reflectance as the basis for estimating the depth of a river channel.

In order to estimate the subject's spectral reflectance, we must account for the spectral power distribution (SPD) of the incident light. The general form of this calibration, known as color constancy [10], has traditionally been difficult. Unlike the general form, we assume that we have a single subject illuminated by a varying lighting conditions. We show that accurate estimates of a subject's spectral reflectance can be derived from images by modeling the possible illumination and relative reflectance spectra. Further, we show how to build these models from experimentally acquired data.

We define our procedure and discuss work related to image formation in Section 2. The experimental setup for building our models is discussed in Section 3. We evalute the accuracy of our predictions in Section 4. Finally, we draw conclusions and suggest future work in Section 5.

2 Device Calibration

The purpose of device calibration is to undo the effect of changing environmental conditions on the image formation process. Formally, image formation is composed of three components: the spectral power distribution (SPD) of the incident light $E(\lambda)$, the relative spectral reflectance of the surface $S(\lambda)$, and the spectral response of the imaging device's sensor $R(\lambda)$. Assuming the surface is matte or Lambertian, the response of the imager's *kth* sensor to a *(lighting, surface)* pair over the spectral range w is defined by Equation 1.

$$r_k = \int_w E(\lambda)S(\lambda)R_k(\lambda)d\lambda \quad (1)$$

For typical visible light imagers, $w = (400nm, 700nm)$ specifying the visible range, and $k = 3$ corresponding to the red, green, and blue sensors in the imager. Since common commercial imagers intend for their output to be consumed by humans, having only three color sensors is reasonable; human color vision was determined to be a 3-dimensional space by color matching experiments [2]. That is, the use of three orthogonal sensors can represent most of the gamut of human color vision.

This formulation is a bit simplistic. In particular it doesn't capture second-order effects attributed to the camera's lens, shutter speed, and aperture. We assume that the lens' distortion is uniform across the image and that the shutter speed an aperture are set such that the sensor is not saturated. An effect we can't ignore is JPEG image compression. This compression algorithm is both lossy and has a spacial component, considering multiple adjacent pixels at a time. We consider the effects of JPEG compression on this model in Section 4.

2.1 Model Formulation

We build a 3-dimensional linear model for the surface reflectance of the subject using principle component analysis (PCA) [13]; this results in a set of basis functions **B** and their weights w. We can write this in matrix notation (Equation 2) if we discretize the spectral range into n bins; **B** is a $n \times 3$ matrix, w is a 3×1 weight vector, and $\hat{S}(\lambda)$ is a $n \times 1$ vector that estimates of the surface's spectral reflectance. Since we are

*Center for Embedded Networked Sensing
University of California, Los Angeles

ImageSense'08, November 4, 2008, Raleigh, North Carolina, USA

considering outdoor ecological applications, we can apply previous work [8] that has similarly defined a 3-dimensional linear model for daylight (Equation 2) using PCA.

$$\hat{S}(\lambda) \approx \mathbf{B_s} w_s \qquad \hat{E}(\lambda) \approx \mathbf{B_e} w_e \tag{2}$$

Once we have models for illumination and relative reflectance, we must mitigate the effect of the camera's shutter speed and aperture on $R_k(\lambda)$. We assume that this effect, causing under- or over- exposure, is uniform across the sensor and that the sensor is never completely saturated (avoiding the loss of information). By using 2-dimensional chromaticity coordinates we can compensate for this uniform change in brightness. The chromaticity space is the projection of the 3-dimensional color space onto a plane of uniform brightness, and thus mitigates the effects of exposure. The chromaticity space we choose is the x and y dimensions of the xyY color space as defined by CIE.

$$r \approx (\mathbf{B_e} w_e)(\mathbf{B_s} w_s)^T \mathbf{R} \tag{3}$$

Our resulting model for image formation (Equation 3), has six unknowns: the w_e and w_s weight vectors. As described, this system is under constrained since we only have two equations as defined by the chromaticity coordinates. We proceed by estimating these unknowns, w_e and w_s, in sequence. First, we estimate w_e to produce the illuminant's spectra. Then, we transform the image to place it under a reference illuminant. From this "registered" image, we estimate w_s resulting in $\hat{S}(\lambda)$, an estimate of the subject's spectral reflectance. We assume that the same camera is used to produce all of the analyzed images, and thus the effect of $R_k(\lambda)$ on the final pixel value is constant across all images.

2.2 Estimating Incident Lighting

There are a number of lighting estimation techniques suggested by the literature [1], each making different assumptions about the lighting and subject present in the image. We assume that our images have a fixed set of possible illuminants (daylight) and a single subject. By using this application specific information, the Color by Correlation algorithm [4] has been shown to slightly out-perform [6] other lighting esimation algorithms.

To train the Color by Correlation algorithm, we must compute a correlation matrix representing the probability that given illuminant was present in a particular image. Each column of the matrix represents possible illuminant E, and each row is the log probability $Pr(c|E)$ that a particular chromaticity coordinate c would be observed for a subject under that particular illuminant. Since the chromaticity space has infinite extent, the space is quantized to make building a correlation matrix feasible.

$$log(Pr(E|C_{im})) \propto \sum_{\forall c \in C_{im}} log(Pr(c|E)) \tag{4}$$

To test this model, we compute the binary chomaticity vector of an example image C_{im}; this vector is 1 for every value that is present in the image, and 0 elsewhere. Multiplying correlation matrix by an image's binary chomaticity vector, we get the log-likelihood that the sample was produced by each of the possible illuminants assuming uniform prior probabilities $Pr(E)$ and $Pr(C_{im})$ (Equation 4). Then, we simply select the most likely illuminant.

2.3 Changing Illumination

After we've estimated the lighting present in a given image, we must transform the images to be under some reference illuminant. We call this operation *re-lighting* the image. Since we are considering outdoor phenomena, we choose D_{65} (an approximation of daylight defined by CIE) as the reference. To build a re-lighting transformation matrix, we assume that the camera's sensors are sufficiently narrow bandwidth such that Equation 1 can be simplified to Equation 5. That is, they can be modeled as impulse functions at some wavelength λ_k, the center wavelength of the camera's sensor.

$$r_k = E(\lambda_k) S(\lambda_k) \tag{5}$$

This assumption is clearly not true of typical cameras. However, it has been shown [15] that when matte objects are subjected to "reasonable" illuminants (such as daylight), it appears to hold. Using Equation 5, we can define the diagonal lighting transformation matrix T_{light}, as Equation 6.

$$\begin{bmatrix} E_1(\lambda_R)S(\lambda_R) \\ E_1(\lambda_G)S(\lambda_G) \\ E_1(\lambda_B)S(\lambda_B) \end{bmatrix} = T_{light} \begin{bmatrix} E_2(\lambda_R)S(\lambda_R) \\ E_2(\lambda_G)S(\lambda_G) \\ E_2(\lambda_B)S(\lambda_B) \end{bmatrix}$$
$$T_{light}[i,i] = E_1(\lambda_i)/E_2(\lambda_i) \tag{6}$$

We choose these center wavelengths to be $\lambda_R = 620nm$, $\lambda_G = 530nm$, and $\lambda_B = 450nm$; these are close to the center wavelength of the sensors on typical digital cameras [3]. The Color by Correlation algorithm produces $E_2(\lambda)$ and we have already assumed that $E_1(\lambda)$ is the standard D_{65} illuminant. To re-light the image, we need only compute the diagonal elements of T_{light} and then transform each of the image's pixels individually.

This formulation only works if we specify $E_1(\lambda_k)$ and $E_2(\lambda_k)$ in absolute terms. However, the spectral power distribution of an illumination is typically normalized such that $E(\lambda_{560}) = 100$ (as is the case for the D_{65} specification). This has the effect of multiplying T_{light} by β as defined in Equation 7.

$$\beta = \frac{E_2(\lambda_{560})}{E_1(\lambda_{560})} \cdot 100 \tag{7}$$

The β term can be factored out of the resulting transformed image if we use chromaticity coordinates instead of absolute color coordinates. This is intuitively true since chromaticity coordinates are designed to be independent of brightness, the effect for which β compensates. Further, using chromaticity coordinates is a reasonable requirement as we have already leveraged chromaticity coordinates to produce a brightness invariant image for lighting estimation.

2.4 Estimating Relative Spectral Reflectance

We can now estimate the weights w_s for the relative spectral reflectance basis functions $\mathbf{B_s}$ (see Equation 2). Unlike lighting estimation, however, we have less insight into the relationship between relative spectral reflectance and the chromaticity coordinates. Accordingly, we choose to estimate

Figure 1: Moss at James Reserve during July 2008 after a long dry period 1(a), left. The experimentally captured images, 1(b), right.

the parameters of our relative spectral reflectance model using non-linear regression. The input to this non-linear regression will be the 2-dimensional chromaticity coordinates. Similar to the Color by Correlation algorithm, we quantize the chromaticity space into $n \times n$ bins, using each as feature in our predictive model. These features are stable between images since we previously corrected for changes in illumination using the re-lighting transform.

Our previous work [7] showed that using this technique produced reasonable results for laboratory data. The dataset used in that work had consistent illumination since all images were taken under controlled laboratory lighting. As a result, it did not require the images to be chromatically registered using a re-lighting transform.

3 Experimental Setup

We acquired six samples of *Tortula princeps*, a drought tolerant moss, from the James Reserve, seen in Figure 1(a). We hydrated the moss and allowed it to dry for approximately 6 hours (from 12pm until 6pm), simulating a rain event in the field. During this drying period, the moss' spectral reflectance is expected to change. We placed the moss outside under natural illumination and collected samples of the incident illumination, the moss' relative spectral reflectance, and images containing the moss and MacBeth Color Checker (seen in Figure 1(b)). Samples were collected every 15 minutes; in total, 23 samples were collected.

In order to measure both the incident illumination as well as the plant's relative spectral reflectance, we used a spectroradiometer (Licor 1800). To measure the absolute spectral power distribution of the incident illumination, we calibrated the response of the spectroradiometer using a reference tungsten illuminant, similar to the CIE A reference. Similarly, the spectral reflectance of the plant was measured with respect to same tungsten illuminant. Samples of both the plant and the incident illumination were taken at 2nm increments from 390nm to 750nm.

Images of moss were taken using two standard consumer-grade cameras with their auto white-balance settings turned off. We used a Canon EOS 450D to capture 10MP images in both RAW format and JPEG format; this camera represents a relatively high-end imager. Additionally, we used a Pentax Optio S5z to capture 5MP images in JPEG format; this camera represents a lower-end imager. Each image contained both the moss sample as well the MacBeth Color Checker reference; this chart contains 24 color swatches of known spectral reflectance.

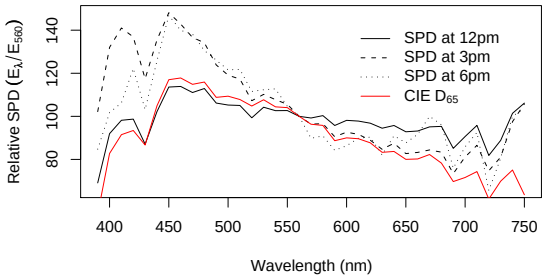

(a) Measured Illumination

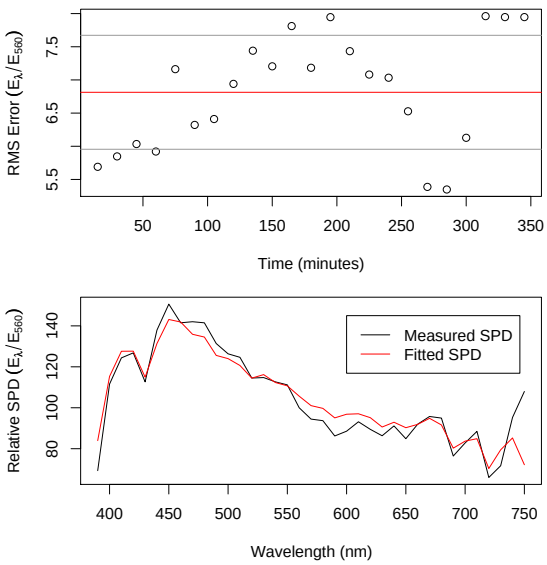

(b) Illumination model accuracy

Figure 2: The spectral power distribution (SPD) of the illumination measured during the course of the experiment as well as the CIE standard D_{65} illuminant is shown in 2(a). The accuracy of the daylight model built by Judd et. al. [8] for our measured illuminants is plotted against time in 2(b), the red line is the mean and the grey lines are the first standard deviation. Below, we show the fit for the sample with the largest RMS error (the 21st sample at 315 minutes).

4 Evaluation

We verified that the illumination we measured using the spectroradiometer was reasonable by comparing it to the CIE standard D_{65} illuminant as seen in Figure 2(a). Each of these spectra have been normalized such that $E(\lambda_{560}) = 100$. We see that our measured spectra have the same characteristic shape as D_{65} although they are slightly bluer late in the day; they contain more power in the 400nm–500nm range than D_{65}. This similarity suggests that our measurements are producing reasonable spectra.

Using the daylight model derived by Judd et. al. [8], we computed the weights (w_e) of the basis functions ($\mathbf{B_e}$), as defined in Equation 2, for our measured illuminants. As shown in Figure 2(b), the RMS error of this model does follow some time dependent trend through the course of the day. Initially, this might suggest that the model is missing some relevant

Figure 3: Error of the Color by Correlation model derived from images of moss and the MacBeth Color Checker taken with the Pentax Optio S5z camera under varying illumination.

information. However, we see that the fit for the example with the largest absolute RMS error still is quite good. This further confirms that our measurements are accurate.

The 2-dimensional chromaticity distributions of the sampled images (quantized into 32×32 bins) are stored in matrices that we convert into row-major ordered vectors. Each vector is normalized by the number of pixels in the image and associated with the illumination measured using the spectroradiometer; these become the columns in the correlation matrix. Recall, to produce the log-likelihood that example image has been illuminated by particular illumination (Equation 4), we simply multiply the correlation matrix by the binary chromaticity vector.

The training set for the Color by Correlation algorithm is selected at random from the set of experimentally obtained samples. We hand segmented the images from both cameras into an images containing only the moss and images only containing the MacBeth Color Checker. Figure 3 shows average RMS residual error between the predicted illumination and the measured illumination as a function of the training set size. Interestingly, for large training set sizes ($n \geq 16$), the moss images had a slightly lower error than the images containing the MacBeth Color Chart. This is odd because the moss' reflectance is changing over time, where as the chart's reflectance is constant. In these cases approximately 70% of samples were used for training, so we believe this is simply an effect of over-training the model.

Though not shown, the model trained using both raw and JPEG images taken from the Canon camera produced similar residual error. This interesting result shows that JPEG compression has a minimal effect on the accuracy of the Color by Correlation algorithm when applied to these data. We can understand this result by considering how JPEG compression works. It computes the discrete 2-dimensional cosine transform of each 8×8 pixel block in the image, producing the spatial frequency of colors within a given image block. Then, it discards some information about the high frequency signals, retaining most information about the lower frequency signals. For both the moss and the MacBeth Color Chart, the spacial frequency in all three color dimensions is relatively low. This suggests that JPEG compression would have minimal effect on the chromaticity-based signals we are using to build our model.

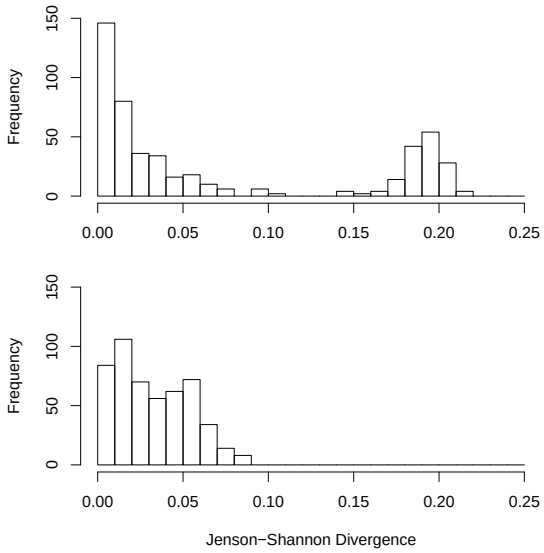

(a) Pairwise divergence before (top) and after (bottom) relighting

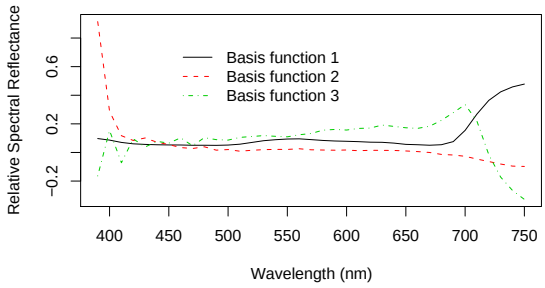

(b) Basis functions for relative spectral reflectance

Figure 4: 4(a) The Jenson-Shannon Divergence, before and after re-lighting, of all pairs of images containing the MacBeth Color Chard under varying daylight illumination. Optimally, all divergences would be zero after the lighting transformation. 4(b) The basis functions, as determined by functional PCA for the relative spectral reflectance of the moss as it dries over time.

Once we obtain an accurate estimate of the image's lighting we can correct for that illuminant using T_{light} (see Equation 6). We test this transform on our segmented images containing the MacBeth Color Chart because its spectral reflectance doesn't change (unlike the moss). We to compute the 2-dimensional Jenson-Shannon Divergence of the discretized chromaticity coordinates for all pairs of examples. We expect the divergences to small since the subjects are identical.

Histograms of these divergences are shown in Figure 4(a). As we can see, the lighting transformation compresses the distribution of divergences towards zero as expected. An unfortunate consequence is that it has also increased the variance among the previously well clustered examples. We hypothesize that this is caused by inherent error in estimating our sensors as impulse functions, a poor choice of center wavelengths, or color alterations resulting from JPEG image

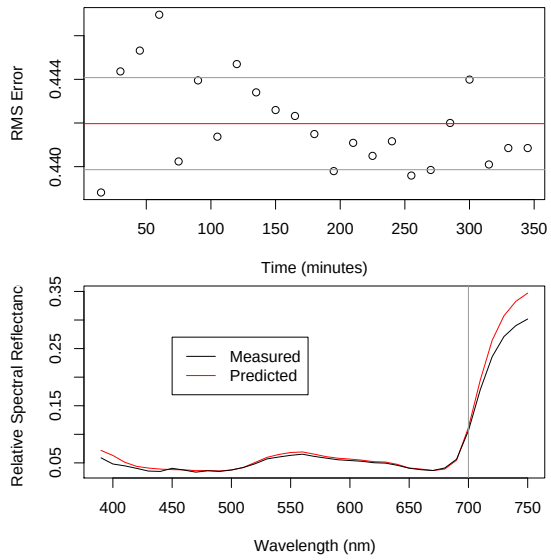

Figure 5: The RMS residual error of the spectral reflectance predicted by our procedure is shown above; the red line is the average error and the grey lines are the first standard deviation. Below, we show the predicted spectral reflectance of the observation with the largest error (observation 4 at time 60 minutes).

compression.

After the images have been transformed we must predict the parameters of the relative spectral reflectance model (shown in Figure 4(b)). We have chosen to use only the first three basis functions for our model because they account for 99.96% of the variance contained in the moss' measured relative spectral reflectance. The first basis function, plotted in black, represents the average spectral reflectance across all samples. The second and third basis functions, plotted in red and green respectively, show the type of variation seen. In particular, we see that there is significant variation in the blue (400nm – 450nm) and red (675nm – 750nm) parts of the spectrum. We expect some variation near 400nm because it is near the minimum wavelength our spectroradiometer can measure. It is not clear what caused the variation around 700nm. We suspect it was due to drift in the spectroradiometer's sensors during the course of the experiment.

Given this model, we must predict weights of these basis functions (w_s from Equation 3). We do this by training three regression-tree based models, one for each parameter, using the 2-dimensional chromaticity coordinates from the images previously registered by re-lighting. We trained this estimation model using 12 samples, a value that produced reasonable results for the lighting estimation. The RMS residual error of this prediction is shown in Figure 5. We can see that the magnitude of the error is rather large and there are some rather significant outliers. In comparison, the best possible values for w_s produce a mean RMS residual error of 0.021, approximately 20 times smaller than the error produced by the spectral reflectance estimation model.

To better understand this error we plot the measured and estimated spectral reflectance for the largest outlier, sample 4 occurring at 60 minutes. As we can see in Figure 5, the fit is quite good. The vast majority of the error comes from wavelengths greater than 700nm. This error is somewhat expected since it is present in the model's basis functions as well as the original measurements. Still, there is certainly significant room for improvement. In particular, the error of this prediction can be reduced by improving the accuracy of the re-lighting transform.

5 Conclusion

In this paper we have described a procedure to estimate the relative spectral reflectance of natural subject using an imager. We have instantiated this procedure to measure the spectral reflectance of *Tortula princeps*, a drought tolerant moss plant, collecting the required data through laboratory experiments. Finally, we have shown that our proposed procedure accurately estimates relative spectral reflectance in the context of this application.

In future work, we plan to use estimated spectral reflectance to predict the underlying natural phenomena; in this case, CO_2 uptake. We plan to refine this procedure by iteratively computing the models for illumination and relative spectral reflectance as suggested by [12]. Since the output of these two models are used together, such iterative modeling would help to make the prediction of each more accurate in the presense of likely variation in the other. Finally, to increase the accuracy of image re-lighting, we intend to apply sensor sharpening techniques to better chromatically register the images.

6 References

[1] K. Barnard, V. Cardei, and B. Funt. A comparison of computational color constancy algorithms. Part I: Methodology and experiments with synthesized data. *IEEE Transactions on Image Processing*, 11(9):972–984, 2002.

[2] R. P. Feynman, R. B. Leighton, and M. Sands. *The Feynman Lectures on Physics*, chapter 35: Color Vision. Addison-Wesley, 1963.

[3] G. D. Finlayson and S. D. Hordley. Color constancy at a pixel. *Journal of the Optical Society of America A*, 18(2):253–264, 2001.

[4] G. D. Finlayson, S. D. Hordley, and P. M. Hubel. Colour by correlation: a simple, unifying approach to colour constancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23:1209–1221, 2001.

[5] A. Gitelson, Y. Gritz, and M. Merzlyak. Relationships between leaf chlorophyll content and spectral reflectance and algorithms for non-destructive chlorophyll assessment in higher plant leaves. *Journal of Plant Physiology*, 160(3):271–282, 2003.

[6] S. D. Hordley and G. D. Finlayson. Reevaluation of color constancy algorithm performance. *Journal of the Optical Society of America A*, 23(5):1008–1020, 2006.

[7] J. Hyman, E. Graham, M. Hansen, and D. Estrin. Imagers as sensors: Correlating plant CO_2 uptake with digital visible-light imagery. *Workshop on Data Management in Sensor Networks*, 2007.

[8] D. Judd, D. MacAdam, G. Wyszecki, et al. Spectral distribution of typical daylight as a function of correlated color temperature. *Journal of the Optical Society of America*, 54(8):1031–1040, 1964.

[9] C. J. Legleiter, D. A. Roberts, A. Marcus, and M. A. Fonstad. Passive optical remote sensing of reiver channel morphology and in-stream habitat: Physical basis and feasibility. *Remote Sensing of the Environment*, 93:493–510, 2004.

[10] L. T. Maloney and B. A. Wandell. Color constancy: a method for recovering surface spectral reflectance. *Readings in Computer Vision: Issues, Problems, Principles, and Paradigms*, 1987.

[11] J. Marchant, N. Tillett, and C. Onyango. Dealing with Color Changes Caused by Natural Illumination in Outdoor Machine Vision. *Cybernetics and Systems*, 35(1):19–33, 2004.

[12] D. H. Marimont and B. A. Wandell. Linear models of surface and illuminant spectra. *Journal of the Optical Society of America A*, 9(11):1905–1913, 1992.

[13] J. Ramsay and B. Silverman. *Functional Data Analysis*. Springer, 1997.

[14] T. Shi, E. Clothiaux, B. Yu, A. Braverman, and D. Groff. Detection of daytime arctic clouds using misr and modis data. *Remote Sensing of Environment*, 2006.

[15] J. Worthey and M. Brill. Heuristic analysis of von Kries color constancy. *Journal of the Optical Society of America A*, 3(10):1708–1712, 1986.