

An Optimisation Approach to the Recovery of Reflection Parameters from a Single Hyperspectral Image

Sejuti Rahman¹ * and Antonio Robles-Kelly^{1,2}

¹School of Engineering, Australian National University, Canberra ACT 0200, Australia

²National ICT Australia (NICTA) [†]Locked Bag 8001, Canberra ACT 2601, Australia

Abstract

In this paper, we present a method to recover the parameters governing the reflection of light from a surface making use of a single hyperspectral image. To do this, we view the image radiance as a combination of specular and diffuse reflection components and present a cost functional which can be used for purposes of iterative least squares optimisation. This optimisation process is quite general in nature and can be applied to a number of reflectance models widely used in the computer vision and graphics communities. We elaborate on the use of these models in our optimisation process and provide a variant of the Beckmann-Kirchhoff model which incorporates the Fresnel reflection term. We show results on synthetic images and illustrate how the recovered photometric parameters can be employed for skin recognition in real world imagery, where our estimated albedo yields a classification rate of $95.09 \pm 4.26\%$ as compared to an alternative, whose classification rate is of $90.94 \pm 6.12\%$. We also show quantitative results on the estimation of the index of refraction, where our method delivers an average per-pixel angular error of 0.15 degrees. This is a considerable improvement with respect to an alternative, which yields an error of 9.9 degrees.

Keywords: Hyperspectral imaging, reflectance modelling, reflection parameter recovery

*Corresponding author. E-mail: sejuti.rahman@nicta.com.au Tel: +61(2)6267-6228

[†]NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

1 Introduction

In this paper, we focus on the recovery of the parameters governing the reflection of light from objects in the scene as captured by imaging spectroscopy sensors. The modeling of surface reflectance is a topic that is of pivotal importance, and has hence attracted considerable effort in both computer vision and computer graphics.

In graphics, the problem is of interest since it allows physically realistic images of synthetic surfaces to be generated. In computer vision, the modeling of surface reflectance has attracted considerable attention due to its relevance to scene analysis and image understanding. For instance, Nayar and Bolle [33] have used photometric invariants derived from the BRDF to recognise objects with different reflectance properties. This work builds on that reported in [32], where a background to foreground reflectance ratio is introduced. In a related development, Dror *et al.* [9] have shown how surfaces may be classified from single images through the use of reflectance properties.

Estimation of reflectance parameters is also closely related to the problem of recovering the shape of an object from its shading information. The classic approaches to shape from shading developed by Ikeuchi and Horn [21], and by Horn and Brooks [15], hinge on the compliance with the image irradiance equation and local surface smoothness. Zheng and Chellappa [55] proposed a gradient consistency constraint that penalises differences between the image intensity gradient and the surface gradient for the recovered surface. Moreover, although shape-from-shading [15] usually relies on the assumption of Lambertian reflectance [25], photometric correction or specular subtraction may be applied as a preprocessing step to improve the results obtained. For instance Brelstaff and Blake [6] used a simple thresholding strategy to identify specularities on moving curved objects. Other lines of research remove specularities by either using additional hardware [33], imposing constraints on the input images [28], requiring color segmentation [24] as post-processing steps, or using reflectance models to account for the distribution of image brightness [39].

Alternatively, the BRDF can be used to compute a reflectance map, to which a number of shape-from-shading methods may be applied [16, 14]. The methods used to model or approximate the BRDF can be divided into those that are physics-based, semi-empirical [37, 53] or empirical [25, 38, 45] in nature. Although the literature from physics is vast, it is perhaps the work of Beckmann on smooth and rough surface reflectance that is better known in the vision and graphics communities [2]. While it is based on physically meaningful surface parameters, the Beckmann

theory relies on the evaluation of the Kirchhoff wave scattering integral and breaks down when the surface roughness or the scattering angle are large. To overcome this problem, Vernold and Harvey [51] have developed a model which accounts for self shadowing on rough surfaces. Ragheb and Hancock [40] have exploited this modification of the Beckman-Kirchoff theory to develop a means of measuring surface roughness parameters using reflectance. Another widely used physics-based model that accounts for specular reflectance by modeling the angular distribution of surface microfacets is that in [49]. A survey of reflectance and shading models can be found in [43].

For purposes of estimating the reflectance properties and light source direction of a scene, Ikeuchi et al. [22] have proposed an iterative least square optimisation based on a simplified Torrance-Sparrow reflectance model. The algorithm in [22] recovers the surface reflectance properties and light source direction from a range and a brightness image. With known geometry of the scene, Sato et al. [41] have utilised brightness values within shadows to solve a system of equations to simultaneously recover reflectance parameters and illumination distribution. Tomimaga et al. [48] have proposed a method for estimating various reflection parameters of the Phong reflectance model [38] along with shape, illumination direction and color from a single image. However, all of these methods assume that the light source and viewing positions are located at infinity. Therefore these methods do not offer a solution for real scenes under light and viewing positions at finite distances which do not follow parallel illumination and orthogonal projection. Later, based on the simplified Torrance-Sparrow reflectance model, Hara et al. [13] have developed a method without the distant illumination and orthogonal projection assumption that estimates reflectance property and light source position from a single image. However, their method still requires prior knowledge of the 3D shape of the scene that narrows down the applicability of this approach. Following a similar approach, the method by Boivin et al. [4] recovers reflectance properties of different surfaces from a single image. With known illumination, camera properties, and 3D geometric model, this method iteratively fits multiple reflection models from the simplest to the most complex one until the error between original and estimated image becomes less than a predefined threshold.

We note that a number of models in the literature employ object shape, index of refraction, reflectance, and a measure of surface roughness to embody the surface reflectance. These models [2, 49, 7, 53] also incorporate a Fresnel term [5]. Moreover, the reflectance and index of refraction are, in general, wavelength dependant. This leads to the natural extension of these techniques to multispectral and hyperspectral image understanding. However, the recovery of reflection parame-

ters from a single hyperspectral image is relatively under researched. The method in [19], based on Dichromatic reflectance model, recovers the illuminant power spectrum and Dichromatic parameters in a structural optimisation setting. However, this method does not consider the wave-length dependency of the specular reflection. In [18], an iterative optimisation approach is presented. The method is applicable to a number of reflectance models [2, 49, 53] and able to recover shape and reflectance parameters from a single image with known illumination direction. Both of these approaches are based on a distant light source assumption and, hence, can break down when applied to real scenes under light and viewing positions at finite distances. Furthermore, in [18], the wave-length dependency of diffuse albedo is ignored and the surface reflectance is not modelled as a composition of both specular and diffuse components but rather the focus is in the optimisation of the model under consideration using Euler-Lagrange equations.

Here, we present a method to recover the reflectance model parameters as a function of wave-length. We do not require prior knowledge on the geometry of the scene nor assume a point light source at infinity. Given a known illumination direction and a single image at input, we depart from a general formulation that hinges in the notion that the light reflected from an object can be deemed to be either specular or diffuse. This permits the recovery of the reflection parameters through an optimisation approach, whose target function is a linear combination of both, the specular and the diffuse reflection components. This, together with the use of Cauchy’s and Sellmeier’s expansions [5] permits the recovery of the index of refraction governing the Fresnel term of models such as Cook and Torrance’s [7], Torrance and Sparrow’s [49], Beckmann-Kirchoff’s [2] and Wolff *et al.*’s [53].

The paper is organised as follows. We commence by expressing the object reflection as a linear combination of a diffuse and a specular component. With this linear combination at hand, we proceed to formulate the problem in an optimisation setting in Section 3. Further discussion is provided in Section 4, where we show how the linear combination presented in Section 2 can be expressed in terms of models elsewhere in the literature. Here, we focus on the use of those models in [7, 49, 2] and [53]. We describe implementation issues in Section 5 and present results for our method on synthetic data and real-world imagery in Section 6, where we also compare against a number of alternatives. Finally, in Section 7, we conclude on the developments presented here.

2 Recovering Reflectance Model Parameters

As mentioned earlier, our aim is to recover the reflectance model parameters. Nonetheless, there are a wide variety of reflectance models, we can depart from a generalist view of the reflection process based upon a physical interpretation, where the geometry of the scene, photogrammetric variables and wavelength dependent Fresnel term are used to describe the image formation process. From this general viewpoint, we can then cast the recovery of the reflection parameters in an optimisation setting.

2.1 General Reflectance Model

Here, we adopt a quite general approach applicable to a number of reflectance models in the existing literature [53, 49, 2, 7]. We formulate the general reflectance model as a linear combination of two reflection components, namely diffuse and specular. This concept of modelling the scene radiance as a linear combination of two components was formally proposed by Shafer through the Dichromatic reflectance model [45]. Furthermore, this is consistent with a number of references elsewhere in the literature [53, 18, 7].

It is worth noting in passing that optimisation approaches have been used in computer vision to recover the illuminant [10] and the surface normals [15]. As per the recovery of the reflectance, Fresnel reflection and other parameters, Weyrich *et al.* [52] have fit a BRDF model to the acquired data. In [18], the authors use Euler-Lagrange equations to recover the reflectance parameters corresponding to a number of reflectance models. These approaches follow the assumption that, given enough data, the reflectance parameters may be recovered accordingly. Note that, in our case, hyperspectral imaging does provide an information-rich representation of the scene which delivers a large amount of data and allows the wavelength dependence to be introduced into the optimisation process. This is similar rationale to that used in Weyrich *et al.*'s [52] work, where large amounts of trichromatic data is acquired using parallel hardware for purposes of fitting.

Here, we further express the diffuse term and specular term as a function of reflection angle variables, photogrammetric variables and wavelength dependent index of refraction. Note that, a number of reflectance models follow a local coordinate system which corresponds to the reflection geometry given in the left-hand panel of Figure 1. This geometry is defined with respect to a local coordinate system whose origin is the surface location and where the z' -axis is aligned to the normalised local surface normal \vec{N} . In the figure, we have denoted the unit incident light direction




Figure 1: Left-hand panel: Geometry of reflection showing the angular variables in a local coordinate system as used in Equation 1; Right-hand panel: Global coordinate system corresponding to the reparameterisation in Equation 2.

vector \vec{L} , with slant θ_i and tilt ϕ_i . Also, following the costume in the literature, we denote the unit view vector as \vec{V} , whose slant is θ_s and tilt is ϕ_s . The unit half-way vector is denoted by \vec{H} and is given by the angular bisector of \vec{L} and \vec{V} , i.e. $\vec{H} = \frac{\vec{L} + \vec{V}}{|\vec{L} + \vec{V}|}$. We denote the angle between \vec{N} and \vec{H} as θ_h .

Making use of the reflection geometry in Figure 1, the illuminant power spectrum $L(\lambda_l)$, the image radiance at pixel u and wavelength λ_l is given by,

$$f(u, \lambda_l) = W_{diff}(u) f_{diff}(L(\lambda_l), \Theta(u), \rho(u, \lambda_l), \eta(u, \lambda_l)) + W_{spec}(u) f_{spec}(L(\lambda_l), \Theta(u), \sigma_m(u), \eta(u, \lambda_l)) \quad (1)$$

In the equation above, the first term on the right-hand-side corresponds to the diffuse reflection component where $f_{diff}(\cdot)$ stands for the diffuse image radiance and W_{diff} denotes the weight of the diffuse term. The second term on the right-hand-side describes the specular component where $f_{spec}(\cdot)$ denotes specular image radiance and W_{spec} denotes the weight of the specular term. In Equation 1, $f_{diff}(\cdot)$ is expressed as a function of the reflection angular variable-set $\Theta(u) = \{\theta_i, \phi_i, \theta_s, \phi_s, \theta_h\}$, wavelength dependent diffuse albedo $\rho(u, \lambda_l)$, and index of refraction $\eta(u, \lambda_l)$. On the other hand, $f_{spec}(\cdot)$ is expressed as a function of the reflection angular variable-set $\Theta(u)$, microfacet slope $\sigma_m(u)$, and wavelength-dependent refracted index $\eta(u, \lambda_l)$.

Further, if the illumination direction is known, we can replace the reflection angle variables by

surface gradients and reparameterise the general reflectance model with respect to surface gradients. The benefit of doing this is that such a reparameterisation implies that the number of geometric variables is two per pixel, i.e. $p(u)$ and $q(u)$. This reduces the number of reflectance equations needed to estimate the surface normal \vec{N} at pixel u . This reparameterisation is effected by making use of the viewer-centered coordinate system shown in the right-hand panel of Figure 1. As a result, the (x, y) -axes correspond to the image row and column directions. This implies that, by definition, the surface normal \vec{N} at u is $[-p(u), -q(u), 1]^T$. This yields the reparameterised image radiance given by

$$f(u, \lambda_l) = W_{diff}(u)f_{diff}(L(\lambda_l), p(u), q(u), \rho(u, \lambda_l), \eta(u, \lambda_l)) \\ + W_{spec}(u)f_{spec}(L(\lambda_l), p(u), q(u), \sigma_m(u), \eta(u, \lambda_l)) \quad (2)$$

In such a coordinate system, the z -axis is lined-up with the optical axis and the positive x -axis points toward the right-hand side of the field of view. It is worth noting that this reparameterisation can be viewed as a rotation of the local coordinate system which aligns the viewer direction with the z -axis making use of the rotation matrix \mathbf{R} as follows

$$z = \mathbf{R}z' \quad (3)$$

Reflecting the notion that, in the local coordinate system, the surface normal is lined up with the z' -axis, whereas in the global coordinate system, the viewer is aligned with the z -axis we can rewrite Equation 3 as follows

$$[0, 0, 1]^T = \frac{1}{\sqrt{p^2 + q^2 + 1}} \mathbf{R}[-p(u), -q(u), 1]^T \quad (4)$$

where we have assumed a normalised surface normal and used the fact that, in our global coordinate system, the viewer direction becomes $[0, 0, 1]^T$.

2.2 Target Function

Given the illumination direction and a single hyperspectral image at input, we aim at recovering the reflectance parameters, i.e. index of refraction, albedo, microfacet slope, and surface normals. In our formulation, we depart from Equation 2 in order to express the diffuse and specular term of the image radiance as a function of the surface gradients, photogrammetric variables (microfacet slope and wavelength dependent diffuse albedo) and wavelength dependent index of refraction. We also recover the weights $W_{diff}(u)$ and $W_{spec}(u)$.

Our method is an iterative one which employs a cost function $C(\mathcal{I})$ defined over the hyperspectral image \mathcal{I} . Here, we have expressed our cost function as the summation of two other cost functionals $C_{diff}(\mathcal{I})$ and $C_{spec}(\mathcal{I})$, where C_{diff} is couched as the squared difference between the image diffuse spectral radiance \mathcal{I}_{diff} and the first term on the right-hand-side of Equation 2. The other cost functional, denoted by C_{spec} , corresponds to the squared difference between the image specular radiance \mathcal{I}_{spec} and the last term on the right-hand-side of Equation 2.

Thus, our cost function is given by

$$C(\mathcal{I}) \triangleq C_{diff}(\mathcal{I}) + C_{spec}(\mathcal{I}) \quad (5)$$

where

$$C_{diff}(\mathcal{I}) \triangleq \sum_{u \in \mathcal{I}} \sum_{\lambda_l \in \mathcal{W}} \left[\mathcal{I}_{diff}(u, \lambda_l) - W_{diff}(u) f_{diff}(L(\lambda_l), p(u), q(u), \rho(u, \lambda_l), \eta(u, \lambda_l)) \right]^2 \quad (6)$$

can be viewed as a data term corresponding to the diffuse image radiance and

$$C_{spec}(\mathcal{I}) \triangleq \sum_{u \in \mathcal{I}} \sum_{\lambda_l \in \mathcal{W}} \left[\mathcal{I}_{spec}(u, \lambda_l) - W_{spec}(u) f_{spec}(L(\lambda_l), p(u), q(u), \sigma_m(u), \eta(u, \lambda_l)) \right]^2 \quad (7)$$

is the converse for the specular component of \mathcal{I} . In the equations above, \mathcal{W} is the wavelength range for the imagery under consideration.

Note that, by making use of the linear combination in Equation 5, we can separate the image radiance into its diffuse and specular components using a specular removal method such as those in [47] or [28]. This separation responds to the observation that, in Equation 2, the albedo $\rho(\cdot)$ and W_{diff} contribute only to the diffuse component, whereas σ_m and W_{spec} contribute only to the specular component. Moreover, the index of refraction $\eta(\cdot)$ and surface gradients $p(u)$, $q(u)$ influence both the specular and diffuse terms.

Therefore, we can adopt an optimisation approach where our aim of computation is the recovery of $\rho(\cdot)$ and W_{diff} by maximising C_{diff} such that

$$(\rho^*(u, \lambda_l), W_{diff}^*(u)) \leftarrow \operatorname{argmin}_{\rho(u, \lambda_l), W_{diff}(u)} C_{diff}(\mathcal{I}) \quad (8)$$

In a similar manner, we note that the microfacet slope $\sigma_m(u)$ and weight $W_{spec}(u)$ can be recovered by optimising $C_{spec}(\mathcal{I})$ as follows

$$(\sigma_m^*(u), W_{spec}^*(u)) \leftarrow \operatorname{argmin}_{\sigma_m(u), W_{spec}(u)} C_{spec}(\mathcal{I}) \quad (9)$$

Finally, if the weights $W_{spec}^*(u)$ and $W_{diff}^*(u)$, albedo and microfacet slope are available, the index of refraction $\eta(\cdot)$ and surface gradients $p(u)$, $q(u)$ can be computed by optimising the combined cost function C . This can be expressed as

$$(\eta^*(u, \lambda_l), p^*(u), q^*(u)) \leftarrow \operatorname{argmin}_{\eta(u, \lambda_l), p(u), q(u)} C(\mathcal{I}) \quad (10)$$

Thus, by expressing Equation 5 in terms of $C_{spec}(\cdot)$ and $C_{diff}(\cdot)$, we obtain twice as many data terms for the estimation of the index of refraction and the surface shape. Furthermore, we can estimate the albedo $\rho(\cdot)$ and the weight W_{diff} by extremising the cost function in Equation 6 alone. Analogously, the microfacet slope σ_m and the weight W_{spec} can be estimated making use of Equation 7. Since these two sets of parameters can be estimated independently, this treatment also opens-up the possibility of effecting these two optimisation steps in parallel and, once the optimal values of σ_m , W_{spec} , $\rho(\cdot)$, and W_{diff} are in hand, the index of refraction $\eta(\cdot)$ and the surface normal $\vec{N} = [-p(u), -q(u), 1]^T$ at pixel u can be estimated.

3 Optimisation Procedure

Here, we opt for a coordinate descent [12] optimisation scheme. This approach is shared with other methods elsewhere in the literature [19, 13, 41], where the complexity of the multivariable optimisation problem is tackled by optimising comparatively simple scalar subproblems.

The idea is to minimise the multivariable cost function $C(\mathcal{I})$ in Equation 5 by minimising it along one coordinate direction at a time while keeping all other coordinates fixed. Thus, the multivariable problem is actually solved by computing a sequence of scalar minimisation subproblems, where each of these is either a linear or a non-linear least square optimisation depending on whether or not the cost function is linear in the coordinate space, i.e. the variable under consideration. This provides the additional advantage that, for the linear least square subproblems, the cost function is variable-wise convex and, thus, a globally optimal solution is guaranteed. For the non-linear least squares, the measure of non-linearity and the ability to chose an initial estimate close to the optimal value effectively determines whether the method will converge to a global minimum.

The algorithm thus proceeds as follows. At each iteration it recovers $\rho(u, \lambda_l)$, $W_{diff}(u)$, $\sigma_m(u)$, $W_{spec}(u)$, $\eta(u, \lambda_l)$ and $\{p(u), q(u)\}$ in interleaved separate minimisation steps. Once it estimates one parameter, the current estimate is used to obtain the remaining parameters. This is, in effect, an update process where the iterations proceed until convergence is reached. The algorithm terminates

when none of the parameters change between two successive iterations by an amount greater or equal to a predefined threshold and/or when maximum number of iteration is completed.

For the least square optimisation involved we have selected the Levenberg–Marquardt method [27, 29]. This method is actually a combination of the gradient descent and the Gauss–Newton method and efficiently utilises the complimentary advantages of these two methods. Being controlled by the value of a damping parameter, the method acts more like a gradient descent method when the parameters are far from their optimal value, whereas, it acts more like the Gauss–Newton method when the parameters are close to their optimal values. The method adaptively controls its own damping and this feature makes it applicable to a wide variety of problems. It increases the damping if a step fails to reduce the error; otherwise it decreases the damping. Thus it defensively navigates a region of parameter space in which the model is highly nonlinear [26]. For updating the damping parameter, we adopt the strategy proposed in [36]. This is as it assures faster convergence than that originally proposed by Marquardt [29]. More importantly, it offers protection against a rank-deficient Jacobian and, therefore, it prevents the errors caused by singular Hessian matrices often found when the parameter space is highly nonlinear.

In the following subsections, we elaborate further on the step sequence of our algorithm, which, for the reader’s reference, we have summarised in the pseudo code shown in Algorithm 1. At input, our algorithm takes a hyperspectral image \mathcal{I} whose pixel values correspond to the spectral radiance $I(u, \lambda_l)$ for each of the wavelength-indexed bands $\{\lambda_1, \dots, \lambda_{|\mathcal{W}|}\}$.

At this point, it is worth noting in passing that, following [15], we can write the image radiance as the product of the light power spectrum and the reflectance $R(u, \lambda_l)$, i.e.

$$I(u, \lambda_l) = L(\lambda_l)R(u, \lambda_l) \quad (11)$$

The expression above has been used widely in the literature [23] and is consistent with reflectance models elsewhere, such as that in [10]. Indeed, this treatment allows for the estimation of the illuminant power spectrum $L(\lambda_l)$ *a priori*, which we estimated using the method in [19], so as to recover the image reflectance $R(u, \lambda_l)$. With the image reflectance in hand, we separate $R(u, \lambda_l)$ into the diffuse and specular reflectance $R_{diff}(u, \lambda_l)$ and $R_{spec}(u, \lambda_l)$ using the method presented in [47]. This can be done without any loss of generality due to the manner in which we have expressed the image radiance in Equation 2, which is consistent with the relation $R(u, \lambda_l) = R_{diff}(u, \lambda_l) + R_{spec}(u, \lambda_l)$. These two steps are listed under the preprocessing tag in Algorithm 1.

Here, we would like to stress that we use one of existing specular removal methods as a

preprocessing tool. This is not exclusive to our method but rather common across other reflectance parameter estimation approaches [13, 31]. Specularity removal is often used to decompose the image radiance into specular and diffuse reflection. Thus, there has been a considerable amount of research on separating reflectance components. This literature is not the focus of the work presented here and, therefore, we refer the interested reader to the survey in [1].

Note that, here, we employ different reflectance models in order to recover the relevant parameters. These are summarised in Table 1. The accuracy of these is expected to depend on the diffuse-specular separation. Nonetheless, our method is not overly affected by errors on the specularity removal. This is as the only parameter which is exclusive to the specular pixels is the microfacet slope $\sigma_m(u)$, which is shared by both, the Torrance-Sparrow [49] and Beckmann-Kirchhoff [2] models. This, de facto, constrains the estimation further. This is an important observation since diffuse pixels tend to be more abundant in images as compared to specular ones. Furthermore, in our experience, the errors in the specularity removal tend to be more evident at grazing angles, which does not overly affect the estimation of the albedo and index of refraction.

Broadly speaking, the reflectance component separation methods can be categorised into two classes, i.e. color-based and polarization-based methods. Color-based methods can not estimate the diffuse component when the object is white whereas polarization-based methods are only applicable to materials whose molecules are isotropically arranged, e.g. plastics, paints, papers, metals, woods, clothes, glasses, or liquids. Moreover, in practice, polarizers often fail to remove strong specularities in objects made of materials such as metals, glasses, or liquids [30]. As a result, here we adopt the color-based component separation technique in [47]. Although this method is based on trichromatic image, it is quite straightforward to apply this to a hyperspectral image. This method requires illumination spectrum normalisation, so the spectra of illumination, $L(\lambda_l)$ need to be known beforehand. For us, it is not a problem as we are estimating $L(\lambda_l)$ at the preprocessing step of our algorithm. We then normalise image radiance by dividing it by illuminant spectrum for all $\lambda_l \in \mathcal{W}$ and provide the method the normalised image intensity value as input. Note that other specularity removal methods elsewhere in the literature may also be used.

To initialise the surface normal parameters $p(u)$ and $q(u)$, we follow Worthington and Hancock [54] and make use of the image gradients. Recall that, in [54], the parameters $p(u)$ and $q(u)$ are initialised making use of a gray-scale image. This is in contrast with our imagery, which is wavelength indexed. To tackle this problem, we note that, from robust statistics [17], it can be shown that, if the estimation errors are normally distributed with zero-mean, the expected value

can be used as an unbiased estimator. Thus, we have

$$p(u) = \frac{1}{|\mathcal{W}|} \sum_{l=1}^{|\mathcal{W}|} \frac{\partial R_{diff}(u, \lambda_l)}{\partial x} \quad \text{and} \quad q(u) = \frac{1}{|\mathcal{W}|} \sum_{l=1}^{|\mathcal{W}|} \frac{\partial R_{diff}(u, \lambda_l)}{\partial y} \quad (12)$$

The rest of the variables are initialised with random values in the interval $[0, 1]$.

3.1 Recovery of Diffuse Reflectance Parameters

As mentioned earlier, the structure of the cost function in Equation 5, permits the recovery of albedo $\rho(u, \lambda_l)$ and diffuse weight W_{diff} , making use of the diffuse radiance devoid of the specular reflection. As a result, in this subsection, we focus on these two parameters, whereas in Subsections 3.2 and 3.3, we turn our attention to other photometric variables and object's shape.

3.1.1 Recovery of Diffuse Albedo

To commence, we recover $\rho(u, \lambda_l)$ at each pixel $u \in \mathcal{I}$ and wavelength $\lambda_l \in \mathcal{W}$ given the current estimate of index of refraction $\eta(u, \lambda_l)$, diffuse weight $W_{diff}(u)$, and surface gradients $p(u), q(u)$ as recovered at the previous iteration. This corresponds to Lines 3-7 of Algorithm 1. As the albedo only contributes to the diffuse component of reflectance, we estimate it making use of the diffuse term in cost function in Equation 6. Moreover, to recover $\rho(u, \lambda_l)$ at each pixel u and wavelength λ_l , we optimise over the neighborhood Ω_u of u , this yields

$$\rho^*(u, \lambda_l) = \underset{\rho(u, \lambda_l)}{\operatorname{argmin}} \sum_{v \in \Omega_u} \left[R_{diff}(v, \lambda_l) - \frac{W_{diff}(v)}{L(\lambda_l)} f_{diff}(L(\lambda_l), p(v), q(v), \rho(v, \lambda_l), \eta(v, \lambda_l)) \right]^2 \quad (13)$$

where we have introduced the illuminant power spectrum $L(\lambda_l)$ so as to express the squared difference above in terms of diffuse reflectance as an alternative to the diffuse image radiance. Hereafter, we will employ this formalism for the sake of consistency.

Furthermore, note that, for each material in the image under consideration, the diffuse albedo is uniform across the scene. Thus, after recovering $\rho^*(u, \lambda_l)$ for all the pixels in the image, we cluster the corresponding diffuse albedo into $\{r_1, \dots, r_c\}$ regions. This implies that each region r represents a unique material with unique diffuse albedo. By assuming normally distributed estimation errors with zero-mean, we can further improve our estimate of $\rho(u, \lambda_l)$ in Lines 8-10 of Algorithm 1 by taking the expected value over all pixel $u \in r$ as given by

$$\rho(u, \lambda_l) = \frac{1}{|r|} \sum_{u \in r} \rho^*(u, \lambda_l) \quad (14)$$

Algorithm 1 Estimating reflectance model parameters

Input: The hyperspectral input image \mathcal{I} whose pixel values correspond to the measured spectral radiance $I(u, \lambda_l)$ indexed to wavelength $\lambda_l \in \mathcal{W}$.

Output: $\{\rho(u, \lambda_l), \eta(u, \lambda_l), \sigma_m(u), W_{spec}(u), W_{diff}(u), p(u), q(u)\} \mid \forall u \in \mathcal{I}, \lambda_l \in \mathcal{W}$

Preprocessing: (i) Estimate the illuminant power spectrum $L(\lambda_l)$ using the method in [19] so as to compute $R(u, \lambda_l)$ from $I(u, \lambda_l)$ using Equation 11.
(ii) Separate $R(u, \lambda_l)$ into $R_{diff}(u, \lambda_l)$ and $R_{spec}(u, \lambda_l)$ using the method in [47].

```
1:  $t \leftarrow 1$ 
2: while true do
3:   for all  $u \in \mathcal{I}$  do
4:     for all  $\lambda_l \in \mathcal{W}$  do
5:        $\rho^*(u, \lambda_l) \leftarrow \operatorname{argmin}_{\rho(u, \lambda_l)} C_{diff}(\mathcal{I}) \mid_{W_{diff}(u)^{t-1}, \eta(u, \lambda_l)^{t-1}, p(u)^{t-1}, q(u)^{t-1}}$ 
6:     end for
7:   end for
8:   for all  $r \in \mathcal{I}$  do
9:      $\rho(u, \lambda_l)^t \leftarrow \frac{1}{|r|} \sum_{u \in r} \rho^*(u, \lambda_l), \forall u \in r, \lambda_l \in \mathcal{W}$ 
10:  end for
11:  for all  $u \in \mathcal{I}$  do
12:     $W_{diff}(u)^t \leftarrow \operatorname{argmin}_{W_{diff}(u)} C_{diff}(\mathcal{I}) \mid_{\rho(u, \lambda_l)^t, \eta(u, \lambda_l)^{t-1}, p(u)^{t-1}, q(u)^{t-1}}$ 
13:  end for
14:  for all  $u \in \mathcal{S}$  do
15:     $\sigma_m(u)^t \leftarrow \operatorname{argmin}_{\sigma_m(u)} C_{spec}(\mathcal{I}) \mid_{W_{spec}(u)^{t-1}, \eta(u, \lambda_l)^{t-1}, p(u)^{t-1}, q(u)^{t-1}}$ 
16:  end for
17:  for all  $u \in \mathcal{S}$  do
18:     $W_{spec}(u)^t \leftarrow \operatorname{argmin}_{W_{spec}(u)} C_{spec}(\mathcal{I}) \mid_{\sigma_m(u)^t, \eta(u, \lambda_l)^{t-1}, p(u)^{t-1}, q(u)^{t-1}}$ 
19:  end for
```

```

20:   for all  $u \in \mathcal{I}$  do
21:       for all  $\lambda_l \in \mathcal{W}$  do
22:            $\eta^*(u, \lambda_l) \leftarrow \operatorname{argmin}_{\eta(u, \lambda_l)} C(\mathcal{I}) \mid_{\rho(u, \lambda_l)^t, W_{diff}(u)^t, \sigma_m(u)^t, W_{spec}(u)^t, p(u)^{t-1}, q(u)^{t-1}}$ 
23:       end for
24:   end for

25:   for all  $u \in \mathcal{I}$  do
26:        $C_k(u)^t = \operatorname{argmin}_{C_k(u)} \sum_{\lambda_l \in \mathcal{W}} \left( \eta^*(u, \lambda_l) - \sum_{k=1}^M C_k(u)^{t-1} \lambda_l^{-2(k-1)} \right)^2$ 
27:       for all  $\lambda_l \in \mathcal{W}$  do
28:            $\hat{\eta}(u, \lambda_l)^t = \sum_{k=1}^M C_k(u)^t \lambda_l^{-2(k-1)}$ 
29:       end for
30:   end for

31:   for all  $r \in \mathcal{I}$  do
32:        $\eta(u, \lambda_l)^t \leftarrow \frac{1}{|r|} \sum_{u \in r} \hat{\eta}(u, \lambda_l), \forall u \in r, \lambda_l \in \mathcal{W}$ 
33:   end for

34:   for all  $u \in \mathcal{I}$  do
35:        $[p(u)^t, q(u)^t] \leftarrow \operatorname{argmin}_{p(u), q(u)} C(\mathcal{I}) \mid_{\rho(u, \lambda_l)^t, W_{diff}(u)^t, \sigma_m(u)^t, W_{spec}(u)^t, \eta(u, \lambda_l)^t}$ 
36:   end for

37:    $converge \leftarrow \sum_{u \in \mathcal{I}} \sum_{\lambda_l \in \mathcal{W}} \delta(\rho(u, \lambda_l)) \leq \varepsilon \text{ and } \sum_{u \in \mathcal{I}} \sum_{\lambda_l \in \mathcal{W}} \delta(\eta(u, \lambda_l)) \leq \varepsilon \text{ and } \sum_{u \in \mathcal{I}} \delta(\sigma_m(u)) \leq \varepsilon$ 
38:    $\text{and } \sum_{u \in \mathcal{I}} \delta(W_{diff}(u)) \leq \varepsilon \text{ and } \sum_{u \in \mathcal{I}} \delta(W_{spec}(u)) \leq \varepsilon \text{ and } \sum_{u \in \mathcal{I}} \delta(\{p(u), q(u)\}) \leq \varepsilon$ 

39:   if  $converge$  or  $t > t_{max}$  then
40:       break
41:   else
42:        $t \leftarrow t + 1$ 
43:   end if
44: end while

45: return  $\rho(u, \lambda_l), \eta(u, \lambda_l), \sigma_m(u), W_{spec}(u), W_{diff}(u), p(u), q(u)$ 

```

over all the material regions in the scene.

3.1.2 Recovery of Diffuse Weight

As mentioned earlier, we also recover the diffuse weight W_{diff} at each pixel $u \in \mathcal{I}$. This is done making use of the optimal value of diffuse albedo $\rho(u, \lambda_l)$ yielded by the optimisation in Equations 13 and 14 and the other parameters, i.e. the index of refraction $\eta(u, \lambda_l)$ and surface gradients $p(u)$ and $q(u)$ recovered at the previous iteration. This process is outlined in Lines 11-13 of Algorithm 1.

The diffuse weight determines the contribution of the diffuse reflectance in the linear combination in Equation 2. We, therefore, estimate this parameter by optimising the cost function in Equation 6 making use of the fact that the diffuse weight should be fixed over the wavelength domain. Thus, our optimisation is effected per-pixel such that

$$W_{diff}^*(u) = \underset{W_{diff}(u)}{\operatorname{argmin}} \sum_{l=1}^{|\mathcal{W}|} \left[R_{diff}(u, \lambda_l) - \frac{W_{diff}(u)}{L(\lambda_l)} f_{diff}(L(\lambda_l), p(u), q(u), \rho(u, \lambda_l), \eta(u, \lambda_l)) \right]^2 \quad (15)$$

3.2 Recovery of Specular Reflectance Parameters

3.2.1 Recovery of Microfacet Slope

With the diffuse albedo and weight, we now focus in the recovery of the microfacet slope and the specular weight. As the microfacet slope only influences the specular reflection, it cannot be estimated for matte pixels, which do not convey any information on the specular reflection parameters. As a result, we only estimate σ_m for those pixels $u \in \mathcal{S}$, where \mathcal{S} is the set of specular pixels in \mathcal{I} .

With the optimal values of specular weight $W_{spec}(u)$ and the index of refraction $\eta(u, \lambda_l)$ computed at the previous iteration, we recover the microfacet slope making use of the optimisation corresponding to

$$\sigma_m^*(u) = \underset{\sigma_m}{\operatorname{argmin}} \sum_{l=1}^{|\mathcal{W}|} \left[\left(R_{spec}(u, \lambda_l) - \frac{W_{spec}(u)}{L(\lambda_l)} f_{spec}(L(\lambda_l), p(u), q(u), \sigma_m(u), \eta(u, \lambda_l)) \right)^2 + \alpha \mathcal{R} \right] \quad (16)$$

this accounts for Lines 14-16 of Algorithm 1 and is consistent with the fact that every wavelength-indexed band for the pixel u share the same σ_m .

Here, we have also introduced a regularisation term \mathcal{R} where α is a constant that determines the contribution of the regulariser to the optimisation process. A number of regularisation schemes can be used, these span from curvature constraints as applied in shape-from-shading [54] to surface integrability [11]. Since, in practice, σ_m should not overly vary over the neighborhood of the pixel u , here we employ a regularisation term which enforces a smoothly varying microfacet slope given by

$$\mathcal{R} = \sum_{v \in \Omega_u} (\sigma_m(u) - \sigma_m(v))^2 \quad (17)$$

3.2.2 Recovery of Specular Weight

This parameter determines the contribution of specular component of reflectance in the linear combination in Equation 2. We, therefore, estimate this parameter by optimising cost function in Equation 7 so as to encourage data closeness between the measured specular reflectance R_{spec} and the specular term in Equation 2. For totally diffuse pixels, the specular weight should be zero. Our method therefore estimates W_{spec} for each pixel $u \in \mathcal{S}$ given the optimal value of microfacet slope $\sigma_m(u)$ estimated using Equation 16 and the current estimates of the refractive index $\eta(u, \lambda_l)$ and surface gradients $p(u), q(u)$. This corresponds to Lines 17-19 in Algorithm 1.

In a manner analogous to the diffuse weight recovered in Subsection 3.1.2, W_{spec} varies per pixel, remaining fixed over the wavelength domain. Thus, we optimise over the wavelength-indexed bands as follows

$$W_{spec}^*(u) = \underset{W_{spec}(u)}{argmin} \sum_{l=1}^{|\mathcal{W}|} \left[R_{spec}(u, \lambda_l) - \frac{W_{spec}(u)}{L(\lambda_l)} f_{spec}(L(\lambda_l), p(u), q(u), \sigma_m(u), \eta(u, \lambda_l)) \right]^2 \quad (18)$$

3.3 Recovery of Index of Refraction and Object's Shape

3.3.1 Recovery of Index of Refraction

Finally, we turn our attention to the two parameters that apply to both, specular and diffuse reflection. Here, we estimate the index of refraction by noting that each material has its own unique refracted index defined over the spectral domain. Thus, our algorithm recovers the index of refraction per material. Further, following [20], we have constrained the index of refraction to follow Cauchy's dispersion equation [5].

Analogous to our approach to recover the diffuse albedo, we first estimate the index of refraction per pixel. Then, we refine our estimated value by updating $\eta(\cdot)$ per material making use of the expected value over all pixels belonging to each material in the scene. Since $\eta(\cdot)$ contributes to both, the specular and diffuse component of reflection, we optimise using the cost function in Equation 5, which comprises the data closeness terms for both, the specular and diffuse reflectance. This is depicted in Lines 20-24 of Algorithm 1. Here, in a manner akin to Subsection 3.2.1, we assume that, at each wavelength λ_l , $\eta(\cdot)$ varies slowly about the neighborhood of pixel u . Therefore, we optimise over Ω_u , i.e. the neighborhood of pixel u , so as to recover $\eta(\cdot)$ at pixel u and wavelength λ_l as follows

$$\eta^*(u, \lambda_l) = \underset{\eta(u, \lambda_l)}{\operatorname{argmin}} \sum_{v \in \Omega_u} \left[\left(R_{diff}(v, \lambda_l) - \frac{W_{diff}(v)}{L(\lambda_l)} f_{diff}(L(\lambda_l), p(v), q(v), \rho(v, \lambda_l), \eta(v, \lambda_l)) \right)^2 + \left(R_{spec}(v, \lambda_l) - \frac{W_{spec}(v)}{L(\lambda_l)} f_{spec}(L(\lambda_l), p(v), q(v), \sigma_m(v), \eta(v, \lambda_l)) \right)^2 \right] \quad (19)$$

and, with the per-pixel estimates at hand, We constrain the refractive indexes to follow Cauchy's dispersion equation.

Recall that Cauchy's equation is an empirical expression relating the refractive index to the wavelength through material specific dispersion coefficients. To enforce Cauchy's dispersion, we first estimate dispersion coefficients by fitting Cauchy's equation to the estimated index of refraction as follows

$$C_k^*(u) = \underset{C_k}{\operatorname{argmin}} \sum_{l=1}^{|\mathcal{W}|} \left(\eta^*(u, \lambda_l) - \sum_{k=1}^M C_k(u) \lambda_l^{-2(k-1)} \right)^2 \quad (20)$$

Here, the coefficients $C_k(u)$, $k \in \{1, \dots, M\}$ can be viewed as a characteristic of the material. In Algorithm 1, Lines 25-30 we recover the M dispersion coefficients at each pixel u by optimising over the corresponding wavelength-indexed bands. With the estimated dispersion coefficients, the index of refraction can be updated, as shown in Line 28 of Algorithm 1, using the expected value for Cauchy's equation given by

$$\eta(u, \lambda_l) = \frac{1}{|r|} \sum_{u \in r} \sum_{k=1}^M C_k^*(u) \lambda_l^{-2(k-1)} \quad (21)$$

where we have used the regions recovered while updating the diffuse albedo.

3.3.2 Recovery of Surface Normal

We now proceed to recover the object surface gradients making use of the viewer-centered global coordinate system presented in Section 2.1. Recall that, on our viewer-centered coordinate system, the surface normal \vec{N} can be expressed in terms of surface gradients. As a result, the vectors $[1, 0, p]^T$ and $[0, 1, q]^T$ are, hence, tangent to the surface. By definition, the surface normal is perpendicular to all vectors in the tangent plane and parallel to the cross-product of these two, that is

$$[1, 0, p(u)]^T \times [0, 1, q(u)]^T = [-p(u), -q(u), 1]^T \quad (22)$$

and, as mentioned earlier, the normalised surface normal is given by

$$\vec{N}(u) = \frac{1}{[p(u)^2 + q(u)^2 + 1]^{\frac{1}{2}}} [-p(u), -q(u), 1]^T \quad (23)$$

The expression above is important since it permits the recovery of the object's shape at each pixel by estimating only 2 variables, i.e. $p(u)$ and $q(u)$, per pixel u . As per the reparameterisation in Equation 2, we perform the optimisation procedure making use of both, the diffuse and specular components. This is reflected in Lines 34-36 of Algorithm 1. This is a straightforward task since the other reflection parameters are in hand. Moreover, we introduce a regularisation term corresponding to the integrability constraint widely used in shape-from-shading [11, 54]. Thus, the optimisation becomes

$$\begin{aligned} [p^*(u), q^*(u)] = \underset{p(u), q(u)}{\operatorname{argmin}} \sum_{v \in \Omega_u} \sum_{l=1}^{|\mathcal{W}|} \left[\left(R_{diff}(v, \lambda_l) - \frac{W_{diff}(v)}{L(\lambda_l)} f_{diff}(L(\lambda_l), p(v), q(v), \rho(v, \lambda_l), \eta(v, \lambda_l)) \right)^2 \right. \\ \left. + \left(R_{spec}(v, \lambda_l) - \frac{W_{spec}(v)}{L(\lambda_l)} f_{spec}(L(\lambda_l), p(v), q(v), \sigma_m(v), \eta(v, \lambda_l)) \right)^2 \right. \\ \left. + \left(\frac{\partial p(v)}{\partial y} - \frac{\partial q(v)}{\partial x} \right)^2 \right] \end{aligned} \quad (24)$$

In the optimisation above, the last term on the right-hand-side corresponds to the regularisation term. Also, note that we optimise the cost functional over the neighbourhood Ω_u of pixel u . This implies that the surface normal \vec{N} varies slowly about the pixel under consideration. This assumption does not imply any strong constraint on the surface shape rather reflects that, over small neighbourhoods of u , the surface can be considered to be locally smooth.

3.4 Stopping Criteria

Finally, we discuss two criteria of the algorithm to stop searching for optimal values of the parameters: (i) when the algorithm is already converged; no improvement is seen at the current cycle of optimisation along different coordinate directions, (ii) when the algorithm already completed maximum number of iterations t_{max} and therefore giving up. Lines 37-38 of Algorithm 1 describe the convergence criteria, where the function $\delta(\cdot)$ is used to denote the absolute difference between the estimates at iterations t and $t - 1$ for any of the reflectance parameters. Therefore, as long as there is a significant change in the current estimate for at least one of the parameters and $t < t_{max}$, the algorithm continues, otherwise, it stops and returns the parameter values estimated so far.

4 Reflectance Models

So far, we have made no assumptions regarding the specific reflectance model to be used for purposes of optimisation. In this section, we examine a number of reflectance models elsewhere in the literature which conform to the general formulation and optimisation approach described in Sections 2 and 3. It is worth noting in passing that, in previous sections, we have used the diffuse and specular image radiance, which can be related to the reflectance in the models hereafter making use of the relations

$$\begin{aligned} R_{diff}(u, \lambda_l) &= \frac{1}{L(\lambda_l)} f_{diff}(L(\lambda_l), p(u), q(u), \rho(u, \lambda_l), \eta(u, \lambda_l)) \\ R_{spec}(u, \lambda_l) &= \frac{1}{L(\lambda_l)} f_{spec}(L(\lambda_l), p(u), q(u), \sigma_m(u), \eta(u, \lambda_l)) \end{aligned} \quad (25)$$

In practice, we can use the relations above so as to substitute the reflectance corresponding to one or more of the models elsewhere in the literature into the cost function in Equation 5. Thus, we elaborate upon four reflectance models elsewhere in the literature and their specific parameters.

4.1 The Torrance-Sparrow Model

We commence by exploring the use of the Torrance-Sparrow model [49], which describes the reflectance from rough surfaces as comprised by two components. These are Lambertian diffuse reflection and the specular lobe. Here, we will focus on the component describing the specular lobe exhibited by smooth and rough surfaces. The main idea in [49] is to model a surface as a

collection of perfectly smooth mirror like microfacets. In the model, these microfacets are perfect mirrors.

For a given light source and viewer directions \vec{L} and \vec{V} , the specular reflection is solely dependent on those microfacets with a normal \vec{N} aligned with the halfway vector \vec{H} . The corresponding equation for the specular reflectance is therefore given by:

$$R_{TS}(u, \lambda_l) = \frac{A_f}{4} F(\theta_i, \eta(u, \lambda_l)) \frac{G(\theta_i, \theta_s, \phi_i, \phi_s)}{\cos \theta_s} D(\theta_h, \sigma_m) \quad (26)$$

where

$$G(\theta_i, \theta_s, \phi_i, \phi_s) = \min \left(1, \frac{2 (\vec{N} \cdot \vec{H}) (\vec{N} \cdot \vec{V})}{(\vec{V} \cdot \vec{H})}, \frac{2 (\vec{N} \cdot \vec{H}) (\vec{N} \cdot \vec{L})}{(\vec{V} \cdot \vec{H})} \right) \quad (27)$$

$$D(\theta_h, \sigma_m) = c \exp\left(-\frac{\theta_h^2}{2\sigma_m^2}\right) \quad (28)$$

the notation \cdot denotes the dot product between vectors, c is a constant and σ_m is the roughness parameter.

Note that, according to Equation 26, the specular reflectance $R_{TS}(\cdot)$ at pixel u and wavelength λ_l is governed by the microfacet area A_f , the wavelength dependent Fresnel term $F(\cdot)$, geometric attenuation factor $G(\cdot)$ and the microfacet slope distribution function $D(\cdot)$. In the model, the geometric attenuation factor $G(\cdot)$, as given in Equation 27, accounts for the microfacet-to-microfacet masking and shadowing. The microfacet distribution function $D(\cdot)$ used here is a rotationally symmetric normal distribution with zero mean, i.e. $\bar{\theta}_h = 0$, and standard deviation σ_m , which captures the probability density of the microfacets pointing in the direction of the halfway vector \vec{H} .

As the surface is not a perfect reflector, only a fraction of the incident light reaching the facets will be reflected. The Fresnel term $F(\cdot)$ determines the fraction of incident light which is reflected by each facet. The computation of the Fresnel term is quite expensive computationally and, therefore, here we apply the Schlick approximation [42] given by

$$F(\theta_i, \eta(u, \lambda_l)) = R_o + (1 - R_o) (1 - \cos \theta_i)^5 \quad (29)$$

where, as before, θ_i is the angle of incidence and $\eta(u, \lambda_l)$ is the refractive index at pixel u and wavelength λ_l . In Schlick's approximation, R_o is the reflectance at normal incidence, defined as $R_o = \left(\frac{1 - \eta(u, \lambda_l)}{1 + \eta(u, \lambda_l)} \right)^2$.

4.2 The Cook-Torrance Model

A model akin to that of Torrance and Sparrow is that developed by Cook and Torrance [7]. This is actually a combination of the Torrance-Sparrow model [49] and the Blinn model [3]. Just like the Torrance-Sparrow model, it considers the Fresnel term, surface self-shadowing and microfacets for reflectance modelling. However, it provides microfacet distribution functions for the Phong model [38], the Trowbridge-Reitz [50] and the Beckmann [2] distributions. Here we describe the reflectance equation of this model using the Beckmann distribution. Therefore, the reflectance equation, as described by this model, is given as,

$$R_{CT}(u, \lambda_l) = F(\theta_i, \eta(u, \lambda_l)) \frac{G(\theta_i, \theta_s, \phi_i, \phi_s)}{\cos\theta_s \cos\theta_i} D_{BK}(\theta_h, \sigma_m) \quad (30)$$

where $D_{BK}(\cdot)$ stands for the Beckmann distribution function given by

$$D_{BK}(\theta_h, \sigma_m) = \frac{1}{\sigma_m^2 \cos^4\theta_h} \exp\left(-\frac{\tan\theta_h}{\sigma_m}\right)^2 \quad (31)$$

4.3 The Beckmann-Kirchhoff Model

We now turn our attention to another model widely cited in the computer vision and graphics literature. The Beckmann-Kirchhoff model [2] describes the reflection from a surface as a sum of two terms. The first of these corresponds to the scattering component in the specular direction and the latter represents the diffuse scattering. The term describing the specular spike given by

$$R_{BK}(u, \lambda_l) = P_0^2(\theta_h, \sigma) \exp(-g(\theta_i, \theta_s, \lambda_l, \sigma_h)) F(\theta_i, \eta(u, \lambda_l)) \quad (32)$$

where,

$$P_0(\theta_h, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{\theta_h^2}{2\sigma^2}\right) \quad (33)$$

and

$$g(\theta_i, \theta_s, \lambda_l, \sigma_h) = \left(2\pi \frac{\sigma_h}{\lambda_l} (\cos\theta_i + \cos\theta_s)\right)^2 \quad (34)$$

The third term in Equation 32 corresponds to the wavelength dependent Fresnel term $F(\cdot)$ as described in Equation 29. Note that, even though there is no Fresnel term in the model as proposed in [2], we have included it without any loss of generality so as to make the specular reflectance component consistent across the models presented in this section. This also reflects the notion that no surface is a perfect reflector and, therefore, the Fresnel term is used to determine the fraction of the incident light reflected from the surface as a function of incident angle and index of refraction.

According to Equation 32, the specular reflectance $R_{BK}(\cdot)$ at pixel u and wavelength λ_l is a product of three terms. The first of these is the magnitude of the specular reflectance P_0 which is a function of θ_h . This function is nearly zero for all scattering directions except a very narrow range around the halfway vector \vec{H} . In the Beckmann-Kirchhoff model, P_0 is expressed as a sinc function. Nonetheless, here, for the sake of computational efficiency, we have followed Nayar et al. [34] and employed a Gaussian function with a very small standard deviation σ . The second term is comprised by the function $g(\cdot)$, which is related to the surface roughness through $\frac{\sigma_h}{\lambda_l}$, the incidence angle θ_i and reflected light angle θ_s . The three cases $g \ll 1$, $g \approx 1$, and $g \gg 1$ correspond to smooth surface, moderately rough and rough surface, respectively.

Note that both, the Torrance-Sparrow and Beckmann-Kirchhoff models incorporate the surface roughness parameter, i.e. σ_m in Equation 26 and σ_h in Equation 32, as a pertinent parameter to influence the reflectance characteristic of the surface. However, these two reflectance models are based on two different surface models. Whereas the Beckmann-Kirchhoff model uses the microfacet height distribution, the Torrance-Sparrow model is based on slope distribution model. Hence, the roughness parameters of these two models are not equivalent. This is important, since our work defines the surface based on slope distribution model and uses a roughness parameter akin to that in the Torrance-Sparrow model. This is not overly restrictive since the surface roughness parameter in the Beckmann-Kirchhoff model can be transformed into that used in our derivation following the approach presented in [34]. We do this making use of the correlation distance T for the height distribution model. We get

$$\sigma_h = \frac{T}{2} \tan(\sqrt{2}\sigma_m) \quad (35)$$

where

$$g(\theta_i, \theta_s, \lambda_l, \sigma_m) = \left(\frac{\pi T}{\lambda_l} \tan(\sqrt{2}\sigma_m) (\cos\theta_i + \cos\theta_s) \right)^2 \quad (36)$$

4.4 The Wolff Model

Finally, we elaborate upon the model proposed by Wolff [53]. This is a diffuse reflection model for smooth dielectric surfaces. When the angle between the illuminant and the viewing direction is large, a significant deviation from Lambert's law [25] is prevalent for many dielectric objects. The Wolff model explains this deviation from Lambert's law by considering the refractive attenuation at the surface-air layer. It multiplies Lambert's cosine law with two Fresnel terms: one for incident

light and the other for reflected light. The diffuse reflectance, as defined by the Wolff model is given by:

$$R_{WF}(u, \lambda_l) = \rho(u, \lambda_l) \cos \theta_i [1 - F(\theta_i, \eta(u, \lambda_l))] \left[1 - F(\theta'_s, \frac{1}{\eta(u, \lambda_l)}) \right] \quad (37)$$

where, as in previous sections, $\rho(u, \lambda_l)$ is the diffuse albedo, $\eta(u, \lambda_l)$ is the index of refraction of the dielectric medium, θ_i is the angle of incidence for the light ray and θ'_s is the internal angel of incidence that relates the reflected angle θ_s with Snell's law, i.e. $\theta'_s = \arcsin(\frac{\sin(\theta_s)}{\eta(u, \lambda_l)})$.

5 Implementation issues

Recall that we have expressed the image reflectance as a linear combination of the specular and diffuse reflection. In our experiments, we use a linear combination of three reflectance models described in the previous section. This is given by

$$R(u, \lambda_l) = W_{diff}(u)R_{WF}(u, \lambda_l) + W_{lobe}(u)R_{TS}(u, \lambda_l) + W_{spike}(u)R_{BK}(u, \lambda_l) \quad (38)$$

In Equation 38, each term represents a different reflection component. We have selected the Wolff model [53] to describe the diffuse reflectance. We have further divided the specular term into the spike and lobe, where we use the Torrance-Sparrow model [49] for the specular lobe and the Beckmann-Kirchhoff model [2] to describe the specular spike. These two terms are governed by the weights W_{lobe} and W_{spike} , which control the contribution of the specular lobe and spike to the image radiance.

The parameters for each of these models are summarised in Table 1. The optimisation procedure described in Section 3 can be readily applied to recover all the parameters except the contributions of the lobe and spike as captured by W_{lobe} and W_{spike} . These can be estimated making use of the optimisation given by

$$[W_{lobe}^*(u), W_{spike}^*(u)] = \underset{[W_{lobe}(u), W_{spike}(u)]}{argmin} \sum_{l=1}^{|W|} \left(R_{spec}(u, \lambda_l) - W_{lobe}(u)R_{TS}(u, \lambda_l) - W_{spike}(u)R_{BK}(u, \lambda_l) \right)^2 \quad (39)$$

as an alternative to Equation 18. In the equation above, we have accounted for the relationship between the specular image radiance and reflectance and opted for a notation consistent with Equation 26 and 32.

As shown in Algorithm 1, our coordinate descent algorithm operates on each coordinate in a cyclic order until convergence. At each cycle, it recovers $\rho(u, \lambda_l)$, $W_{diff}(u)$, $\sigma_m(u)$, $\{W_{lobe}(u), W_{spike}(u)\}$, $\eta(u, \lambda_l)$ and $\{p(u), q(u)\}$ in interleaved separate minimisation steps. Note that the recovery of $\rho(u, \lambda_l)$, $W_{diff}(u)$ and $\{W_{lobe}(u), W_{spike}(u)\}$ can be formulated as linear least squares minimisation problems. These are therefore variable-wise convex and have a globally optimal unique solution.

First, we turn our attention to the recovery of diffuse albedo using the following equation, which is in fact, a model specific form of Equation 13, where we place the Wolff reflectance model to account for the diffuse reflectance.

$$\rho^*(u, \lambda_l) = \underset{\rho(u, \lambda_l)}{\operatorname{argmin}} \sum_{v \in \Omega_u} \left[R_{diff}(v, \lambda_l) - W_{diff}(v) R_{WF}(v, \lambda_l) \right]^2 \quad (40)$$

The second order derivative of Equation 40 with respect to albedo is $2 \sum_{v \in \Omega_u} C(v, \lambda_l)^2$ where, $C(v, \lambda_l) = W_{diff}(v) \cos \theta_i [1 - F(\theta_i, \eta(v, \lambda_l))] \left[1 - F(\theta'_s, \frac{1}{\eta(v, \lambda_l)}) \right]$. Note that the second order derivative is always positive and, thus, Equation 40 is convex with respect to the diffuse albedo. Analogously, we can derive the second order derivative of model specific form of Equation 15 with respect to diffuse weight, which is also always positive, $2 \sum_{l=1}^{|\mathcal{W}|} R_{WF}(u, \lambda_l)^2$ and hence guarantees global minima.

For the two specular weights, $W_{lobe}(u)$ and $W_{spike}(u)$, we use Equation 39. The determinant of corresponding Hessian matrix is:

$$\det(H) = 4 \sum_{l=1}^{|\mathcal{W}|} R_{TS}(u, \lambda_l)^2 \sum_{l=1}^{|\mathcal{W}|} R_{BK}(u, \lambda_l)^2 - 4 \left(\sum_{l=1}^{|\mathcal{W}|} R_{TS}(u, \lambda_l) R_{BK}(u, \lambda_l) \right)^2 \quad (41)$$

which is also positive and therefore confirms the unique minima for this function.

The recovery of $\sigma_m(u)$, $\eta(u, \lambda_l)$, and $\{p(u), q(u)\}$ are nonlinear least squares problems. Note that, as mentioned earlier, the measure of non-linearity and the ability to chose initial estimates close to the real parameter determines whether our method will converge to a minimum. In all our experiments, the algorithm converges with our provided initial values, as described in Section 3.

By calculating the remaining second order derivatives for the corresponding cost functions, we can further comment on the non-linearity of the estimation task in hand. For example, for the microfacet slope $\sigma_m(u)$, the cost function is “flat” when the angle $\theta_h(u)$ has large values; i.e. when the halfway vector significantly deviates from the surface normal. According to specular

reflectance models [2, 49, 7], in this case, u is not a specular pixel and therefore the algorithm can not estimate $\sigma_m(u)$ at pixel u that does not contribute to the specular reflectance. That is why we estimate microfacet slope only at specular pixels.

We have already seen that the cost function for index of refraction has contribution from both specular and diffuse component. If we are to estimate $\eta(u, \lambda_l)$ where local neighborhood of u is a purely diffuse region, just like $\sigma_m(u)$, the specular part of the cost function becomes flat for larger value of θ_h . In this case, $\eta(u, \lambda_l)$ can still be recovered from the diffuse part. On the other hand, the diffuse part becomes flat for pixels with poor illumination and low image reflectance.

For all our least square optimisation, we have used the Levmar C++ library available for download from FORTH-ICS ¹. This library provides a Levenburg-Marquardt method which can employ both, an analytic Jacobian or finite difference approximated Jacobians. For computing the gradients $\{p(u), q(u)\}$, we avoid the possibly complex analytical derivation of the Jacobian and use finite difference approximated Jacobians. For all the other parameters, we provide the solver with analytic Jacobians. This has the effect of a reduced computational complexity.

In our experiments, the Levmar solver has been set to terminate when at least one of the following conditions is met: when the gradient of the cost drops below a threshold ϵ_1 , when the change in parameter values has fallen to some threshold ϵ_2 , when cost itself has reached some acceptable value ϵ_3 , and when specified number of iterations it_{max} is completed. Therefore, the library requires four user defined threshold values; We provide 10^{-15} for ϵ_1 and ϵ_2 , 10^{-20} for ϵ_3 , and 100 for it_{max} . The library also requires a user input τ to set the initial value for the damping term. According to the damping strategy of the algorithm, τ should be smaller only if we believe that our provided initial value of the unknown parameters are close to optimal value. Therefore, we set τ to 10^{-6} for $\{p(u), q(u)\}$ as we get the initial value from image gradients, whereas, we set 10^{-1} for the rest.

6 Experiments

In this section, we provide example results for our method and compare with alternatives. In all our experiments, we have used the method in [19] to recover the power spectrum of the illuminant and expressed the image reflectance as a linear combination as given in Equation 38.

¹<http://www.ics.forth.gr/~lourakis/levmar>

Table 1: Parameters to be recovered for each of the reflectance models under consideration.

Reflectance Model	Corresponding Equation	Parameters to estimate
TS	$R_{TS}(u, \lambda_l) = \frac{A_f}{4} F(\theta_i, \eta(u, \lambda_l))$ $\frac{G(\theta_i, \theta_s, \phi_i, \phi_s)}{\cos \theta_s} D(\theta_h, \sigma_m)$	$\sigma_m(u),$ $\eta(u, \lambda_l), \{p(u), q(u)\}$
BK	$R_{BK}(u, \lambda_l) = P_0^2(\theta_h, \sigma)$ $\exp(-g(\theta_i, \theta_s, \lambda_l, \sigma_m))$ $F(\theta_i, \eta(u, \lambda_l))$	$\sigma_m(u),$ $\eta(u, \lambda_l), \{p(u), q(u)\}$
Wolff	$R_{WF}(u, \lambda_l) = \rho(u, \lambda_l) \cos \theta_i$ $[1 - F(\theta_i, \eta(u, \lambda_l))]$ $\left[1 - F(\theta'_s, \frac{1}{\eta(u, \lambda_l)})\right],$	$\rho(u, \lambda_l)$ $\eta(u, \lambda_l), \{p(u), q(u)\}$ $\{p(u), q(u)\}$

6.1 Synthetic Data

In this section, we perform experiments on a number of synthetic imagery so as to verify the accuracy of the recovered photometric parameters and object shape. Our synthetic data comprises of the *Stanford Bunny* and the *Happy Buddha* whose mesh files can be downloaded from the Stanford University Computer Graphics Laboratory ². Making use of these meshes, we have rendered the images using the reflectance model given in Equation 38, making use of two indexes of refraction, microfacet slope values and diffuse albedos. Note that this results in two images per mesh, which we name *Stanford Bunny A*, *Stanford Bunny B*, *Happy Buddha A* and *Happy Buddha B*. The

²<http://graphics.stanford.edu/data/3Dscanrep/>

Reflectance model parameters	Image Name			
	Happy Buddha A	Happy Buddha B	Stanford Bunny A	Stanford Bunny B
Microfacet slope (mean, std)	(4.0, 0.01)	(8.0, 0.01)	(4.0, 0.01)	(8.0, 0.01)
Albedo	yellow paper	yellow paper	orange gel box	orange gel box
Index of refraction	BK7	BK7	Si	Si

Table 2: Parameter values for our synthetic images

parameter values for each of these images is summarised in Table 2.

It is worth noting in passing that, for the two microfacet slope values, we have used a normal distribution, whereas for the ground truth value of albedo, we have used spectra acquired in-house making use of a StellarNet spectrometer for both, a yellow paper and an orange gel box. For the indexes of refraction, we use those corresponding to BK7, i.e. a high quality optical glass, and Silicon (Si). The ground truth value for the indexes of refraction are available at Filmetrics³. All our imagery is comprised of 30 bands in the range of $[430nm - 720nm]$, in $10nm$ steps.

With these images in hand, we first turn our attention to the capacity of our method to recover the diffuse albedo. To this end, we show, in Figure 2 the diffuse albedo plots for both, the ground truth and that recovered by our method. In the bottom row of the figure we show the plot for the average average estimated normalised albedo. In the figures, the error bars denote the standard deviation across the estimates for all pixels. Note that for both, the *Stanford Bunny* and the *Happy Buddha* the mean albedos recovered by our method are in close accordance with the ground truth. Moreover, the standard deviation is quite small.

We now provide a qualitative illustration of the capacity of our method to recover the reflection parameters which correspond to the specular and diffuse reflection components in the imagery. To this end, we show, in Figure 3 and 4, the specular, diffuse and composite reflection for the *Happy Buddha A* and *Happy Buddha B*. In the figures, from left-to-right, we show the composite, the specular lobe, the specular spike and the diffuse reflectance for both, the ground truth value and our estimated value. The top row shows images rendered with all ground truth values whereas the bottom row shows images rendered with our estimated parameters. It is worth noting that, in the figures, we show pseudocolour imagery obtained making use of the CIE Colour Matching functions as proposed by Stiles and Burch [46].

Similarly, in Figures 5 and 6, we show the field of surface normals, i.e. needlemaps, for the *Happy Buddha A* and the *Stanford Bunny A* images recovered by our method and two alternatives. These are the method in [54] and that presented in [18]. The first of these is a shape-from-shading one which imposes a hard constraint on the surface normals. Since this method assumes Lambertian reflectance, we have used, as input, the diffuse component delivered by the method in [47]. The method in [18] is an optimisation one which employs Euler-Lagrange equations to recover the shape and index of refraction.

³<http://www.filmetrics.com/refractive-index-database>.

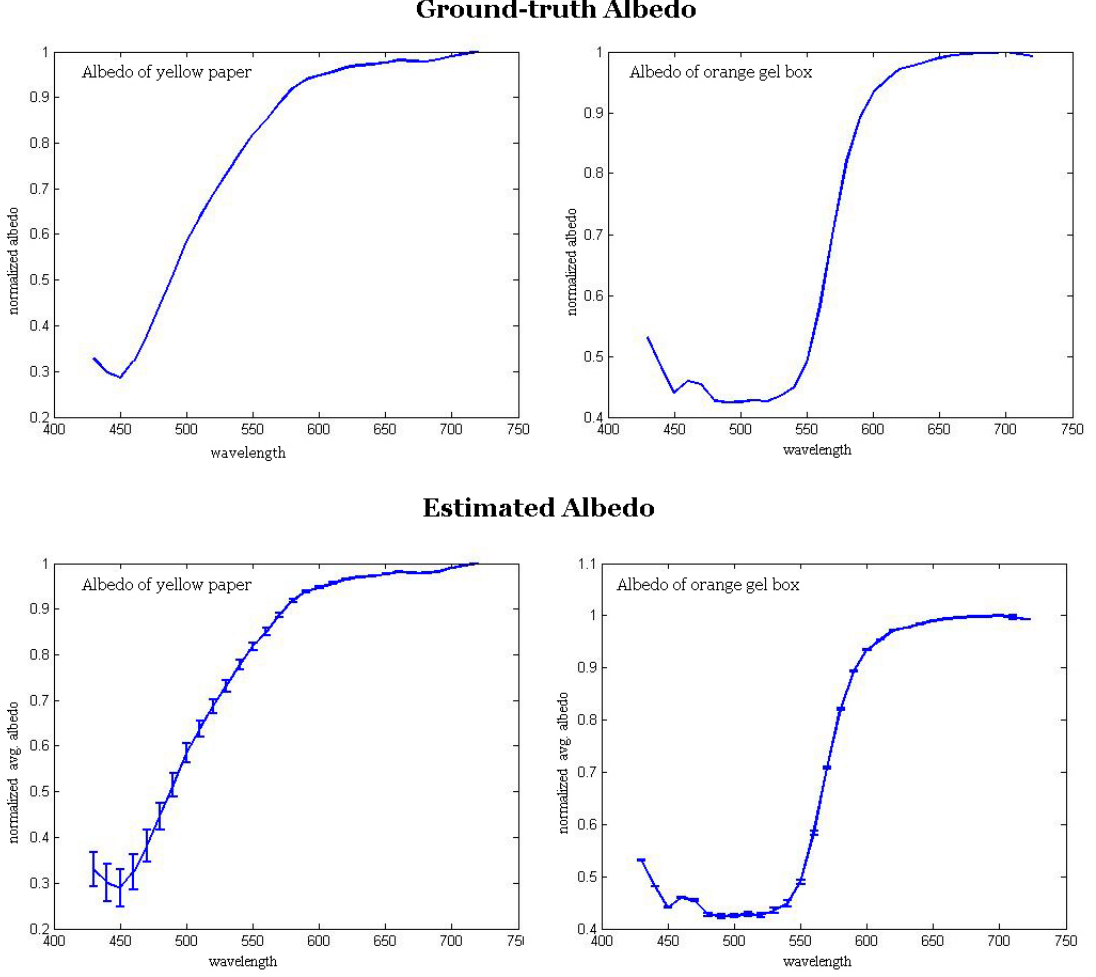


Figure 2: Top row: Plots of the normalised ground truth albedo used for the *Happy Buddha* and the *Stanford Bunny* images; Bottom row: Plots of the average estimated normalised albedo over all pixels in the *Happy Buddha* and *Stanford Bunny* images. Here, the error bars account for the standard deviation of the estimated albedo over all pixels in the image.

In the figures, our method delivers the surface normal field which better reflects the shape of the object. Indeed, for both images, the needlemap yielded by our method is the one which better captures the details of the object, with less over smoothing and abrupt bogus changes on the direction of the surface normals. This is despite the challenging nature of both object's shape.

We provide a quantitative analysis on the index of refraction recovered by our method as compared to that delivered by the alternative in [18]. In Table 3, we show the average per-pixel angular error, in degrees, for the index of refraction yielded by our method and the alternative. Note that our method delivers a considerable improvement over the method in [18]. As mentioned earlier, we impose the index of refraction to follow Cauchy's dispersion equation. This constraint signif-




Figure 3: A qualitative illustration of the capacity of our method to recover the reflection model parameters for the *Happy Buddha A*. From left-to-right, we show the composite, the specular lobe, the specular spike and the diffuse reflectance for both, the ground truth (top row) and the estimated images rendered with the parameters recovered by our method (bottom row).

icantly improves our results as it provides a further constraint on the optimisation process. This can be appreciated in the four images, where the index of refraction is quite robust to the shape variations between the two meshes and the combinations of microfacet slope values and diffuse albedos.

Finally, we present the computation time for our algorithm in Table 4. Note that, the compu-




Figure 4: A qualitative illustration of the capacity of our method to recover the reflection model parameters for *Happy Buddha B*. From left-to-right, we show the composite, the specular lobe, the specular spike and the diffuse reflectance for both, the ground truth (top row) and the estimated images rendered with parameters recovered by our method (bottom row).




Figure 5: (a) Ground truth surface normals for the *Happy Buddha*; (b) Needlemap delivered by our method; (c) Needlemap delivered by the method in [54]; (d) The surface normal field delivered by the algorithm in [18].

tational cost of our method depends on the image size, number of wavelength-indexed bands in the image, neighborhood size for each pixel and the reflectance models involved. Therefore, as a reference, we provide the per pixel, per band time for the *Happy Buddha A* image in Figure 3 for each of the reflectance parameters under consideration. Note that, in our implementation, we

	<i>Happy Buddha A</i>	<i>Happy Buddha B</i>	<i>Stanford Bunny A</i>	<i>Stanford Bunny B</i>
Our method	0.1555	0.1552	0.1549	0.1561
Method in [18]	6.4066	6.4001	9.9730	9.9910

Table 3: Average per-pixel angular error (in degrees) with respect to ground truth for the index of refraction recovered by our method and the alternative in [18].

discard dark and achromatic pixels as these convey no information about reflectance. We do this based on a threshold value on the image intensity. This yields 72326 valid pixels for the Happy Buddha A image. Moreover, the specular removal method in [47] yields 5381 specular pixels. In our experiments, we consider a 3x3 neighborhood size for each pixel while estimating ρ , η , and p, q . This altogether gives us the total computation time of 19.78 minutes on a workstation with a




Figure 6: (a) Ground truth surface normals for the *Stanford Bunny*; (b) Needlemap delivered by our method; (c) Needlemap delivered by the method in [54]; (d) The surface normal field delivered by the algorithm in [18].

2.0 GHz Intel Xeon 7500 series processor and 32GB of RAM.

Note that, in our method, the estimation of the index of refraction requires the longest time as compared to other parameters. The reasons for this are twofold. Firstly, the refractive index is estimated per pixel and per band. Secondly, its computation involves both, the diffuse and the specular terms of our cost function.

To reduce the computational complexity, recall that, in Section 3.1.1, we segment the image into $\{r_1, \dots, r_c\}$ regions, where each region r corresponds to a specific material. We can take advantage of this segmentation so as to reduce the computation time for the refractive index. This is as each material has its own unique refractive index defined over the spectral domain. Thus, we can replace the pixel-based computation in Equation 19 with the following region-based estimation.

$$\eta^*(r, \lambda_l) = \underset{\eta(r, \lambda_l)}{\operatorname{argmin}} \sum_{v \in r} \left[\left(R_{diff}(v, \lambda_l) - \frac{W_{diff}(v)}{L(\lambda_l)} f_{diff}(L(\lambda_l), p(v), q(v), \rho(v, \lambda_l), \eta(v, \lambda_l)) \right)^2 + \left(R_{spec}(v, \lambda_l) - \frac{W_{spec}(v)}{L(\lambda_l)} f_{spec}(L(\lambda_l), p(v), q(v), \sigma_m(v), \eta(v, \lambda_l)) \right)^2 \right] \quad (42)$$

It is also worth noting that, as explained in Section 2.2, our formulation of the problem allows for the optimisation on the diffuse and specular reflection components to be effected in parallel. Parallel execution of our method can significantly speed up the process.

6.2 Real-world Imagery

[In this section, we provide results on real world data.](#) Our real world hyperspectral imagery comprises of the database reported in [19], which was acquired using an OKSI Turnkey Hyperspectral Camera System. The images correspond to 51 human subjects, each captured under one of 10 directional light sources with varying directions and spectral power. The light sources are divided into two rows. The first of these is placed above the camera system and the second one at the same height as the camera. The main directions of the lights are adjusted so as to point towards the centre of the scene.

We commence by showing results on the recovered fields of surface normals. In Figure 7, we show needlemap plots for two sample images in the database. In the Figure, we show the fields

Estimated Parameter	Notation	Reflectance Model Involved	Computation Time (in sec)
Diffuse albedo	$\rho(u, \lambda)$	Wolff	9.21×10^{-05} per pixel, per band
Diffuse weight	$W_{diff}(u)$	Wolff	1.28×10^{-06} per pixel
Microfacet slope	$\sigma_m(u)$	TS, BK	0.051 per specular pixel
Specular weight	$W_{spec}(u)$	TS, BK	3.71×10^{-05} per specular pixel
Index of refrac.	$\eta(u, \lambda)$	Wolff, TS, BK	2.31×10^{-04} per pixel, per band
Surface gradients	$\{p(u), q(u)\}$	Wolff, TS, BK	2.92×10^{-03} per pixel

Table 4: Computation time per pixel, band and reflection parameter for the *Happy Buddha A* image

of surface normals delivered by our method and two other alternatives. These are the method in [54] and that presented in [18]. The first of these is a shape-from-shading one which is based on Lambertian reflectance. Therefore, we provide, as input to this method, the diffuse component

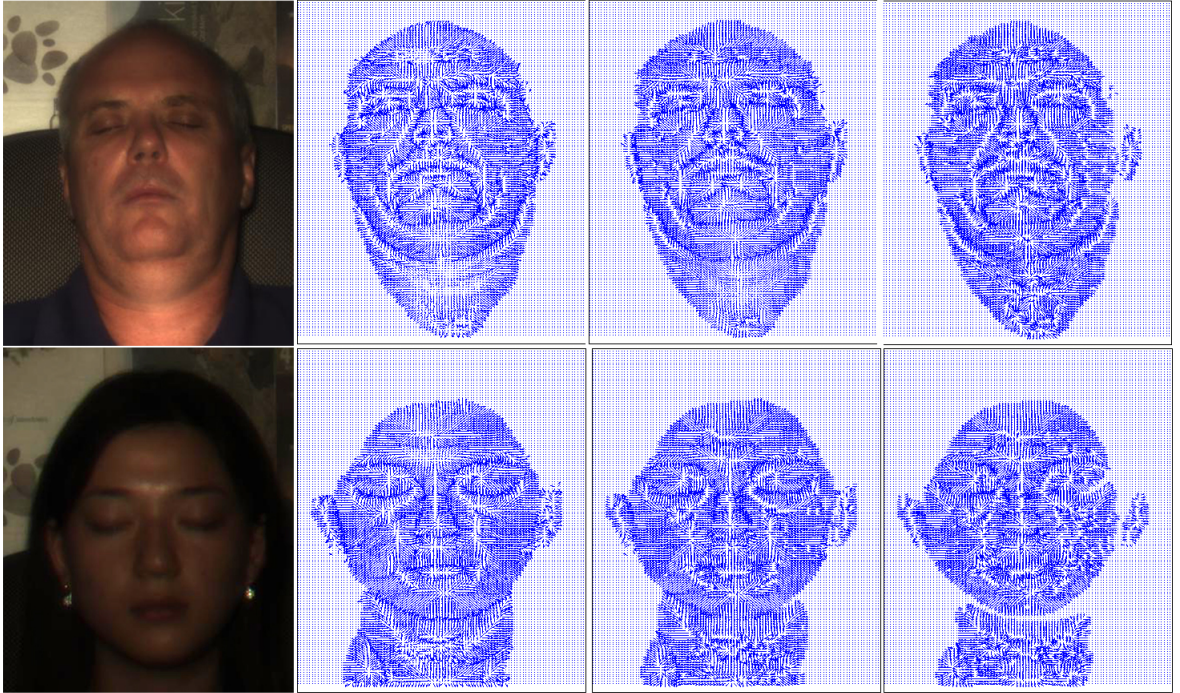


Figure 7: Fields of surface normals, i.e. needlemaps, for sample faces in our database. First column: input images; Second column: Needlemap delivered by our method; Third column: Needlemap recovered by the method in [18]; Fourth column: The surface normal field yielded by the algorithm in [54].

delivered by the method in [47]. The method in [18] is an optimisation one which takes a hyper-spectral image at input and employs Euler-Lagrange equations to recover the shape. Note that, in the figure, our method delivers the surface normals that better capture the details of the face, with less over smoothing and abrupt bogus changes at specular regions.

Next, we present results on skin recognition. We do this following the notion that an accurate recovery of the reflectance, i.e. diffuse albedo, should yield a high recognition rate. Moreover, there is evidence that the Torrance-Sparrow model [49] and the Beckmann micro-facet distribution [2] can be used to model the skin reflectance [52]. The Torrance-Sparrow model [49] has also been used for purposes of rendering human skin [8], whereas the Schlick approximation for the Fresnel term has been employed to model skin reflection [35].

Notice that our diffuse albedo is expected to be invariant to changes in illumination direction and surface geometry. Here, we pose this recognition task as a classification problem where the skin and non skin reflectance spectra comprise of positive and negative classes, respectively. To obtain a training data set, we select skin and non skin regions from a single image captured under a light source placed in a high oblique position in front of the subject. On an average, there are 856 skin pixels and 7796 non-skin pixels selected from several regions in each image as training data. Subsequently, the albedo is used as input to a Support Vector Machine (SVM) classifier [44] with a Radial Basis Function (RBF) kernel. We have effected ten trials where the parameters of the classifier are selected using 5-fold cross validation at training time. For testing, the resulting SVM classifier is then applied to the testing images. Note that the test images are acquired under different illuminant conditions than those corresponding to the training imagery, which further complicates the recognition task.

In Figure 8, we present sample skin probability maps obtained using different features as input to the classifier. In the figure, the first row shows sample training image. The second row shows the sample testing images. Note that the illuminant directions and power spectra differ between the training and testing imagery. This is evident by the difference in shading and shadows. The four remaining rows show the skin probability maps generated by the SVM classifier trained using different reflectance-based features. The third row shows skin probability maps obtained using the spectral diffuse reflectance recovered by our method whereas the fourth row shows that yielded by the diffuse reflectance estimated by the alternative in [19]. The fifth row shows the maps obtained by using the reflectance computed via the normalisation of the image radiance by the illuminant. The bottom row shows the probability maps for the raw radiance.

From Figure 8, we note that the skin probability maps in the third and fourth rows depict more visually accurate results. This is due to our method and the alternative in [19] being less prone to error due to changes in illumination direction and power spectra between the training and testing




Figure 8: Skin probability maps obtained using different reflectance-based feature as input for classification. Top row: sample training images; Second row: sample test images of the corresponding subjects, captured under a different illumination condition. The third to last rows show the skin probability maps obtained using the reflectance delivered by our method (third row), the reflectance yielded by the alternative in [19] (fourth row), the reflectance obtained via normalisation by the illuminant (fifth row) and the raw radiance (bottom row).

Feature for Classification	CDR(%)	FDR(%)	CR(%)
Estimated reflectance by our method	92.89 ± 9.48	4.10 ± 2.77	95.09 ± 4.26
Estimated reflectance by method in [19]	85.12 ± 13.36	5.10 ± 6.30	90.94 ± 6.12
Reflectance by illuminant normalisation	70.63 ± 16.95	5.50 ± 5.69	84.75 ± 8.04
Raw radiance	47.27 ± 20.58	12.48 ± 11.64	71.23 ± 9.17

Table 5: Skin detection rates, false detection rates and classification rates yielded by different features. Top row: results yielded by the reflectance recovered by our method; Second row: those yielded when the reflectance estimated by the alternative in [19] is used; Third row: recognition rates corresponding to the reflectance obtained via illuminant normalisation of the image radiance; Bottom row: recognition rates for the raw radiance.

images. Further, nonetheless the probability maps in the third and fourth rows are comparable, our method still provides a margin of improvement over the alternative. This is particularly noticeable on the subjects in the middle columns, where the eyebrows, lips and beard have been better outlined by our method.

As mentioned earlier, for the probability maps on the fifth row, we have use the reflectance computed by normalising the image radiance with respect to the illuminant spectrum. Although being illuminant invariant, this classification feature gives poor results at pixels near the face boundary and the specular regions. This is due to the fact that, normalising the image radiance by the illuminant power spectrum does not achieve surface shading independence. Moreover, it disregards the specular component inherent in the dichromatic reflection model. In contrast, our method estimates the reflectance taking into account the shading and specularities. Therefore, still our method delivers a reflectance feature that allows for classification even at grazing angles and specular spikes and lobes. In the other hand, the raw radiance spectra is not illuminant invariant. Therefore, as expected, this feature generates poor results when the illumination direction and power spectra change between the training and testing images. As a result, the false negative rate is high in skin areas and false positives are rife in non-skin regions.

Finally, we provide a quantitative analysis in Table 5. In the table, we compare the performance achieved using the reflectance yielded by our method with that yielded by the alternatives. The table shows the classification rate (CR), correct detection rate (CDR), and false detection rate (FDR). Here, CDR stands for the percentage of skin pixels correctly classified, whereas FDR

corresponds to the percentage of non-skin pixels falsely classified as skin. The CR is the total percentage of skin and non-skin pixels classified accurately. For our experiments, we have used the ground truth skin data provided with the dataset. As reported in [19], this has been obtained by manually labelling skin and non-skin pixels for all subjects in the dataset. From the table, we conclude that the reflectance delivered by our method outperforms the alternatives in all the three measures used for analysis, with a better average performance and a lower variance. This is consistent with the qualitative result shown in Figure 8.

7 Conclusions

In this paper, we have presented a method to recover photometric parameters and object’s shape from a single hyperspectral image. The method is quite general in nature, stemming from the use of a general formulation for the image radiance based upon a linear combination of the diffuse and specular reflection components. This treatment permits, in turn, the use of a coordinate descent least-square optimisation approach to recover the parameters governing the reflection process. Moreover, the formulation provided here allows for the optimisation on the diffuse and specular reflection to be effected in parallel. We have illustrated how this optimisation process can be used in combination with the Beckmann-Kirchhoff, the Torrance-Sparrow, and the Wolff reflectance model to recover the index of refraction, microfacet slope, diffuse albedo, and surface normal from a single image. We have also shown results on skin recognition and compared with an alternative elsewhere in the literature.

8 Acknowledgements

The authors would like to thank Dr. Cong Phuoc Huynh for providing the reflectance computed via illuminant normalisation and the raw radiance used as classification features for the skin recognition results shown in Section 6.2.

References

- [1] A. Artusi, F. Banterle, and D. Chetverikov. A survey of specular removal methods. *Computer Graphics Forum*, 30(8):2208–2230, 2011.

- [2] P. Beckmann and A. Spizzichino. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon, New York, 1963.
- [3] J. M. Blinn. Models of light reflection for computer synthesized pictures. In *ACM SIGGRAPH*, pages 273–282, 1977.
- [4] S. Boivin and A. Gagalowicz. Image-based rendering of diffuse, specular and glossy surfaces from a single image. In *Proceedings of SIGGRAPH*, pages 107–116, 2001.
- [5] M. Born, E. Wolf, and A. B. Bhatia. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 2003.
- [6] G. Brelstaff and A. Blake. Detecting specular reflection using Lambertian constraints. In *International Conf. on Computer Vision*, pages 297–302, 1988.
- [7] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Trans. on Graphics*, 1(1):7–24, 1982.
- [8] C. Donner, T. Weyrich, E. d’Eon, R. Ramamoorthi, and S. Rusinkiewicz. A layered, heterogeneous reflectance model for acquiring and rendering human skin. *ACM Trans. Graph.*, 27(5):1–12, 2008.
- [9] R. O. Dror, E. H. Adelson, and A. S. Willsky. Recognition of surface reflectance properties from a single image under unknown real-world illumination. In *Proc. of the IEEE Workshop on Identifying Objects Across Variations in Lighting*, 2001.
- [10] G. D. Finlayson and G. Schaefer. Solving for colour constancy using a constrained Dichromatic reflection model. *International Journal of Computer Vision*, 42(3):127–144, 2001.
- [11] R. T. Frankot and R. Chellappa. A method of enforcing integrability in shape from shading algorithms. *IEEE Trans. on PAMI*, 4(10):439–451, 1988.
- [12] J. Friedman, T. Hastie, H. Hööfling, and R. Tibshirani. Pathwise coordinate optimization. *The Annals of Applied Statistics*, 1:302–332, 2007.
- [13] K. Hara, K. Nishino, and K. Ikeuchi. Light source position and reflectance estimation from a single view without the distant illumination assumption. *IEEE Trans. on PAMI*, 27(4):493–505, 2005.

- [14] B. K. P. Horn. Hill shading and the reflectance map. *Proceedings of the IEEE*, 69(1):14–47, 1981.
- [15] B. K. P. Horn and M. J. Brooks. The variational approach to shape from shading. *CVGIP*, 33(2):174–208, 1986.
- [16] B. K. P. Horn and R. W. Sjoberg. Calculating the reflectance map. *Applied Optics*, 18(11):1770–1779, 1979.
- [17] P. Huber. *Robust Statistics*. Wiley, 1981.
- [18] C. P. Huynh and A. Robles-Kelly. Simultaneous photometric invariance and shape recovery. In *IEEE International Conf. on Computer Vision*, 2009.
- [19] C. P. Huynh and A. Robles-Kelly. A solution of the Dichromatic model for multispectral photometric invariance. *International Journal of Computer Vision*, 90(1):1–27, 2010.
- [20] C. P. Huynh, A. Robles-Kelly, and E. R. Hancock. Shape and refractive index recovery from single-view polarisation images. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2010.
- [21] K. Ikeuchi and B.K.P. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1-3):141–184, 1981.
- [22] K. Ikeuchi and K. Sato. Determining reflectance properties of an object using range and brightness images. *IEEE Trans. on PAMI*, 13(11):1139–1153, 1991.
- [23] R. Kimmel, M. Elad, D. Shaked, R. Keshet, and I. Sobel. A variational framework for retinex. *International Journal of Computer Vision*, 52(1):7–23, 2003.
- [24] G. Klinker, S. A. Shafer, and T. Kanade. A physical approach to color image understanding. *International Journal of Computer Vision*, 4(1):7–38, 1990.
- [25] J. H. Lambert. *The Photometria, sive, De mensura et gradibus luminis, colorum et umbrae*. Eberhard Klett, 1760.
- [26] M. Lampton. Damping-undamping strategies for the Levenberg-Marquardt nonlinear least-squares method. *Computers in Physics*, 11(1):110–115, 1997.

- [27] K. Levenberg. A method for the solution of certain non-linear problems in least squares. *The Quarterly of Applied Mathematics*, (2):164–168, 1944.
- [28] S. Lin and H. Shum. Separation of diffuse and specular reflection in color images. In *IEEE Conf. on Computer Vision and Pattern Recognition*, 2001.
- [29] D. Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *SIAM Journal on Applied Mathematics*, 11:431–441, 1963.
- [30] D. Miyazaki, T. Shibata, and K. Ikeuchi. Wavelet texture wavelet-method: Appearance compression by polarization, parametric reflection model, and daubechies wavelet. *International Journal of Computer Vision*, 86(2-3):171–191, 2010.
- [31] D. Miyazaki, R.T. Tan, K. Hara, and K. Ikeuchi. Polarization-based inverse rendering from a single view. In *International Conf. on Computer Vision*, volume 2, pages 982 –987, 2003.
- [32] S. K. Nayar and R. M. Bolle. Computing reflectance ratios from an image. *Pattern Recognition*, 26:1529–1542, 1993.
- [33] S. K. Nayar and R. M. Bolle. Reflectance based object recognition. *International Journal of Computer Vision*, 17(3):219–240, 1996.
- [34] S. K. Nayar, K. Ikeuchi, and T. Kanade. Surface reflection: Physical and geometrical perspectives. *IEEE Trans. on PAMI*, 13(7):611–634, 1991.
- [35] C. S.-L. Ng and L. Li. A multi-layered reflection model of natural human skin. *Computer Graphics International Conference*, 2001.
- [36] H. B. Nielsen. Damping parameter in Marquardt’s method. Technical report, Informatics and Mathematical Modelling, Technical University of Denmark, DTU, 1999.
- [37] M. Oren and S. K. Nayar. Generalization of the Lambertian model and implications for machine vision. *International Journal of Computer Vision*, 14(3):227–251, 1995.
- [38] B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.
- [39] H. Ragheb and E. R. Hancock. A probabilistic framework for specular shape-from-shading. *Pattern Recognition*, 36(2):407–427, 2003.

- [40] H. Ragheb and E. R. Hancock. The modified Beckmann-Kirchhoff scattering theory for rough surface analysis. *Pattern Recognition*, 40(7):2004–2020, 2007.
- [41] I. Sato, Y. Sato, and K. Ikeuchi. Illumination from shadows. *IEEE Trans. on PAMI*, 25(3):290–300, 2003.
- [42] C. Schlick. A fast alternative to Phong’s specular model. *Graphics Gems*, 4:363–366, 1994.
- [43] C. Schlick. A survey of shading and reflectance models. *Computer Graphics Forum*, 13(2):121–131, 1994.
- [44] B. Scholkopf and A. J. Smola. *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. MIT Press, 2001.
- [45] S. A. Shafer. Using color to separate reflection components. *Color Research and Applications*, 10(4):210–218, 1985.
- [46] W. S. Stiles and J. M. Burch. Interim report to the Commission Internationale de l’Éclairage Zurich, 1955, on the National Physical Laboratory’s investigation of colour-matching. *Optica Acta*, 2:168–181, 1955.
- [47] R. T. Tan and K. Ikeuchi. Separating reflection components of textured surfaces using a single image. *IEEE Trans. on PAMI*, 27(2):178–193, 2005.
- [48] S. Tominaga and N. Tanaka. Estimating reflection parameters from a single color image. *IEEE Computer Graphics and Applications*, 20(5):58–66, 2000.
- [49] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America*, 57(9):1105–1114, 1967.
- [50] T. S. Trowbridge and K. P. Reitz. Average irregularity representation of a rough surface for ray reflection. *Journal of Optical Society of America*, 65(5):531–536, 1975.
- [51] C. L. Vernold and J. E. Harvey. A modified Beckmann-Kirchoff scattering theory for non-paraxial angles. In *Scattering and Surface Roughness II*, Proc. SPIE 3426, pages 51–56, 1998.

- [52] T. Weyrich, W. Matusik, H. Pfister, B. Bickel, C. Donner, C. Tu, J. McAndless, J. Lee, A. Ngan, H. Wann Jensen, and M. H. Gross. Analysis of human faces using a measurement-based skin reflectance model. *ACM Trans. Graph.*, 25(3):1013–1024, 2006.
- [53] L. B. Wolff. Diffuse-reflectance model for smooth dielectric surfaces. In *IEEE Conf. on Computer Vision and Pattern Recognition*, pages 2956–2968, 1994.
- [54] P. L. Worthington and E. R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE Trans. on PAMI*, 21(12):1250–1267, 1999.
- [55] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Trans. on PAMI*, 13(7):680–702, 1991.