Vision, Video, and Graphics (2003)
P. Hall, P. Willis (Editors)

Lambertian Correction for Rough and Specular Surfaces

Antonio Robles-Kelly[1] and Edwin R. Hancock[2]

Department of Computer Science
The University of York
York, Y01 5DD, UK
[1]arobkell@cs.york.ac.uk [2]erh@cs.york.ac.uk

Abstract

This paper describes a method for performing Lambertian reflectance for rough and specular surfaces. Rather than using an existing reflectance model, we present a method for estimating the reflectance function from image data. The method makes use of the Gauss map between a surface and a unit sphere. 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 the polar function. A simple analysis shows how the tabular representation of the reflectance function can be obtained using the accumulative distribution of image gradients. By modifying the reflectance function and back-projecting, we can render the surface with alternative lighting models. Here, we choose to back-project a Lambertian reflectance model. This allows us to be remove specularities from shiny surfaces and compensate from boundary "flattening" for rough surfaces. We illustrate the utility of the method on a variety of real world imagery.

Categories and Subject Descriptors (according to ACM CCS): I.4.8 [Image Processing and Computer Vision]: Lambertian correction, BRDF approximation

1. Introduction

The modelling 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 bi-directional reflectance distribution function (BRDF) is to hand then a number of surface analysis tasks may be addressed. For instance Nayar and Bolle [1] have used photometric invariants derived from the BRDF to recognise objects with different reflectance properties. In a related development, Dror *et al.* [2] 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 [3], if a BRDF is to hand then photometric correction or specularity subtraction may be applied as a preprocessing step to improve the results obtained. It is interesting to note that there have been several attempts to remove specularities from images of non-Lambertian objects [4]. The main limitation of these methods is that they rely on the use of the BRDF to characterise the specular spike and limb. As a result, they lack the generality required to process real-world imagery in an unsupervised or automatic way. Finally, there has been recent interest in describing texture as a surface relief phenomenon process using ideas from physics to model the BRDF [5].

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 [6]. Although, it is based on physically meaningful surface parameters, the Beckmann theory is both intractable for analysis problems and breaks down when the surface roughness is large or the scattering angle is large. However, recently, Vernold and Harvey [7] have overcome this latter problem by developing a model which accounts for self shadowing on rough surfaces. 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 [8, 9]. One of the most popular models is that developed by Phong [9]. A survey of reflectance and shading models can be found in [10].

However, neither the models developed in physics nor the

empirical models developed in graphics are well suited for surface analysis tasks in computer vision. It is for this reason that Wolff [11] and Nayar and Oren [12] have developed semi-empirical models that account for departures from Lambertian reflectance. 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 modelling and estimation of the bidirectional reflectance function remains an elusive task. The main problem is that the BRDF has four degrees of freedom. Ward [13] has shown how to simplify matters by approximating the BRDF using a Gaussian lobe. The BRDF may also be approximated using spherical harmonics [14], splines [15] and wavelets [16]. Lafortune and co-workers [17] at Cornell have developed an approximation method that employs a set of reciprocal, energy-preserving functions with non-linear parameter dependance. Dana and Nayar [18] have provided a method for collecting empirical BRDFs and have catalogued the BRDFs for a large number of surfaces of different physical properties.

In this paper, we present a method for estimating the reflectance distribution from image data that avoids using basis functions or using a predetermined BRDF 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 intensity values from the surface to corresponding locations on the Gauss sphere. 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 Gauss sphere. We pose the problem of recovering the reflectance function as that of estimating a tabular representation of the polar function. To overcome the problem that we do not have surface normal correspondences to hand, we show how to estimate the polar angles on the Gauss sphere using image intensity gradients. A simple analysis shows how the tabular representation of the reflectance function can be obtained using the accumulative distribution of image gradients.

With the reflectance distribution to hand, then a number of different analysis and synthesis tasks may be addressed. For instance, the acquired reflectance models may be used to render synthetic surfaces or may be modified and back-projected to perform view synthesis. Here as an illustration we focus on the latter task and illustrate how the method can be be used for Lambertian intensity 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. 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 [19]. If specularities can be subtracted from images, then shape-from-shading may be applied more effectively to recover surface shape [20]. Similarly, if boundary flattening effects can be compensated, then shape-from-shading may be applied to rough surfaces. Here we show how to backproject a Lambertian reflectance model onto the imaged surface using the inverse mapping between the Gauss sphere and the image.

2. Preliminaries

In this section, we provide the background for our method. To commence, let the surface under study be denoted by $S \in \Re^3$. In practice, we will be working with brightness images formed on the image plane. Hence, we commence by projecting the surface S onto the image plane Π. We also construct a Gauss map for the surface, by projecting surface normals onto a unit sphere $\hat{S} \in \Re^3$. Corresponding locations on the surface and the sphere are such that the surface normal directions are identical. The mapping is such that the surface normals at corresponding locations on the surface and the sphere have identical directions. This correspondance mapping between the surface S and the sphere $\hat{S}$ allows the intensity values of the rendered sphere $\hat{S}$ to be mapped onto a viewer plane $\hat{\Pi}$ analogous to the image plane Π. The plane $\hat{\Pi}$ is chosen so that the viewer direction vector $\vec{\mathbf{V}}$ is equal to the light-source direction vector $\vec{\mathbf{L}}$. This geometry is illustrated in Figure 1.

Next, we consider to show how the BDRF may be expressed in terms of unit vectors in the directions of the surface normal, the viewer direction and the light source direction. To be more formal, let the unit normal vector to the surface at the location s with pixel coordinates (j,k) be $\vec{\mathbf{N}}_s$. The light source, viewer and surface normal vectors can be expressed in terms of the elevation and azimuth angles (θ_L, α_L), (θ_V, α_L) and (θ_N, α_{N_s}) thus

$$\begin{aligned} \vec{\mathbf{L}} &= [\sin(\theta_L)\cos(\alpha_L), \sin(\theta_L)\sin(\alpha_L), \cos(\theta_L)]^T \\ \vec{\mathbf{V}} &= [\sin(\theta_V)\cos(\alpha_V), \sin(\theta_V)\sin(\alpha_V), \cos(\theta_V)]^T \\ \vec{\mathbf{N}}_s &= [\sin(\theta_{N_s})\cos(\alpha_{N_s}), \sin(\theta_{N_s})\sin(\alpha_{N_s}), \\ &\quad \cos(\theta_{N_s})]^T \end{aligned} \tag{1}$$

We use the vector $\vec{\mathbf{N}}_s$ as a reference and define the following elevation and azimuth angle offsets for the light-source and viewer directions $\alpha_{L,N_s} = \alpha_{N_s} - \alpha_L$, $\alpha_{V,N_s} = \alpha_{N_s} - \alpha_V$, $\theta_{L,N_s} = \theta_{N_s} - \theta_L$ and $\theta_{V,N_s} = \theta_{N_s} - \theta_V$.

For objects illuminated by a single light source, the total reflectance from a visible point on the surface can be assumed to be a linear combination of the ambient, diffuse and specular components. Hence, making use of the variables defined above and normalising to unity the total reflectance is given by

$$\begin{aligned} B(\alpha_{L,N_s}, \alpha_{V,N_s}, \theta_{L,N_s}, \theta_{V,N_s})_s = (1 - I_d - I_s) + \\ I_d f_d(\theta_{L,N_s}) + I_s f_s(\alpha_{L,N_s}, \alpha_{V,N_s}, \theta_{L,N_s}, \theta_{V,N_s}) \end{aligned} \tag{2}$$

where I_d and I_s are positive constants satisfying the condition, $I_d + I_s \leq 1$ and, $f_d(\theta_{L,N_s})$, $f_s(\alpha_{L,N_s}, \alpha_{V,N_s}, \theta_{L,N_s}, \theta_{V,N_s})$ are the reflectance functions for the diffuse and specular components.

3. Orientable Surfaces

As noted in the previous section, the directional dependence of the specular term in the reflectance function complicates the approximation of the intensity function if the light-source and viewer directions are different from one another, i.e. $\vec{\mathbf{L}} \neq \vec{\mathbf{V}}$. Furthermore, although the assumption that the diffuse term is independent of the viewer direction reduces the dimensionality of the problem, it severely restricts its real-world utility. As a result, when more complex diffuse reflectance models (for instance the Oren and Nayar model [12]) are employed or the directional dependence of the specular term is taken into account, then the problem of separating the reflectance components from a single image is underconstrained. Even when multiple images are collected from different viewpoints or are acquired under different illumination conditions, then the recovery of the combined specular, diffuse and ambient reflectance components may not be tractable in a closed form.

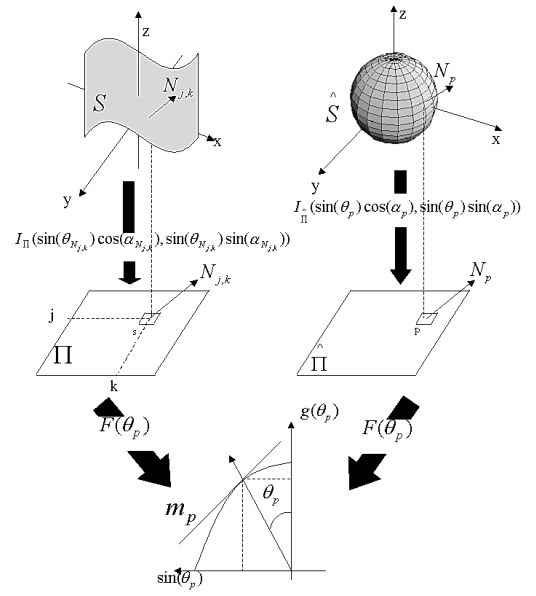

List of symbols

S	Surface under study
$\hat{S}$	Sphere in $\Re^3$
Π, $\hat{\Pi}$	Planes orthogonal to $\vec{\mathbf{V}}$
$I_\Pi(\sin(\theta_{N_S})\cos(\alpha_{N_S}),$ $\sin(\theta_{N_S})\sin(\alpha_{N_S}))$	Mapping of the surface S onto the plane Π
$I_{\hat{\Pi}}(\sin(\theta_p)\cos(\alpha_p),$ $\sin(\theta_p)\sin(\alpha_p))$	Mapping of the sphere $\hat{S}$ onto the plane $\hat{\Pi}$
$\vec{\mathbf{N}}_p$, $\vec{\mathbf{N}}_s$	Surface normals
$F(\theta_p)$	Parametric function
m_p	Slope of the line tangent to $F(\theta_p)$

Figure 1: *Structure of the mapping process*

To overcome this problem, in this paper we exploit the differential geometry of orientable surfaces to approximate the reflectance function using data acquired from a single image. We make use of the fact that given an orientable surface $S \in \Re^3$, there exists a function $F : \Re^3 \mapsto \Re^3$ that maps the normal-vector at any point on the surface S to a corresponding normal-vector at a point on a unit Gauss sphere $\hat{S} \in \Re^3$. As noted above, this mapping is referred to as the Gauss-map of the surface S. Since the vectors $\vec{\mathbf{L}}$ and $\vec{\mathbf{V}}$ remain fixed over all locations on the image plane, it follows that all the intensities present in the image can be mapped completely onto a hypothetical unit Gauss sphere $\hat{S}$ viewed from the direction $\vec{\mathbf{V}}$. The illuminated Gauss sphere together with its field of surface normals can be projected onto a image plane $\hat{\Pi}$ which is perpendicular to the viewer direction $\vec{\mathbf{V}}$. This process is illustrated in Figure 1.

Let the surface under study be $S \in \Re^3$. The surface normal vector $\vec{\mathbf{N}}_s$ at the point s with coordinates j,k on the surface S can be associated with a corresponding normal-vector $\vec{\mathbf{N}}_p$ at a point p on the Gauss sphere $\hat{S}$. On the image plane Π, the two vectors can be written in polar form as

$$\vec{\mathbf{N}}_s = \vec{\mathbf{N}}_p = [\sin(\theta_p)\cos(\alpha_p), \sin(\theta_p)\sin(\alpha_p), \cos(\theta_p)]^T \tag{3}$$

In the above expression, $\alpha_p \in [0, 2\pi)$ and $\theta_p \in [0, \frac{\pi}{2})$ are respectively the azimuth and elevation angles of the projection of surface normal from the corresponding point p on the Gauss sphere $\hat{S}$ to the plane $\hat{\Pi}$. Provided that the surface normal correspondence can be established, then the intensity mapping of the Gauss sphere $\hat{S}$ to the plane $\hat{\Pi}$ can be represented by the function

$$\begin{aligned} I_{\hat{\Pi}}(\sin(\theta_p)\cos(\alpha_p), \sin(\theta_p)\sin(\alpha_p)) \\ = I_\Pi(\sin(\theta_{N_s})\cos(\alpha_{N_s}), \sin(\theta_{N_s})\sin(\alpha_{N_s})) \\ = B(\alpha_{L,p}, \alpha_{V,p}, \theta_{L,p}, \theta_{V,p}) \end{aligned} \tag{4}$$

where $I_\Pi(\sin(\theta_{N_s})\cos(\alpha_{N_s}), \sin(\theta_{N_s})\sin(\alpha_{N_s}))$ is the mapping of the surface S to the image plane Π, $B(\alpha_{L,p}, \alpha_{V,p}, \theta_{L,p}, \theta_{V,p})$ is the BRDF for the surface S and $\alpha_{L,p}$, $\alpha_{V,p}$, $\theta_{L,p}$ and $\theta_{V,p}$ are angular variables given by $\alpha_{L,p} = \alpha_p - \alpha_L$, $\alpha_{V,p} = \alpha_p - \alpha_V$, $\theta_{L,p} = \theta_p - \theta_L$ and $\theta_{V,p} = \theta_p - \theta_V$.

As a result, the mapping of the Gauss sphere $\hat{S}$ to the plane $\hat{\Pi}$ is given by the BRDF of a sphere with the same optical and reflectance properties as the surface S when viewed from the direction $\vec{\mathbf{V}}$ and illuminated by a light source in the direction $\vec{\mathbf{L}}$. Moreover, since it is formed by orthographic projection from a sphere, the image can be represented as a polar function with four degrees of freedom.

To take our analysis further, we note that the problem

of recovering the reflectance distribution function from a single image is an ill-posed one. Since the BRDF is four-dimensional, we require additional restrictions to approximate it. At this point, we note that the angular variables $\alpha_{L,p}$,$\alpha_{V,p}$, $\theta_{L,p}$ and $\theta_{V,p}$ are determined by the angles θ_p and α_p. To simplify matters, we choose to view the Gauss sphere $\hat{S}$ in the light source direction, i.e. we set $\vec{\mathbf{L}} = \vec{\mathbf{V}}$. As a result, we can write $\alpha_{L,p} = \alpha_{V_p}$ and $\theta_{L,p} = \theta_{V_p}$. Further, without any loss of generality we can set $\alpha_L = \theta_L = 0$.

By viewing the Gauss sphere in the light source direction, we simplify the mapping from the Gauss sphere $\hat{S}$ onto the plane $\hat{\Pi}$. The simplified polar function has two degrees of freedom and is given by

$$I_{\hat{\Pi}}(\sin(\theta_p)\cos(\alpha_p), \sin(\theta_p)\sin(\alpha_p)) = f(\theta_p, \alpha_p) \qquad (5)$$

We now take advantage of the fact that the expression above is defined only in terms of the polar coordinates of the point p on the Gauss sphere $\hat{S}$. We also note the following. First, all surface normals on the same circle of lattitude have the same elevation angle θ_p irrespective of the azimuth angle α_p. Second, the assumption $\vec{\mathbf{V}} = \vec{\mathbf{L}}$ implies that the BRDF on the Gauss sphere is not dependant on the angle α_p. Hence, the observed brightness values on the Gauss sphere can be generated by revolving the function

$$g(\theta_p) = f(\theta_p, 0) \qquad (6)$$

in azimuth angle α_p about the axis defined by the viewer and light source directions. As a result, the problem of describing the observed brightness distribution over the Gauss sphere reduces to that of approximating the function $g(\theta_p)$ and computing its trace of revolution.

4. Parametric Representation

Our idea in this paper is to approximate the BRDF from a single image. We show how this can be performed by using the differential structure of the observed brightness on the image plane Π to estimate $g(\theta_p)$. Hence, we commence rewriting $g(\theta_p)$ as the integral of the partial derivative of the observed brightness with respect to the angular variable θ_p. As a result, $g(\theta_p)$ is given by

$$g(\theta_p) = \frac{1}{2\pi}\int_0^{2\pi}\int_0^{\theta_p}\frac{\partial f(\theta,\alpha)}{\partial\theta}d\theta d\alpha \qquad (7)$$

In other words, the BRDF on the unit Gauss sphere can be expressed in term of the cumulative distribution of the derivatives of the reflectance on the unit Gauss sphere.

Our aim is find a polar parametric function $F \in \Re^2$ whose trace can be used to approximate that of the generator of the BRDF on the unit Gauss sphere $\hat{S}$, i.e. $g(\theta_p)$. The radial coordinate of the function is the Euclidean distance between the point p and the center-point of the Gauss sphere $\hat{S}$ on the viewer plane $\hat{\Pi}$, i.e.

$$r_p = \sqrt{(\sin(\theta_p)\cos(\alpha_p))^2 + (\sin(\theta_p)\sin(\alpha_p))^2} = \sin\theta_p \qquad (8)$$

Hence

$$F = \begin{bmatrix} r_p \\ g(\theta_p) \end{bmatrix} = \begin{bmatrix} \sin(\theta_p) \\ \frac{1}{2\pi}\int_0^{2\pi}\int_0^{\theta_p}\frac{\partial f(\theta,\alpha)}{\partial\theta}d\theta d\alpha \end{bmatrix} \qquad (9)$$

Unfortunately, since the surface normals are not to hand, the correspondences between locations on the surface and the Gauss sphere are not available. Hence, the quantity θ_p is unkown. In other words the function F only allows the surface S to be mapped onto the Gauss sphere $\hat{S}$ in an implicit manner. In the next section we show how the lack of correspondence information can be overcome by equating the image intensity gradient and the slope of the function $g(\theta_p)$. This allows us to approximate θ_p by performing numerical integration using the cumulative distribution of the inverse image gradients. In this way the BRDF can be approximated without the need to recover the surface normals.

5. BRDF Approximation

Our aim in this section is to show how the distribution of intensity gradients on the image plane can be used to estimate the BRDF. We do this by using image gradients to determine the mapping onto the Gauss sphere. To commence, we note that the derivative of the functions F at the point p is

$$m_p = \frac{\Delta g(\theta_p)}{\Delta\sin(\theta_p)} \qquad (10)$$

Since the intensity of the point s on the image, i.e. $\nu(s)$, is mapped onto the point p on the Gauss sphere, we can equate m_p with the magnitude of the image intensity gradient, i.e. $m_p = \mid \nabla\nu(s) \mid$. The image gradient $\nabla\nu(s)$ can be computed using the formula

$$\nabla\nu(s) = \begin{bmatrix} \Delta\nu_x(s) \\ \Delta\nu_y(s) \end{bmatrix} = \frac{1}{r}\begin{bmatrix} \nu(j+1,k) - \nu(j-1,k) \\ \nu(j,k+1) - \nu(j,k-1) \end{bmatrix} \qquad (11)$$

where r is the spacing of sites on the pixel lattice. Furthermore, on the Gauss sphere $\hat{S}$, it is always possible to choose points to be sampled so that the difference in brightness is a constant τ. As a result, we can write

$$\mid \nabla\nu(s) \mid \Delta\sin(\theta_p) \approx \Delta g(\theta_p) \approx \tau \qquad (12)$$

Re-arranging terms, we get

$$\Delta\sin(\theta_p) \approx \frac{\tau}{\mid \nabla\nu(s) \mid} \qquad (13)$$

To recover θ_p from the expression above we perform numerical integration. To do this we order the image intensity gradients such that $\nu(1) < \nu(2) < \ldots < \nu(s)$ and $\nu(s) = g(\theta_p)$. The numerical estimate of $\sin(\theta_p)$ is

$$\sin(\theta_p) \approx \sum_{i=1}^{s}\frac{\tau}{\mid \nabla\nu(i) \mid} + \kappa \qquad (14)$$

where κ is the integration constant. Hence, we can use the accumulative distribution of inverse gradients to index the azimuthal angle on the unit Gauss sphere. This indexation

property means that we can approximate the function F or equivalently $g(.)$ by tabulation.

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

$$\hat{h}(l) = \frac{\sum_{(s)\in H(l)} \mid \nabla\nu(s) \mid}{\mid H(l) \mid} \tag{15}$$

With the average image intensity gradient to hand, we define the tabular approximation $\hat{F}$ to $F(\theta_p)$ as the set of Cartesian pairs

$$\hat{F} = \{((\tau \sum_{i=1}^{l} \hat{h}(i)^{-1} + \kappa), l); l = 1, 2, \ldots, n_{max}\} \tag{16}$$

where n_{max} is the maximum brightness value in the image. All that remains is to compute the constants τ and κ. We do this by making use of the maximum and minimum values of $\sin(\theta_p)$. Since the maximum and minimum values of $\sin(\theta_p)$ are unity and zero when $\theta_p = \frac{\pi}{2}$ and $\theta_p = 0$, we can set κ to unity. Evaluating the numerical integral for $l = n_{max}$ (i. e. $\sin(0) = 0$), we get

$$\tau = -\sum_{i=1}^{n_{max}} \hat{h}(i)^{-1} \tag{17}$$

In practice the inverse of the average intensity gradient is stored as a floating-point vector J whose elements indexed $l = \{1, 2, \ldots, n_{max}\}$ are given by

$$J_l = \begin{cases} \hat{h}(l)^{-1} & \text{if } \mid H(l) \mid \neq 0 \\ 0 & \text{otherwise} \end{cases} \tag{18}$$

The elements of the vector J corresponding to intensity values l not present in the image are set to zero. To overcome divide-by-zero problems, we interpolate the zero elements of the vector using piecewise-linear interpolation of the adjacent non-zero elements to compute a vector $\hat{J}$.

6. Lambertian Re-mapping

In this paper, we illustrate the utility of our method for acquiring the reflectance distribution function. In particular, we illustrate how it may be used for re-mapping a Lambertian reflectance model onto an imaged surface. This re-mapping may be used to both remove specularities from shiny surfaces and to correct for reduced boundary contrast for rough surfaces.

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

Our aim is to use the tabular representation of the reflectance function $\hat{J}$ to retrieve the Lambertian intensity at a given point on the surface S illuminated from a light source with direction vector $(0, 0, 1)$. To do this we note that the tabular function $\hat{F}$ is a list of Cartesian pairs in which the first element is the sine of the tilt angle of the surface normals, i.e. $\sin\theta_p$ while the second element is the associated image brightness , i.e. l at a point indexed on the surface S. For Lambertian reflectance, the observed reflectance is propertional to the cosine of the angle of light incidence, i.e. to $\cos\theta_p$. Hence, we can perform Lambertian re-illumination by noting the observed brightness l at a pixel and identifying the associated value of $\sin\theta_p$; the corresponding corrected Lambertian intensity is $\cos\theta_p$.

This intensity re-mapping can be effected using the measured image gradient. Suppose that S_s^* is a neighbourhood of points of area ε centred at the pixel location indexed s. We compute corrected Lambertian intensity by averaging $\cos\theta_p$ over the neighbourhood S_s^*. Since the angle θ_p is defined on the Gaussian sphere, while the intensity is required on the image plane, we weight the average using the appropriate Jacobian. From the analysis presented in the follwoing section it follows that the weighting factor is proportional to the image gradient. Thus, the corrected Lambertian intensity at the pixel s is given by

$$\hat{\rho}_s = \frac{1}{\mu(s)} \sum_{s^* \in S_s^*} \left\{ T(s^*) \mid \nabla\nu(s^*) \mid \right\} \tag{19}$$

where

$$\mu(s) = \sum_{s^* \in S_s^*} \mid \nabla\nu(s^*) \mid \tag{20}$$

is the average intensity gradient and

$$T(s^*) = \cos\left(\arcsin\left(\tau \sum_{i=1}^{\nu(s^*)} \hat{h}(i)^{-1} + \kappa \right) \right) \tag{21}$$

where $\nu(s^*)$ is the raw image intensity at the point $s^* \in S^*$ on S. This averaging process effectively smooths the estimate of the Lambertian reflectance.

7. Experiments

In this section we report our experimental evaluation of the new method for estimating the reflectance distribution function and using it for specularity removal and limb correction. The study is divided into two parts. We commence with a study of synthetic imagery aimed at evaluating the performance of the method on data with known ground truth. The second part of the study focuses on real world data and aims to demonstrate the utility of the method on objects composed of a variety of materials.

7.1. Synthetic Images

Our goal in this section is to evaluate the effectiveness of the specularity removal process on synthetic data. The synthetic data used for our experiments consists of 54 images of spheres rendered using the Phong and Oren-Nayar reflectance models under different lighting conditions. For our experiments we have set the material color to $(0.5, 0.5, 0.5)$. To ground-truth the method, we have compared the results of specularity removal with a sphere rendered using Lambert's law. For our experiments, we have set the viewer direction to be $\vec{\mathbf{V}} = (0, 0, 1)$. We have varied the angle between the illuminant $\vec{\mathbf{L}}$ and the viewer direction in increments of 0.157 radians between 0 radians and $\frac{\pi}{4}$ radians. In the case of the Phong model, we have used the three different shininess parameter settings $\eta = 10$, $\eta = 100$ and $\eta = 200$. In the case of the 36 spheres rendered using the Oren-Nayar model, we have varied the roughness parameter σ over the interval $[0, 1]$ in steps of 0.2.

We commence by illustrating some of the qualitative properties of the specularity subtraction method. In Figure 2 from left-to-right the different panels show an example image of a Phong sphere, the output image after Lambertian remapping and the Lambertian ground-truth image. The original image contains a clear specularity. The Lambertian remapping cleanly removes the specularity, and there is no residual image structure in the proximity of specularity. The re-mapped image and the ground truth Lambertian image are in good agreement.

In Figure 3, we show the plots of the tabular approximation to the reflectance function for the Phong sphere in Figure 2. The panels in the top row show are from left-to-right, the raw vector J and its piecewise linear interpolation $\hat{J}$. The process may be applied identically to each of the three colour-channels comprising the image. Furthermore, since the color material for our synthetic imagery has been set to $(0.5, 0.5, 0.5)$, the image brightness is the same for each of the three RGB colour-channels. As a result, the plots for the three colour bands are identical. Hence, we only show a single plot that applies equally to any of the three colour-channels in the image. In the bottom row of the figure , we show the plots of the estimated function $\hat{F}(\theta_p)$ and the approximated mapping $I_{\hat{\Pi}}(\sin(\theta_p)\cos(\alpha_p), \sin(\theta_p)\sin(\alpha_p)) = f(\theta_p, \alpha_p)$. In these two plots, the specular structure of the reflectance function is clearly visible as a spike near the origin.

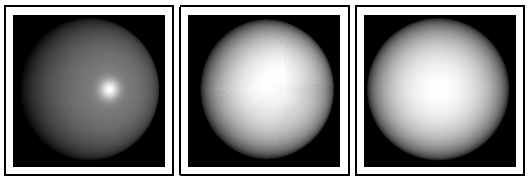

Figure 2: *Input image (sphere rendered using the Phong model), output image and ground-truth Lambertian sphere.*

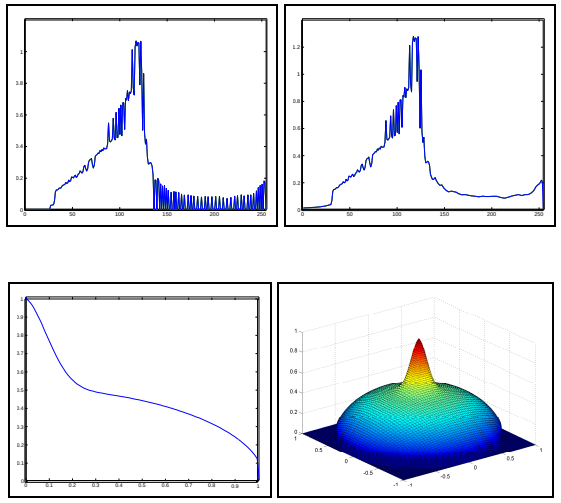

Figure 3: *Top row: plots of the vectors J and $\hat{J}$; Bottom row: plots of $\hat{F}(\theta_p)$ and $f(\theta_p, \alpha_p)$ computed from the Phong sphere in Figure 2*

We now turn our attention to the errors associated with the reflectance distribution estimates. In the left-hand side of Figure 4, we show the mean squared error of the approximated reflectance function $f(\theta_p, \alpha_p)$ with respect to the ground-truth as a function of both the angle between the viewer and light source directions, and the shininess parameter η for 18 Phong spheres. The right-hand plot in Figure 4 shows the error plot for the 36 Oren-Nayar spheres as a function of the parameter σ and the angle between the light source and viewing directions. From these plots we draw the following conclusions. First, we note that our method behaves better with shiny surfaces than it does with rough ones. Second, the errors are greatest when $\theta_p \simeq \frac{\pi}{2}$, i.e. near the occluding boundary. This is because the gradient distortion is greatest where the intensity is zero.

7.2. Real-World Images

Our real-world imagery consists of pictures of objects composed of white porcelain and of terracota. The porcelain is shiny and hence exhibits well collimated specularities. The terracotta is rough. However, it does have a weak specu-

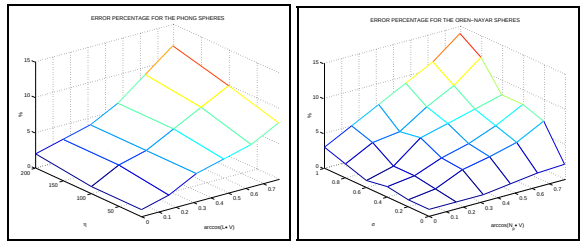

Figure 4: *Error plots.*

lar component and this results in a more diffuse specular structure. We have performed experiments on real-world images of two shiny porcelain objects and one rough terracota object. The rough object is a bear. 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 collimated tungsten light source of known direction. For our experiments, the estimated light source vector is $\vec{\mathbf{L}} = (0.43, 0, 0.9)^T$. 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 crossed, then the specularities are extinguished since they correspond to coherent reflectance from the object surface.

In Figure 5, we show the results obtained with the three objects used in our experiments. In the top row From left to right the panels show the original images of the terracotta bear and the two porcelain objects (i.e. the vase 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 polaroids. 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 object. Because of its intrinsic roughness, it appears relatively flat. There is also some surface brightening due to a weak and dispersed specular component; this is evident on the belly and snout of the bear. After Lambertian re-mapping, the intensity gradients across the the object are enhanced (i.e. it appears less "flat") and the weak specular structure is removed. In the case of the two 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 specularity subtraction is accurate.

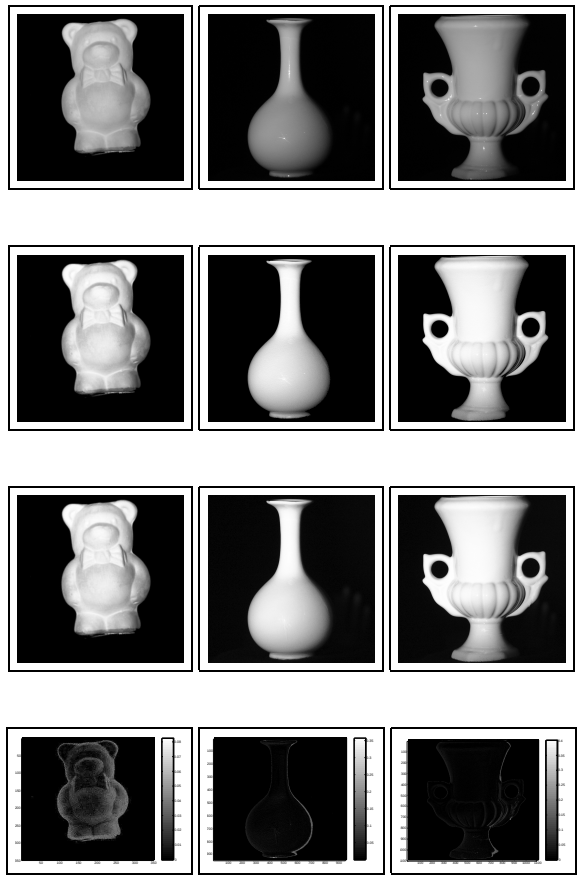

Figure 5: *Results on real-world images.*

8. Conclusions

In this paper, we have presented a novel approach for approximating the BRDF of possibly specular objects from a single image. Although the new method is applicable only when the light source and viewer directions are approximately equal, it can be used to estimate the BRDF 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. To illustrate the utility of the method, we have used the estimated reflectance 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.

References

1. S. Nayar and R. Bolle. Reflectance based object recognition. *International Journal of Computer Vision*, 17(3):219–240, 1996.
2. 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 Worksho on Identifying Objects Across Variations in Lighting*, 2001.
3. B. K. P. Horn and M. J. Brooks. The variational approach to shape from shading. *CVGIP*, 33(2):174–208, 1986.
4. G. Brelstaff and A. Blake. Detecting specular reflection using lambertian constraints. In *Int. Conf. on Comp. Vision*, pages 297–302, 1988.
5. B. van Ginneken, M. Starvridi, and J. J. Koenderink. Diffuse and specular reflectance from rough surfaces. *Applied Optics*, 37(1):130–139, 1998.

6. P. Beckmann and A. Spizzochino. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon, New York, 1963.

7. 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.

8. R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. volume 15, pages 307–316, 1981.

9. B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.

10. C. Schlick. A survey of shading and reflectance models. *Computer Graphics Forum*, 13(2):121–131, 94.

11. L. B. Wolff. On the relative brightness of specular and diffuse reflection. In *Int. Conf. on Comp. Vision and Patt. Recognition*, pages 369–376, 1994.

12. S. K. Nayar and M. Oren. Visual appearance of matte surfaces. *SCIENCE*, 267:1153–1156, 1995.

13. G. J. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(2):265–272, 1992.

14. S. Westin, J. Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. In *SIGGRAPH 92 Conference Proceedings*, pages 255–264, 1992.

15. X. He, P. Heynen, R. Phillips, K. Torrance, D. Salesin, and D. Greenberg. A fast and accurate light reflection model. In *Siggraph 92 Conference Proceedings*, volume 26, pages 253–254, 1992.

16. P. Schröder and W. Sweldens. Spherical wavelets: Efficiently representing functions on the sphere. In *Siggraph 95 Conference Proceedings*, pages 161–172, 1995.

17. E. P.F. Lafortune, Sing-Choong Foo, K. E. Torrance, and D. P. Greenberg. Non-linear approximation of reflectance functions. In *SIGGRAPH 97 Conference Proceedings*, pages 117–126, 1997.

18. K. J. Dana and S. K. Nayar. Correlation model for 3d texture. In *Int. Conf. on Comp. Vision*, pages 1061–1066, 1999.

19. A. Zisserman.and P. Giblin and A. Blake. The information available to a moving observer from specularities. *Image and Vision Computing*, (7):38–42, 1989.

20. H. Ragheb and E. R. Hancock. Highlight removal using shape-from-shading. In *European Conf. on Comp. Vision*, number 2351 in LNCS, pages 626–641, 2002.