

Simultaneous Photometric Invariance and Shape Recovery

Cong Phuoc Huynh¹

huynh@rsise.anu.edu.au

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

Antonio Robles-Kelly^{1,2}

antonio.robles-kelly@nicta.com.au

²National ICT Australia, Locked Bag 8001
Canberra ACT 2601, Australia

Abstract

In this paper, we identify the constraints under which the generally ill-posed problem of simultaneous recovery of surface shape and its photometric invariants can be rendered tractable. We examine the cases where a single or more images are acquired using different lighting directions with known illuminant power. Given these conditions, we state the constraints upon which the recovery of the surface geometry and its photometric parameters can be estimated. With these constraints, we then show how the recovery process may be formulated as an optimisation algorithm which aims to fit the reflectance models under study to the image reflectance. The approach presented here is general and can be applied to a family of reflectance models that are based on the Fresnel reflection theory. Thus, we provide a theoretical and computational background for recovering shape, material index of refraction and microscopic roughness from multi-spectral images.

1. Introduction

Photometric invariance has found applications in computer vision and pattern recognition for purposes of recognition and shape recovery. For instance, Nayar and Bolle [9] have used photometric invariants derived from the BRDF to recognise objects with different reflectance properties. In a related development, Dror *et al.* [3] have shown how surfaces may be classified from single images through the use of reflectance properties.

The intrinsic relation between photometric invariance and shape recovery is due to the fact that the reflectance of an object is determined not only by the light source and viewing directions, but also by the material properties of the surface under study. This is 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 [6], and by Horn and Brooks [5], hinge on the compliance with the image irradiance equation and local surface smoothness. Zheng and

Chellappa [19] proposed a gradient consistency constraint that penalises differences between the image intensity gradient and the surface gradient for the recovered surface. Worthington and Hancock [18] impose the Lambertian radiance constraint in a hard manner. Prados and Faugeras [12] have presented a shape from shading approach applicable to the case where the light source and viewing directions are no longer co-incident. Kimmel and Bruckstein have used the apparatus of level sets methods to recover solutions to the Eikonal equation [8].

The non-collinearity of the viewing and light directions on non-Lambertian surfaces using the continuous solutions to the image irradiance equation suggests a generalisation of integral shape-from-shading schemes for purposes of photometric invariance. Along these lines, most of the effort related to photometric invariance with respect to shape recovery is devoted to modeling the effects encountered on shiny or rough surfaces. For shiny surfaces, there are specular spikes and lobes which must be modelled. For instance, Brelstaff and Blake [2] use a thresholding strategy to identify specularities on moving curved objects. Ragheb and Hancock [13] have developed a probabilistic framework for specular subtraction which uses the Torrance and Sparrow model [14] to account for the distribution of image brightness.

The methods above often rely on the use of a reflectance model to capture the photometric invariance of non-Lambertian surfaces. Along these lines, it is perhaps the work of Beckmann [1] on smooth and rough surface reflectance that is the best known in the vision and graphics communities. While it is based on physically meaningful surface parameters, the Beckmann theory relies on the evaluation of the Kirchoff wave scattering integral and breaks down when the incident or the scattering angle is large. Vernold and Harvey [15] have overcome this latter problem by developing a model which accounts for self shadowing on rough surfaces. Torrance and Sparrow [14] have developed a physically realistic account of specular reflectance that models the angular distribution of surface microfacets.




Figure 1. Reflection Geometry.

2. Motivation

Here, we focus on the constraints for simultaneous recovery of object shape and photometric parameters from single view images. This is not only practically useful, but theoretically relevant since these constraints provide a unified setting and understanding of photometric invariant and shape recovery methods based upon variational techniques. Moreover, the differential equations presented here, which govern the relation between shape and reflectance are consistent with prior literature in the fields of shape from shading, photometric stereo and regularisation theory.

Thus, the paper is organised as follows, in Section 3, we provide the background on the relationship between photometric invariants and shape recovery. In this section, we provide a general characterisation for a number of reflectance models based upon Fresnel theory and provide a formulation for the cost function which we aim to minimise. In Section 4, we examine in more detail a number of reflectance models for which the general formulation in Section 3.1 applies. In Section 5, we provide experimental results on shape and photometric variable recovery.

3. Photometric Invariants and Shape Recovery

In this section, we provide a unifying formulation that applies to a family of reflectance models based on the Fresnel reflection theory. This family is comprised by those models which depend on three sets of variables related to the reflection geometry, the material's index of refraction and the microfacet structure. More specifically, these reflectance models, in general, involve a wavelength dependent Fresnel term that accounts for the reflection, transmission and refraction of light through the boundary between different object media. This term, in turn, depends on the incident angle and index of refraction.

As a result of the wavelength dependence of the Fresnel terms, we cast the problem in a general setting so as to recover the photometric invariants of materials and surface shape from multispectral imagery. It is worth stressing that, the theory here is equally applicable to monochromatic

or trichromatic imagery by fixing the discrete wavelength-indexed terms accordingly. Moreover, we also state the constraints upon which the problem is well-defined in terms of the number of illumination directions needed to fit the reflectance model and recover the shape of the object under study.

In the reflectance models herein, the reflection geometry is defined with respect to a local coordinate system whose origin is the surface location and whose z-axis is aligned to the normalised local surface normal \vec{N} . The incident light direction \vec{L} is defined by the zenith and azimuth angles θ_i and ϕ_i , respectively. Accordingly, the zenith and azimuth angles of the outgoing direction \vec{V} are θ_s and ϕ_s . For the sake of simplicity, we assume that the incident light is always in the xz -plane, *i.e.* $\phi_i = \pi$. Alternatively, reflectance models can be parameterised with another set of angles, making use of the half-way vector $\vec{H} = \vec{L} + \vec{V}$, which is the sum of the unit-length vectors in the light and viewing directions. Note that the reflection geometry can be equivalently represented by the angular difference θ_d between \vec{L} and \vec{H} , the half angle θ_h between \vec{N} and \vec{H} and the incident angle θ_i . The geometry above is illustrated in Figure 1.

3.1. General Reflectance Model

The reflectance equation at a surface location u and wavelength λ can be, in general, written as a product of two functions $\Gamma(\cdot)$ and $\Lambda(\cdot)$ as follows

$$f(u, \lambda) = \Gamma(\Theta(u), n(u, \lambda))\Lambda(\Theta(u), \Omega(u)) \quad (1)$$

where $\Theta(u)$ is the set of reflection-angle variables describing the incident light, viewing and local surface normal directions, $n(u, \lambda)$ is the wavelength-dependent index of refraction of the surface material under monochromatic light, $\Omega(u)$ is the set of photogrammetric variables, such as the local microfacet slope and its distribution, as well as the masking and shadowing factors. In the model above, $\Theta(u)$ and $\Omega(u)$ are wavelength-independent.

In the general formulation above, the function $\Gamma(\cdot)$ in Equation 1 involves a Fresnel term, which is directly related to the incident angle θ_i and the index of refraction $n(u, \lambda)$. The function $\Lambda(\cdot)$ depends solely on the light source and viewer direction, the surface orientation and the microfacet scattering parameters. In Section 4, we present a number of reflectance models that comply with the general formulation above by showing their correspondences to the general reflectance equation 1. For now, we focus on the constraints governing the optimisation of the cost function corresponding to the reflectance equation above.

3.2. Constraints for Shape and Photometric Variable Recovery

At input, we have a set of M multispectral images $\mathcal{R}_1, \mathcal{R}_2, \dots, \mathcal{R}_M$, where each image \mathcal{R}_l is taken under a different illuminant direction with a known power spectrum. In addition, all the images are observed from the same view point. Each of these images is indexed to the wavelength $\lambda \in \{\lambda_1, \dots, \lambda_K\}$, where $R_l(u, \lambda)$ is the measured spectral reflectance at the pixel-site u on the l^{th} image after being normalised by the respective illumination power spectrum.

From the reflectance images, we aim to recover the model parameters satisfying Equation 1. To solve this problem, we commence by noting that the reflection angles depend on the illumination direction \vec{L} , the viewing direction \vec{V} and the surface orientation \vec{N} . While the surface normal and viewing direction are fixed for all the input images, the angles $\Theta_l(u)$ at every pixel site in the l^{th} image vary with respect to the illumination direction. On the other hand, the parameters $n(u, \lambda)$ and $\Omega(u)$ are invariant to illuminant power and direction.

Let us consider a global coordinate system where the origin is located at the view point, the positive z -axis is in the opposite direction to the line of sight and the positive x -axis points toward the right-hand side of the field of view. In this coordinate system, the function $z(x, y)$ is, in effect, the height map at the surface location (x, y) that corresponds to the pixel location u . The surface normal at (x, y) is, by definition, $\vec{N} = [-p(u), -q(u), 1]^T$, where $p(u)$ and $q(u)$ are the surface gradients, i.e. $p(u) = \frac{\partial z(x, y)}{\partial x} = z_x$ and $q(u) = \frac{\partial z(x, y)}{\partial y} = z_y$.

Therefore, with known illuminant directions, we can reparameterise the general reflectance equation 1 with respect to the surface gradients $p(u)$, $q(u)$ and index it with respect to the image number l as follows

$$f_l(u, \lambda) = \Phi_l(p(u), q(u), n(u, \lambda))\Psi_l(p(u), q(u), \Omega(u))$$

This representation offers the advantage of replacing the reflection angle-variables with the surface gradients. These are invariant with respect to the illuminant direction. More importantly, this formulation implies that the number of geometric variables is two per pixel-site u , which, in turn constraints the number of reflectance equations needed to recover the surface shape. With the new parameterisation, we aim to solve the following system of equations

$$f_l(u, \lambda) = R_l(u, \lambda), \forall l = 1, \dots, M$$

Assuming that the number of pixels in each image is N , the system above consists of $M \times N \times K$ equations with $(|\Omega| + K + 2) \times N$ independent variables. These include K wavelength dependent refractive indices at each pixel and

the number of microfacet scattering variables $|\Omega|$. Without further constraints, this system is only well-defined if and only if the number of equations at least equals the number of variables. In other words, the problem is only solvable given at least $M \geq (|\Omega| + K + 2)/K$ images. For all reflectance models in Section 4, this number is lower-bounded at 2. In summary, this is the theoretical lower bound on the number of illuminants needed in order to recover the surface shape and photometric invariants when the illuminant directions are known.

Note that this lower bound is consistent with the literature of photometric stereo methods for grayscale images with the Lambertian model, where $K = 1$ and $|\Omega| = 0$. In fact, the formulation above can be viewed as a generalisation of Woodham's photometric stereo problem [17] for arbitrary reflectance models.

3.3. Optimisation Approach

When only a single image is at hand, the shape and photometric parameter recovery becomes ill-posed. In this case, we resort to constraints based on implicit assumptions on the spectral and local surface variations. Specifically, when the surface under study is smooth, we can enforce the surface integrability constraint, which has been utilised in the field of Shape-From-Shading by several authors [4, 5]. This constraint states that the cross-partial derivatives of $p(u)$ and $q(u)$ must be equal, i.e. $p_y(u) = q_x(u)$. Furthermore, one can assume smoothness on the spectral variation of the refractive index. This assumption implies that the surface scattering characteristics should vary smoothly across the spatial domain. As we will show later in this section, these assumptions permit the recovery of the shape and photometric parameters making use of a line-search approach.

To take our analysis further, we commence by noting that the parameters satisfying the reflectance equations are the minimisers of the following cost function

$$\begin{aligned} C = & \sum_{l=1}^M \int_{\mathcal{I}} \int_{\mathcal{W}} (R_l(u, \lambda) - f_l(u, \lambda))^2 dud\lambda \\ & + \alpha \int_{\mathcal{I}} \left(\frac{\partial p(u)}{\partial y} - \frac{\partial q(u)}{\partial x} \right)^2 du \\ & + \beta \int_{\mathcal{I}} \int_{\mathcal{W}} \left(\frac{\partial n(u, \lambda)}{\partial \lambda} \right)^2 dud\lambda \\ & + \gamma \int_{\mathcal{I}} \left[\left(\frac{\partial \Omega(u)}{\partial x} \right)^2 + \left(\frac{\partial \Omega(u)}{\partial y} \right)^2 \right] du \quad (2) \end{aligned}$$

where \mathcal{I} is the image spatial domain and \mathcal{W} is the wavelength range.

The arguments of the cost function C are the surface gradients $p(u)$ and $q(u)$, the index of refraction $n(u, \lambda)$ and

$$\begin{aligned}
\sum_{l=1}^M \int_{\mathcal{W}} (R_l(u, \lambda) - f_l(u, \lambda)) \frac{\partial f_l(u, \lambda)}{\partial p(u)} d\lambda + \alpha(p_{yy}(u) - q_{xy}(u)) &= 0 \\
\sum_{l=1}^M \int_{\mathcal{W}} (R_l(u, \lambda) - f_l(u, \lambda)) \frac{\partial f_l(u, \lambda)}{\partial q(u)} d\lambda + \alpha(q_{xx}(u) - p_{yx}(u)) &= 0 \\
\sum_{l=1}^M \int_{\mathcal{I}} (R_l(u, \lambda) - f_l(u, \lambda)) \frac{\partial f_l(u, \lambda)}{\partial n(\lambda)} du + \beta \frac{\partial^2 n(\lambda)}{d\lambda^2} &= 0 \\
\sum_{l=1}^M \int_{\mathcal{W}} (R_l(u, \lambda) - f_l(u, \lambda)) \frac{\partial f_l(u, \lambda)}{\partial \Omega(u)} d\lambda + \gamma (\Omega_{xx}(u) + \Omega_{yy}(u)) &= 0
\end{aligned}$$

Figure 2. Euler-Lagrange equations for the cost function C .

$$\begin{aligned}
p^{(t+1)}(u) &= \bar{p}^{(t)}(u) - \frac{1}{2} \hat{q}^{(t)}(u) + \frac{\epsilon_s^2}{2\alpha} \times \sum_{l=1}^M \sum_{k=1}^K (R_l(u, \lambda_k) - f_l^{(t)}(u, \lambda_k)) \frac{\partial f_l^{(t)}(u, \lambda_k)}{\partial p^{(t)}(u)} \\
q^{(t+1)}(u) &= \bar{q}^{(t)}(u) - \frac{1}{2} \hat{p}^{(t)}(u) + \frac{\epsilon_s^2}{2\alpha} \times \sum_{l=1}^M \sum_{k=1}^K (R_l(u, \lambda_k) - f_l^{(t)}(u, \lambda_k)) \frac{\partial f_l^{(t)}(u, \lambda_k)}{\partial q^{(t)}(u)} \\
n^{(t+1)}(\lambda_k) &= \bar{n}^{(t)}(\lambda_k) + \frac{\epsilon_w^2}{2\beta} \times \sum_{l=1}^M \sum_{u \in \mathcal{I}} (R_l(u, \lambda_k) - f_l^{(t)}(u, \lambda_k)) \frac{\partial f_l^{(t)}(u, \lambda_k)}{\partial n^{(t)}(\lambda_k)} \\
\Omega^{(t+1)}(u) &= \bar{\Omega}^{(t)}(u) + \frac{\epsilon_s^2}{4\gamma} \times \sum_{l=1}^M \sum_{k=1}^K (R_l(u, \lambda_k) - f_l^{(t)}(u, \lambda_k)) \frac{\partial f_l^{(t)}(u, \lambda_k)}{\partial \Omega^{(t)}(u)}
\end{aligned}$$

Figure 3. Line-search update equations for the optimisation with respect to the shape and photometric parameters.

the photogrammetric parameter-set $\Omega(u)$. The weights α , β and γ control the contribution to the cost function of the integrability constraint, the spectral smoothness constraint on the refractive index and spatial smoothness constraint on the surface scattering variables, respectively.

With the cost function C at hand, we derive an iterative scheme based upon the Euler-Lagrange equations so as to minimise the functional above. The resulting Euler-Lagrange equations with respect to the argument functions are shown in Figure 2. In the equations, the x, y subscripts imply partial derivatives with respect to the corresponding axis-variable.

Moreover, we can employ the discrete approximation of the higher order derivatives in the equations shown in Figure 2. To this end, let the spatial domain be discretised into a lattice with a spacing of ϵ_s between adjacent pixel-sites and the wavelength domain in steps of ϵ_w . We index $p(u)$, $q(u)$ and $\Omega(u)$ according to the pixel coordinates (i, j) and $n(u, \lambda)$ according to the wavelength index k . With these ingredients, the partial derivatives can be approximated using finite differences. By substituting the finite-differences into the Euler-Lagrange equations, we obtain a set of update equations for the model parameters with respect to the iteration number t . In Figure 3, we show the set of resulting update equations, with the superscripts denoting the iteration number. The shorthands are defined as follows, in which $\hat{p}(u)$ and $\hat{q}(u)$ are the cross-derivatives of $p(u)$ and $q(u)$ approximated by finite difference (times ϵ_s^2).

$$\begin{aligned}
\bar{p}(u)|_{u=(i,j)} &\triangleq (p_{i,j-1} + p_{i,j+1})/2 \\
\bar{q}(u)|_{u=(i,j)} &\triangleq (q_{i-1,j} + q_{i+1,j})/2 \\
\hat{p}(u)|_{u=(i,j)} &\triangleq (p_{i+1,j+1} + p_{i-1,j-1} \\
&\quad - p_{i-1,j+1} - p_{i+1,j-1})/4 \\
\hat{q}(u)|_{u=(i,j)} &\triangleq (q_{i+1,j+1} + q_{i-1,j-1} \\
&\quad - q_{i-1,j+1} - q_{i+1,j-1})/4 \\
\bar{n}(\lambda_k) &\triangleq (n(\lambda_{k-1}) + n(\lambda_{k+1}))/2 \\
\bar{\Omega}(u)|_{u=(i,j)} &\triangleq (\Omega_{i,j-1} + \Omega_{i,j+1} + \Omega_{i-1,j} + \Omega_{i+1,j})/4
\end{aligned}$$

Note that the second right-hand terms of the update equations are the closed-form negative partial derivatives of the data closeness term with respect to each model parameter. These formulae are in fact instances of line-search optimisation where a search from the current iterate is performed in the steepest gradient descent direction. Here, it is revealed that the Euler-Lagrange equation in the function space is equivalent to gradient descent optimisation in the parameter space.

To enforce numerical stability on the update of parameter values over iterations, we introduce a step length along the steepest descent direction. To this end, we employ the Wolfe condition [10] to ensure that the step length delivers a sufficient decrease in the target function. For each update of the model parameters, we perform a backtracking line search approach by starting with an initial long step length and contracting it upon insufficient decrease of the target function.

4. Reflectance Models Based Upon Fresnel Theory

Note that, so far, we have formulated the constraints in the previous sections in such a manner that the reflectance model under consideration is general in nature. In this section, we show how the general reflectance model presented in Equation 1 captures a family of existing reflectance models in the literature. We do this by establishing a correspondence between the generic parameter sets in the general model and those specific to some of the models used by the computer vision and graphics communities. This is important since it provides a proof of concept that the process of model parameter recovery presented above can be performed on each of these reflectance models at hand. It also provides an explicit link between the equations above and the reflectance models in the literature.

General model	Θ	Ω	$\Gamma(\Theta, n)$	$\Lambda(\Theta, \Omega)$
Beckmann-Kirchoff	θ_i, θ_h	m	$\lambda^2 F(\theta_i, n(\lambda))$	$\frac{m}{\cos^2 \theta_i \cos^4 \theta_h} e^{-\frac{m}{4} \tan^2 \theta_h}$
Vernold-Harvey	$\theta_i, \theta_d, \theta_h$	m	$\lambda^2 F(\theta_i, n(\lambda))$	$\frac{m \cos \theta_i}{\cos^2 \theta_d \cos^2 \theta_h} e^{-\frac{m}{4} \tan^2 \theta_h}$
Torrance-Sparrow	$\theta_i, \theta_s, \phi_s$	$A_f, \vartheta, \pi_{\vartheta}, \sigma_{\vartheta}$	$F(\theta_i, n(\lambda))$	$\frac{A_f}{4} \frac{G_{TS}(\theta_{ip}, \theta_{sp})}{\cos \theta_s} P(\vartheta)$
Wolff	θ_i, θ_s	ρ_W	$[1 - F(\theta_i, n(\lambda))] \times [1 - F(\theta_s, 1/n(\lambda))]$	$\rho_W \cos \theta_i$

Table 1. Parameter correspondences between specific reflectance models and the general model in Section 3.1.

The Fresnel theory has been used extensively in the optics, computer vision and graphics literature to derive reflectance models. Among the physics-based models, the Beckmann-Kirchoff model [1] originated from Kirchoff's theory on the scattering of electromagnetic waves. Torrance and Sparrow [14] employed the Fresnel reflection coefficient to model specular reflection. Wolff [16] derived a diffuse reflectance model for layered dielectrics by analysing the reflection, transmission and refraction of light at the surface boundary.

The models above all have parameters corresponding to surface scattering, reflection geometry and Fresnel reflection coefficients. The parameter correspondences between the general model presented earlier and the following specific models are summarised in Table 1. In the following subsections, we elaborate further on the parameters in Table 1.

4.1. Beckmann-Kirchoff Model

With the reflection angles as described in Figure 1, the Beckmann-Kirchoff model [1] predicts the mean scattered power from a surface point u at wavelength λ as a summation of two terms. The first of these represents the scattering component in the specular direction. The second term corresponds to the diffuse scattering component.

Here we focus our attention on the diffuse scattering component for very rough surfaces. Under normalised illuminant power, the surface reflectance is the same as the diffused scattered power. By far, the two most popular approximations of the diffuse reflectance are formulated in the cases of Gaussian and exponential surface correlation functions [1]. When the surface is very rough and the correlation function is Gaussian, the diffuse reflectance at a given wavelength λ of incident light from a surface patch of area A is approximated by

$$f_{BK}(u, \lambda) = F(\theta_i, n(\lambda)) \frac{\pi T^2 G_{BK}^2}{A \sigma^2 v_z^2} \exp\left(-\frac{T^2 v_{xy}^2}{4 \sigma^2 v_z^2}\right) \quad (3)$$

where, as before, the Fresnel reflection coefficient $F(\theta_i, n(u, \lambda))$ is wavelength dependent via the index of refraction $n(u, \lambda)$ and we have $v_x = k(\sin \theta_i - \sin \theta_s \cos \phi_s)$, $v_y = -k \sin \theta_s \sin \phi_s$, $v_z = -k(\cos \theta_i + \cos \theta_s)$, $v_{xy}^2 =$

$v_x^2 + v_y^2$, $g = \sigma^2 v_z^2$. Here, k is the propagation rate of the incident light, related to its wavelength λ through the equation $k = \frac{2\pi}{\lambda}$.

In Equation 3, σ is the standard deviation of the height variation with respect to the mean surface level and the surface correlation length T gives the relative horizontal spacing between the micro-surface extrema. Note that, the surface slope parameter $\frac{\sigma}{T}$ controls the scattering behaviour for various degrees of roughness. Therefore, it is sufficient to estimate $\frac{\sigma}{T}$ from reflectance data rather than each parameter σ and T separately. In other words, this is equivalent to estimating the parameter $m = \left(\frac{T}{\sigma}\right)^2$, which is the square inverse of the surface slope.

In Equations 3 the geometric factor G_{BK} explains the attenuation of emitted light by the surface orientation with respect to illuminant and viewing directions. The geometric factor is defined as

$$G_{BK}(\theta_i, \theta_s, \phi_s) = \frac{1 + \cos \theta_i \cos \theta_s - \sin \theta_i \sin \theta_s \cos \phi_s}{\cos \theta_i (\cos \theta_i + \cos \theta_s)} \quad (4)$$

To obtain the set of parameters as per Equation 1, we reparameterise the reflection geometry with respect to the incident angle θ_i , the difference angle θ_d and the half angle θ_h , as follows

$$f_{BK}(u, \lambda) = \frac{\lambda^2 F(\theta_i, n(\lambda)) m}{16\pi A \cos^2 \theta_i \cos^4 \theta_h} \exp\left(-\frac{m}{4} \tan^2 \theta_h\right) \quad (5)$$

4.2. Vernold-Harvey Model

It has been noted that the Beckmann-Kirchoff model commonly breaks down at large incident and scattering angles [11, 15]. This is since the geometric factor G_{BK} tends to infinity near the grazing angle. Vernold and Harvey [15] have proposed a variant of the Beckmann-Kirchoff model which can cope well with a wide range of angles. In their work, Vernold and Harvey presented a modified geometric factor $G_{VH}^2 = \cos \theta_i$ instead of G_{BK}^2 .

The modified Beckmann-Kirchoff model proposed by Vernold and Harvey can also be parameterised with respect to the half-vector angles as

$$f_{VH}(u, \lambda) = \frac{\lambda^2 F(\theta_i, n(\lambda)) m \cos \theta_i}{16\pi A \cos^2 \theta_d \cos^2 \theta_h} \exp\left(-\frac{m}{4} \tan^2 \theta_h\right) \quad (6)$$

$$\begin{aligned}
p^{(t+1)}(u) &= \bar{p}^{(t)}(u) - \frac{1}{2} \hat{q}^{(t)}(u) + \frac{\epsilon_s^2}{2\alpha} \times \sum_{l=1}^M \sum_{k=1}^K (\log R_l(u, \lambda_k) - h_l^{(t)}(u, \lambda_k)) \times \\
&\quad \left[\left(\frac{1}{F(\theta_i^{(t)}(u), n^{(t)}(\lambda_k))} \frac{\partial F(\theta_i^{(t)}(u), n^{(t)}(\lambda_k))}{\partial \cos \theta_i^{(t)}(u)} + \frac{1}{\cos \theta_i^{(t)}(u)} \right) \frac{\partial \cos \theta_i^{(t)}(u)}{\partial p^{(t)}(u)} + \left(\frac{m^{(t)}(u)}{2 \cos^3 \theta_h^{(t)}(u)} - \frac{2}{\cos \theta_h^{(t)}(u)} \right) \frac{\partial \cos \theta_h^{(t)}(u)}{\partial p^{(t)}(u)} \right] \\
q^{(t+1)}(u) &= \bar{q}^{(t)}(u) - \frac{1}{2} \hat{p}^{(t)}(u) + \frac{\epsilon_s^2}{2\alpha} \times \sum_{l=1}^M \sum_{k=1}^K (\log R_l(u, \lambda_k) - h_l^{(t)}(u, \lambda_k)) \times \\
&\quad \left[\left(\frac{1}{F(\theta_i^{(t)}(u), n^{(t)}(\lambda_k))} \frac{\partial F(\theta_i^{(t)}(u), n^{(t)}(\lambda_k))}{\partial \cos \theta_i^{(t)}(u)} + \frac{1}{\cos \theta_i^{(t)}(u)} \right) \frac{\partial \cos \theta_i^{(t)}(u)}{\partial q^{(t)}(u)} + \left(\frac{m^{(t)}(u)}{2 \cos^3 \theta_h^{(t)}(u)} - \frac{2}{\cos \theta_h^{(t)}(u)} \right) \frac{\partial \cos \theta_h^{(t)}(u)}{\partial q^{(t)}(u)} \right] \\
n^{(t+1)}(\lambda_k) &= \bar{n}^{(t)}(\lambda_k) + \frac{\epsilon_w}{2\beta} \times \sum_{l=1}^M \sum_{u \in \mathcal{I}} (\log R_l(u, \lambda_k) - h_l^{(t)}(u, \lambda_k)) \frac{1}{F(\theta_i^{(t)}(u), n^{(t)}(\lambda_k))} \frac{\partial F(\theta_i^{(t)}(u), n^{(t)}(\lambda_k))}{\partial n^{(t)}(\lambda_k)} \\
m^{(t+1)}(u) &= \bar{m}^{(t)}(u) + \frac{\epsilon_s^2}{4\gamma} \times \sum_{l=1}^M \sum_{k=1}^K (\log R_l(u, \lambda_k) - h_l^{(t)}(u, \lambda_k)) \left(\frac{1}{m^{(t)}(u)} - \frac{1}{4} \tan^2 \theta_h^{(t)}(u) \right)
\end{aligned}$$

Figure 4. Line-search update equations for the Vernold-Harvey model.

4.3. Torrance-Sparrow Model

Torrance and Sparrow’s model [14] provides an analytical equation of the reflected radiance from mirror-like microfacets whose slopes are randomly distributed. According to the model, the total reflectance from a differential area dA is given by

$$f_{TS}(u, \lambda) = w_d \cos \theta_i + \frac{A_f}{4} F(\theta_i, n(\lambda)) \frac{G_{TS}(\theta_{ip}, \theta_{sp})}{\cos \theta_s} P(\vartheta) \quad (7)$$

In Equation 7, the first term is the diffuse reflectance component that obeys Lambert’s cosine law and is assigned a weight w_d . The latter term is the specular reflectance component. Firstly, A_f is the microfacet’s area. In addition, $G_{TS}(\theta_{ip}, \theta_{sp})$ is the geometric attenuation factor which depends on the projections θ_{ip} and θ_{sp} of the angular variables θ_i and θ_s onto the plane spanned by the facet normal and the mean surface normal \vec{N} . Lastly, ϑ denotes the angle between the facet normal and the mean surface normal.

For isotropic surfaces, the angle ϑ may assume a Gaussian distribution that is rotationally symmetric about the mean surface normal, $P(\vartheta) \sim \pi_{\vartheta} \mathcal{N}(0, \sigma_{\vartheta}^2)$, where $\mathcal{N}(0, \sigma_{\vartheta}^2)$ is a Gaussian distribution with zero mean and a standard deviation σ_{ϑ} .

4.4. Wolff Model

The model proposed by Wolff [16] is derived from the theory of radiative transfer through layered dielectric surface boundaries. To predict the departure in behaviour from the Lambertian model at large angles between the illuminant and viewing directions, Wolff viewed the energy flux emerging through subsurface as a result of refractive and reflective phenomena inside the dielectric body. This model is hence explained through the use of Snell’s law of refraction and the Fresnel attenuation factor. The diffuse subsurface scattering, as defined by the Wolff model, is given by

$$f_W(u, \lambda) = \rho_W \cos \theta_i [1 - F(\theta_i, n(\lambda))] [1 - F(\theta'_s, 1/n(\lambda))] \quad (8)$$

In Equation 8, θ'_s is the zenith angle of light incident on dielectric-air surface boundary before it is refracted and

re-emerges from the surface. This angle is related to the reflection one through Snell’s law making use of the expression $\theta'_s = \arcsin(\sin(\theta_s)/n(\lambda))$. In the equation above, ρ_W is the total diffuse albedo after multiple diffuse subsurface scattering.

5. Experiments

In the following experiments, we fit the Vernold-Harvey model described in Section 4.2 to multispectral images. Since it is analytically more tractable to deal with the Vernold-Harvey model in its log form, we use the following log version

$$\begin{aligned}
h_{VH}(u, \lambda) &= \log \left(\frac{\lambda^2}{16\pi A} \right) + \log F(\theta_i, n(\lambda)) + \log m \\
&+ \log(\cos \theta_i) - 2 \log(\cos \theta_d) - 2 \log(\cos \theta_h) - \frac{m}{4} \tan^2 \theta_h
\end{aligned}$$

The log reflectance equation above yields the update equations in Figure 4 when being substituted into the general form in Figure 3.

5.1. Shape Recovery

In this experiment, we present results on shape recovery from images of synthetically generated surfaces. To obtain the synthetic data, we render images of synthetic shapes using the spectra of refractive index of real-world materials reported in [7]. The images were synthesised using the Vernold-Harvey model with various values of the surface roughness m between 2.5 and 4. Furthermore, we compare our recovered surface normals with those obtained by Worthington and Hancock’s method [18]. This is a regularisation method which takes grayscale Lambertian images at input and delivers surface normals at output.

Table 2 reports the overall accuracy of the needle maps recovered by our method and that of Worthington and Hancock (W&H), across all the materials and surface roughness under study. The error is expressed as the deviation angle per-pixel from the ground-truth needle map, in degrees. Here, our algorithm is comparable to the alternative method in recovering the shape of the Ridge and the Volcano. Note that our algorithm is performed on multispectral images synthesised using Vernold-Harvey model, which is




Figure 5. Needle maps of a Ridge (top) and a Volcano (bottom). First column: Ground truth. Second column: Recovered by our method using a single illuminant. Third column: Recovered by our method using two illuminants. Fourth column: Recovered by Worthington and Hancock’s method.




Figure 6. Skin segmentation results: (a) Input single-band face image at $670nm$, (b)–(d) Skin segmentation maps using: (b) index of refraction recovered from an image captured under a single illuminant, (c) index of refraction recovered from images captured under two illuminants, (d) raw spectral reflectance.

more complex than the Lambertian images input to Worthington and Hancock’s algorithm. Also note that the needle map accuracy is consistent in both the single illuminant and two illuminant cases. This demonstrates the robustness of the algorithm in cases where shadows appear under a light source direction, but not the other.

Figure 5 illustrates the needle maps recovered from the synthetic images. Here, we present the results in the single and two illuminant cases. Qualitatively, our algorithm is able to recover a similar curvature to the ground truth. In addition, surface smoothness as well as surface extrema and discontinuities are also visible in our needle map. Meanwhile, Worthington and Hancock’s algorithm seems to distort the overall shape of the Ridge and creates undesirable curvature near the boundary of the Volcano image.

5.2. Photometric Variable Recovery

In this experiment, we first report the error on the refractive indices recovered from the synthetic images mentioned earlier. Columns 5–6 of Table 2 show the mean and standard deviation of the error on the value of the estimated refractive indices, as compared to their ground truth. These figures imply that the relative error of this estimation is in the order of less than 7% for all the materials under study with an average refractive index of around 1.48. Since we use the same reflectance model to render the simulated data, the optimal index of refraction is supposed to vary consis-

tently across different shapes, as shown in columns 5–6 of Table 2. This result suggests that our algorithm is a potential passive method of refractive index measurement.

Next, we verify the accuracy of the recovered model parameters through skin segmentation on real-world multi-spectral images. To achieve this, we recover the refractive index of materials from multi-spectral images and utilise it as a descriptor for recognition purposes. In addition, we also compare the results of performing our recovery algorithm on single images and pairs of images of the same scene taken under two different light directions. The degree of similarity in the performance in both cases shows that the additional constraint imposed in the single-image case is valid. The experiment is performed on an image database of 50 human subjects, each captured under one of 10 light sources with varying directions and spectral power.

We focus our attention to the photometric invariance of the index of refraction recovered by our algorithm and its applications to recognition tasks. Specifically, we treat the index of refraction as a feature vector for skin segmentation. The segmentation task is viewed as a classification problem where the skin and non skin spectra comprise positive and negative classes. Here we also compare the segmentation accuracy when using the index of refraction with the raw reflectance spectra which has been normalised by the illuminant power.

To obtain the training examples, we select skin and non-skin rectangular regions, of sizes 25×17 and 24×16 respectively, from 10% of the images in the database, from which the index of refraction and raw reflectance are extracted and used as training features. Subsequently, we train a Support Vector Machine (SVM) classifier with a second order polynomial kernel on the training set. The resulting SVM model is applied for classifying skin versus non skin pixels in the remaining images of the subjects.

In Figure 6, we show the original image of a human face captured at the wavelength of $670nm$ in the first column. The skin segmentation maps making use of the index refraction recovered under one and two light source directions are shown in columns 2 and 3. The right most column shows the skin map recovered using the raw reflectance. The brightness of the pixel corresponds to the likelihood of being skin. Note that the first two segmentation maps are similar. This confirms the effectiveness of our regularisers in enforcing additional constraints in the case of single light direction. These constraints, as can be seen, achieve a performance close to that for two light source directions. On the other hand, the raw reflectance spectra result in more false positives and negatives than the index of refraction, which demonstrates that the refractive index is a better photometric invariant for skin recognition purposes.

To support our qualitative results, in Table 3, we quantify the skin segmentation performance with the recovered

	Shape error			Refractive index error	
	One illuminant	Two illuminants	W&H	One illuminant	Two illuminants
Ridge	12.0 ± 12.3	12.8 ± 13.9	11.9 ± 9.4	0.0990 ± 0.0555	0.0990 ± 0.0555
Volcano	18.0 ± 12.44	18.0 ± 12.4	18.17 ± 14.6	0.0990 ± 0.0555	0.0990 ± 0.0555

Table 2. The accuracy of the recovered parameters. Columns 2–3 show the error (in degrees) of the needle maps recovered by our method compared to Worthington and Hancock’s [18] (in column 4). Columns 5–6 show the error of the estimated refractive indices compared to the ground truth. Note that Worthington and Hancock’s method does not recover photometric variables.

	CR(%)	CDR(%)	FDR(%)
Refractive index (1 Illuminant)	86.23 ± 4.27	87.06 ± 10.63	14.05 ± 3.87
Refractive index (2 Illuminants)	87.35 ± 1.59	88.11 ± 4.75	12.72 ± 2.98
Raw reflectance	74.43 ± 0.9	70.67 ± 2.24	21.21 ± 0.44

Table 3. The classification rate (CR), correct detection rate (CDR) and false detection rate (FDR) on facial images using the recovered index of refraction in the one-illuminant and two-illuminant settings and using raw reflectance spectra as features for skin recognition.

index of refraction and raw spectral reflectance as descriptors in terms of the classification rate (CR), the correct detection rate (CDR) and false detection rate (FDR). The correct detection rate is the percentage of skin pixels correctly classified. The false detection rate is the percentage of nonskin pixels incorrectly classified. The classification rate is the overall percentage of skin and nonskin pixels correctly classified. To obtain the result, we randomly select 10% of the image database as training data and compute the average performance over 20 random tests. The results are consistent with the skin maps shown in Figure 6 in the sense that the average segmentation accuracy using the refractive index in the one-illuminant and two-illuminant settings are similar. Furthermore, the overall performance in the two-illuminant setting is more stable due to the additional reflectance equations introduced by the second illuminant direction. On the other hand, the raw reflectance spectra, as before, yields lower performance in both the overall classification rate and correct detection rate. It also suffers from a high number of false positives.

6. Conclusions

In this paper, we have provided a principled link between shape recovery and photometric invariance. We have done this by providing the constraints governing the recovery of the reflectance and object shape parameters. The setting presented here applies to a number of reflectance models used by the computer vision and graphics communities and is consistent with results in shape-from-shading, photometric stereo and regularisation theory. We have shown the utility of this general set of optimisation constraints for purposes of material recognition and shape recovery.

Acknowledgements

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.

References

- [1] P. Beckmann and A. Spizzichino. *The Scattering of Electromagnetic Waves from rough Surfaces*. Pergamon Press, 1963.
- [2] G. Brelstaff and A. Blake. Detecting Specular Reflection Using Lambertian Constraints. In *Int. Conf. on Comp. Vision*, pages 297–302, 1988.
- [3] 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.
- [4] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. In B. K. P. Horn and M. J. Brooks, editors, *Shape from Shading*, pages 89–122. MIT Press, Cambridge, MA, 1989.
- [5] B. K. P. Horn and M. J. Brooks. The variational approach to shape from shading. *CVGIP*, 33(2):174–208, 1986.
- [6] K. Ikeuchi and B. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17(1-3):141–184, 1981.
- [7] S. N. Kasarova, N. G. Sultanova, C. D. Ivanov, and I. D. Nikolov. Analysis of the dispersion of optical plastic materials. *Optical Materials*, 29(11):1481 – 1490, 2007.
- [8] R. Kimmel and A. M. Bruckstein. Tracking level sets by level sets: a method for solving the shape from shading problem. *Computer vision and Image Understanding*, 62(2):47–48, July 1995.
- [9] S. Nayar and R. Bolle. Reflectance based object recognition. *International Journal of Computer Vision*, 17(3):219–240, 1996.
- [10] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, New York, 2nd edition, 2006.
- [11] J. A. Ogilvy. *Theory of wave scattering from random rough surfaces*. Adam Hilger, Bristol, 1991.
- [12] E. Prados and O. Faugeras. Perspective shape from shading and viscosity solutions. In *IEEE Int. Conf. on Computer Vision*, pages II:826–831, 2003.
- [13] H. Ragheb and E. R. Hancock. A probabilistic framework for specular shape-from-shading. *Pattern Recognition*, 36(2):407–427, 2003.
- [14] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9):1105–1114, 1967.
- [15] C. L. Vernold and J. E. Harvey. A Modified Beckmann-Kirchoff Scattering Theory for Non-paraxial Angles. In *Scattering and Surface Roughness*, number 3426 in Proc. of the SPIE, pages 51–56, 1998.
- [16] L. B. Wolff. Diffuse-reflectance model for smooth dielectric surfaces. *Journal of the Optical Society of America*, 11(11):2956–2968, November 1994.
- [17] R. J. Woodham. Photometric method for determining surface orientation from multiple images. In B. K. P. Horn and M. J. Brooks, editors, *Shape from Shading*, pages 513–532. MIT Press, Cambridge, MA, 1989.
- [18] P. L. Worthington and E. R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(12):1250–1267, 1999.
- [19] Q. Zheng and R. Chellappa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 13(7):680–702, 1991.