

Estimating the Surface Radiance Function from Single Images

Antonio Robles-Kelly¹*, Edwin R. Hancock²

¹*NICTA, Building 115, ANU, 0200 ACT, Australia*

²*Department of Computer Science, University of York, York YO1 5DD, UK*

Abstract

This paper describes a simple method for estimating the surface radiance function from single images of smooth surfaces made of materials whose reflectance function is isotropic and monotonic. The method makes implicit use of the Gauss map between the surface and a unit sphere. We assume that the material brightness is monotonic with respect to the angle between the illuminant direction and the surface normal. Under conditions in which the light source and the viewer directions are identical, we show how a tabular representation of the surface radiance function can be estimated using the cumulative distribution of image gradients. Using this tabular representation of the radiance function, surfaces may be rendered under varying light source direction by rotating the corresponding reflectance map on the Gauss sphere about the specular spike direction. We present a sensitivity study on synthetic and real-world imagery. We also present two applications which make use of the estimated radiance function. The first of these illustrates how the radiance function estimates can be used to render objects when the light and viewer directions are no longer coincident. The second application involves applying corrected Lambertian radiance to rough and shiny surfaces.

1 Introduction

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, if a bidirectional reflectance distribution function (BRDF) is at hand then a number of surface analysis tasks may

* Corresponding author.

Email addresses: Antonio.RoblesKelly@nicta.com.au (Antonio Robles-Kelly¹), erh@cs.york.ac.uk (Edwin R. Hancock²).

be addressed. For instance, Nayar and Bolle [1] have used photometric invariants derived from the BRDF to recognise objects with different reflectance properties. This work builds on that reported in [2], where a background to foreground reflectance ratio is introduced. In a related development, Dror *et al.* [3] have shown how surfaces may be classified from single images through the use of reflectance properties. Moreover, although shape-from-shading usually relies on the assumption of Lambertian reflectance [4], if a BRDF is at hand then photometric correction or specular subtraction may be applied as a preprocessing step to improve the results obtained. Alternatively, the BRDF can be used to compute a reflectance map, to which a number of shape-from-shading methods may be applied [5,6].

The reflectance distribution is determined not only by the light source and viewer directions, but also by the material properties of the surface under study. There are a number of ways in which reflectance can depart from the Lambertian case. However, most of the effort in the literature 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. There have been several attempts to remove specularities from images of non-Lambertian objects. For instance Brelstaff and Blake [7] use a simple thresholding strategy to identify specularities on moving curved objects. Ragheb and Hancock [8] have developed a probabilistic framework for specular subtraction which uses the Torrance and Sparrow model to account for the distribution of image brightness. The main limitation of these methods is that they rely on the use of the closed form of the BRDF to characterise the specular spike and lobe. As a result, they lack the generality required to process real-world imagery in an unsupervised or automatic way. Turning our attention to rough surface reflectance, here the main departure from the Lambertian case is due to limb brightening. This results in a loss of contrast at the object limbs and an apparent flattening of perceived relief. There has been recent interest in describing texture as a surface relief phenomenon process using ideas from physics to model the BRDF [9,10]. However, the accurate modeling of rough surface reflectance has proved to be extremely elusive, and is still an active area of research in applied optics [9,11,12].

1.1 Related Literature

The methods used to model or approximate the BRDF can be divided into those that are physics-based, semi-empirical or empirical in nature. Although the literature from physics is vast, it is perhaps the work of Beckmann on smooth and rough surface reflectance that is the best known in the vision and graphics communities [13]. While it is based on physically meaningful surface parameters, the Beckmann theory is both, intractable for analysis problems since it relies on the evaluation of the Kirchhoff wave scattering integral and breaks down when the surface roughness is large or the scattering angle is large. However, recently, Vernold and Harvey [14]

have overcome this latter problem by developing a model which accounts for self shadowing on rough surfaces. Ragheb and Hancock [15] have exploited this modification of the Beckman-Kirchoff theory to develop a means of measuring surface roughness parameters using reflectance measurements. By contrast, in the graphics community it is the development of computationally efficient tools for the purposes of realistic surface rendering that is of primary interest, and, hence, it is empirical models that have been the focus of activity [16,17]. One of the most popular models is that developed by Phong [17]. Torrance and Sparrow [18] have developed a more physically realistic account of specular reflectance that models the angular distribution of surface microfacets. Lewis [19] has also derived a more physically plausible variant of the Phong model. Additionally, Schlick [20] has presented a useful numerical simplification of Phong's model. A survey of reflectance and shading models can be found in [21].

However, neither the models developed in physics nor the computational models developed in graphics are well suited for surface analysis tasks in computer vision. It is for this reason that Wolff [22], and, Nayar and Oren[23] have developed phenomenological, or semi-empirical, models that account for departures from Lambertian reflectance. The contribution of Nayar and Oren [23] has been to develop a semi-empirical model for rough surface reflectance that results in limb brightening. Wolff, on the other hand, has a physically-based model for reflectance from shiny surfaces which uses the Fresnel term to model refractive attenuation for incident light [22]. Recently, Wolff et al. [24] have suggested a combination of these two separate models for surfaces with intermediate roughness. However, although these models provide a more accurate account of reflectance from rough and shiny surfaces, there is no real methodology for estimating their underlying parameters. Moreover, due to their dependence on both, viewer and light source direction, the task of approximating the BRDF from a single image is under-constrained.

Despite these efforts, the physical and phenomenological modeling of the BRDF remains an elusive task. An alternative is to make empirical estimates, or learn, the BRDF under controlled lighting and viewing conditions. There have also been attempts to model the reflectance properties of human skin from real-world imagery [25,26]. Hertzmann and Seitz [27] have shown how the BRDF can be recovered making use of a spherical reference object and multiple views of the scene.

The main obstacle to estimating the BRDF is that it has four degrees of freedom, that correspond to the zenith and azimuth angles for the light source and the viewer. As a result, the tabulation of empirical BRDFs can be slow and labour intensive, and the results are not easily reproduced. Notwithstanding these difficulties, the distribution of reflectance with respect to the four angles can be parameterized or represented in a number of alternative ways. These alternatives include spherical harmonics [28], splines [29] and wavelets [30]. The parameterisation can be chosen to reflect the anticipated shape of the BRDF and can embody some simple physics. For instance, Ward [31] has shown how to simplify matters by approxi-

imating the BRDF using a Gaussian lobe. Lafortune and co-workers [32] at Cornell have developed an approximation method that employs a set of reciprocal, energy-preserving functions with non-linear parameter dependence. Finally, Dana and Nayar [33] have provided a method for collecting and cataloguing empirical BRDFs for a large number of surfaces of different physical properties.

1.2 Contribution

In this paper, we confine our attention to the simpler problem of estimating the surface radiance function from single images. The surface radiance function captures the reflectance properties of the surface for fixed illuminant direction and variable viewer direction, and may be viewed as a “slice” of the more complicated BRDF. Our procedure does not involve complex calibration procedures, and, hence, provides a computationally cheap alternative to the complicated setups employed by measurement-intensive approaches. Of course, acquiring a reflectance model for the purposes of photorealistic rendering from a single image is experimentally difficult and time consuming.

We present an essentially non-parametric method for estimating the surface radiance function from image data that avoids using basis functions or using a pre-determined model to characterise the specular spike and limb. Our method makes implicit use of the Gauss map, i.e. the projections of the surface normals onto a unit sphere. We map image brightness values from the surface to locations on the unit sphere which have identical surface normal directions. Under conditions in which the light source direction and the viewer direction are identical, we show how the reflectance function can be represented by a polar function on the unit sphere. We pose the problem of recovering the reflectance function as that of estimating a tabular representation of this polar function. To overcome the problem that we do not have the surface normal correspondences at hand, under conditions of isotropic reflectance, we show how to estimate the zenith angles on the unit sphere using image brightness gradients. A simple analysis shows how the tabular representation of the reflectance function can be obtained using the cumulative distribution of image gradients.

We also provide a sensitivity study and investigate two applications of the estimated reflectance distribution function. The first of these is to show the method can be used for purposes of forward rendering. The second application involves Lambertian radiance re-mapping. This is an important task since it allows both specularities to be removed from shiny surfaces and the compensation of boundary “flattening” effects for rough surfaces. As mentioned earlier, in computer vision, the identification of specularities plays an important role in shape-analysis. For instance, the apparent movement of specularities provides an important shape-cue [34]. If specularities can be subtracted from images, then shape-from-shading may

be applied more effectively to recover accurate surface shape [35]. In particular, if not removed then specularities manifest themselves as high curvature “pimples” on the surface. Similarly, if boundary flattening effects can be compensated, then shape-from-shading may be applied to rough surfaces and more accurate boundary shape recovered. Here we show how to back project a Lambertian reflectance model onto the imaged surface using the inverse mapping between the Gauss sphere and the image.

The outline of the paper is as follows. Section 2 introduces some of the background to the study by defining the BRDF and the geometry of the reflectance process. In Section 3, we describe the critical component of the work reported in this paper, namely how to estimate the zenith angle on the unit sphere using the distribution of image brightness gradients. We examine the practical considerations underpinning the algorithm in Section 5. Section 6 provides an experimental evaluation of the method. This includes an algorithm sensitivity study and results which show how the estimated radiance function can be used for the purposes of forward rendering and extracting corrected Lambertian reflectance. Finally, Section 7 offers conclusions and suggests directions for future investigation.

2 Preliminaries

In this section, we provide the background for our method. Our overall aim is to make an estimate of the reflectance distribution function from a single image of a piecewise smooth surface. We limit our study to those surfaces whose reflectance function is isotropic and monotonically decreasing with respect to the angle between the viewer and the light source. Surfaces of this kind are those of shiny and moderately rough objects. Examples of the materials that satisfy this requirement are porcelain, smooth terracotta and plastics.

We work with images collected using a digital camera. We are interested in a surface illuminated by a parallel beam of light and viewed by a camera which is sufficiently distant from the surface so that perspective effects may be ignored. Let the illuminated surface under study be denoted by $S \in \mathbb{R}^3$. Consider the orthographic projection of the surface S onto the image plane Π . We are interested in the interaction of the light with the surface at a location s . The geometry of the reflectance process can be expressed in terms of unit vectors in the light source \vec{L} , viewer \vec{V} and surface normal \vec{N} directions. The unit vectors have elevation and azimuth angles (θ_L, α_L) , (θ_V, α_V) and (θ_N, α_N) respectively. We use the surface normal vector \vec{N} as a reference and define the elevation and azimuth angle offsets for the light-source and viewer directions, i.e. $\alpha_i = \alpha_N - \alpha_L$, $\alpha_r = \alpha_N - \alpha_V$, $\theta_i = \theta_N - \theta_L$ and $\theta_r = \theta_N - \theta_V$. Suppose that the incident irradiance at the point s on the surface is $f_I(\theta_i, \alpha_i)$ and that the outgoing radiance from the point is $f_O(\theta_r, \alpha_r)$. The BRDF is

defined to be the ratio of the emitted surface radiance to the incident irradiance, i.e.

$$\rho(\theta_i, \alpha_i, \theta_r, \alpha_r) = \frac{f_O(\theta_r, \alpha_r)}{f_I(\theta_i, \alpha_i) \cos(\theta_i) d\omega} \quad (1)$$

where $d\omega$ is the element of solid angle in the viewer direction. The geometry of the reflection process is shown in Figure 1. Suppose that the surface is imaged with a camera whose lens has diameter D , and that the image plane is at a distance z' behind the lens. Further, suppose that the location s with area A on the surface is imaged by a sensor element u of area A' , and that the angle subtended by the sensor element to the axis of the camera is β_u . Based on the conservation of radiant flux, the irradiance at the sensor element is $I_u = f_o(\theta_r, \alpha_r) \frac{\pi}{4} \left(\frac{D}{z'}\right)^2 A' \cos^4 \beta_u$. Hence, the brightness of the pixels in the image are proportional to the radiance from the corresponding surface location. When the image covers only a narrow angle of the scene, one can assume that $\cos \beta_u = 1$ and we can write

$$I_u = \tau f_O(\theta_r, \alpha_r) = \tau \rho(\theta_i, \alpha_i, \theta_r, \alpha_r) \cos \theta_i d\omega$$

where $\tau = \frac{\pi}{4} \left(\frac{D}{z'}\right)^2 A'$.

Our aim in this paper is to present a method for estimating the surface radiance function $f_O(\theta_r, \alpha_r)$ from single images of an illuminated surface. We confine our attention to the case where the light source direction is fixed to be the same as that of the viewer, and we have that $\theta = \theta_i = \theta_r$ and $\alpha = \alpha_i = \alpha_r$. Under these conditions, the angular dependence of the BRDF is determined solely by the direction of the surface normal. Moreover, we assume that the light source is a point at infinity, and hence the irradiance is constant, i.e. $f_I(\theta, \alpha) = k$. As a result, we can estimate the reduced reflectance function

$$\hat{\rho}(\theta, \alpha) = k \frac{f_O(\theta, \alpha)}{\cos(\theta) d\omega} \quad (2)$$

As a result the measured image brightness is proportional to

$$f_O(\theta, \alpha) = \frac{1}{k} \rho(\theta, \alpha) \cos(\theta) \quad (3)$$

Hence, estimating the radiance function $f_O(\theta, \alpha)$ of the surface illuminated under conditions in which the light source and viewer directions are identical provides a means of recovering the reflectance function for the object under study.

If the surface normal directions are known, then the estimation of the radiance function $f_O(\theta, \alpha)$ can be easily effected. If I is the brightness at the surface location where the surface normal has elevation and azimuth angles θ and α , then $f_O(\theta, \alpha) = \frac{I}{\tau}$. The radiance function estimation process can be conveniently represented by the

Gauss map of the surface under study onto a unit sphere. For an orientable surface $S \in \mathbb{R}^3$, the Gauss map $G : S \mapsto \hat{S}$ maps points on the surface S onto locations on the unit sphere \hat{S} which have identical surface normal directions. The surface normal directions can hence be used to map brightness values from the image onto the unit sphere. The normalised polar distribution of brightness values on the unit sphere \hat{S} is the radiance function $f_O(\theta, \alpha)$ for the surface. To avoid ambiguities, we assume that points on the surface with identical surface normal directions have identical brightness values.

Here we are interested in the situation where only a single image is at hand, and the surface normal directions are unknown. Now the mapping of the brightness values from the image of the surface onto the unit sphere is not straightforward. In fact, the task of estimating surface normal directions from measured brightness values is an under-constrained one, which has preoccupied researchers in the field of shape-from-shading for several decades. Even for surfaces which exhibit simple Lambertian reflectance, the problem is not tractable in closed-form. However, it is interesting to note that in the case of Lambertian reflectance, the angle between the light source direction and the surface normal vector is recoverable [36]. If the viewer and the light source directions are identical, then the zenith angle θ of the surface normal is recoverable in closed form (it is simply the inverse cosine of the normalised image brightness). The Gauss map of the brightness onto a unit sphere is also very simple in this case. If the axis of the unit sphere is taken to be parallel to the viewer direction, then circles of latitude have constant brightness, and the “latitude” itself is determined by a function of the brightness. Hence, on the unit sphere, the brightness falls away from its maximum value at the visible pole to zero at the equator.




Fig. 1. Geometry of the reflection process

However, for non-Lambertian reflectance the situation is more complex. In the case of non-Lambertian reflectance, provided that the reflectance properties of the surface are isotropic and homogeneous, the problem is again simplified considerably if the viewer and light source directions are identical. The isotropy assumption will ensure that circles of latitude on the unit sphere will still have constant brightness. However, if the surface normal directions are not at hand then, if we wish to estimate the radiance function $f_O(\theta, \alpha)$, i.e. the distribution of the brightness with latitude, we require a means of estimating the zenith angle θ .

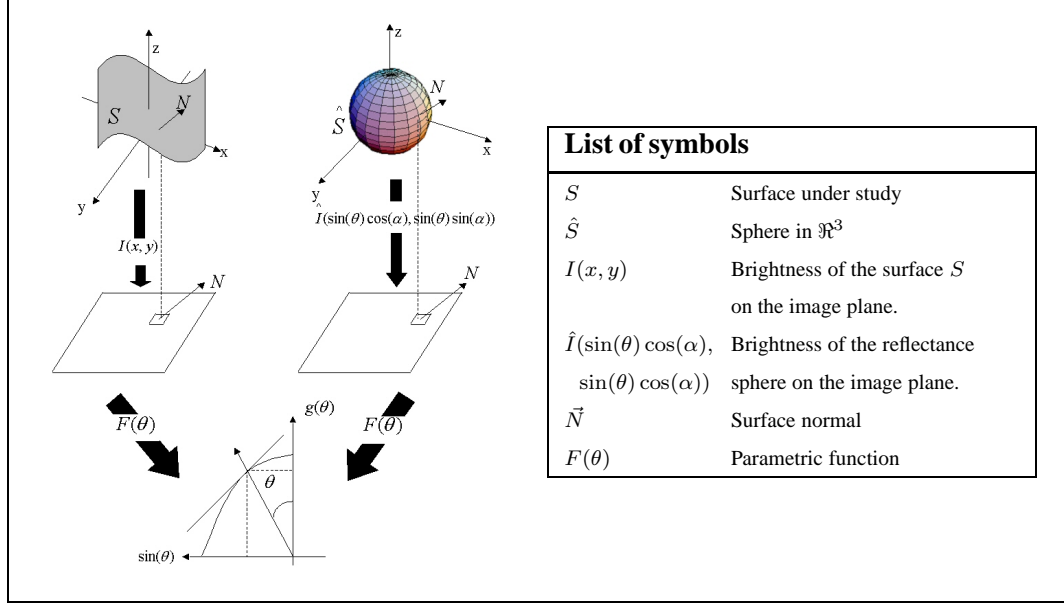


Fig. 2. Structure of the mapping process

3 Radiance Function Estimation

Our main contribution in this paper is to show how the distribution of image-gradients can be used to estimate the zenith angle θ and hence recover the radiance function $f_O(\theta, \alpha)$.

3.1 Theory

To commence, consider a unit sphere where the brightness at the point with elevation angle θ and zenith angle α is $f_O(\theta, \alpha)$. We construct the image of the unit-sphere under projection onto the image plane. The point p on the unit sphere with coordinates (θ, α) maps to the location with Cartesian coordinates $(x, y) = (\sin(\theta) \cos(\alpha), \sin(\theta) \sin(\alpha))$ on the image plane. The brightness at the point (x, y) in the image is

$$\hat{I}(\sin(\theta) \cos(\alpha), \sin(\theta) \sin(\alpha)) = \tau f_O(\theta, \alpha)$$

At the location (x, y) , the image-gradient of the projection is related to the radiance function $f_O(\theta, \alpha)$ by

$$\nabla \hat{I}(\sin(\theta) \cos(\alpha), \sin(\theta) \sin(\alpha)) = \begin{bmatrix} \frac{\partial \hat{I}}{\partial x} \\ \frac{\partial \hat{I}}{\partial y} \end{bmatrix} = \frac{\tau}{\text{Det}(J)} \begin{bmatrix} \frac{\partial f_O(\theta, \alpha)}{\partial \theta} \\ \frac{\partial f_O(\theta, \alpha)}{\partial \alpha} \end{bmatrix} \quad (4)$$

where the Jacobian matrix for the transformation from the unit sphere to the image plane is

$$J = \begin{bmatrix} \frac{\partial x}{\partial \theta} & \frac{\partial x}{\partial \alpha} \\ \frac{\partial y}{\partial \theta} & \frac{\partial y}{\partial \alpha} \end{bmatrix} = \begin{bmatrix} \frac{\partial(\sin(\theta) \cos(\alpha))}{\partial(\theta)} & \frac{\partial(\sin(\theta) \cos(\alpha))}{\partial(\alpha)} \\ \frac{\partial(\sin(\theta) \sin(\alpha))}{\partial(\theta)} & \frac{\partial(\sin(\theta) \sin(\alpha))}{\partial(\alpha)} \end{bmatrix} \quad (5)$$

We exploit this relationship between the image gradient and the derivatives of the radiance function to establish mappings between the brightness values of the imaged surface and the unit sphere. To do this we assume that locations on the surface can be locally approximated by a sphere that has been rendered using the radiance function $f_O(\theta, \alpha)$. Points on the surface and the sphere that are in correspondence under the Gauss map will have identical image gradients. The measured image gradient ∇I can be computed using the formula

$$\nabla I(x, y) = \frac{1}{2\delta} \begin{bmatrix} I(x + \delta, y) - I(x - \delta, y) \\ I(x, y + \delta) - I(x, y - \delta) \end{bmatrix} \quad (6)$$

where δ is the spacing of sites on the pixel lattice. Correspondences between the unit sphere and the observed surface can be established by identifying the location (x, y) in the projected image of the unit sphere where $\nabla I(x, y) = \nabla \hat{I}(\sin(\theta) \cos(\alpha), \sin(\theta) \sin(\alpha))$. The radiance function can be estimated by mapping the measured intensity $I(x, y)$ onto the required location (θ, α) on the unit-sphere, and $f_O(\theta, \alpha) = I(x, y)/\tau$.

However, here we consider a simplified version of this process and restrict ourselves to the case of recovering the zenith angle dependence alone. The reason for this is that accurate estimation of image gradient direction is an elusive task, that is not easily accomplished on the discrete pixel lattice unless either support masks of large extent or significant image smoothing is used. Hence, we restrict ourselves to the case where the radiance function is assumed isotropic, i.e.

$$\frac{\partial f_O(\theta, \alpha)}{\partial \alpha} = 0$$

To make this simplification explicit we consider the generating function $g(\theta)$ which describes the zenith angle dependence of the radiance function. Provided that the radiance function $f_O(\theta, \alpha)$ is monotonically decreasing for $\theta \in [0, \frac{\pi}{2}]$, the generating function is given by

$$g(\theta) = \frac{1}{2\pi} \int_0^{2\pi} \left(f_O(0, \alpha) + \int_0^\theta \frac{\partial f_O(\theta, \alpha)}{\partial \theta} d\theta \right) d\alpha \quad (7)$$

In other words, the generating function $g(\cdot)$ on the unit sphere can be expressed in terms of the cumulative distribution of the derivatives of the surface radiance

function. In particular, the magnitude of the image gradient is given by

$$|\nabla \hat{I}| = \frac{\tau}{\cos(\theta)} \frac{\partial g(\theta)}{\partial \theta} = \tau \frac{\partial g(\theta)}{\partial \sin(\theta)} \quad (8)$$

Solving the differential equation, we find that

$$\sin \theta = \int \frac{\tau}{|\nabla \hat{I}|} dg(\theta) \quad (9)$$

Since $\hat{I} = \tau g(\theta)$, we can write

$$\sin \theta = \int_0^{\hat{I}} \frac{1}{|\nabla \hat{I}|} d\hat{I} \quad (10)$$

Hence, we can use the cumulative distribution of reciprocal gradients to index the zenith angle on the unit sphere.

To recover θ from the expression above we perform numerical integration. To do this, we sort the image gradients according to the associated image brightness values. Accordingly, let ∇I_l be the image gradient associated with the brightness value l . This allows us to compute the function

$$F(l) = \sin(\theta(l)) = \sum_{m=0}^l \frac{\tau}{|\nabla I_m|} + \kappa \quad (11)$$

that returns an estimate of the value of $\sin(\theta(l))$ corresponding to the measured brightness value l , where κ is the integration constant. Provided that the brightness decreases monotonically with zenith angle, then the inverse function $F^{-1}(\sin(\theta(l))) = l$ returns the brightness associated with the zenith angle $\theta(l)$. Hence, the estimate of the generating function $g(\theta)$ for the surface radiance function $f_O(\theta, \alpha)$ is

$$g(\theta) = \frac{1}{\tau} F^{-1}(\sin \theta)$$

3.2 Numerical Implementation

To pursue this idea, in principle, we only require a single image gradient corresponding to each of the distinct brightness levels in the image. In practice, we make use of the cumulative distribution of image gradients in order to minimise the approximation error by averaging. Let $Q_l = \{(x, y) \mid I(x, y) = l\}$ be the set of pixels with brightness value l . For the brightness value l , the average gradient is given by

$$h(l) = \frac{\sum_{(x,y) \in Q_l} |\nabla I(x, y)|}{|Q_l|} \quad (12)$$

The distribution of average gradients is then stored as a vector $\vec{h} = (h(0), \dots, h(\nu))^T$ where ν is the maximum brightness value. Zero entries of the vector, which correspond to brightness values that are not sampled in the image, can cause divide-by-zero errors when the radiance function is computed. To overcome this problem, we smooth the components of the vector by performing piecewise linear interpolation of the adjacent non-zero elements. The resulting vector is denoted by \hat{h} . With the average image gradient at hand, the value of $\sin \theta$ corresponding to the image brightness value l is given by the function

$$\hat{F}(l) = \tau \sum_{i=0}^l \hat{h}(i)^{-1} + \kappa \quad (13)$$

All that remains is to compute the constants τ and κ . Since the minimum value of $\sin(\theta)$ is zero, $\kappa = 1$, and hence

$$\tau = - \left(\sum_{i=0}^{\nu} \hat{h}(i)^{-1} \right)^{-1} \quad (14)$$

Since the brightness values are quantised, our representation of the generator function will also be discretely sampled. In practice, the sampled representation is stored as a vector whose element indexed j is given by

$$U(j) = \hat{F}^{-1} \left(\sin(j \frac{\pi}{2\nu}) \right)$$

Provided that the normalisation constant τ is known, the vector U can be used as a look-up table for the generator function $g(\theta)$.

4 Applications of the Radiance Function

There are clearly a number of applications which can be addressed using an estimate of the surface radiance function. Here we focus on the problems of rendering objects under variable light source direction and mapping rectified Lambertian reflectance onto rough or shiny surfaces. In related work we have investigated more complex tasks including surface material classification [37] and over-exposed image correction [38].

4.1 Variable Illumination Direction

In practice, we will be interested in rendering objects observed from a variety of viewpoints and lighting directions for which the condition $\vec{V} = \vec{L}$ does not hold. Hence, we need to relax the restrictions imposed when estimating the radiance function from the surface S . To do this, we note that the viewer direction \vec{V} is related to the specular spike direction \vec{L}' using the rotation

$$\vec{V} = R_{xyz} \vec{L}' \quad (15)$$

where R_{xyz} is a rotation matrix. The rotation matrix is given by

$$R_{xyz} = \begin{bmatrix} L'_z + L'^2_x c & -L'_x s + L'_x L'_y c & L'_y s + L'_x L'_z c \\ L'_z s + L'_x L'_y c & L'_z + L'^2_y c & -L'_x s + L'_y L'_z c \\ -L'_y s + L'_x L'_z c & L'_x s + L'_y L'_z c & L'_z (1 + L'_z c) \end{bmatrix}$$

where $\vec{L}' = [L'_x, L'_y, L'_z]^T$ is the unit direction-vector for the specular spike, $c = 1 - L'_z$ and $s = \sin(\arccos(L'_z))$. With these ingredients, we can use the matrix R_{xyz} to rotate the surface normal \vec{N} as follows

$$\vec{N}_t = R_{xyz} \vec{N} \quad (16)$$

Hence, when the point t is illuminated from the direction $\vec{L}' \neq \vec{V}$ then the brightness at the point t is given by

$$I_t = \tau g(\arccos(R_{xyz} \vec{N} \cdot \vec{V})) \quad (17)$$

This rendering process can be achieved by performing simple vector operations. Further, since the function $g(\theta)$ can be stored in memory as a floating-point vector of order n_{max} , the rendering can be done in a fast, memory-efficient way.

4.2 Lambertian Re-mapping

In this subsection, we take our study one step further by showing how our method can be used for re-mapping a Lambertian reflectance model onto an imaged surface. This re-mapping may be used either to remove specularities from shiny surfaces, or to correct for reduced boundary contrast for rough surfaces.

The idea underpinning our method is to re-map the image brightness using the inverse Gauss map from the unit sphere onto the original image. The radiance function may be modified in a number of ways. For instance, we could exchange the

```

1  For  $j = 1, 2, \dots, n_{cols} - 1$ 
2      For  $k = 1, 2, \dots, n_{rows} - 1$ 
3           $|\nabla I(j, k)| = \sqrt{(I(j+1, k) - I(j-1, k))^2 + (I(j, k+1) - I(j, k-1))^2}$ ;
4  Blank vectors  $Q$  and  $h$ ;
5  For  $j = 1, 2, \dots, n_{cols} - 1$ 
6      For  $k = 1, 2, \dots, n_{rows} - 1$ 
7           $h(I(j, k)n_{max}) = h(I(j, k)n_{max}) + |\nabla I(j, k)|$ ;
8           $Q(I(j, k)n_{max}) = Q(I(j, k)n_{max}) + 1$ ;
9  For  $j = 0, 1, \dots, n_{max}$ 
10     If  $Q(j) \neq 0$  then  $h(j) = \frac{h(j)}{Q(j)}$ ;
11     else  $h(j) = 0$ ;
12   $h(0) = \epsilon$ ;
13   $\hat{h} =$ Piecewise-linear interpolation using vector  $h$ ;
14   $\hat{F}(0) = 0$ ;
15  For  $j = 1, 2, \dots, n_{max}$ 
16     If  $\hat{h}(j) \geq \varsigma$  then  $h_{old} = \hat{h}(j)$ ;
17      $\hat{F}(j) = \hat{F}(j-1) + h_{old}^{-1}$ ;
18  For  $j = 0, 1, \dots, n_{max}$ 
19      $\hat{F}(j) = 1 - \frac{\hat{F}(j)}{\hat{F}(n_{max})}$ ;
20  Blank vector  $U$ ;
21  For  $j = 0, 1, \dots, \nu$ 
22      $U(\text{floor}(\frac{2\nu}{\pi} \arcsin(\hat{F}(j)))) = j$ ;
23   $U =$ Piecewise-linear interpolation using vector  $U$ ;

```

Fig. 3. Pseudocode for our method

tabular representations acquired from images of different surfaces. Alternatively, an analytic reflectance model can be back-projected. However, here we confine our attention to a simple Lambertian reflectance model.

Formally, our aim is to use the function \hat{F} to retrieve the Lambertian radiance at a given point on the surface S illuminated from a light source with direction vector $(0, 0, 1)^T$. For a given value of image brightness, the function returns the sine of the elevation angle of the surface normals, i.e. $\sin(\theta)$. For Lambertian reflectance, the observed radiance is proportional to the cosine of the angle of light incidence, i.e. to $\cos(\theta)$. Hence, we can perform Lambertian re-illumination by noting the observed brightness l_p at a pixel p and identifying the associated value of $\sin(\theta)$; the corresponding corrected Lambertian radiance is $\cos(\theta)$. Hence, if the pixel brightness at a pixel is l_p then the remapped Lambertian brightness is

$$T(l_p) = \cos\left(\arcsin\left(\tau \sum_{i=1}^{l_p} \hat{h}(i)^{-1} + \kappa\right)\right) \quad (18)$$

5 Practical Considerations

Having presented the theoretical basis for our method, in this section we note a number of important practical considerations underpinning the implementation of the algorithm.

In Figure 3, we present the pseudocode for the algorithm. There are a number of features of the pseudocode that deserve comment. First, the sequence of operations has been slightly altered in order to reduce the number of floating point operations. The complexity of the algorithm, when processing an image of n_{rows} by n_{cols} , whose format can hold n_{max} distinct brightness values, is $3\nu + 8n_{rows}n_{cols} + 5n_{max}$. The algorithm requires $\nu + n_{rows}n_{cols} + 4n_{max}$ memory locations. Second, recall that for $\theta = \frac{\pi}{2}$ the quantity ∇I is, ideally, infinity. Hence, in Line 12 of the pseudocode, we have set $h(0)$ to $\epsilon \gg 1$. Third, as the brightness tends to unity, the inverse of the image intensity gradient tends to infinity. Numerically, this is undesirable since it leads to an unstable estimate of $\sin(\theta)$. To overcome this problem, we require the values $h(\cdot)$ (i.e. the magnitude of the image intensity gradient) to be greater than or equal to the threshold ς . For the cases in which $h(\cdot) < \varsigma$, we use the adjacent values of the vector h so as to satisfy the condition in Line 16 of our pseudocode. In our implementation, the length of the vector U has been fixed to 158, which implies a resolution in θ of approximately 0.56° . We have set the variables ϵ and ς to $3 \max(h)$ and 0.1 respectively.

Finally, we note the following. The vector U is used a lookup-table representation of the inverse of $\tau g(\theta)$. As a consequence of the discrete nature of the approximation used to compute \hat{F} , there is no guarantee that all the values of $l = \{0, 1, \dots, \nu\}$, which correspond to values of θ are captured by the instruction $\text{floor}(\frac{2\nu}{\pi} \arcsin(\hat{F}(j)))$, which returns the largest integer not greater than $\frac{2\nu}{\pi} \arcsin(\hat{F}(j))$. To overcome this problem, we interpolate the missing coefficients of the vector U making use of its non-zero elements.

6 Experiments

In this section, we provide experimental evaluation of the new method for estimating the radiance function. This study is divided into four parts. We commence by illustrating the qualitative properties of the method on synthetic and real-world data. Second, we explore the accuracy of the method against ground truth on both synthetic and real-world imagery ground-truthed using range data. Third, we show results for forward rendering. Finally, we present results on the use of Lambertian correction for specularity removal from shiny surfaces and limb-correction for rough surfaces.

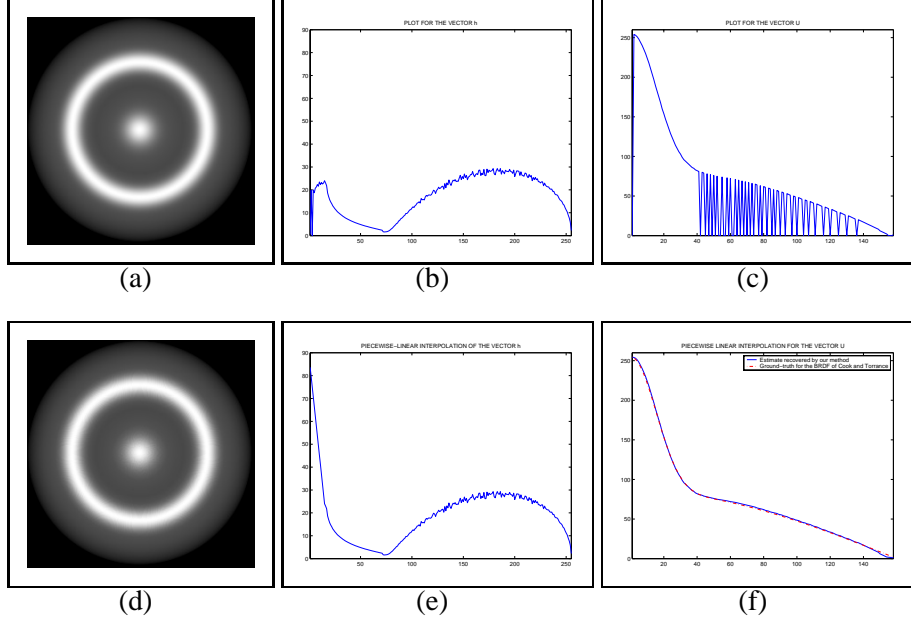


Fig. 4. Step-sequence of the method. Top row, from left-to-right: Input image; plot of the vector h ; vector U before being interpolated. Bottom row, from left-to-right: rendering of the surface using the radiance function recovered by our method; plot of the vector \hat{h} ; vector U delivered by our pseudocode.

6.1 Radiance Function Estimation

In this section we demonstrate some of the features of the radiance function estimation. To this end, we commence by illustrating how the algorithm behaves on synthetic data. Next, we illustrate the method on a variety of real world data.

6.1.1 Algorithm Behaviour

In this section, we illustrate the step-sequence of the algorithm making use of an example synthetic surface. We have generated a surface of a torus with an elliptic dome superimposed. We have rendered the surface making use of the BRDF of Cook and Torrance [39] and set the average microfacet-slope to 0.3 with a diffuse component of Lambertian nature whose magnitude is 0.5. The rendering of the surface, which we have used as input image to our algorithm, is shown in Figure 5(a). Our input image is 510×510 pixels and is quantised into 256 gray levels.

In Figure 5(b), we show the elements of the vector h used for the piecewise-linear interpolation. The elements of the vector \hat{h} returned by the piecewise-linear interpolation routine are shown in Figure 5(e). It is clear that the discontinuities in the vector h have been corrected by the piecewise-linear interpolation step and the estimate of \hat{h} is smooth. The elements of the vector U are shown in Figure 5(c). We have plotted its piecewise-linear interpolation, together with the ground-truth

curve for the Cook and Torrance’s BRDF, in Figure 5(f). There are a number of features to note from this plot. First, the specular structure of the radiance function is clearly visible as a spike near the origin. Second, there are small differences when $\sin(\theta) \simeq 0$, and $\sin(\theta) \simeq 1$, i.e. as the coefficient indexes tend to the extrema, but these are not significant. In Figure 5(d), we show the rendering of the surface making use of the radiance function recovered by the algorithm.

6.1.2 Qualitative Properties of the Radiance Function

Using imagery from a variety of surfaces, we now explore some of the qualitative properties of the radiance functions estimated when our algorithm is applied to a variety of real-world images. The images used in our study are shown in the top row of Figures 5, 6 and 7. The objects studied a porcelain mug, a terracotta urn, a gilt mask, a machined steel part, a pile of billiard balls, a single billiard ball, a terracotta bear, a terracotta teapot and a rough cylinder. They hence have a variety of surface characteristics, and include both shiny and rough surfaces, and also metals and dielectrics. Our images are 24-bit color of size 340 by 340 pixels. We have edited the images so that the background is black and only the object under study is visible. This is a preprocessing step that has no serious implications for the applicability of the method.

In the second row of the figures we show plots of the vector of the average image gradient $\hat{h}(l)$ versus the brightness l . The third row of each figure shows the interpolated vector \hat{h} used to compute the tabular representation of the radiance function. Finally, the function \hat{F}^{-1} is shown as a function of $\sin \theta$ in the fourth row of each figure. The process is applied identically for the three RGB colour-channels comprising the image. Hence, the different curves in the plots in the second, third and fourth rows show the results corresponding to the RGB channels (coloured according to band). The black curve in each plot is the Lambertian radiance curve.

There are a number of features to note from these plots. First, it is clear that the discontinuities in the vector \hat{h} have been corrected by the piecewise-linear interpolation step and the estimate of \hat{F} is smooth. Secondly, the shapes of the reflectance functions are very different for the different objects. Moreover, the radiance curves are different for the different colour bands. Of all the objects studied, it is the red band, for those composed of terracotta (the urn, the bear and the teapot), that produces the radiance curve that is closest to the predictions of Lambert’s law. In the case of the porcelain mug, the curves are close to the predictions of Lambert’s law, however there is a clear specular spike where $\sin(\theta) \simeq 0$. For the metallic objects (e.g. the gilt mask and the machined steel part) the departures from Lambert’s law are most significant.

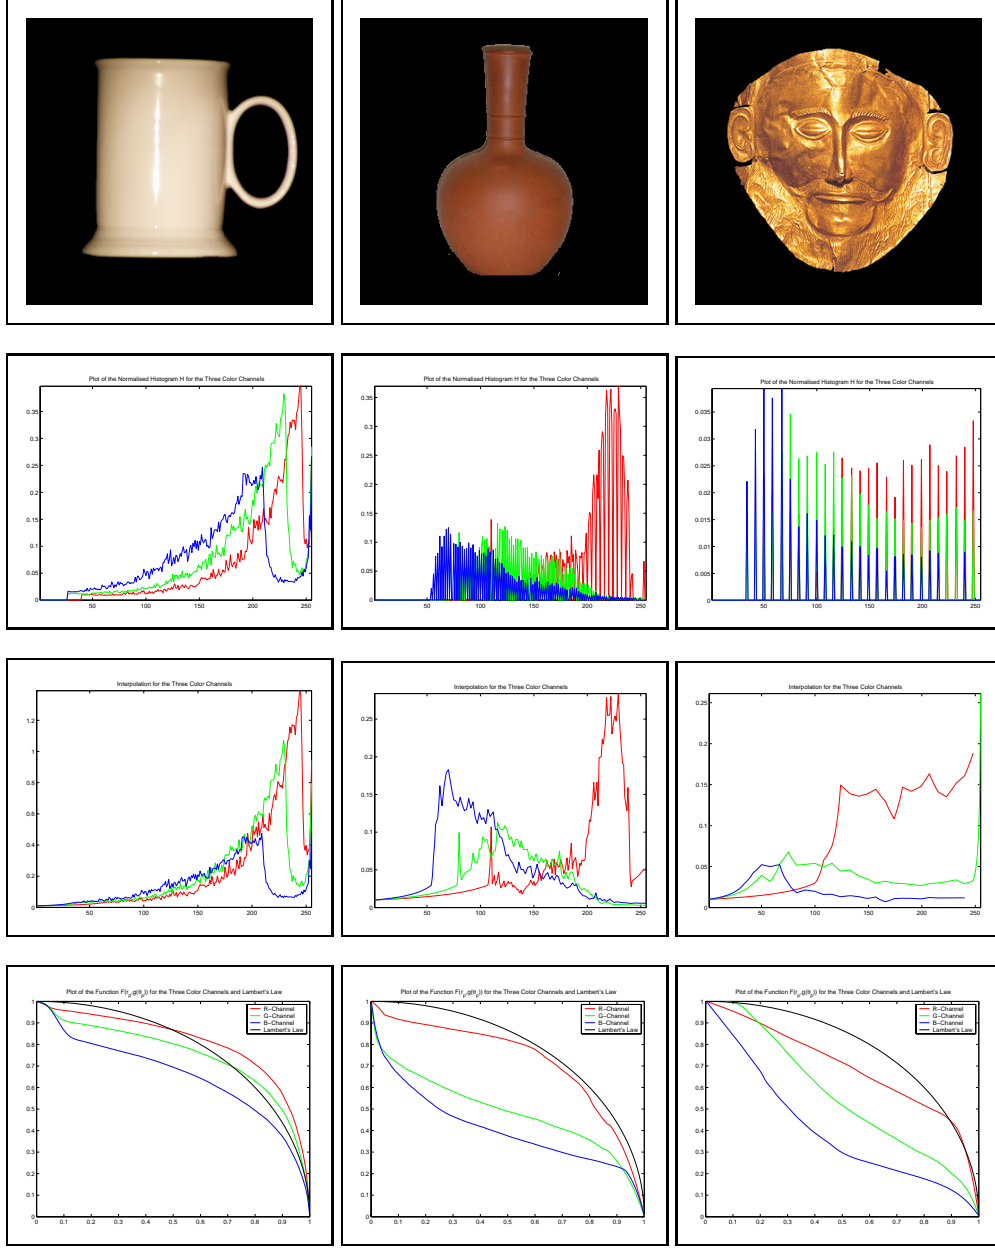


Fig. 5. Input images and plots of the variables computed through the radiance function approximation process.

6.2 Sensitivity Study

We now turn our attention to the accuracy of our radiance function estimation method. For this sensitivity study, we have used synthetic and real-world imagery. Our synthetic data consists of a series of torii and teapots rendered using the BRDFs of Phong [17], Oren and Nayar [40] and Cook and Torrance [39]. Our real-world data is provided by images of samples of sandpaper rolls of different grades, and images of rough terracotta objects and shiny porcelain objects. Our experimen-

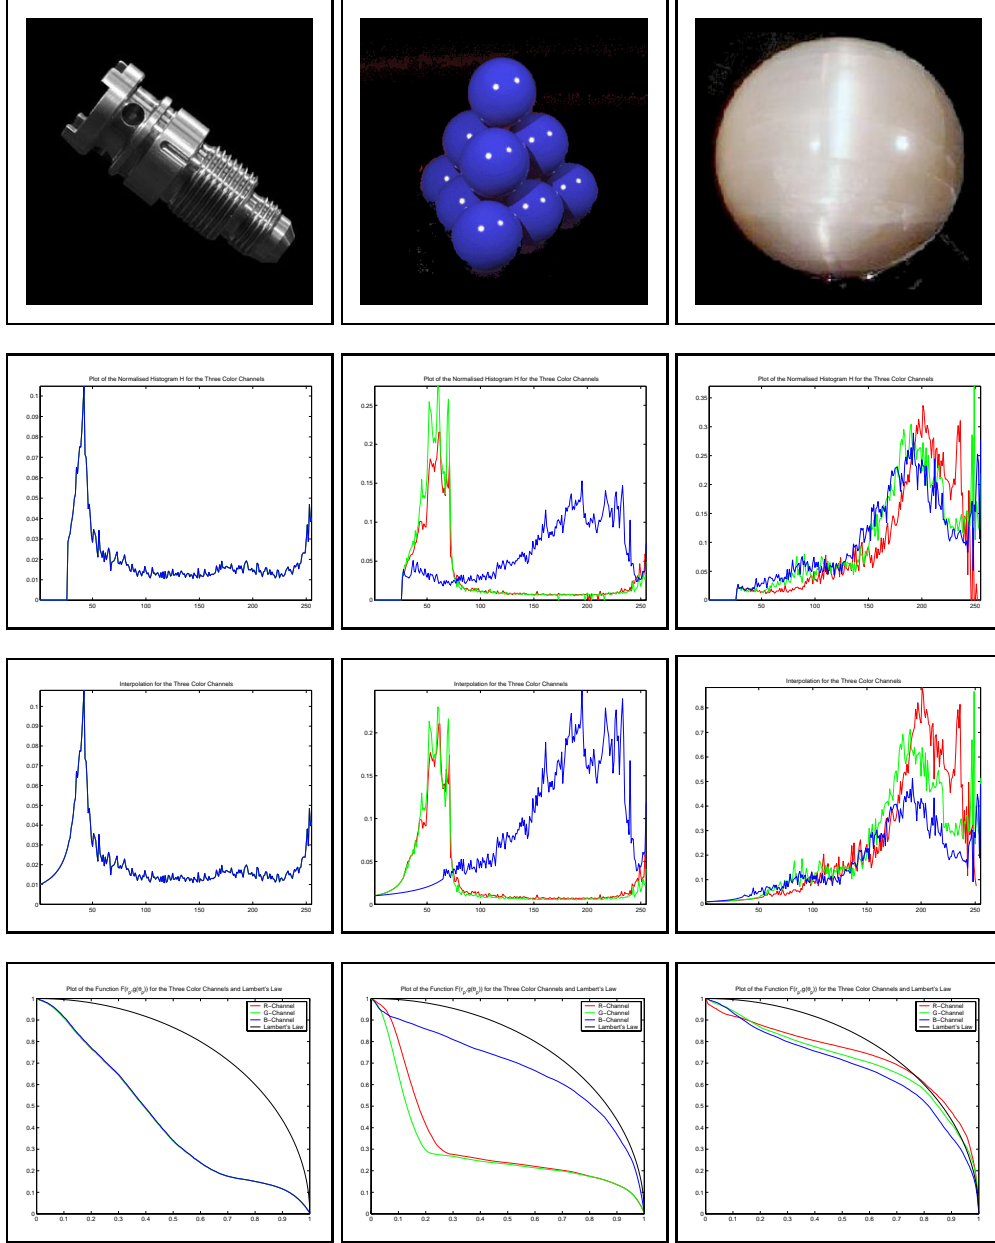


Fig. 6. Input images and plots of the variables computed through the radiance function approximation process.

tal procedure is as follows. Using as input a single image that satisfies the condition $\vec{V} \approx \vec{L}$, we have estimated the surface radiance function for each object. For the different objects under study, we have investigated the effect of variable light source direction. To do this we have acquired surface height data. In the case of the synthetic tori and teapots, the images were generated by rendering a synthetic surface mesh. In the case of the real-world objects, the height data was acquired using a range-sensor and a surface mesh fitted to the range data. Using the estimated radiance function, we have rendered the synthetic and real world surfaces under variable illuminant directions. We have compared these synthesised images

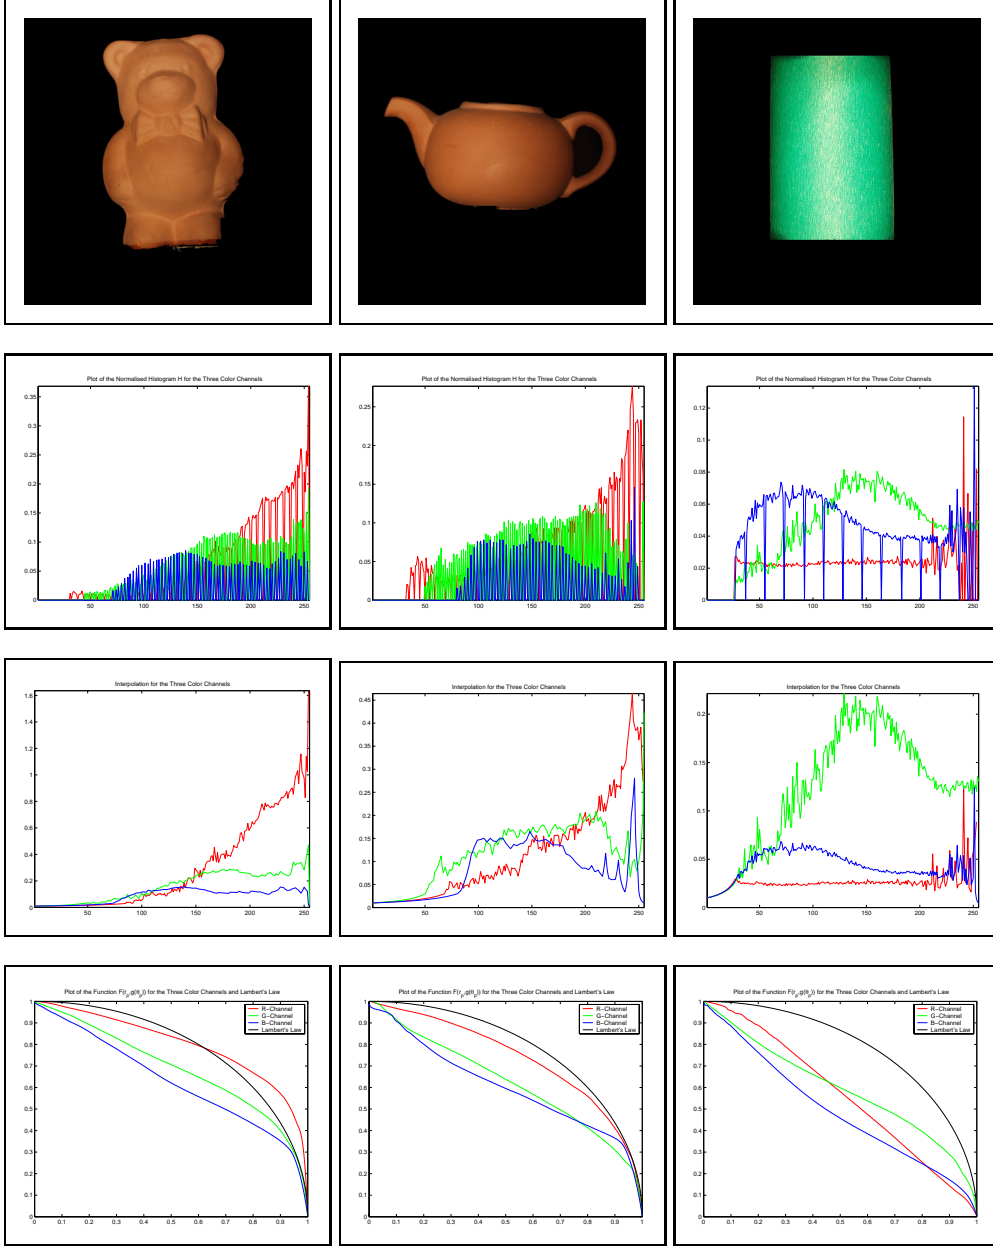


Fig. 7. Input images and plots of the variables computed through the radiance function approximation process.

with real-world ones acquired under similar illumination conditions. We align the synthetic and acquired real-world images, and compute the mean-absolute error between the rendered surface meshes and the real-world imagery. The error for the pair of images with dimensions n_{cols} by n_{rows} with a set of foreground pixels Ω is given by

$$e = \frac{1}{|\Omega|} \sum_{q \in \Omega} |I_M(u) - I_D(u)| \quad (19)$$

where $I_M(u)$ and $I_D(u)$ are the respective brightness values at the pixel u in the real-world and synthesised images.




Fig. 8. Top row: Input images; Bottom row: Reflectance map on the Gauss sphere computed using the radiance function recovered by our algorithm.




Fig. 9. Top row: Input images; Bottom row: Reflectance map on the Gauss sphere computed using the radiance function recovered by our algorithm.

6.2.1 Synthetic Data

For our synthetic data, we have used a series of tori and teapots rendered using isotropic and monotonic BRDFs. We have generated the images by setting $\vec{L} = [0, 0, 1]^T$ and with the viewer direction $\vec{V} = [-\sin(\theta_V), 0, \cos(\theta_V)]^T$, varying so that the angle θ_V increments from 0° to 50° in 10° steps. The first model considered is that of Phong, where the radiance function is

$$g_{Phong}(\theta) = 0.5 \cos(\theta_i) + 0.5 \cos(\theta_i)^\eta$$

where $\theta_i = \arccos(\vec{N} \cdot \vec{L})$. For each of the different viewing directions we have generated synthetic images with the three values of shininess $\eta = 10, 100$ and 1000 . The second model used is that of Oren and Nayar whose radiance function is

$$g_{ON}(\theta) = E_0 \cos(\theta_i) \left(A + B \max\{0, \cos(\theta_r - \theta_i)\} \sin(\max\{\theta_i, \theta_r\}) \tan(\min\{\theta_i, \theta_r\}) \right)$$

where E_0 is a constant, $\theta_r = \arccos(\vec{N} \cdot \vec{V})$ and the dimensionless parameters A and B depend on the surface roughness σ . These are given by $A = 1.0 - 0.5\sigma^2/(\sigma^2 + 0.33)$ and $B = 0.45\sigma^2/(\sigma^2 + 0.09)$. We have generated synthetic images using the roughness values $\sigma = 0, 0.05, 0.10, 0.15, 0.20$ and 0.25 . The third model studied is that comprised by the linear combination of a diffuse and Cook and Torrance's specular components. Hence, the model becomes

$$g_{CT} = 0.5(\vec{N} \cdot \vec{H}) + 0.5 \frac{\mathcal{D}\mathcal{G}\mathcal{F}}{\pi(\vec{N} \cdot \vec{V})(\vec{N} \cdot \vec{L})}$$

where \vec{H} is the half-way vector and the terms \mathcal{D} , \mathcal{G} and \mathcal{F} are the microfacet-slope distribution, geometric factor and Fresnel term, respectively, and are given by




Fig. 10. Top row: Ground-truth renderings of the torus; Bottom row: Tori rendered making use of the radiance function recovered by our algorithm.




Fig. 11. Top row: Ground-truth renderings of the teapot; Bottom row: Teapots rendered making use of the radiance function recovered by our algorithm.

$$\mathcal{D} = \frac{\exp\left\{-\tan\left(\frac{\arccos(\vec{N} \cdot \vec{H})}{m_p}\right)^2\right\}}{4m_p^2(\vec{N} \cdot \vec{H})^4}$$

$$\mathcal{G} = \min\left\{1, \frac{2(\vec{N} \cdot \vec{H})(\vec{N} \cdot \vec{V})}{(\vec{V} \cdot \vec{H})}, \frac{2(\vec{N} \cdot \vec{H})(\vec{N} \cdot \vec{L})}{(\vec{V} \cdot \vec{H})}\right\}$$

$$\mathcal{F} = 0.5 \frac{\sin\left((\vec{L} \cdot \vec{H}) - \arcsin\left(\frac{\sin(\vec{L} \cdot \vec{H})}{m_p}\right)\right)^2}{\sin\left((\vec{L} \cdot \vec{H}) + \arcsin\left(\frac{\sin(\vec{L} \cdot \vec{H})}{m_p}\right)\right)^2}$$

where we have denoted the microfacet-slope by m_p . Here, we have generated images by varying the microfacet-slope parameter m_p from 0.20 to 0.40 in increments of 0.05. In our experiments, we have set $\vec{L}' = \vec{V}$ when synthesising images using the Phong and Oren and Nayar model. For the series of images rendered using the Cook and Torrance's BRDF, we have set the specular spike direction in the direction of the half-way vector \vec{H} between the light source and viewer directions, i.e. $\vec{L}' = \vec{H}$.

In the top row of Figures 8 and 9, from left-to-right, we show three example input images of tori and teapots rendered by setting $\theta_V = 0^\circ$ and using the Phong model for $\eta = 1000$, the Oren and Nayar model for $\sigma = 0.25$ and the Cook and Torrance model for $m_p = 0.40$. The bottom row shows unit spheres rendered using the generator function $g(\theta)$ estimated using our algorithm. In the top row of Figures 10 and 11, we show the ground-truth for the input images in the top row of Figures 8 and 9 rendered by setting $\theta_V = 50^\circ$. The bottom row of the figures shows the tori and teapots rendered using the reflectance functions recovered by our algorithm for the input images in Figures 8 and 9. From the images in the figure its clear that the renderings of the 3D data using the recovered reflectance functions are in close accordance with those obtained using the relevant model.

We present more quantitative results in Figure 12. Here, we have plotted the mean-absolute error between the images rendered using the three reflectance models and

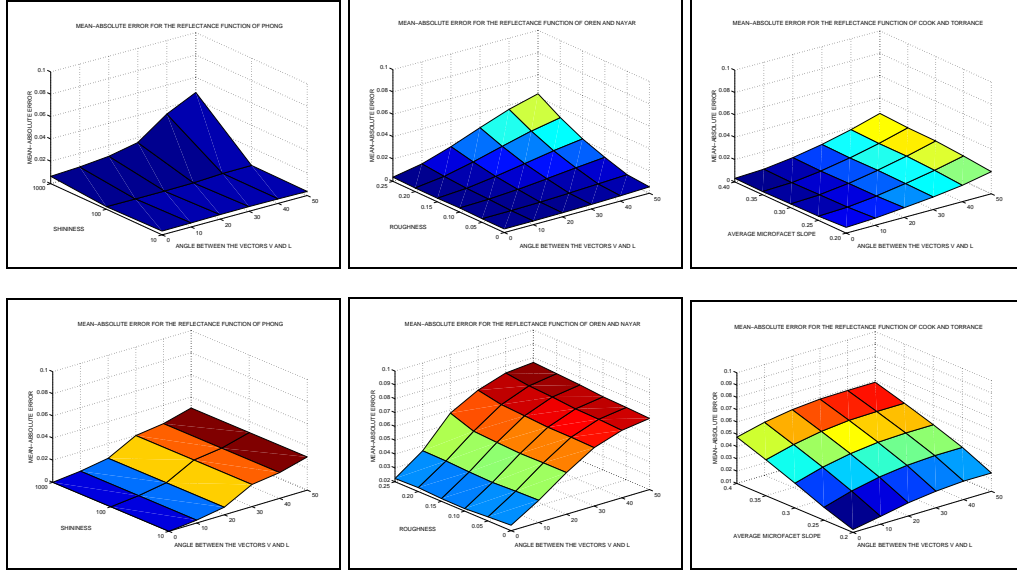


Fig. 12. Error plots for the reflectance models of Phong, Oren-Nayar and Cook-Torrance. Top row: Plots for the tori. Bottom row: Plots for the teapots.

those rendered using the reflectance functions recovered from the input images. From these plots, we draw the following conclusions. First, we note that our method behaves better with the Cook and Torrance model than it does with those of Phong or Oren and Nayar. Second, the errors are greater as the magnitude of the parameters and the angle θ_V increase, i.e. near the occluding boundary and extreme points of the plots. This is because the gradient distortion is greatest where the intensity is zero.

6.2.2 Real-world Data

Our real-world data consists of imagery of sandpaper rolls of different grades, and terracotta and porcelain objects. The imagery was captured using the procedure described in the previous subsection. The viewer direction is offset from the light source direction by about 2° . For each object, we have kept fixed the illuminant direction $\vec{L} = [0, 0, 1]^T$ and rotated the camera by 0° , 25° and 50° with respect to the light source direction \vec{L} . We have also set the vector \vec{L}' to the half-way vector \vec{H} . We have edited the images so that the background is black and only the object under study is visible. This is a preprocessing step that has no serious implications for the applicability of the method. All our imagery is true-color and has been captured at a resolution of 2240 by 1680 pixels and resampled to 840 by 630 pixels. A diagram illustrating the imaging setup is shown in Figure 13.

Our first set of images are of four sandpaper rolls with grades: 100, 150, 180 and 240. In the top row of Figure 14 we show, from left-to-right, in increasing grade, the input images for the four sandpaper rolls. In the bottom row of Figure 14, we show the projection of the reflectance functions onto Gaussian spheres. In the top row of




Fig. 13. Left-hand panel: Side-view diagram of our capture setup; Right-hand panel: Top-view diagram.

Figure 15, we show the images for $\theta_V = 50^\circ$. In the bottom row, we show the result of rendering the surface height data of the sandpaper rolls using the reflectance functions approximated from the input images in the top row of Figure 14. The images are in good agreement with one another. This implies that our method has accurately estimated the structure of the reflectance function from a single image.

Next, we turn our attention to images of rough terracotta and shiny porcelain objects. The top row of Figure 16 shows the input images. The second row shows the surface height data for each object. The unit sphere representations of the surface radiance function $g(\theta)$ are shown in the bottom row of the figure. In the top row of Figure 17, we show the images obtained when $\theta_V = 50^\circ$. In the bottom row of the figure we show the images obtained when the surface height data is rendered using the recovered surface radiance function. Again, the two sets of images are very similar in appearance.

Finally, we provide a quantitative analysis of the image data. In Tables 1 and 2 we show the mean-absolute error between the real-world and synthesised images. The error is greatest for the extreme values of θ_V (i.e. $\theta_V = 50^\circ$). We also note that the error is in average larger for the sandpaper rolls. The reason for this is due to the specular limb, which makes the object appear “flatter”.




Fig. 14. Top row: Input images; Bottom row: Recovered reflectance spheres.




Fig. 15. Top row: Images captured with a $\theta_V = 50^\circ$; Bottom row: Renderings of the 3D meshes using the reflectance function delivered by our algorithm.

6.3 Forward Rendering

In this subsection, we show how our method can be used for the purposes of forward rendering. Here, we have taken computer-generated surface models and have rendered them using radiance functions acquired from real-world imagery.

In Figure 18 we display the radiance functions for the sample objects as rendered spheres viewed from a number of different directions. In the left-hand column of the figure, we show the raw images. The second column shows the view of the sphere in the light source direction. The remaining columns show views of the spheres from different directions. Again, there are a number of features worth noting. Perhaps the most important of these is the difference of the specular component. In the case of the shiny dielectrics (the billiard balls and the porcelain mug), the specular structure

Table 1
Mean-Absolute Error for the Sandpaper Rolls

Roll Grade	Error as a Function of the Angle Between the Illuminant and the Viewer		
	0°	25°	50°
100	0.0376	0.0397	0.0418
150	0.0308	0.0385	0.0394
180	0.0354	0.0369	0.0419
240	0.0348	0.0376	0.0396

Table 2
Mean-Absolute Error for the Real-world Objects

Object Name	Error as a function of the Angle Between the Illuminant and the Viewer		
	0°	25°	50°
Porcelain Hand	0.0267	0.0299	0.0319
Porcelain Bear	0.0100	0.0187	0.0195
Porcelain Mug	0.0259	0.0396	0.0429
Porcelain Vase	0.0110	0.0217	0.0304
Terracotta Bear	0.0418	0.0419	0.0422
Terracotta Teapot	0.0197	0.0211	0.0213
Terracotta Jug	0.0161	0.0192	0.0199
Terracotta Watering Can	0.0150	0.0130	0.0223

is highly localised. In the case of the rough terracotta objects, the specular structure is weak and dispersed. In the case of the metallic objects, the specular structure is strong and dispersed. Of the objects studied, it is the blue billiard ball which gives the most realistic appearance.

We now take this study one step further by rendering a more complex object. In the top row of Figure 19, we again show the real-world images used to estimate the radiance function. The second row shows a view of a sphere rendered using the radiance functions estimated from the different objects. In the two bottom rows we show the results of performing forward rendering of a computer model of




Fig. 16. Top row: Input images; Middle row: Renderings of the ground-truth depth data acquired using a range scanner; Bottom row: Projection of the recovered reflectance maps onto Gaussian spheres.




Fig. 17. Top row: Images captured with a $\theta_V = 50^\circ$; Bottom row: Renderings of the 3D meshes using the reflectance function delivered by our algorithm.

“Yoda” downloaded from the Avalon public 3D-model archive and a 3D-model of a “Teapot” generated in-house. The radiance function estimation process took in average 0.2 secs on a 2.0 GHz., dual-Xeon workstation.

Finally, to illustrate the utility of the method for purposes of forward rendering, we have rendered a torus, a teapot and a sphere using the radiance functions shown in Figures 5 and 6. The renderings for these objects, when positioned in a simple environment with two omnidirectional lights, is shown in Figure 20.

From our experiments, it is clear that the rendered surfaces have an appearance that resembles closely that of the input object material. Furthermore, the algorithm has also shown to be able to cope well with small angular differences between the viewer and light source directions by recovering the radiance function of complex objects under a variety of reflectivity conditions.

6.4 Lambertian Re-mapping

We now turn our attention to the removal of specularities and the correction of limb brightening effects for images of real-world objects. Our real-world imagery consists of pictures of objects composed of white porcelain and of terracotta. The porcelain is shiny and hence exhibits well colimated specularities. The terracotta is rough. However, it does have a weak specular component and this results in a more diffuse specular structure. We have performed experiments on real-world images of three shiny porcelain objects and two rough terracotta objects. The rough objects are a bear and a teapot. The shiny objects are two vases and an urn. We have acquired the images under controlled lighting conditions in a dark room using an E10 Olympus digital camera. The objects are illuminated by a single colimated tungsten light source of known direction. For our experiments, the estimated light source vector is $\vec{L} = (0.43, 0, 0.9)^T$. To ground-truth the specularity removal method, we have used two polaroid filters. One filter is placed between the object and the light source. A second filter is placed between the object and the camera. For each object we collect two images. The first of these is with the polaroids aligned. The second image is collected with the polaroids crossed. When the polaroid filters are




Fig. 18. Left hand column: Input images; Right hand columns: Views of spheres rendered using the recovered radiance functions for the input images.




Fig. 19. Top row: input images for the radiance function approximation process; Second, third and fourth rows: spheres, “Yoda” and “Teapot” rendered making use of the radiance functions approximated from the input images.

crossed, then the specularities are extinguished since they correspond to coherent reflectance from the object surface.

In Figure 21, we show the results obtained with the five objects used in our experiments. In the top row from left to right the panels show the original images of the two terracotta objects (i.e. the terracotta bear and the teapot) and the three porcelain objects (i.e. the two vases and the urn). These images are collected with aligned polaroids. In the second row we show the result of Lambertian re-mapping. For comparison, in the third row, we show the images obtained using crossed po-




Fig. 20. Renderings of a torus, a teapot and a sphere using the radiance functions in Figures 5 and 6.




Fig. 21. Results on real-world images.

laroids. In the bottom row of the figure, we show the gray-scale difference between the cross-polaroid images and the images resulting from Lambertian re-mapping. There are a number of observations that can be drawn from these examples. We turn our attention first to the terracotta objects. Because of their intrinsic roughness, these objects appear relatively flat. There is also some surface brightening due to a weak and dispersed specular component. This effect is evident on the belly and snout of the bear, and, the body of the teapot. After Lambertian re-mapping, the image gradients across the two objects are enhanced, i.e. they appear less “flat”, and the weak specular structure is removed. In the case of the three porcelain objects, there are sharp specularities on the surfaces. These are cleanly removed, with no evident residual structure. In the case of both the terracotta and the porcelain objects, the agreement with the crossed polaroid images is good. Finally, from the difference images it is clear that the specular subtraction is accurate.

7 Conclusions

In this paper, we have presented a novel approach for approximating the radiance functions and hence the BRDF of possibly specular objects from a single image. Although the new method is applicable only when the light source and viewer di-

rections are approximately equal, it can be used to estimate the radiance function so that inverse rendering may be performed. Furthermore, the method is a computationally efficient alternative to the approximation of the BRDF via basis functions, which is the standard approach in the literature. We have also used the estimated radiance function to perform Lambertian re-mapping. This allows us to remove specularities from shiny objects and also to compensate for boundary “flattening” effects for rough surfaces.

There are a number of advantages of the approach over competing ones. Firstly, it provides a method for recovering the radiance function from a single image of a rough or specular object without prior knowledge of its surface structure. Secondly, the tabular representation of the recovered BRDF allows the use of floating-point vectors for storing, in a memory efficient way, the approximated radiance function. Thirdly, the computational cost of the algorithm is low since it makes use of a non-iterative approach which is devoid of free parameters. Fourthly, the radiance function estimated using the method can be used to render objects when the light source and viewer directions are no longer identical. Finally, the rendering process can be performed using simple vectorial operations in a computationally efficient way.

There are a number of ways in which the work developed in this paper may be further extended. Since the method provides a means of locating the Gauss-map, it may also provide a route to shape-from-shading for non-Lambertian surfaces. Second, the recovered radiance functions may be compared with theoretical models. This may provide a means of determining surface parameters. For instance, by fitting the modified Beckman-Kirchoff model to our radiance functions, roughness and correlation measurements could be made. Thirdly, it would be interesting to extend the method to multiple light source or viewer directions and to determine whether the method provides a new route to photometric stereo.

8 Acknowledgment

The work presented here was done while Antonio Robles-Kelly was with The University of York, York, United Kingdom.

References

- [1] S. Nayar, R. Bolle, Reflectance based object recognition, *International Journal of Computer Vision* 17 (3) (1996) 219–240.
- [2] S. K. Nayar, R. M. Bolle, Computing reflectance ratios from an image, *Pattern Recognition* 26 (1993) 1529–1542.

- [3] R. O. Dror, E. H. Adelson, 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] B. K. P. Horn, M. J. Brooks, The variational approach to shape from shading, CVGIP 33 (2) (1986) 174–208.
- [5] B. K. P. Horn, R. W. Sjöberg, Calculating the reflectance map, Applied Optics 18 (11) (1979) 1770–1779.
- [6] B. K. P. Horn, Hill shading and the reflectance map, Proceedings of the IEEE 69 (1) (1981) 14–47.
- [7] G. Brelstaff, A. Blake, Detecting specular reflection using lambertian constraints, in: Int. Conf. on Comp. Vision, 1988, pp. 297–302.
- [8] H. Ragheb, E. R. Hancock, A probabilistic framework for specular shape-from-shading, Pattern Recognition 36 (2) (2003) 407–427.
- [9] B. van Ginneken, M. Starvridi, J. J. Koenderink, Diffuse and specular reflectance from rough surfaces, Applied Optics 37 (1) (1998) 130–139.
- [10] T. K. Leung, J. Malik, Recognizing surfaces using three-dimensional textons, in: Proc. of the IEEE Int. Conf. on Compl. Vis., 1999, pp. 1010–1017.
- [11] S. K. Nayar, K. Ikeuchi, T. Kanade, Surface reflection: Physical and geometrical perspectives, IEEE Trans. on Pattern Analysis and Machine Intelligence 13 (7) (1991) 611–634.
- [12] J. E. Harvey, C. L. Vernold, A. Krywonos, P. L. Thompson, Diffracted radiance: A fundamental quantity in a non-paraxial scalar diffraction theory, Applied Optics 38 (31) (1999) 6469–6481.
- [13] P. Beckmann, A. Spizzochino, The Scattering of Electromagnetic Waves from Rough Surfaces, Pergamon, New York, 1963.
- [14] C. L. Vernold, J. E. Harvey, A modified beckmann-kirchoff scattering theory for non-paraxial angles, in: Scattering and Surface Roughness, no. 3426 in Proc. of the SPIE, 1998, pp. 51–56.
- [15] H. Ragheb, E. R. Hancock, Estimating surface characteristics using physical reflectance models, in: Proc. of the IEEE Conference on Computer Vision and Pattern Recognition, 2003, pp. II:177–183.
- [16] R. L. Cook, K. E. Torrance, A reflectance model for computer graphics, Vol. 15, 1981, pp. 307–316.
- [17] B. T. Phong, Illumination for computer generated pictures, Communications of the ACM 18 (6) (1975) 311–317.
- [18] K. E. Torrance, E. M. Sparrow, Theory for off-specular reflection from roughened surfaces, Journal of Optical Society of America 57 (9) (1967) 1105–1114.

- [19] R. R. Lewis, Making shaders more physically plausible, in: *Proceedings of the Fourth Eurographics Workshop on Rendering*, 1993, pp. 47–62.
- [20] C. Schlick, A fast alternative to phong’s specular model, *Graphics Gems 4* (1994) 363–366.
- [21] C. Schlick, A survey of shading and reflectance models, *Computer Graphics Forum* 13 (2) (94) 121–131.
- [22] L. B. Wolff, On the relative brightness of specular and diffuse reflection, in: *Int. Conf. on Comp. Vision and Patt. Recognition*, 1994, pp. 369–376.
- [23] S. K. Nayar, M. Oren, Visual appearance of matte surfaces, *SCIENCE* 267 (1995) 1153–1156.
- [24] L. B. Wolff, S. K. Nayar, M. Oren, Improved diffuse reflection models for computer vision, *Int. Journal of Comp. Vision* 30 (1) (1998) 55–71.
- [25] S. R. Marschner, S. H. Westin, E. P. F. Lafortune, K. E. Torrance, D. P. Greenberg, Image-based brdf measurement including human skin, in: *10th Eurographics Rendering Workshop*, 1999.
- [26] P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, M. Sagar, Acquiring the reflectance field of a human face, in: *SIGGRAPH 2000*, 2000, pp. 145–156.
- [27] A. Hertzmann, S. M. Seitz, Shape and materials by example: A photometric stereo approach, in: *Int. Conf. on Comp. Vision and Patt. Recognition*, 2003, pp. 533–540.
- [28] S. Westin, J. Arvo, K. Torrance, Predicting reflectance functions from complex surfaces, in: *SIGGRAPH 92 Conference Proceedings*, 1992, pp. 255–264.
- [29] X. He, P. Heynen, R. Phillips, K. Torrance, D. Salesin, D. Greenberg, A fast and accurate light reflection model, in: *Siggraph 92 Conference Proceedings*, Vol. 26, 1992, pp. 253–254.
- [30] P. Schröder, W. Sweldens, Spherical wavelets: Efficiently representing functions on the sphere, in: *Siggraph 95 Conference Proceedings*, 1995, pp. 161–172.
- [31] G. J. Ward, Measuring and modeling anisotropic reflection, *Computer Graphics* 26 (2) (1992) 265–272.
- [32] E. P. Lafortune, Sing-Choong Foo, K. E. Torrance, D. P. Greenberg, Non-linear approximation of reflectance functions, in: *SIGGRAPH 97 Conference Proceedings*, 1997, pp. 117–126.
- [33] K. J. Dana, S. K. Nayar, Correlation model for 3d texture, in: *Int. Conf. on Comp. Vision*, 1999, pp. 1061–1066.
- [34] A. Z. P. Giblin, A. Blake, The information available to a moving observer from specularities, *Image and Vision Computing* (7) (1989) 38–42.
- [35] H. Ragheb, E. R. Hancock, Highlight removal using shape-from-shading, in: *European Conf. on Comp. Vision*, no. 2351 in LNCS, 2002, pp. 626–641.

- [36] P. L. Worthington, 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) (1999) 1250–1267.
- [37] A. Robles-Kelly, E. R. Hancock, Radiance function estimation for object classification, in: *CIARP 2004*, 2004, p. To appear.
- [38] A. Robles-Kelly, E. R. Hancock, Correction of underexposed images using scene radiance estimation, in: *Proc. of the Int. Conf. on Image Processing*, 2004, to appear.
- [39] R. L. Cook, K. E. Torrance, A reflectance model for computer graphics, *ACM Trans. on Graphics* 1 (1) (1982) 7–24.
- [40] M. Oren, S. K. Nayar, Generalization of the lambertian model and implications for machine vision, *Int. Journal of Comp. Vision* 14 (3) (1995) 227–251.