Highlight Removal Using Shape-from-Shading

Hossein Ragheb and Edwin R. Hancock

Department of Computer Science,
University of York, York YO10 5DD, UK.
{hossein,erh}@minster.cs.york.ac.uk

Abstract. One of the problems that hinders the application of conventional methods for shape-from-shading to the analysis of shiny objects is the presence of local highlights. The first of these are specularities which appear at locations on the viewed object where the local surface normal is the bisector of the light source and viewing directions. Highlights also occur at the occluding limb of the object where roughness results in backscattering from microfacets which protrude above the surface. In this paper, we consider how to subtract both types of highlight from shiny surfaces in order to improve the quality of surface normal information recoverable using shape-from-shading.

1 Introduction

Shape-from-shading is concerned with recovering surface orientation from local variations in measured brightness. There is strong psychophysical evidence for its role in surface perception and recognition [12]. Some of the pioneering work in the area was performed by Horn and his co-workers [10]. However, despite considerable effort over the past two decades, reliable shape recovery from shading information has proved an elusive goal [10]. The reasons for this are twofold. Firstly, the recovery of surface orientation from the image irradiance equation is an under-constrained process which requires the provision of boundary conditions and constraints on surface smoothness to be rendered tractable. Secondly, real-world imagery rarely satisfies these constraints. Several authors have attempted to develop shape-from-shading methods which overcome these shortcomings. For instance, Oliensis and Dupuis [18], and Bichsel and Pentland [3] have developed solutions for which shape-from-shading is not under-constrained, but which require prior knowledge of the heights of singular points of the surface. Meanwhile, Kimmel and Brookstein have shown how the apparatus of level-set theory can be used to solve the image irradiance equation as a boundary value problem [11]. Frankot and Chellappa [8] have focused on the differential geometry of the recovered surface and develop a Fourier domain approach for imposing integrability constraints. Ferrie and Lagrarde [7] have used the Darboux-frame smoothing method of Sander and Zucker [21] to impose constraints from differential geometry on the recovered needle-map. A detailed comparative review of these and other related methods for shape-from-shading can be found in the recent comprehensive survey paper of Zhang, Tsai, Cryer and Shah [27].

A. Heyden et al. (Eds.): ECCV 2002, LNCS 2351, pp. 626–641, 2002.

The observation underpinning this paper is that although considerable effort has gone into the recovery of accurate surface geometry, existing shape-from-shading methods are confined to situations in which the reflectance is predominantly Lambertian. When the surface under study is shiny, then the estimated geometry may be subject to error. The main problem that can occur is that surface intensity highlights may lead to misestimation of surface curvature. The most familiar example here is that of surface specularities. These occur at locations on the surface where the local surface normal direction is the bisector of the light source and viewing directions For this reason, if specular hihglights can be accurately located, then they can provide important cues that can be used to constrain the recovery of surface shape. However, there is a less well known effect that results in limb brightening. This is due to surface roughness and results from oblique scattering from microfacets that protrude above the limb perpendicular to the line of sight.

The problem of non-Labmertian and specular reflectance has been widely studied [13,22]. For instance, Healey and Binford [9] have shown how to simplify the Beckmann distribution [1] using a Gaussian approximation to the distribution of specular angle. This simplification can be used in conjunction with the Torrance and Sparrow model [23] to model intensity variations in the analysis of surface curvature. In a comprehensive treatment of specular shape-from-shading Brelstaff and Blake [5] have analysed the geometric constraints provided by specularities, and have shown how to detect specularities using Lambertian irradiance constraints. Drawing on psychophysics, Blake and Bulthoff [4] have developed a computational model of the shape information available to a moving observer from the apparent movement of specularities. Several authors have looked critically at the physics underlying specular reflectance. For instance, Nayar, Ikeuchi and Kanade [16] have shown that the Torrance and Sparrow model [23] is applicable to the modelling of the specular lobe rather than the specular spike. Wolff [24] also has a model which combines diffuse and specular reflectance components, in which the parameters are chosen on the basis of the known physical properties of particular surfaces. In a series of recent papers, Lin and Lee have shown how specular reflections due to multiple light-sources can be located in multi-band imagery [14]. Finally, Nayar, Fang and Boult [17] have used polarisation filters to detect specular reflection.

There has also been a considerable body of work devoted to reflectance from rough surfaces. As noted above, this process is responsible for limb brightening. Oren and Nayar [19] have developed a model which can be used to account for reflectance from surfaces with a rough microfacet structure. Dana, Nayar, Van Ginneken and Koenderink [6] have catalogued the BRDF's for 3D surface textures. Recently, Magda, Kriegman, Zickler and Belhumeur [15] have commented on how shape can be recovered from surfaces with arbitrary BRDF's. Finally, Wolff [24] has shown how the Fresnel term can be used to model reflectance from a variety of surfaces.

In this paper our aim is to incorporate both specular and rough limb reflectance into the shape-from-shading process. This is a two-step process. First, we make estimates of the local surface normals using geometric constraints on

the directions of Lambertian and specular reflectance to recover surface normal directions. The approach is a probabilistic one, which uses a mixture model to estimate the posterior mean direction of Lambertian and specular reflectance.

Once the posterior mean surface normals are to hand, then we can perform photometric correction on the original image. This is again a two-step process. First, we subtract specularities using the Torrance and Sparrow model. Second, we correct the residual non-specular component using the Oren and Nayar model. The result is a corrected Lambertian image from which both local specularities and limb-brightening effects are removed. By applying a Lamertian shape-from-shading algorithm to the corrected image, we obtain an improved estimate of the surface normal directions.

2 Reflectance Geometry

In this section we outline the geometry of the reflectance processes which underpin our shape-from-shading model. We adopt a two-component model in which the predominantly Lambertian surface reflectance exhibits local specular highlights.

2.1 Specular Reflectance

The first component of our reflectance process is concerned with modelling local specular highlights on the observed surface. For specular reflection the surface normal, the light source direction and the viewing direction are coplanar. The incidence angle is equal to the angle of specular reflectance. Hence, for specular reflection, the direction of the surface normal $\boldsymbol{N}_S^{(n)}$ is the bisector of the light source ($\boldsymbol{L}$) and the viewing ($\boldsymbol{V}$) directions and the unit-vector is $\boldsymbol{N}_S^{(n)} = \frac{(\boldsymbol{L}+\boldsymbol{V})}{\|\boldsymbol{L}+\boldsymbol{V}\|}$. It is important to stress that the surface normal for the specular reflectance component is fully constrained if the light source direction and the viewing direction are known. We therefore keep the specular surface normal direction fixed throughout our iterative recovery of the needle-map. The geometry of the specular reflectance process is illustrated in Figure 1a.

2.2 Lambertian Reflectance

In the case of Lambertian reflectance from a matte surface of constant albedo illuminated with a single collimated light-source, the observed intensity is independent of the viewing direction. The observed intensity depends only on the quantity of absorbed light, and this in turn is proportional to the cosine of the incidence angle. Suppose that $\boldsymbol{L}$ is the unit-vector in the direction of the light source and that $\boldsymbol{N}_L(i,j)$ is the unit-vector in the surface normal direction for Lambertian reflectance at the pixel (i,j). According to Lambert's law, the observed image intensity at the pixel with coordinates (i,j) is $E(i,j) = \boldsymbol{N}_L(i,j) \cdot \boldsymbol{L}$.

Lambert's equation provides insufficient information to uniquely determine the surface normal direction. However, as recently observed by Worthington and

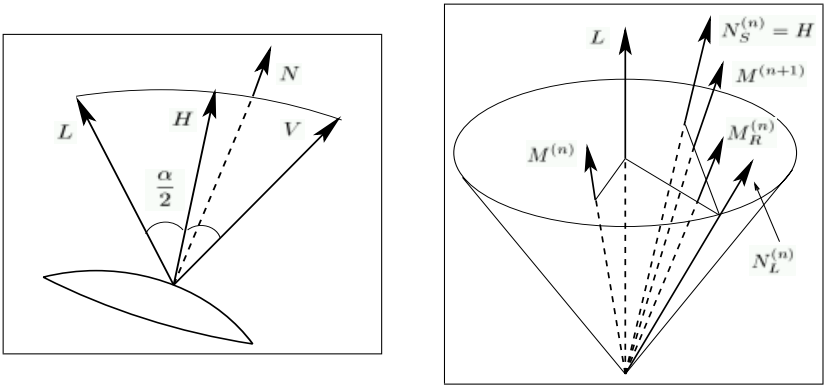

Fig. 1. Geometry of the specular reflectance and the needle-map update process.

Hancock [26], the equation does have a simple geometric interpretation which can be used to constrain the direction of the surface normal. The equation specifies that the surface normal must fall on the surface of a right-cone whose axis is aligned in the light-source direction $\boldsymbol{L}$ and whose apex angle is $\cos^{-1}(E)$.

Worthington and Hancock [26] exploit this property to develop a two-step iterative process for shape-from-shading. The process commences from a configuration in which the surface normals are placed on the position on the irradiance cone where their projections onto the image plane are aligned in the direction of the local (Canny) image gradient. This geometry is illustrated in figure 1b.

In the first step, the surface normal directions are subjected to smoothing in such a way as to satisfy curvature consistency constraints. The resulting smoothed surface normal $\hat{\boldsymbol{N}}_L^{(n)}$ will not fall on the irradiance cone and will hence not satisfy Lambert's law. To overcome this problem, in the second step of the process the smoothed surface normal is rotated onto the nearest location on the irradiance cone. The resulting surface normal, which satisfies Lambert's law, is

$$\boldsymbol{N}_L^{(n+1)} = \hat{\boldsymbol{N}}_L^{(n)} \cos(\theta) + (\boldsymbol{R} \times \hat{\boldsymbol{N}}_L^{(n)}) \sin(\theta) \tag{1}$$

where the rotation axis and the rotation angle are given by

$$\boldsymbol{R} = \frac{\hat{\boldsymbol{N}}_L^{(n)} \times \boldsymbol{L}}{\|\hat{\boldsymbol{N}}_L^{(n)} \times \boldsymbol{L}\|} , \quad \theta = \cos^{-1}(\hat{\boldsymbol{N}}_L^{(n)} \cdot \boldsymbol{L}) - \cos^{-1}(E) \tag{2}$$

In Worthington and Hancock's shape-from-shading method, which deals with matte surfaces that are free from specularities, these smoothing and back-projection steps may be interleaved and iterated until convergence, i.e. a stable needle-map is obtained.

3 Probabilistic Framework

The aim in this paper is to use a previously reported Bayes-decision scheme for separating the two reflectance modes [20]. We compute the *a posteriori* prob-

abilities of specular or Lambertian reflectance. This is done using the iterated conditional modes algorithm of Besag [2]. Although the method has notoriously poor global optimisation properties, we use it here because it is simple and relatively efficient.

The aim is to label pixels according to the reflectance mode from which they originated. The class identity for the pixel (i, j) at iteration n is denoted by $\omega_{i,j}^{(n)}$. The class-identity may be drawn from the set $\Omega = \{S, L\}$ where S is the specular reflectance label and L is the Lambertian reflectance label. For each image location, we maintain a specular surface normal and a Lambertian surface normal which satisfy the geometric constraints outlined in Section 2. At iteration n of the algorithm the currently available estimates of the two surface normals are respectively $\boldsymbol{N}_L^{(n)}(i, j)$ and $\boldsymbol{N}_S^{(n)}(i, j)$. In the case of the specular component, the normal direction is in the direction of local specular reflection, and does not change with iteration number. In the case of Lambertian reflectance, the surface normal direction varies with iteration number, but is always projected to be positioned on the irradiance cone.

To develop our decision process, we require two probabilistic modelling ingredients. The first of these are separate probability density functions which can be used to represent the distributions of surface normals for the two reflectance components. We evaluate these densities at the posterior mean surface normal $\boldsymbol{M}^{(n)}(i, j)$ computed at iteration n. The reason for doing this is that the current values of the two normals are guaranteed to satisfy the geometric constraints outlined in Section 2. As a result, they will be associated with vanishing angular error. Accordingly, we let $q_{i,j}^{(n)}(L) = p(\boldsymbol{M}^{(n)}(i, j)|\omega_{i,j}^{(n)} = L)$ be the probability distribution for the posterior mean surface normal under the Lambertian reflectance model. Similarly, we let $q_{i,j}^{(n)}(S) = p(\boldsymbol{M}^{(n)}(i, j)|\omega_{i,j}^{(n)} = S)$ denote the distribution function for the posterior mean surface normal for the specular reflectance component.

The second probabilistic ingredient is a smoothness prior for the selected surface normal. This component of the model incorporates contextual information. Indexing the surface normals according to their pixel locations, suppose that $\Gamma_{i,j}^{(n)} = \{\boldsymbol{M}^{(n)}(k, l)|(k, l) \in G_{i,j}\}$ is the set of posterior mean surface normals in the neighbourhood $G_{i,j}$ of the pixel (i, j). We let $P_{i,j}^{(n)}(L) = P(\boldsymbol{N}_L^{(n)}(i, j)|\Gamma_{i,j}^{(n)})$ be the conditional probability (or smoothness prior) of the Lambertian surface normal at the location (i, j) given the field of surrounding posterior mean surface normals. With these ingredients, then according to the iterated conditional modes, the probability that the pixel (i, j) belongs to the Lambertian class at iteration n is

$$P(\omega_{i,j}^{(n)} = L|\boldsymbol{M}^{(n)}(i, j)) = \frac{q_{i,j}^{(n)}(L)P_{i,j}^{(n)}(L)}{\sum_{\Lambda\in\Omega} q_{i,j}^{(n)}(\Lambda)P_{i,j}^{(n)}(\Lambda)} \tag{3}$$

The probability that the surface normal belongs to the specular class is the complement, These probabilities can be used to separate the two reflectance modes. With these probabilities to hand, we can update the estimate of the

posterior mean surface normal in the following manner

$$\begin{aligned}\boldsymbol{M}^{(n+1)}(i,j) = \boldsymbol{N}_S^{(n)}(i,j)P(\omega_{i,j}^{(n)} = S|\boldsymbol{M}^{(n)}(i,j)) \\ + \boldsymbol{N}_L^{(n)}(i,j)P(\omega_{i,j}^{(n)} = L|\boldsymbol{M}^{(n)}(i,j))\end{aligned} \tag{4}$$

4 Probability Distributions

To apply the Bayes framework for the posterior mean surface normal estimation, we require probability distributions for Lambertian and specular reflectance, together with a smoothness prior for the surface normal directions.

Specular Reflection. The modelling of specular reflectance and specular highlights has attracted considerable attention in the computer vision and computer graphics communities [23,1,9]. Here we are interested in two approaches to the problem. The first of these is concerned with modelling specular intensities. Here the model of Torrance and Sparrow [23] captures the physics of scattering by the micro-facet structure of a surface. The second approach is to model the angular distribution associated with the reflected light in the proximity of specularities. Here the Beckmann distribution [1] provides a relatively simple model which captures the angular shape of the specular spike. Healey and Binford [9] have an alternative model which can be used to model the distribution of specular intensities for regions of high surface curvature.

We use the Torrance and Sparrow model to distribution of specular intensities. According to this model, the specular intensity is given by

$$I_S^{(n)}(i,j) = \left(K\frac{F}{\pi}\right)\frac{G}{(\boldsymbol{M}^{(n)}\cdot\boldsymbol{V})}\frac{D(\alpha)}{(\boldsymbol{M}^{(n)}\cdot\boldsymbol{L})} \tag{5}$$

The model is controlled by four terms. The first of these is the Fresnel term which is close to unity, i.e. $F \cong 1.0$. Secondly, there is the geometrical attenuation factor

$$G = min\left[1, 2\frac{(\boldsymbol{M}^{(n)}\cdot\boldsymbol{N}_S)(\boldsymbol{M}^{(n)}\cdot\boldsymbol{V})}{(\boldsymbol{V}\cdot\boldsymbol{N}_S)}, 2\frac{(\boldsymbol{M}^{(n)}\cdot\boldsymbol{N}_S)(\boldsymbol{M}^{(n)}\cdot\boldsymbol{L})}{(\boldsymbol{V}\cdot\boldsymbol{N}_S)}\right] \tag{6}$$

Thirdly, there is the facet slope function which we model using the Beckmann distribution [1] to model the distribution of the angle $\alpha = \cos^{-1}(\boldsymbol{M}^{(n)}(i,j)\cdot\boldsymbol{N}_S^{(n)}(i,j))$ between the posterior mean surface normal $\boldsymbol{M}^{(n)}(i,j)$ and the predicted direction of the specular spike $\boldsymbol{N}_S^{(n)}$. The distribution is

$$D(\alpha) = \frac{1}{\sigma_S^2\cos^4\alpha}\exp\left[-\left(\frac{\tan(\alpha)}{\sigma_S}\right)^2\right] \tag{7}$$

where σ_S is a parameter which controls the angular shape of the distribution. This distribution can be used to model the shape of both the specular spike and the specular lobe. It makes no attempt to model the distribution of specular intensities.

Fourthly, and finally, K is a constant which normalizes the equation. We assume that the observed specular intensities follow a Gaussian distribution with variance σ_S^2. Under these assumptions we can write

$$q_{i,j}^{(n)}(S) = \frac{1}{\sqrt{2\pi}\sigma_S} \exp\left[-\frac{1}{2}\left(\frac{E(i,j) - I_S^{(n)}(i,j)}{\sigma_S}\right)^2\right] \tag{8}$$

This distribution models the variation in specular intensity resulting from the physical variation in specular reflectance direction.

Lambertian Reflectance. Our model of the Lambertian reflectance process assumes that the observed intensity values follow a Gaussian distribution with variance σ_L^2. The mean intensity is $(\boldsymbol{M}^{(n)} \cdot \boldsymbol{L})$. Under these assumptions we can write

$$q_{i,j}^{(n)}(L) = \frac{1}{\sqrt{2\pi}\sigma_L} \exp\left[-\frac{1}{2}\left(\frac{E(i,j) - \boldsymbol{M}^{(n)}(i,j) \cdot \boldsymbol{L})}{\sigma_L}\right)^2\right] \tag{9}$$

Smoothness Prior. Our model for the surface normal smoothness prior is based on the average value of the inner product of the surface normal at the location (i,j) with the surrounding field of surface normals. We write

$$P_{i,j}^{(n)}(\Lambda) = \frac{1}{2|G_{i,j}|}\left[|G_{i,j}| + \sum_{(k,l)\in G_{i,j}} \boldsymbol{N}_{\Lambda}^{(n)}(i,j) \cdot \boldsymbol{M}^{(n)}(k,l)\right] \tag{10}$$

When the posterior mean surface normals from the neighbourhood $G_{i,j}$ are aligned in the direction of $\boldsymbol{N}_{\Lambda}^{(n)}(i,j)$, then $P_{i,j}^{(n)}(\Lambda) = 1$, the larger the misalignment then the smaller the value of smoothness prior.

5 Specularity Subtraction

Having described the Bayes framework and the associated two-mode reflectance model, we are now in a position to perform specularity removal. In this section we describe a shape-from-shading algorithm which leads to images free of specularities.

We commence by initialising the algorithm. The initial Lambertian surface normal $\boldsymbol{N}_L^{(0)}(i,j)$ is constrained to lay on the irradiance cone in the direction of the image gradient. The subsequent iterative steps of the algorithm are as follows:

– 1: The field of posterior mean surface normals (initially equal to $\boldsymbol{N}_L^{(0)}$) is subjected to local smoothing. Here we use the curvature sensitive smoothing method [26]. The smoothed surface normal is denoted by $\boldsymbol{M}_R^{(n)}(i,j)$.

- 2: We update the current estimate of the Lambertian surface normal by projecting the smoothed posterior mean surface normal onto the nearest location on the irradiance cone. This gives us the revised surface normal $\boldsymbol{N}_L^{(n)}$.
- 3: With the $\boldsymbol{M}_R^{(n)}(i,j)$ to hand we compute the conditional measurement densities $q_{i,j}^{(N)}(L)$ and $q_{i,j}^{(N)}(S)$ for the two reflectance modes. Taking $\boldsymbol{M}_R^{(n)}(i,j)$, $\boldsymbol{N}_L^{(n)}$ and $\boldsymbol{N}_S^{(n)}$, we compute the smoothness priors $P_{i,j}^{(n)}(L)$ and $P_{i,j}^{(n)}(L)$. Then, we compute the updated *a posteriori* probabilities for both reflectance modes.
- 4: Using $\boldsymbol{N}_L^{(n)}$ and $\boldsymbol{N}_S^{(n)}$ and the updated *a posteriori* probabilities, we compute the new posterior mean surface normal $\boldsymbol{M}^{(n+1)}(i,j)$ and we return to step 1.

The steps of the algorithm are summarised in Figure 1b. The posterior mean surface normals delivered by our shape-from-shading algorithm can be used for the purposes of reconstructing the specular intensity component I_S using the Torrance-Sparrow model given in equation 5. With the reconstructed specular intensity to hand, we can compute the matte reflectance component $I_M(i,j) = E(i,j) - I_S(i,j)$.

6 Correcting for Limb-Brightening

As mentioned earlier, there may also surface brightness anomalies due to rough reflectance from the limbs of objects. Our aim in this section is to show how the Nayar and Oren model can be used to further correct the images obtained by specular subtraction for limb-brightening.

It is well-known that there are a large number of situations where Lambert's law is significantly in error. These include locations near the occluding contour of objects under any illumination conditions. It also applies to situations where the angle of illumination incidence is greater than 50° relative to the viewing direction. Here there will be significant departures from Lambertian reflectance both near the occluding boundary and over a large portion of object area bounded on one side by the shadow boundary with respect to illumination [25]. These non-Lambertian effects are observable for both shiny and rough surfaces.

Oren and Nayar have a qualitative reflectance model for rough surfaces [19]. For a point on a rough surface with illuminant incidence angle θ_i and viewing, or reflectance angle, θ_r the reflectance functions is

$$L_r(\theta_i, \theta_r, \phi_r - \phi_i; \sigma) = \frac{\rho}{\pi} E_0 \cos(\theta_i)(A + B \max[0, \cos(\phi_r - \phi_i)] \sin(\alpha) \tan(\beta)) \tag{11}$$

where $A = 1.0 - 0.5\frac{\sigma^2}{\sigma^2+0.33}$, $B = 0.45\frac{\sigma^2}{\sigma^2+0.09}$ and $\alpha = \max[\theta_i, \theta_r]$, $\beta = \min[\theta_i, \theta_r]$. It is important to note that the model reduces to the Lambertian case when $\sigma = 0$. Here, we aim to utilize this model to deduce a corrected Lambertian reflectance image from the matte component delivered by our specular

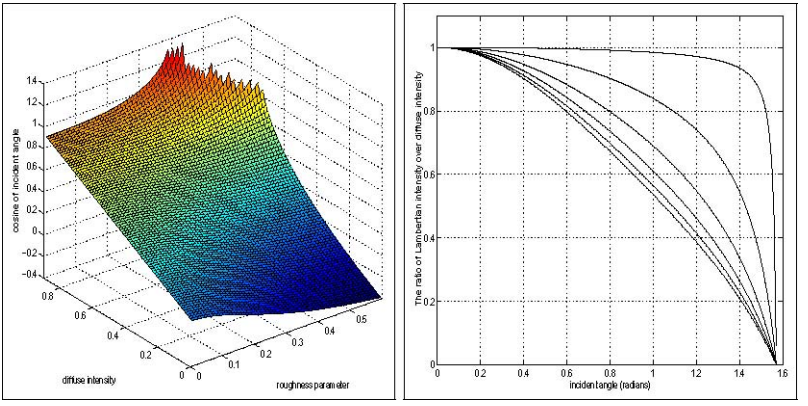

Fig. 2. Plots showing the behaviour of the Oren and Nayar model for rough surfaces.

subtraction method. To do this, we assume that the surface roughness σ is almost constant and the reflectance measurements are obtained in the plane of incidence i.e. $(\phi_r = \phi_i = 0)$. We also confine our attention to the case where the angle between the light source and the viewing directions is small, i.e. $\theta_r = \theta_i = \theta$. With these two restrictions, we can write $\cos(\phi_r - \phi_i) = 1$ and $\alpha = \beta = \theta$. Hence, the non-specular (or diffuse) intensity predicted by the simplified Oren and Nayar model is

$$I_M(i,j) = A\cos\theta + B\sin^2\theta \tag{12}$$

Hence, the matte intensity consists of two components. The first of these is a Lambertian component $A\cos\theta$. The second is the non-Lambertian component $B\sin^2\theta$ which takes on it maximum value where $\theta = \frac{\pi}{2}$, i.e. close to the occluding boundary. To perform Lambertian correction, we proceed as follows. At every pixel location, we use Equation (12) to estimate the angle θ using the subtracted matte intensity and solving the resulting quadratic equation in $\cos\theta$. The solution is

$$\cos\theta = \frac{A \mp \sqrt{A^2 - 4B(I_M(i,j) - B)}}{2B} \tag{13}$$

We take the sign above which results in a value of $A\cos\theta$ which is closest to the matte intensity I_M (in the majority of cases this involves taking the solution associated with the minus sign). This hence allows us to reconstruct the corrected Lambertian reflectance image $I_L = A\cos\theta$. It also gives us an estimate of the opening angle of the Lambertian reflectance cone. This can then be used in the Worthington and Hancock shape-from-shading scheme which assumes the Lambertian reflectance model to recover improved surface normal estimates.

In Figure 2a we show the Lambertian reflectance $\cos\theta$ (equation 13) as a function of the roughness parameter σ and the matte intensity I_M. When the roughness is zero, then the Lambertian and matte intensities are equal to one another. When the roughness increases, then the departures from Lambertian reflectance become more marked.

In Figure 2b we plot the ratio $\frac{I_L}{I_M}$ as a function of the incidence angle θ. The different curves are for different values of the roughness parameter σ. For zero roughness, the ratio is flat, i.e. the reflectance is purely Lambertian. As the roughness increases, then so the value of the ratio decreases by increasing the incidence angle. For normal incidence, the ratio is always unity, i.e. the reflectance is indistinguishable from the Lambertian case, whatever the value of the roughness parameter is.

7 Experiments

The images used in our experiments have been captured using an Olympus 10E camera. The objects studied are made of white porcelain and are hence shiny. Each object has been imaged under controlled lighting conditions in a darkroom. The objects have been illuminated using a single collimated tungsten light source. The light source direction is recorded at the time the images are captured.

To ground-truth the surface highlight removal process, we have used a pair of polaroid filters. We have placed the first filter between the light source and the illuminated object. The second filter was placed between the illuminated object and the camera. For each object we have collected a pair of images. The first of these is captured when the second filter (i.e. the one between the camera and the object) is rotated until there is maximal extinction of the observed specularities. The second image is obtained when the polaroid is rotated through 90 degrees, i.e. there is minimal extinction of the specularities. We refer to the polarisation conditions of the former image as "uncrossed" and of the latter as "crossed".

In Figure 4 we show the results obtained for three of the objects used in our study. The objects are a porcelain bear, a porcelain vase and a porcelain urn. The top row of the figure shows the images obtained with uncrossed polaroids while the second row shows the images obtained with crossed polaroids. The third row shows the difference between the crossed and uncrossed polaroid images. The strongest differences occur at two different locations. Firstly, there are the positions of specularities. From the uncrossed polaroid images it is clear that there are several quite small specular reflections across the surface of the bear. The vase has larger specularities on the neck and the centre of the bulb. The urn has a complex pattern of specularities around the handles. From the crossed polaroid images it is clear that most of the specular structure is removed. The second feature in the difference images are the locations of occluding object limbs, where oblique scattering occurs.

In the fourth row of Figure 4, we show the matte images I_M obtained after specularity subtraction using the Torrance and Sparrow model, The fifth row shows the reconstructed specular intensity obtained using the Torrance and Sparrow model, i.e. I_S. The sixth row shows the difference between the corrected matte images in the fourth row and the uncrossed polaroid images in the top row. Turning our attention to the matte images and the specular images, it is clear that for each of the objects the specular structure is cleanly removed and the matte appearance is close to that obtained with the crossed polaroids. Also,

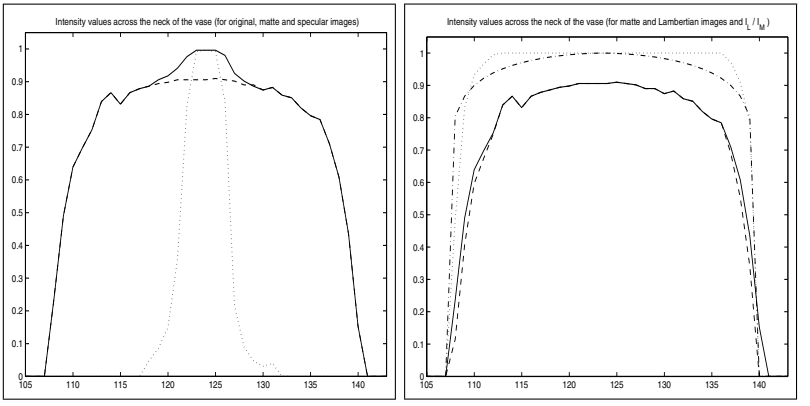

Fig. 3. Intensity plots for different reflectance components across the neck of the vase.

the pattern of specularities obtained in each case corresponds to that obtained by subtracting the crossed polaroid images from the uncrossed polaroid images.

In Figure 5 we investigate the shape information recoverable. The top row shows the Lambertian images after correction for rough limb reflectance using the simplified Oren and Nayar model, i.e. I_L. The second row shows the difference between the corrected Lambertian images and the uncrossed polaroid images.

In the third row we show the needle-maps obtained when we apply shape-from-shading to the images obtained only by specular subtraction, i.e. I_M. In the fourth row of the figure we show the needle maps obtained when the shape-from-shading is applied to the corrected Lambertian images (I_L) appearing in the top row of this Figure. The fifth row of Figure 5 shows the difference in needle-map directions for the matte (I_M) and Lambertian images (I_L). Here the main differences occur at the limbs of the objects.

The sixth row of Figure 5 show the curvedness estimated using the surface normals delivered by the corrected Lambertian images. In the case of the urn the ribbed structure emerges well. The complex surface structure of the bear, in particular the boundaries of the arms and legs, is clearly visible. For the urn the symmetric structure of the neck and the bulb is nicely preserved.

In Figure 6 we show some image reconstructions obtained from the Lambertian image surface normals. The images in the figure are organised into three pairs of rows. In each pair the top row shows the image reconstructions. The lower row shows the differences between the reconstructed images and the subtracted matte images (the fourth row of Figure 4). The first pair of rows show the results obtained using a simple Lambertian model, the second pair are the results obtained with the Oren and Nayer method, and the final pair show the results obtained using Wolff's Fresnel model. The images in the second pair show almost no difference with the matte components. However, the remaining two pairs show stronger patterns of difference. This suggests that the correction process using the Oren and Nayar model has been performed successfully. The reason for this is that the reconstructed images are almost identical to the matte

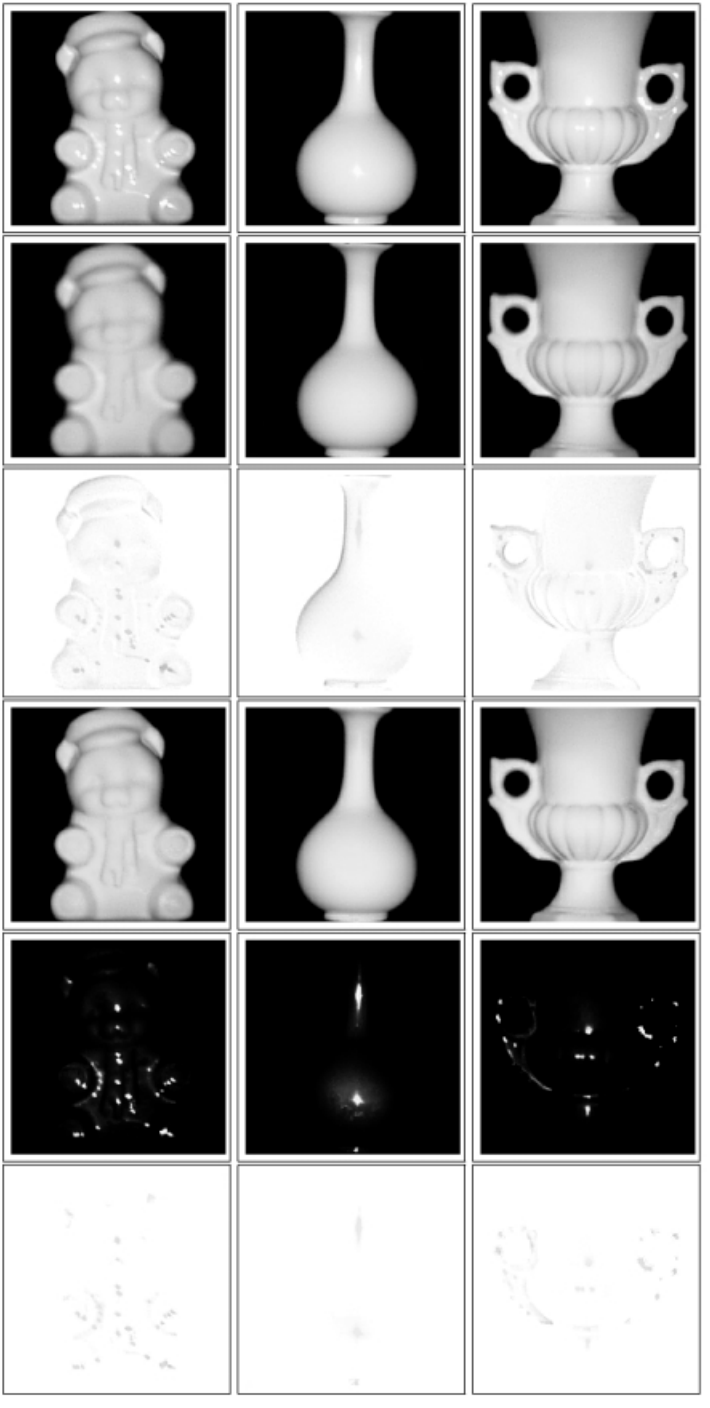

Fig. 4. Applying our specular SFS to separate the specular and matte components.

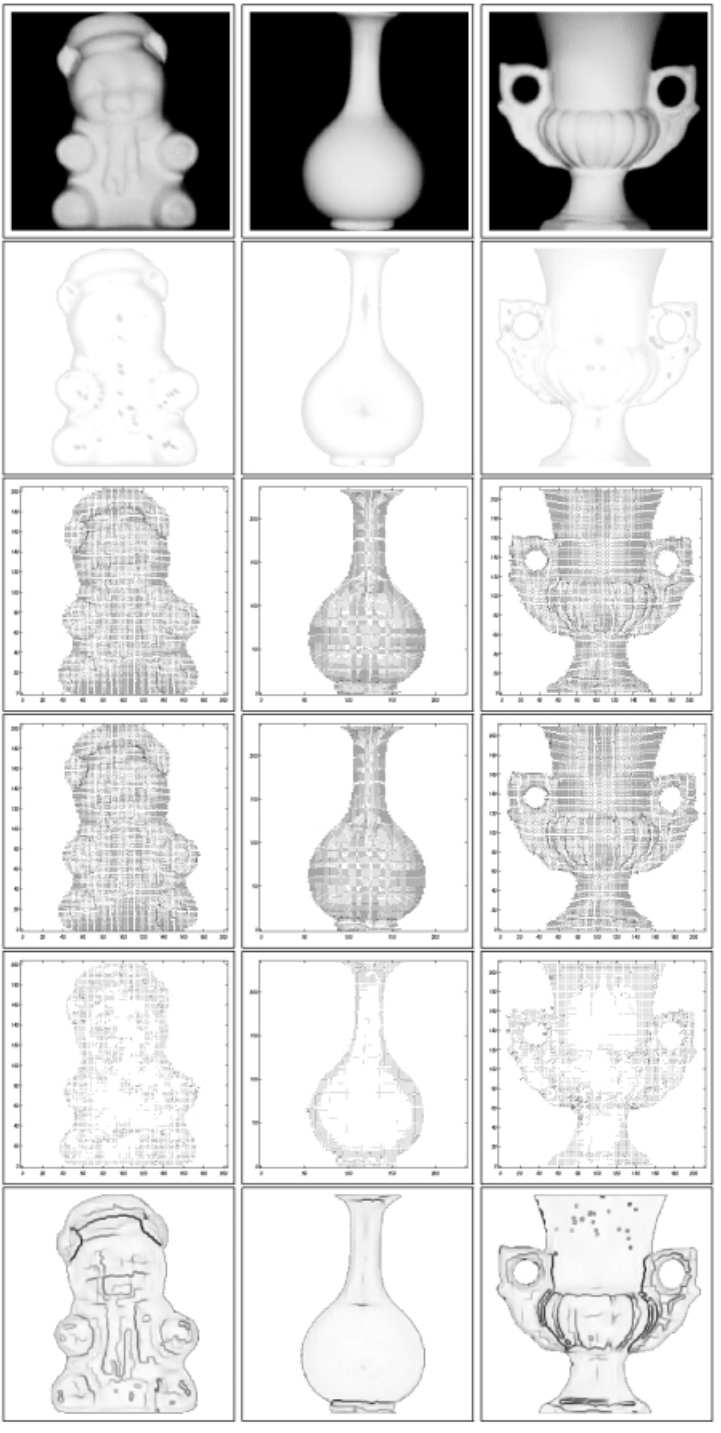

Fig. 5. Surface normals obtained by running SFS over matte and Lambertian images.

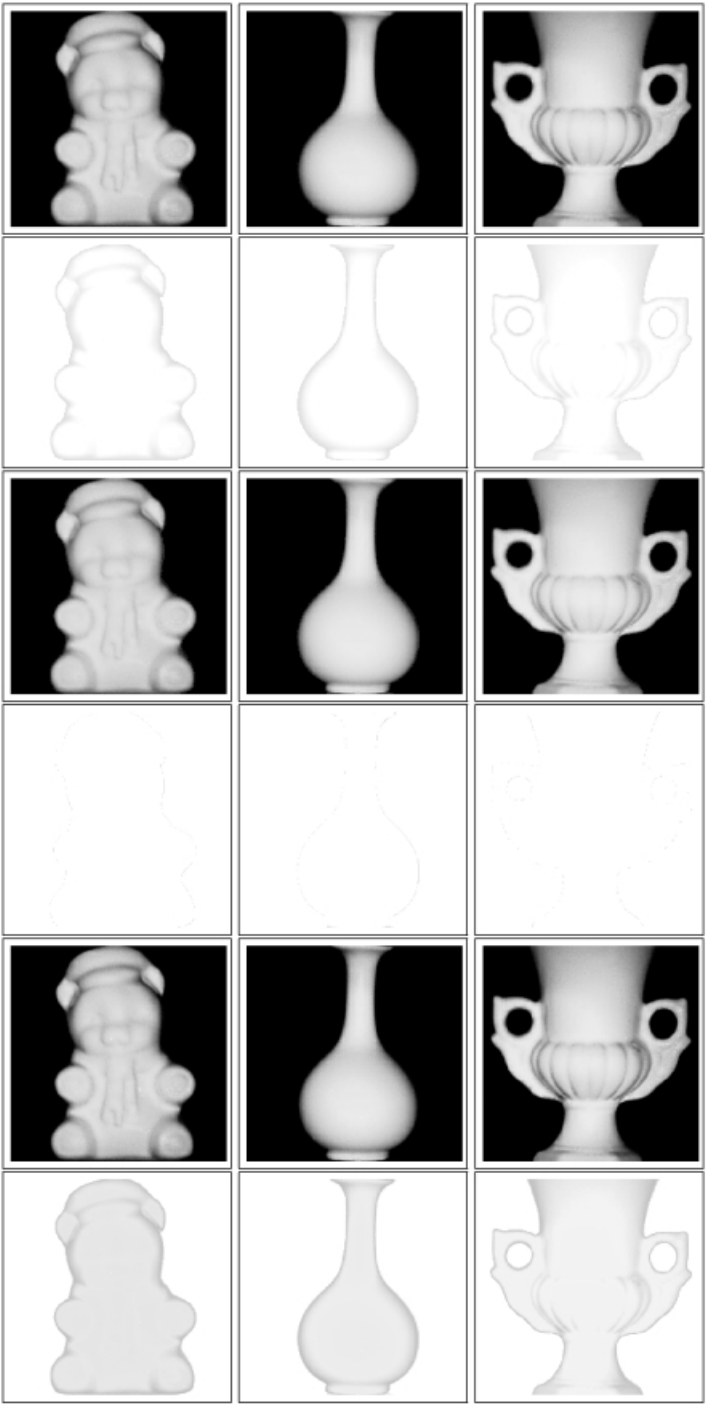

Fig. 6. Reconstructed images using Lambertian, Oren-Nayar and Wolff models.

images. It is also clear that the images reconstructed using the Oren and Nayar model are brighter than the ones reconstructed using the Lambertian model, whereas, the images reconstructed using the Wolff model are darker than both of them.

In Figure 3 we provide some analysis of the different reflectance models used in our experiments. In the left hand panel of the figure, the solid curve is the intensity cross-section along a horizontal line crossing the uncrossed image of the neck of the vase shown in Figure 4. The dashed-curve shows the matte image I_M resulting from specular subtraction, while the dotted curve is the specular component I_S reconstructed using the Torrance and Sparrow model. The specularity on the neck is clearly visible as a peak in the solid curve. This peak is cleanly subtracted in the matte (dashed) curve. In the right-hand panel we focus on the corrected Lambertian image. Here the solid curve is the matte reflectance I_M obtained by specular subtraction. The dashed curve is the corrected Lambertian reflectance I_L. The differences between the two curves are small except at the limbs of the object. To examine the effect of the model in more detail, the dotted curve shows the ratio of corrected Lambertian and matte reflectance $\rho = \frac{I_L}{I_M}$. The ratio drops rapidly towards zero as the limbs are approached. Also shown on the plot as a dash-dot curve is the predicted value of the ratio based on the assumption that the object has a circular cross-section. If x is the distance from the centre and r is the radius of the circle, then value of the ratio at a distance x from the centre is $\rho(x) = \frac{A\sqrt{1-(\frac{x}{r})^2}}{A\sqrt{1-(\frac{x}{r})^2}+B(\frac{x}{r})^2}$. This simple model is in reasonable agreement with the empirical data.

8 Conclusions

In this paper we have shown how to use shape-from-shading to perform photometric correction of images of shiny objects. Our approach is to use estimated surface normal directions together with reflectance models for specular and rough reflectance to perform specularity removal and rough limb-correction. Specularities are modelled using the Torrance and Sparrrow model while the rough limb brightening is modelled using the Oren and Nayar model. We commence by using an iterated conditional modes algorithm to extract surface normals using a mixture of specular and matte reflectance directions. The resulting surface normal directions are used to perform specularity subtraction. Finally, we correct the residual matte reflectance component for rough limb scattering using the Oren and Nayar model. The resulting corrected Lambertian images can be used as input to a conventional shape-from-shading algorithm and result in improved recovery of object-geometry.

References

1. P. Beckmann and A. Spizzochino, *The Scattering of Electromagnetic Waves from Rough Surfaces,* Pergamon, New York, 1963.
2. J. Besag, "On the statistical analysis of dirty pictures," *J. R. Statis. Soc. Lond. B*, Vol. 48, pp. 259-302, 1986.

3. M. Bichsel and A.P. Pentland, "A Simple Algorithm for Shape from Shading," *CVPR*, pp. 459-465, 1992.
4. A. Blake and H. Bulthoff, "Shape from Specularities: computation and psychophysics," *Phil Trans R. Soc. Lond. B*, Vol. 331, pp. 237-252, 1991.
5. G. Brelstaff and A. Blake, "Detecting Specular Reflection Using Lambertian Constraints," *ICCV*, pp. 297-302, 1988.
6. K. Dana, S. Nayar, B. Van Ginneken and J. Koenderink, "Reflectance and Texture of Real-World Surfaces," *CVPR*, pp. 151-157, 1997.
7. F.P. Ferrie and J. Lagarde, "Curvature Consistency Improves Local Shading Analysis," *ICPR*, Vol. I, pp. 70-76, 1990.
8. R.T. Frankot and R. Chellappa, "A Method for Enforcing Integrability in Shape from Shading Algorithms," *ICCV*, pp. 118-125, 1987.
9. G.H. Healey and T.O. Binford, "Local shape from specularity," *ICCV*, pp. 151-160, 1987.
10. B.K.P. Horn and M.J. Brooks, "The Variational Approach to Shape from Shading," *CVGIP*, Vol. 33, No. 2, pp. 174-208, 1986.
11. R. Kimmel and A.M. Bruckstein, "Tracking Level-sets by Level-sets: A Method for Solving the Shape from Shading Problem," *CVIU*, Vol. 62, No. 1, pp. 47-58, 1995.
12. J.J. Koenderink, A.J. van Doorn and A.M.L. Kappers, "Surface Perception in Pictures," *Perception and Psychophysics*, Vol. 52, No. 5, pp. 487-496, 1992.
13. C.C.J. Kuo and K.M. Lee, "Shape from Shading With a Generalized Reflectance Map Model," *CVIU*, Vol. 67, No. 2, pp. 143-160, 1997.
14. S. Lin and S.W. Lee, "Estimation of Diffuse and Specular Appearance," *ICCV*, pp. 855-860, 1999.
15. S. Magda, D. Kriegman, T. Zickler and P. Belhumeur, "Beyond Lambert: Reconstructing Surfaces with Arbitrary BRDFs," *ICCV*, Vol. 2, pp. 391-399, 2001.
16. S.K. Nayar, K. Ikeuchi and T. Kanade, "Surface Reflection: Physical and Geometrical Perspectives," *PAMI*, Vol. 13, No. 7, pp. 611-634, 1991.
17. S.K. Nayar, X. Fang and T. Boult, "Removal of specularities using color and polarization," *CVPR*, pp. 583-590, 1993.
18. J. Oliensis and P. Dupuis, "A Global Algorithm for Shape from Shading," *CVPR*, pp. 692-701, 1993.
19. M. Oren and S.K. Nayar, "Generalization of the Lambertian Model and Implications for Machine Vision," *IJCV*, vol. 14, No. 3, pp. 227-251, 1995.
20. H. Ragheb, and E.R. Hancock, "Separating Lambertian and Specular Reflectance Components using Iterated Conditional Modes," *BMVC*, pp. 541-552, 2001.
21. P.T. Sander and S.W. Zucker, "Inferring surface trace and differential structure from 3D images," *PAMI*, Vol. 12, No. 9, pp 833-854, 1990.
22. H.D. Tagare and R.J.P. deFigueiredo, "A Theory of Photometric Stereo for a Class of Diffuse Non-Lambertian Surfaces," *PAMI*, Vol. 13, No. 2, pp. 133-151, 1991.
23. K. Torrance and E. Sparrow, "Theory for Off-Specular Reflection from Roughened Surfaces," *JOSA*, Vol. 57, pp. 1105-1114, 1967.
24. L.B. Wolff, "On The Relative Brightness of Specular and Diffuse Reflection," *CVPR*, pp. 369-376, 1994.
25. L.B. Wolff, S.K. Nayar and M. Oren, "Improved Diffuse Reflection Models for Computer Vision," *IJCV*, Vol. 30, No. 1, pp. 55-71, 1998.
26. P.L. Worthington and E.R. Hancock, "New Constraints on Data-closeness and Needle-map consistency for SFS," *PAMI*, Vol. 21, No. 11, pp. 1250-1267, 1999.
27. R. Zhang, P. Tsai, J.E. Cryer and M. Shah, "Shape from Shading: A Survey," *PAMI*, Vol. 21, No. 8, pp. 690-706, 1999.