

## Verification of Sky Models for Image Calibration

Rishi Ramakrishnan, Juan Nieto and Steve Scheduling  
University of Sydney, Australia

{r.ramakrishnan, j.nieto, s.scheduling}@acfr.usyd.edu.au

### Abstract

*Perception systems operating in outdoor scenarios face challenges due to the high dynamic range of the image, as different regions are illuminated by varying amounts of sunlight and skylight. A pre-processing step like image calibration can be used to convert the pixel values to an illumination independent domain such as reflectance. Each pixel is therefore represented by a characteristic material description, instead of an illumination and viewpoint dependent pixel colour. This assists object identification, segmentation and classification algorithms [20].*

*This paper investigates modelling the sky colour through a number of parametric approaches typically used in the computer graphics community for rendering purposes, namely those developed by Preetham et al. [17] and Hosek-Wilkie [6]. The models are compared in terms of chromaticity with observations taken from a camera and are used to develop an environment map for the application of inverse reflectometry of diffuse objects in an outdoor environment. This is of particular importance for applications involving imaging objects whose primary illumination source is skylight. It was found that the Hosek-Wilkie [6] model produced more robust estimations and was less sensitive to changes in azimuth, while both models had similar reconstruction results with angular errors of approximately 0.15 radians.*

### 1. Introduction

Computer vision algorithms that operate consistently in outdoor situations are important for both the robotics and remote sensing communities. Algorithms such as feature extraction and description, and object identification, classification and tracking, are used in robotic applications for a number of purposes including mapping and navigation. In remote sensing, it is important to correct for variations in observations due to geometry and illumination prior to the application of any high level processing such as classification.

In outdoor imaging scenarios, achieving such a consis-

tency is challenging due to the influence of the illumination source. The main illumination source in these contexts is the sun and this can be split into three components; sunlight, skylight and indirect illumination. The sun itself can be considered as a light source emitting light rays in a (roughly) spherical pattern, but due to the large distances involved, the rays that strike the earth can be approximated as being parallel to one another. As these rays pass through the atmosphere they undergo scattering due to interactions with particles of different sizes. Those which do not scatter are called sunlight and have high luminance values, thereby casting sharp shadows in the scene. Light rays that scatter off different particles have their frequencies altered thereby generating the blue colour of the sky, with their luminance at a significantly lower value. This is known as skylight and casts soft shadows in the scene. Indirect illumination occurs when light reflects off multiple objects and is highly dependent on the material properties.

Knowledge of the incident illumination, geometry of the scene and the captured image, can be used to determine the characteristic spectra of an object in the scene through a process known as inverse reflectometry. Once this occurs, algorithms such as object identification and classification can operate independently of the conditions. Our main application for the use of sky models in outdoor inverse reflectometry scenarios is for the identification and classification of mineral ores on a mine face. A standard method is to scan the face with a hyperspectral camera and convert the radiance measurements to reflectance through the use of a small reflectance panel placed in the scene in a region illuminated by sunlight [19]. However, this is only valid for one specific point and geometric orientation, and fails when the illumination varies due to occlusions. By performing inverse reflectometry, a per-pixel calibration of hyperspectral images can take place allowing pixels to be described by spectra independent of the illumination. Following this, a number of classification techniques can be utilised to identify the minerals present in the scene.

The contributions of this paper are the verifications of various sky models in terms of their colour and the preliminary application of inverse rendering in outdoor environ-

ments. The focus is on recovering reflectance spectra of objects whose primary source of illumination is skylight. Section 2 presents the different sky models that have been used in the computer graphics community. Section 3 describes the methodology to perform a verification of the models through camera calibration, comparison with observations, and an application that uses these models. The results of the experiments are shown in Section 4.

## 2. Related Work

Inverse reflectometry is the technique of inferring material reflectance given the image, geometry and lighting conditions. The geometry can be obtained in outdoor scenarios through the use of sensors such as laser scanners to develop three dimensional models [3][11], while the illumination can be captured through techniques such as environment mapping. This involves placing reflective spheres in the scene [3][4] or taking six images and projecting them onto a cube. These mappings encompass the entire scene and can be sampled from using informed methods in order to calculate the illumination at different points in the scene [1][10]. An issue with techniques such as sphere and cube mapping is the fact that they require either an object in the scene or rotating the camera, which for our future application of utilising hyperspectral cameras, is not feasible as the camera only rotates about the vertical axis. A potential solution to these problems is through the use of an illumination model.

Sky model approaches involve estimating the luminance and chromaticity of the sky through a series of equations. They are typically used in the computer graphics community for rendering outdoor scenes and can be split into two categories based on whether they estimate colour or not. The luminance only sky model by Perez *et al.* [15] has served as the basis for many other models, including those developed by Preetham *et al.* [17], and recently Hosek-Wilkie [6]. This model is a parametrised equation based on measurements of the sky, with the relative illuminance being a function of zenith angle, and the angle between the sun and a sky element. It contains variable coefficients that can be altered based on the conditions.

This model was extended by Preetham *et al.* [17] to also estimate chromaticity by simulating the sky using an analytical model developed by Nishita *et al.* [14]. Following the simulation of the sky, an optimisation process was run to generate an equation as a function of the sun position and turbidity. The luminance values for this model were analysed in [21] and it was found that there were several problematic cases where negative luminance could be achieved depending on the magnitude of the turbidity and elevation angle. Nevertheless, the model has become widely used due to its simple implementation and computational speed [7].

A model recently proposed by Hosek-Wilkie [6] adds ad-

ditional terms to the Perez *et al.* [15] model in an attempt to fix gradient issues around the zenith and a glow around the sun position. A simulation of the sky was generated from the scattering equations to build reference datasets and a model was fitted and Bezier curves are used to estimate the parameters. The models presented in [17] and [6] (hereby referred to as the Preetham and Hosek-Wilkie models) are the most common and recent sky models, and will therefore be utilised in this paper to determine their suitability for outdoor inverse rendering.

## 3. Methodology

In order to verify the sky models, we compare the expected chromaticity values at specific angles with those from a calibrated camera. After a comparison between the two is made, we utilise these models in a proof of concept application of inverse reflectometry for diffuse materials in an outdoor scenario. The system we aim to develop is shown in Figure 1, which shows a hyperspectral camera and laser scanner being used to develop a geometric model and spectral description of each point in the scene. Once these two data sources are registered together, a Global Positioning System (GPS) receiver and compass gives the required location information to estimate the illumination at each point, calibrate the hyperspectral image, and perform thematic mapping on the mine face. This system is an extension of those developed in Murphy *et al.* [12] and Nieto *et al.* [13], and should allow more robust classification to occur.

### 3.1. Camera Characterisation

In cameras containing a Charge-Coupled Device (CCD) sensor array, the incident light from the environment is measured by a sensor after it has passed through a colour filter array, often in the form of a Bayer pattern. The image that a user receives is typically a non-linearly transformed representation of the radiance in the scene due to the spectral sensitivities of the sensor, and the application of algorithms

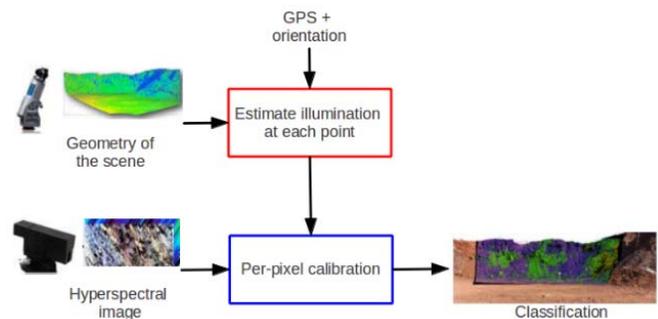


Figure 1: Overview of the proposed image calibration system for hyperspectral imagery.

such as white balancing and gamma correction. These are applied to produce visually pleasing images, but are inconvenient when comparing the sky models and observations as they are in different colour spaces and cannot be compared directly unless we account for these. Therefore, we use the raw image captured by the camera, which is proportional to the scene radiance. The camera used for the experiments is a Canon Powershot A720 IS and can be made to store raw images by loading the Canon Hacker Development Kit (CHDK). Once the images are captured, they are decoded and converted to a portable graymap format prior to demosaicing. Therefore, for each pixel, we have a three dimensional vector  $P_{raw}$  that represents the measurement in the raw colour domain and is nominally linearly related to the luminance [9].

The camera saves the measurements of the sky in the raw colour domain while the two sky models generate colours in the  $XYZ$  colour space with reference to the  $D65$  reference illuminant ( $XYZ_{D65}$ ). In order to compare the two data we transform them to the  $XYZ_{D65}$  colour space, which is a tristimulus, device independent model [8][9]. The raw colour domain measurement, where each channel is characterised by the sensor response function of the camera, are converted by generating a device dependent transformation matrix. These response functions are found by imaging a X-Rite ColorChecker board under a known illuminant which in this case, is a lamp. The raw values for each patch  $P_{raw}$  are obtained by manually selecting a region of interest and calculating the mean for each channel. The spectra of the lamp  $I_{lamp}(\lambda)$  is found by placing a reflectance panel in the scene and measuring the reflected spectra with an ASD spectrometer. This spectrometer is also used to measure the reflectance spectra of each colour patch  $R(\lambda)$ . The spectral response functions can be approximated in a number of ways [2][8], and we use the quadratic fitting method developed in [16]. The pixel intensity in a raw image is calculated as:

$$P_{j,raw}[k] = \int_{380nm}^{830nm} I(\lambda)R_j(\lambda)S_k(\lambda)d\lambda, \quad (1)$$

$$k : \{R, G, B\}, j \in [1, 24],$$

where  $j$  is the index of the colour patch and  $k$  is the colour channel [8]. We simulate the appearance of the ColorChecker board under a  $D65$  reference illuminant light source by substituting the illumination spectra and camera response functions into Equation 1. The tristimulus values for each colour patch in the  $XYZ_{D65}$  are also calculated as:

$$P_{j,xyz}[k] = \int_{380nm}^{830nm} I(\lambda)R_j(\lambda)\bar{m}_k(\lambda)d\lambda, \quad (2)$$

$$k : \{X, Y, Z\}, j \in [1, 24],$$

where  $m$  are the colour matching functions [9]. A  $3 \times 5$  matrix is generated that allows transformation between

the  $XYZ_{D65}$  colour space and the raw colour domain  $M_{XYZ_{D65} \rightarrow raw}$ .

$$M_{XYZ_{D65} \rightarrow raw} = P_{raw}P_{XYZ_{D65}}^+, \quad (3)$$

$$M_{raw \rightarrow XYZ_{D65}} = P_{XYZ_{D65}}P_{raw}^+, \quad (4)$$

where  $P_c$  is a  $3 \times 24$  matrix containing the tristimulus values for each colour patch in the specified colour space  $c$ , and  $P^+$  is the pseudo-inverse of matrix  $P = [r\ g\ b\ r\ g\ b\ 1]^T$  or  $P = [x\ y\ z\ xyz\ 1]^T$ . This polynomial transformation method was shown to decrease the colour differences between the two colour spaces [5]. The transformation matrices are utilised for comparing the sky model estimates with the observations, as well as calculating the total illumination in the scene.

### 3.2. Sky Dataset Collection

Comparing the modelled chromaticity of the sky with observational data involves capturing images of the sky in raw format. The input parameters into the sky models are location, orientation, time and turbidity, so these are the variables we can alter. Sky models fail when clouds are present in the sky, unless the location and height of these clouds are known. Unfortunately, we do not have this data so for this work we focus on clear sky conditions. The images of the sky are taken from the horizon level to zenith in ten degree increments at a number of different orientations and times of day. As Figure 2 shows, images were captured facing away from the sun in order to minimise artefacts brought about by inscattering of sunlight.

After raw images of the sky are captured, a minimal amount of pre-processing is applied so as not to impact the measurements. The raw file is converted to a readable image format and saved in 16 bit form. The saturation limit of the raw image is 10 bits, so no scaling of the data has occurred. The image is demosaiced to produce an image with three layers in the camera sensor domain. The measurements for a specific angle are taken by sampling a square region around the centre of the image, which reduces the influence of vignetting in the image. In order to compare the observations against the models we first convert the raw data to the  $XYZ_{D65}$  colour space using the transformation matrix  $M_{raw \rightarrow XYZ_{D65}}$ . This is an approximation of the required transform as ideally, the conversion matrix should have been developed using the same illuminant sources [9]. However, this requires the sky spectral information from each point on the hemisphere and it is therefore not feasible to compute for each pixel. The  $D65$  standard illuminant has a broad band spectrum and approximates daylight, so this is a reasonable illumination model to use for the estimation.

### 3.3. Diffuse Inverse Reflectometry

As a proof of concept for an outdoor image calibration system, we investigate the use of sky models to perform

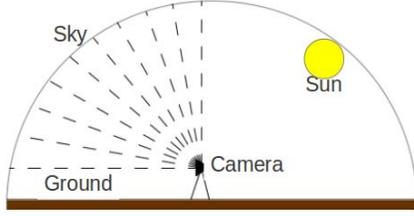


Figure 2: Collection of sky data is done by imaging at ten degree increments in elevation at a fixed azimuth. To avoid inscattering, datasets are captured facing away from the sun.

inverse reflectometry for diffuse materials. In order for this to occur we need information about the illumination source and scene structure, as well as an observation of the scene.

In this context the illumination source is described by first finding the location and orientation of the scene using a GPS receiver and a compass respectively. The latitude, longitude, time and time-zone of the the experiment are used as inputs into a solar position algorithm [18] that calculates the azimuth and elevation angle of the sun. The location of the sun is used in the sky model, along with the turbidity factor which we assume to be between one and two for a clear sky scenario [17]. The models are then used to generate a hemispherical environment map that can be used to illuminate the scene. In this experiment we uniformly sample the hemisphere 100 times to generate the incident illumination at a point in the scene:

$$\langle L(x) \rangle = \frac{1}{N} \sum_{i=1}^N \frac{L_m(x \leftarrow \psi_i) V(x, \psi_i) \cos(\psi_i, N_x)}{p(\psi_i)} \quad (5)$$

where  $L(x)$  is the illumination at a point  $x$  in the scene,  $N$  is the number of samples,  $L_m(x \leftarrow \psi_i)$  is the illumination from direction  $\psi_i$  determined from sampling the sky models,  $V(x, \psi_i)$  is the binary visibility function that determines whether point  $x$  is occluded from the light sample,  $N_x$  is the normal of the object and  $p(\psi_i)$  is the probability density function for the environment map.

This paper is focused on utilising the sky models and their chromaticity predictions, so we minimise indirect illumination by imaging inside a cube with one side left open as shown in Figure 3. The inside of the cube is lined with a thick black fabric in order to minimise light from reflecting off the sides, thereby re-illuminating the object. The diffuse object that will be imaged is the ColorChecker board and this is placed at the bottom of the cube. The camera is situated directly above it and the positions of the object, camera and cube are measured and modelled. This is analogous to obtaining the geometry of the scene through the use of a laser scanner as shown in Figure 1. In order to obtain only skylight as the illumination source, we orientate the cube

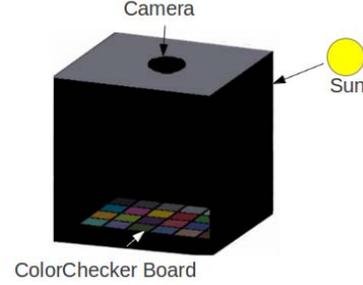


Figure 3: An image is taken of a ColorChecker board inside a black box that is exposed to skylight. Inverse rendering is then used to calculate material properties using the sky models as the illumination source.

such that the opening side faces away from the sun and then take a picture. After imaging the scene, the solar position is calculated and we are able to perform inverse rendering in the raw colour domain. This domain is chosen as other colour spaces such as  $XYZ$  or  $sRGB$  are non-linearly sensitive to a number of frequencies and the observations are already captured in the raw format. To verify consistency amongst the sky models, we compare the material characteristics they infer to ground truth values found by integrating the reflectance spectra with the spectral response functions for the three channels.

The inverse reflectometry method generates an approximation to the material reflectance spectra as the raw pixel intensity  $P$  is calculated according to Equation 1. Through discretisation:

$$P[k] = \sum_{i=380nm}^{830nm} I(i) R(i) S_k(i), k : R, G, B \quad (6)$$

where  $I$  is the illumination spectra,  $R$  is the reflectance spectra of a material and  $S$  is the sensor sensitivity function. For a particular sensor, the sensor response shows that not all frequencies influence the measurement and so we may approximate the raw intensity by sampling a smaller number of wavelengths with lower bound  $l$  and upper bound  $u$ .

$$P[k] = \sum_{j=l}^u I(j) R(j) S_k(j) \quad (7)$$

Assuming that the reflectance is constant over the region of interest between  $l$  and  $u$ , then:

$$P[k] = R[k] \sum_{j=l}^u I(j) S_k(j), \quad (8)$$

$$R[k] = \frac{P[k]}{\sum_{j=l}^u I(j) S_k(j)}. \quad (9)$$

This means that the reflectance can be approximated by dividing the raw pixel intensity by the illumination in the raw colour domain. The approximation is highly dependent on how wide band the sensor response is and is more suited for hyperspectral sensors as the spectral response is narrow band, allowing the reflectance spectra to be sampled at a higher degree.

Therefore, in order to perform inverse reflectometry for the different colour patches in the scene, we generate a mask that captures the pixels that are associated with a particular object. This mask is used to generate the associated illumination values in the raw colour domain. We then divide the pixel values by the illumination to get a number of reflectance estimates, of which we take the median value. This is because the sampling strategy induces a large amount of noise in the illumination estimates, while there is also noise present in the image itself.

## 4. Results

### 4.1. Camera Characterisation

Figure 4a shows the spectrum of the illumination source as measured by the ASD spectrometer. The two sharp peaks at approximately 546nm and 611nm are indicative of the presence of both mercury and europium in either the lamp used for the experiment, or the external fluorescent lights that were illuminating the area.

All camera images used in these experiments are taken with a shutter speed of  $\frac{1}{200}$  seconds in the raw colour domain. The image of the ColorChecker board is used to generate mean raw values for each colour patch and together with the measured reflectance spectra, we can approximate the sensor response functions. These are shown in Figure 4b and have not been normalised in order to avoid a channel dependent scaling factor being introduced into the calculations. The functions show that the red channel has a wider band than the blue and green channels, and is also less sensitive.

The transformation matrix between *raw* and  $XYZ_{D65}$  is generated by estimating the raw values under a  $D65$  illuminant and calculating the tristimulus  $XYZ_{D65}$  values for each patch as discussed in Section 3.1. The camera dependent transformation matrix  $M_{XYZ_{D65} \rightarrow raw}$  were found to be:

$$\begin{bmatrix} 39.61 & -14.54 & -4.79 & -8.24 & 0.89 \\ -0.05 & 0.19 & 0.02 & -10.11 & 0.92 \\ 0.00 & 0.02 & 0.13 & -2.36 & 0.62 \end{bmatrix}, \quad (10)$$

while  $M_{raw \rightarrow XYZ_{D65}}$  was calculated as:

$$\begin{bmatrix} 0.02 & 0.00 & 0.05 & 0.00 & -0.02 \\ -0.01 & 0.03 & 0.03 & 0.00 & -0.02 \\ 0.00 & -0.03 & 0.25 & 0.00 & -0.09 \end{bmatrix}. \quad (11)$$

### 4.2. Sky Dataset Collection

A total of 18 datasets were captured between June and August at times ranging from mid-morning to afternoon. The azimuth angles were varied between  $12^\circ$  and  $292^\circ$  with all images taken facing away from the sun. The images were taken on what were classified as ‘clear skies’ and an example of this can be seen in Figure 5.

As we are comparing the chromaticity values generated by the sky models with the observations from the experiments, we need to convert the images from *raw* to the  $XYZ_{D65}$  colour space using the empirically derived transformation matrix  $M_{XYZ_{D65} \rightarrow raw}$ . This is further transformed into the  $xyY$  colour space in order to isolate the luminosity  $Y$ , from the two chromaticity channels  $x$  and  $y$ .

Figure 6 shows the errors when comparing the models with the observations. The datasets were grouped based on their azimuth angles and are represented by different markers. The results reveal that the Preetham model errors are highly dependent on the azimuth angle as shown by the varying error curve shapes. On the other hand, the Hosek-Wilkie model is more robust to these changes, with all error curves tightly grouped together. In terms of  $y$  chromaticity, the Hosek-Wilkie model has a non-linear relationship with the elevation angle. Based on these results, the Hosek-Wilkie model has been shown to produce more stable and reliable estimations when compared to the Preetham model, though further improvement could be achieved by focusing on the influence of elevation angle upon  $y$ -chromaticity.

### 4.3. Diffuse Inverse Reflectometry

The ColorChecker Board was imaged inside a dark box with one side left open and facing the sky as shown in Figure 3. The illumination at various points in the scene is calculated by utilising the sky models as an environment map and sampling it 100 times. This simple sampling strategy is not optimal and yields a high amount of noise in the illumination map. As a result of this, we use the median value of the inferred reflectance calculated using Equation

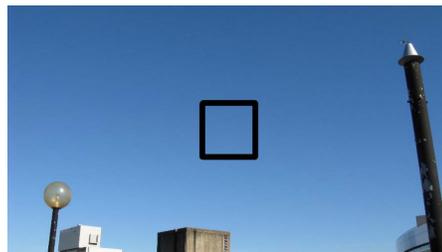
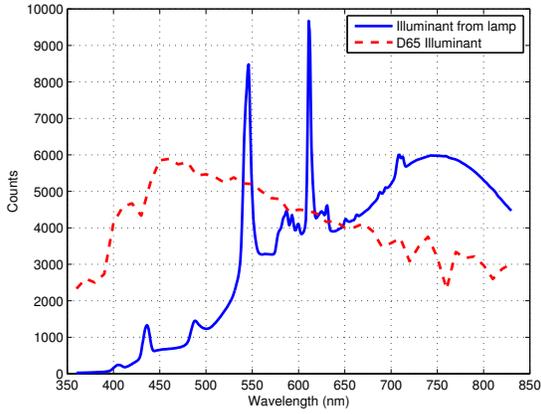
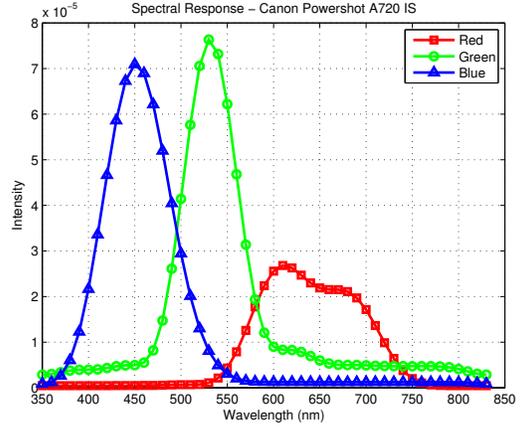


Figure 5: Sample image from experiment 6, elevation 20 degrees. The square indicates the area used to calculate chromaticity of the sky.

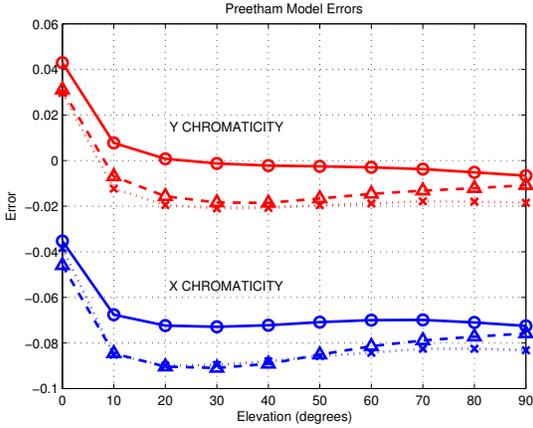


(a) Illumination source and D65 reference illuminant.

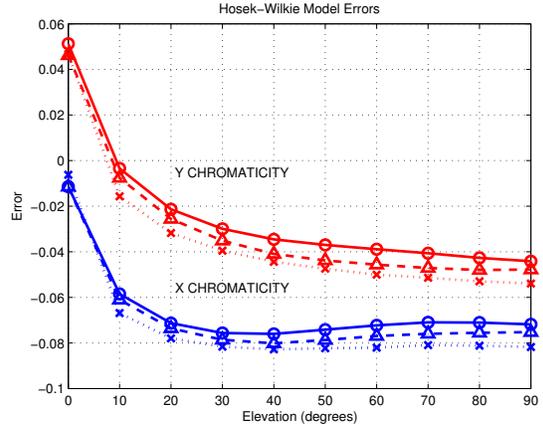


(b) Approximated sensor response functions.

Figure 4: Illumination source spectra and sensor response functions.



(a) Preetham model errors.



(b) Hosek-Wilkie model errors.

Figure 6: Mean errors for each elevation angle between the sky models and the observations. Errors are calculated by subtracting the sky model from the observations. Each marker represents an azimuth group - circles ( $< 82^\circ$ ), triangles ( $82-194^\circ$ ) and crosses ( $> 194^\circ$ ), and the Preetham  $y$ -chromaticity error has been offset by 0.04 for clarity.

9. Figure 7 shows the inferred reflectance for both models for each channel. The inferred reflectance has been scaled by a channel dependent factor  $\alpha_k, k = R, G, B$  calculated by:

$$\alpha_k = [R_{1,t}[k], \dots, R_{24,t}[k]][R_{1,i}[k], \dots, R_{24,i}[24]]^+, \quad (12)$$

where  $R_{j,t}[k]$  and  $R_{j,i}[k]$  are the ground truth and inferred reflectance values respectively for colour patch  $j$  in channel  $k$ . The per-channel scaling values were found to be (1.41, 0.60, 0.40) for the Preetham model, and (1.91, 0.70, 0.44) for the Hosek-Wilkie model.

The scaling values are a combination of multiple factors including a constant scaling factor that exists across all the channels due to the exposure function of the camera. The main error though, exists due to the radiance calibration of the camera, and the luminance estimations made in the models.

The two models provide similar results for all channels which is not unexpected as they both use the Perez *et al.* [15] model as a basis. The inferred reflectance for the red channel shows that the Preetham model outperforms the Hosek-Wilkie model for colour patch 19 (white), but does

not perform as well for colour patch 12 (orange-yellow). The results show that both models fail to detect the drop in

intensity from colour patch 5 to 6 (blue flower and bluish green).

Both the green and blue channels showed extremely similar inferred reflectance values for both the Preetham and Hosek-Wilkie models. Across the channels, it is clear that the last row of the ColorChecker Board (patches 19-24), which represents a decrease in intensity from white to black, is only correctly inferred by both the green and blue channels.

The results show that despite the models not capturing the exact chromaticity of the sky, they can still be used to extract pseudo-reflectance for each channel. One reason for this is the fact that we are using an object parallel to the ground, so illumination from the sky is weighted according to the cosine of the angle between the normal of the object and the angle of incidence, as shown in Equation 5. This gives a higher impact to large elevation angle samples from the environment map.

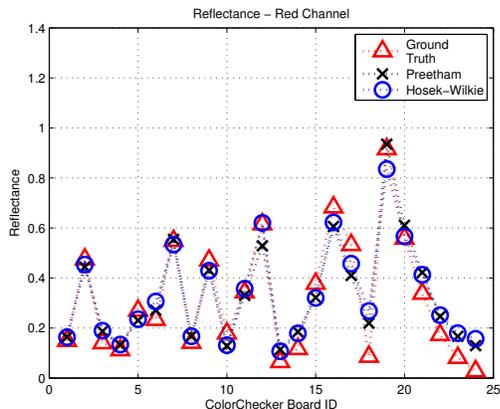
To compare the three dimensional reflectance spectra inferred by the models we use the Spectral Angle Mapper (SAM) metric [19]. This metric is utilised in the remote sensing community for comparing spectral data and calculates the angle between two vectors by evaluating the dot product between them. For the current application of inverse reflectometry, we are generating reflectance spectra that should retain the relative magnitude differences between the colour patches. Therefore, we analyse the SAM error metric across all the colour patches for each channel:

$$SAM[k] = \arccos \frac{R_T[k] \cdot R_I[k]}{|R_T[k]| |R_I[k]|}, \quad (13)$$

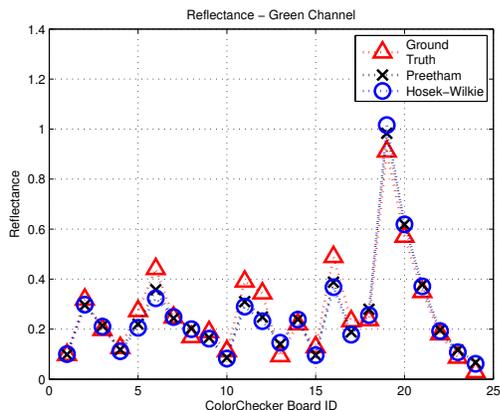
where  $R_T$  is the ground truth reflectance,  $R_I$  is the inferred reflectance and  $k$  is the channel. The results for the SAM error metric are shown in Table 1 and reveal that despite the Hosek-Wilkie model attempting to improve the sky model developed by Preetham, the inferred reflectance errors are slightly higher in the green channel.

## 5. Conclusion

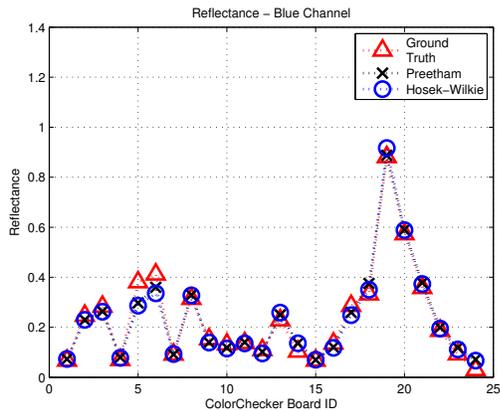
In this paper, a chromaticity analysis of two sky models commonly used in the computer graphic community was presented. The purpose of this was to determine whether



(a) Red channel.



(b) Green channel.



(c) Blue channel.

Figure 7: Comparison between the inferred reflectance for each colour patch using the two sky models and the ground truth.

Model	$SAM_R$	$SAM_G$	$SAM_B$	$\mu$
Preetham	0.17	0.15	0.09	0.14
Hosek-Wilkie	0.17	0.18	0.10	0.15

Table 1: Error between the inferred reflectance and the ground truth measured using the Spectral Angle Mapper measured in radians.

such sky models could be used for image calibration in order to utilise computer vision algorithms using reflectance instead of pixel values. This is extremely important as pixels are dependent on illumination and geometry, while reflectance is characteristic of the material. This has a number of useful applications, especially in remote sensing where thematic mapping of mine faces requires a camera calibration method to take place. We have shown that the Hosek-Wilkie model improves upon the Preetham model by producing robust chromaticity estimates in the presence of azimuth changes. However, there is an angular dependency between the  $y$ -chromaticity estimate and the elevation angles which needs to be investigated in order for the model to be used accurately. Despite these errors, inverse reflectometry for a flat diffuse object in the outdoor environment has been shown to work well for both models in that it is able to retain the relative intensity variations amongst many different colours. Current work includes development of an iterative and spatial propagative system that allows inverse reflectometry to occur for objects ranging from horizontal to vertical orientations, using both RGB and hyperspectral cameras.

## Acknowledgements

This work has been supported by the Rio Tinto Centre for Mine Automation and the Australian Centre for Field Robotics, University of Sydney.

## References

- [1] S. Agarwal, R. Ramamoorthi, S. Belongie, and H. W. Jensen. Structured Importance Sampling of Environment Maps. *ACM Transactions on Graphics (TOG)*, 22(3):605–612, 2003.
- [2] D. L. Bongiorno, M. Bryson, D. G. Dansereau, and S. B. Williams. Spectral characterization of COTS RGB cameras using a linear variable edge filter. In N. Sampat and S. Battiato, editors, *IS&T/SPIE Electronic Imaging. International Society for Optics and Photonics*, volume 8660, Jan. 2013.
- [3] P. Debevec, C. Tchou, A. Gardner, T. Hawkins, C. Poullis, J. Stumpfel, A. Jones, N. Yun, P. Einarsson, T. Lundgren, M. Fajardo, and Philippe Martinez. Estimating surface reflectance properties of a complex scene under captured natural illumination. *ACM Transactions on Graphics*, pages 1–11, 2004.
- [4] P. Dutre, K. Bala, P. Bekaert, and P. Shirley. *Advanced Global Illumination*. AK Peters Ltd, 2006.
- [5] G. Hong, M. R. Luo, and P. A. Rhodes. Colorimetric Characterization Based on Polynomial Modeling. *Color Research & Application*, 26(1):76–84, 2000.
- [6] L. Hosek and A. Wilkie. An analytic model for full spectral sky-dome radiance. *ACM Transactions on Graphics (TOG)*, 31(4):1–9, 2012.
- [7] L. Hosek and A. Wilkie. Adding a Solar-Radiance Function to the Hošek-Wilkie Skylight Model. *Computer Graphics and Applications, IEEE*, 33(3):44–52, 2013.
- [8] J. Jiang, D. Liu, J. Gu, and S. Susstrunk. What is the space of spectral sensitivity functions for digital color cameras? *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, pages 168–179, Jan. 2013.
- [9] M. H. Kim and J. Kautz. Characterization for High Dynamic Range Imaging. *Computer Graphics Forum*, 27(2):691–697, Apr. 2008.
- [10] T. Kollig and A. Keller. Efficient illumination by high dynamic range images. In *Proceedings of the 14th Eurographics workshop on Rendering. Eurographics Association*, number x, pages 45–50. Eurographics Association, 2003.
- [11] T. Machida, H. Takemura, and N. Yokoya. Inverse reflectometry for real objects with diffuse and specular interreflections. *Electronics and Communications in Japan (Part II: Electronics)*, 90(1):50–60, Jan. 2007.
- [12] R. J. Murphy, S. T. Monteiro, and S. Schneider. Evaluating Classification Techniques for Mapping Vertical Geology Using Field-Based Hyperspectral Sensors. *IEEE Transactions on Geoscience and Remote Sensing*, 50(8):3066–3080, Aug. 2012.
- [13] J. Nieto, S. Monteiro, and D. Viejo. 3D geological modelling using laser and hyperspectral data. In *Geoscience and Remote Sensing Symposium (IGARSS), 2010 IEEE International*, pages 4568–4571, 2010.
- [14] T. Nishita, Y. Dobashi, and E. Nakamae. Display of clouds taking into account multiple anisotropic scattering and sky light. In *Proceedings of the 23rd annual conference on Computer graphics and interactive techniques*, pages 379–386. ACM, 1996.
- [15] R. Perez, R. Seals, and J. Michalsky. All-weather model for sky luminance distribution preliminary configuration and validation. *Solar energy*, 50:235–245, 1993.
- [16] T. W. Pike. Using digital cameras to investigate animal colouration: estimating sensor sensitivity functions. *Behavioral Ecology and Sociobiology*, 65(4):849–858, 2010.
- [17] a. J. Preetham, P. Shirley, and B. Smits. A practical analytic model for daylight. *Proceedings of the 26th annual conference on Computer graphics and interactive techniques - SIGGRAPH '99*, pages 91–100, 1999.
- [18] I. Reda and A. Andreas. Solar position algorithm for solar radiation applications. *Solar energy*, 76.5:577–589, 2004.
- [19] R. A. Schowengerdt. *Remote Sensing: Models and Methods for Image Processing*. Academic Press, third edition, 2006.
- [20] O. Wang, P. Gunawardane, S. Scher, and J. Davis. Material classification using BRDF slices. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2805–2811, June 2009.
- [21] G. Zotti, A. Wilkie, and W. Purgathofer. A critical review of the Preetham skylight model. *WSCG 2007 Short Communications Proceedings I*, pages 23–30, 2007.