Shape and Refractive Index Recovery from Single-View Polarisation Images

Cong Phuoc Huynh[1]   Antonio Robles-Kelly[1,2]   Edwin R. Hancock[3]

[1]RSISE, Bldg. 115, Australian National University, Canberra ACT 0200, Australia

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

[3] Dept. of Comp. Science, The University of York, York YO105DD, UK

Abstract

In this paper, we propose an approach to the problem of simultaneous shape and refractive index recovery from multispectral polarisation imagery captured from a single viewpoint. The focus of this paper is on dielectric surfaces which diffusely polarise light transmitting through the surface boundary into the air. The diffuse polarisation of the reflection process is modelled using a Transmitted Radiance Sinusoid curve and the Fresnel transmission theory. We provide a method of estimating the azimuth angle of surface normals from the spectral variation of the phase of polarisation. Moreover, to render the problem of simultaneous estimation of surface orientation and index of refraction well-posed, we enforce a generative model on the material dispersion equations for the index of refraction. This generative model, together with the Fresnel transmission ratio permits the recovery of the index of refraction and the zenith angle simultaneously. We show results on shape recovery and rendering for real world and synthetic imagery.

1. Introduction

The polarisation of light is a property that describes the preferential orientation of oscillation of its electromagnetic field. Although the human vision system is oblivious to polarisation, their effects can be captured by devices such as polarimeters and more recently, polarisation cameras [25]. Such developments have broadened the applications of polarisation to areas such as target detection [9] and material property recovery [26].

Shape and material properties are usually co-existing factors that influence the appearance of an object to an observer and the polarisation properties of the emitted radiation. In an early discovery, Torrance, Sparrow and Birkebak [20] were able to measure the specular reflectance distribution of rough surfaces for different polarisation orientations. Reflectance models such as the Torrance-Sparrow model [19] and the Wolff model [24] can also be decomposed into polarisation components and are applicable to polarised light sources. In these models, the material properties and the geometry of the reflection process is expressed in a single equation with multiple degrees of freedom. As a result, the simultaneous recovery of the photometric and shape parameters becomes an underconstrained problem. To remedy the ill-posedness nature of the problem, Miyazaki *et al.* [13] assumed that the histogram of zenith angles of a given object is similar to that of a sphere so as to map the degree of polarisation to zenith angle. Using a similar experimental setup, which involves a spherical optical diffuser, Saito *et al.* [16] were able to recover the shape of transparent objects. Later, Atkinson and Hancock [1] published their work concerning the recovery of surface orientation from diffuse polarisation applied to smooth dielectric surfaces.

To overcome the difficulties in shape recovery from single images, the vision community has turned its attention to the use of multiple images for the recovery task. Rahmann and Canterakis [15] proposed a method based on polarization imaging for shape recovery of specular surfaces. Atkinson and Hancock [3] also made use of the correspondences between the phase and degree of polarisation in two views for shape recovery. Aiming to disambiguate the two possible zenith angles recovered from the degree of specular polarisation, Miyazaki *et al.* [11] analysed the derivative of the degree of polarisation while tilting the observed object by a small angle. As an alternative option to the use of multiple views, polarisation information can also be extracted from photometric stereo images for shape recovery. Using a similar experimental setup to photometric stereo, Thilak *et al.* [18] presented a nonlinear least-squares estimation algorithm to extract the complex index of refraction and the zenith angle of surface normals from multiple images illuminated by unpolarised light sources.

Its somewhat surprising that there is little work on the use of polarisation methods for the simultaneous estimation of surface orientation and photometric invariants, including

index of refraction. Whereas material classification methods [6, 22] have been mainly interested in material intrinsic properties, they have often neglected the wavelength dependency of the index of refraction. In this paper, we propose a new approach to make this problem well-posed by using multispectral information from polarisation images taken from a single view point. In contrast to prior literature on polarisation methods for shape and material recovery, our work concentrates on the use of the Transmitted Radiance Sinusoid and Fresnel transmission theory so as to estimate the shape and refractive index simultaneously from polarisation images observed from a single viewpoint. This framework assumes little knowledge of the illumination conditions, as well as the material under study and the 3D surface orientation.

2. Shape and Refractive Index Recovery

In this section, we present our method for the recovery of surface orientation and index of refraction from a set of polarisation images captured in a single view. Firstly, we commence by decomposing the given set of images into polarisation components. Next, we present an approach to estimating the azimuth angle of surface normals from the spectral phase angles. To complete the section, we show how to jointly estimate the zenith angle of the object surface normals and the refractive index of materials.

The polarisation of light reflected from a surface can be measured by mounting a linear polariser in front of the camera's optics. By rotating the polariser, one can capture polarisation components that are oriented at different angles in a plane orthogonal to the light propagation direction. The intensity captured at each image pixel varies sinusoidally with respect to the angle of rotation of the polariser. Specifically, at each pixel u and wavelength λ, the variation of the image intensity can be described by the following Transmitted Radiance Sinusoid (TRS) curve

$$I(u, \lambda, \vartheta) = \frac{I_{max} + I_{min}}{2} + \frac{I_{max} - I_{min}}{2} \cos(2\vartheta - 2\phi) \tag{1}$$

where ϕ is the phase angle of polarisation of the reflected light, and I_{min} and I_{max} are the minimum and maximum intensities on the sinusoid at the pixel and wavelength of reference, respectively.

Thus, at input, we have a set of N multispectral image cubes $\mathcal{I}(\vartheta_1), \mathcal{I}(\vartheta_2), \ldots, \mathcal{I}(\vartheta_N)$, where the image $\mathcal{I}(\vartheta_i)$ has been acquired with the polariser's transmission axis oriented at the i^{th} step-angle with respect to a reference axis. We can index each of these images to wavelength $\lambda \in \{\lambda_1, \ldots \lambda_K\}$ by denoting $I_i(u, \lambda) \triangleq I(u, \lambda, \vartheta_i)$ as the measured intensity at the pixel-site u and wavelength λ on the image $\mathcal{I}(\vartheta_i)$ corresponding to the i^{th} polariser angle ϑ_i.

2.1. Decomposing Polarisation Images

There are several alternative methods of recovering the I_{min}, I_{max} and ϕ in Equation 1 from a successive sequence of polarisation images captured at several polariser angles. Note that Equation 1 contains three unknowns I_{min}, I_{max} and ϕ. Thus, by capturing the same scene with three or more polariser orientations, one can fit a sinusoidal curve through the intensity-polariser angle pairs using a numerical nonlinear least square fitting algorithm. This method is, however, not efficient since the optimisation has to be performed per pixel. Alternatively, these parameters can be obtained making use of the method in [25], where three images at $0°$, $45°$ and $90°$ with respect to a reference axis are acquired so as to compute the phase, intensity and degree of polarisation.

However, the method in [25] is susceptible to noise corruption since it employs only three images. Here, we employ an alternative that yields a stable estimation of the intensity, phase and degree of polarization by solving an over-determined linear system of equations. The method is akin to that described by Nayar *et al.* in [14]. We commence by rewriting Equation 1 for each pixel-site u and wavelength λ in the following vector form

$$
\begin{aligned}
I_i(u, \lambda) &= \begin{bmatrix} 1 \\ \cos(2\vartheta_i) \\ \sin(2\vartheta_i) \end{bmatrix}^T \begin{bmatrix} \frac{I_{max} + I_{min}}{2} \\ \frac{I_{max} - I_{min}}{2} \cos(2\phi) \\ \frac{I_{max} - I_{min}}{2} \sin(2\phi) \end{bmatrix} \\
&= f_i^T x
\end{aligned} \tag{2}
$$

After collecting $N \geq 3$ measurements at three or more polariser orientations, one arrives at the folllowing over-determined linear system

$$\mathbb{I} = \mathbb{A}x \tag{3}$$

where $\mathbb{I} = \begin{bmatrix} I_1(u, \lambda) \\ I_2(u, \lambda) \\ \cdots \\ I_N(u, \lambda) \end{bmatrix}$ and $\mathbb{A} = \begin{bmatrix} f_1^T \\ f_2^T \\ \cdots \\ f_N^T \end{bmatrix}$

Equation 3 is well-constrained since the number of equations $N \geq 3$ is no less than the number of unknowns. Moreover, the coefficient matrix $\mathbb{A}$ depends solely on the polariser angle with respect to a direction in the image plane and, thus, allowing for an efficient solution of Equation 3 over all the image pixels and wavelengths simultaneously.

Having obtained the solution for $x = [x_1, x_2, x_3]^T$, one can recover the maximal and minimal intensities on the sinusoid and the phase of polarisation at each pixel u and wavelength λ as

$$
\begin{aligned}
I_{max} &= x_1 + \sqrt{x_2^2 + x_3^2} \\
I_{min} &= x_1 - \sqrt{x_2^2 + x_3^2} \\
\phi &= \frac{1}{2} \arctan \frac{x_3}{x_2}
\end{aligned}
$$

2.2. Estimation of the Azimuth Angle

Now, we turn our attention to the estimation of the azimuth angle of the surface normals from the phase angles recovered above. Let us consider a diffuse reflection process from dielectric surfaces. In such a process, diffuse reflection takes place due to incicident light penetrating the surface and scattering within the dielectric body. We can safely assume that the internal scattering within the dielectric body is largely unpolarised, and that the diffusely emitted light from the surface occurs due to refraction through the material-air boundary. Further, the polarisation component oscillating parallel to the plane of reflection is the one that is refracted through the media-air surface boundary with the maximal efficiency. This can be explained by the fact that, the internally scattered light incident upon the material-air boundary is unpolarised and the Fresnel transmission component parallel to the reflection plane is always equal to or higher than the perpendicular component. In addition, the diffuse polarisation oriented at an arbitrary angle is a convex combination of its parallel and perpendicular oscillation components [21].

Furthermore, the maximal intensity on the TRS curve at pixel u and wavelength λ is reached at the phase angle ϕ. Thus, the azimuth angle α at pixel u is either the same as the phase of polarisation ϕ or differs from it by π [1], *i.e.* $\alpha(u) = \phi$ or $\alpha(u) = \phi \pm \pi$. To disambiguate the azimuth angle between the two possibilities, ϕ and $\phi \pm \pi$, we assume convexity on the surface under observation. Under this assumption, surface normals point in the opposite direction to that of image gradients. One can determine the azimuth angle as the one between ϕ and $\phi \pm \pi$ that yields the closer orientation to the negative gradient direction.

However, across a broad spectrum, the estimated phase angle ϕ at a given pixel may vary widely. This is due to weak polarisation that causes a drift in phase angle estimation. Hence, the azimuth angle can be estimated as the weighted mean of the spectral phase of polarisation. The weights can be chosen to reflect the fitting error of the TRS curve to the intensity-polariser angle pairs at each wavelength. Here, the fitting error $\epsilon(u, \lambda)$ is quantified as the L_2-norm of the residual $\epsilon(u, \lambda)$ of the Equation 3, i.e. $\epsilon(u, \lambda) = \|\mathbb{I} - \mathbb{A}x\|_2$, where $\mathbb{I}$, $\mathbb{A}$ and x have been computed per pixel and wavelength. The weight $\omega(u, \lambda)$ associated with the phase of polarisation at each wavelength is defined via a kernel weighting function. Here we employ the Epanechnikov kernel such that

$$\omega(u, \lambda) = \begin{cases} 1 - g(\epsilon(u, \lambda))^2 & \text{if } \|g(.)\| < 1 \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

where $g(\epsilon(u, \lambda))^2 = \frac{\epsilon^2(u, \lambda)}{h}$ and h is a bandwidth parameter.

Since the azimuth angle is a directional quantity, instead of averaging the spectral phase angles directly, we estimate the mean of the sines and cosines of the spectral phase angles for each pixel-site as follows

$$\langle \sin(\phi) \rangle_\lambda = \frac{\sum_\lambda \sin(\phi)\omega(u, \lambda)}{\sum_\lambda \omega(u, \lambda)}$$
$$\langle \cos(\phi) \rangle_\lambda = \frac{\sum_\lambda \cos(\phi)\omega(u, \lambda)}{\sum_\lambda \omega(u, \lambda)}$$

where $\langle \cdot \rangle_\lambda$ denotes the mean value across wavelengths.

The estimated azimuth angle at pixel u then becomes

$$\alpha^*(u) = \begin{cases} \arctan(\frac{\langle \sin(\phi) \rangle_\lambda}{\langle \cos(\phi) \rangle_\lambda}) & \text{if } \langle \cos(\phi) \rangle_\lambda > 0 \\ \arctan(\frac{\langle \sin(\phi) \rangle_\lambda}{\langle \cos(\phi) \rangle_\lambda}) + \pi & \text{if } \langle \cos(\phi) \rangle_\lambda < 0 \\ \frac{\pi}{2} & \text{if } \langle \cos(\phi) \rangle_\lambda = 0 \end{cases} \quad (5)$$

2.3. Joint Estimation of the Zenith Angle and Refractive Index

In Section 2.2, we have shown how to estimate the azimuth angle of surface normals from the spectral phase of polarisation, and how to disambiguate between the azimuth candidate angles. To fully determine the surface shape, one needs to estimate the zenith angle of the surface normals with respect to the viewing direction. In this section, we will show how both the zenith angle and the material index of refraction can be recovered from the Fresnel ratio of diffuse polarisation.

Note that, in prior literature, the degree of polarisation has been heavily exploited for the purpose of recovering the zenith angle of surface normals and the index of refraction [1, 2, 4, 11, 12, 18, 23]. However, there are a number of limitations that prevent the application of these methods in general uncontrolled and unpolarised lighting conditions for unknown shapes and materials. The main reason for this is that, these methods usually rely on either known index of refraction [1, 2, 4, 11, 12] or require many images to be captured in multiple views [11] or under several known light source directions [18]. The need for multiple measurements and complicated setups make them impractical for shape and material analysis on real-world data.

Moreover, existing methods are limited by the kind of data at hand. With a single-band or trichromatic image, there is only a single equation relating the zenith angle and index of refraction through the degree of polarisation. Therefore, one is unable to estimate both of these quantities without further prior knowledge, such as known illumination directions or index of refraction. With an increasing number of wavelengths, *i.e.* image bands, we can utilise the correlation of the index of refraction across the whole spectrum making use of material dispersion equations. There are several approximating functions of dispersion in the physics literature. Of these, perhaps Cauchy's and Sellmeier's [5]

are the most popular. These equations allow the representation of the index of refraction with a few coefficients making use of a linear combination of rational functions of wavelength. As more samples are acquired at an increasing number of wavelengths, the number of equations exceeds the number of variables, rendering the problem solvable. Therefore, the recovery problem becomes well-constrained and can be formulated as a non-linear least square minimisation problem. In brief, the rationale behind our method lies in the use of dispersion equations, which, as we will show in our experiments, enforces constraints that replace the need for the prior knowledge on the index of refraction or illumination conditions.

We focus on showing how the zenith angle and the material index of refraction are jointly estimated. Suppose that the imagery is acquired at K discrete sampled wavelengths $\lambda \in \{\lambda_1, \ldots, \lambda_K\}$. As shown in Section 2.1, at each pixel location u and wavelength λ, we have already obtained the quantities I_{max} and I_{min}. Given unpolarised incident light, the maximal diffuse polarisation is observed in the direction parallel to the reflection plane, which also contains the reflected ray. The minimal diffuse polarisation is observed in the direction perpendicular to the reflected ray. This is a consequence of the developments presented in Section 2.2. Since the unpolarised incident light can be decomposed into incoherent polarisation components oscillating in the parallel and perpendicular planes with equal magnitudes, the polarisation components of the reflected light are proportional to the Fresnel transmission coefficients in the corresponding planes [5], *i.e.*

$$
\begin{aligned}
\frac{I_{min}}{I_{max}} &= \frac{1 - F_\perp(u, \lambda)}{1 - F_\parallel(u, \lambda)} \\
&= \left[\frac{\cos\theta(u)\sqrt{n(u,\lambda)^2 - \sin^2\theta(u)} + \sin^2\theta(u)}{n(u,\lambda)} \right]^2
\end{aligned}
\tag{6}
$$

where $\theta(u)$ is the zenith angle of the surface normal at location u, $n(u, \lambda)$ is the refractive index of the same location at wavelength λ, and $F_\perp(u, \lambda)$ and $F_\parallel(u, \lambda)$ are the perpendicular and parallel Fresnel reflection coefficients.

Let us denote the known quantity $\sqrt{\frac{I_{min}}{I_{max}}}$ as $r(u, \lambda)$. At each pixel location u, we solve the system of Equations 6 with $\lambda \in \{\lambda_1, \ldots, \lambda_K\}$ by converting it to a non-linear least square problem with the following cost function

$$
\begin{aligned}
E(u) = \sum_{i=1}^{K} &\left[\frac{\cos\theta(u)\sqrt{n(u,\lambda_i)^2 - \sin^2\theta(u)} + \sin^2\theta(u)}{n(u,\lambda_i)} \right. \\
&\left. - r(u, \lambda_i) \vphantom{\frac{\sqrt{n}}{n}} \right]^2
\end{aligned}
\tag{7}
$$

With the equation above in hand, we can provide further constraints on the index of refraction according to the Cauchy's dispersion equation, as follows

$$
n(u, \lambda) = \sum_{k=1}^{M} C_1(u)\lambda^{-2(k-1)}
\tag{8}
$$

Here, the dispersion equation relates the index of refraction solely to the wavelength, with coefficients $C_k(u), k \in \{1, \ldots, M\}$ that are characteristic of the material. With this representation of $n(u, \lambda)$, the cost function $E(u)$ in Equation 7 can be reparameterised as a function of $M + 1$ variables, including $\theta(u)$ and $C_k(u), k = 1, \ldots, M$. If the number of M dispersion coefficients is chosen such that $M + 1 \leq K$, where K is the number of wavelengths, then the non-linear least square problem can be solved numerically by standard line-search or trust-region methods. Having obtained the azimuth and zenith angle of the surface normals, one can recover the shape under study by means of a surface integration method such as that in [8]. On the other hand, the index of refraction can be recovered from the dispersion equation making use of the dispersion coefficients $C_k(u)$.

3. Experiments

In this section, we perform experiments using our method for recovering surface orientation and refractive index on synthetic and real-world polarisation images. On the synthetic dataset, we provide the accuracy of our estimated parameters compared to their ground-truth. On the real-world images, we verify the accuracy of the recovered surface orientation and index of refraction through images rendered under novel lighting conditions. For all our experiments, we have set the bandwidth of the Epanechnikov kernel in Section 2.2 to unity. In addition, the form of Cauchy's dispersion equation that we employ approximates the refractive index as an eighth-order rational function of wavelength, with five coefficients. The non-linear least square minimisation problem posed in Section 2.3 is solved by a subspace trust-region method that is based on the interior-reflective Newton method described in [7].

3.1. Synthetic Data

We commence by performing experiments on a synthetic dataset. To build the dataset, we generate the 3D surface of a dome, a ridge, a torus, a volcano and a two-dome test shape. We render images of these synthetic shapes using the refractive index of 24 plastic and liquid materials reported in [10]. The image dataset is synthesized using the Wolff diffuse reflectance model for dielectrics [24] with the total diffuse albedo parameter being the spectral reflectance of Polyvinyl Chloride under ten illumination conditions. These amount

(a) L_3 (b) L_4 (c) L_5 (d) $L_2 + L_4$ (e) $L_1 + L_5$

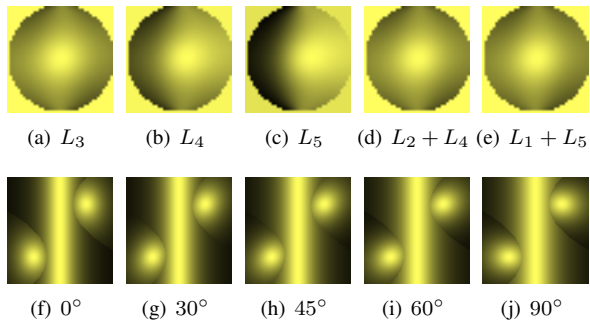(f) $0°$ (g) $30°$ (h) $45°$ (i) $60°$ (j) $90°$

Figure 1. First row: Sample synthetic images of a dome rendered with the refractive index of Polystyrene under five different natural sunlight directions. Second row: Sample synthetic images of a test shape at five different angles of polarisation, under the frontal light source direction.

to a total of 1200 combinations of shape and photometric parameters. For each combination, five polarisation images are generated corresponding to five polariser orientations at $0°$, $30°$, $45°$, $60°$ and $90°$ with respect to the vertical orientation in increments on the clockwise direction. All the multispectral images in our dataset are 30 bands in length, spanning the $430 - 720nm$ range, with a $10nm$ step between successive bands.

The illuminants involved in the simulation lie in the same horizontal plane as the viewing direction. Each illuminant assumes one of five different directions pointing towards the surface under study. We have denoted these illuminants as L_1, L_2, L_3, L_4 and L_5. These directions form angles of $-26.5°$, $-14°$, $0°$, $14°$ and $26.5°$ with respect to the viewing direction, where a negative angle means the light direction is on the left-hand side of the view point and a positive angle means otherwise. In addition, each light assumes the power spectrum of either the natural sunlight or an incandescent light. The power spectrum of the two illuminant types have been acquired from sample light sources using a StellarNet spectrometer. The illuminant combinations used for this dataset include three instances of single illuminant directions, which are L_3, L_4 and L_5, and two instances of two simultaneous illuminant directions which are $L_2 + L_4$ and $L_1 + L_5$.

In Figure 1, we show sample synthetic images generated for two shapes in the dataset. The first row shows the variation of shading with respect to the illumination direction, whereas the second row shows the variation of shading with respect to the angle of polarisation. The pseudocolour RGB images have been synthesized from the generated multispectral images using the Stiles and Burch colour matching function [17]. The images in the figure have been rendered with the refractive index of Polystyrene under the power spectrum of natural sunlight.

In Figure 2, we present the shading and needle maps of

all the simulated shapes under the oblique light direction originating from the right-hand side of the viewing direction, at an angle of $26.5°$ (L_5). The top row shows the input images, corresponding to the vertical polarisation direction. The middle and bottom rows show the recovered diffuse maps and needle-maps of the shapes on the top panels. Note that the surface geometry has been successfully recovered everywhere except for totally dark shadows, irrespective of shading variations. Furthermore, the figure shows that our method results in needle maps highly consistent with the test shapes. Here, the horizontal symmetry of the recovered shading and needle maps also shows that our method relies on polarisation rather than shading so as to reveal the surface geometry. The results also imply that our method is insensitive to changes in illumination power and direction.

To support the qualitative results above, in Table 1, we show the accuracy of the recovered shapes and refractive index spectra. The shape accuracy is quantified as the Euclidean angle, in degrees, between the estimated surface normal direction and the corresponding ground truth per pixel. In columns 2–6, we report the mean of this angular error for each shape and light direction. The first observation is that the shape errors are greatest for flatter shapes such as the Ridge and Volcano. This is since polarisation weakly occurs at surfaces almost perpendicular to the viewing direction. In most cases, the frontal and the combined light directions yield the lowest shape error. In addition, the shape error increases by at most three degrees when the illuminant shifts to the most oblique direction (at an angle of $26.5°$ from the camera axis). In fact, the variation of the mean shape error is less than three degrees across different lighting directions. This observation, again, supports the claim that polarisation is a good cue to surface orientation because it is robust to changes in illumination direction. This is in contrast to Shape from Shading and Photometric Stereo methods, which attribute geometric cues to image shading. Moreover, the standard deviation of the shape error in our results is negligible being never greater than 0.0027 degrees. This means that the recovered shape is stable across different materials and illuminant power spectra.

In Table 1, we also report the accuracy of the recovered refractive index. For each synthetic image, the accuracy is quantified as the Euclidean angle in degrees between the mean of the recovered refractive index spectra across all the pixels and the ground truth refractive index used for synthesizing the image. This accuracy measure is meaningful because our method aims at recovering an apparent refractive index, which may be a scalar multiple of the actual one. Furthermore, for the purposes of recognition, the scaling factor is a less important issue than the spectral variation of refractive index. The reported angular errors in columns 7–11 show that we can recover the refractive index at a high

accuracy level and that the estimation is stable across all the shapes and lighting directions. In particular, the change in lighting direction hardly affects the recovery process. This observation is consistent with the shape recovery results. Therefore, these quantitative results demonstrate that our recovery method, to an extent, is robust against change in illumination and material.

3.2. Real-world Imagery

In this section, we turn our attention to real-world multispectral images acquired in house using a hyperspectral camera with an acousto-optic tunable filter. To acquire the imagery in our experiments, we have tuned the filter to wavelengths in the range of $400 - 650nm$, with a spectral resolution of $5nm$. We have acquired multispectral images of four different objects made of matte plastic and porcelain. To measure polarisation, we mount a linear polariser in front of the camera lens and capture the polarised images when the polariser transmission axis is oriented at each of seven different angles, i.e. 45, 60, 75, 90, 105, 120 and 135 degrees, with respect to the vertical axis in the clock-wise direction.

The input images to the recovery algorithm are captured under two unpolarised artificial sunlights simultaneously illuminating the left and right hand side of the objects. Note that the algorithm does not require prior knowledge of the illuminant power spectrum and direction. Here, we have used two illuminants to ensure that there is no dark shadow in the captured images. To produce realistic rendering results, there is a need for obtaining the material reflectance. We achieve this by capturing an unpolarised multispectral image of the same object under an illuminant direction aligned to the viewing direction.

Our method delivers, at output, the surface orientation and material index of refraction taking at input the polarisation images. For the recovery of the reflectance, we utilise the fact that the incident and reflection angles are the same. Suppose that the illuminant power spectrum is known, we can use the Wolff model equation [24] to solve for the material reflectance since the other parameters involved in the model, i.e. incoming and outgoing light direction and index of refraction, are recovered by our algorithm.

Figure 3 shows sample input images of the four objects under study. The images shown here are the pseudocolour RGB version of the input multispectral images, simulated with the Stiles and Burch colour matching functions [17]. In the figure, the first row shows the polarised images of the test objects with the polariser's transmission axis forming an angle of $45°$ to the vertical direction. The second and third rows show the recovered shading and needle-maps of the objects on the top row. Note that the needle maps show clear overall surface contours and correct surface normal orientation along occlusion boundaries. The shading maps

are intuitively correct, being more prone to error along material boundaries. This is due to the fact that our method is based solely on polarisation information and, therefore, the changes at material boundaries can be interpreted as a variation in object geometry.

Having obtained the surface orientation and material properties of the objects under study, we re-render their images under novel illuminants. For validation, we have captured ground truth images of these objects under directional incandescent lights placed in the horizontal plane going through the camera's position. Specifically, from left to right, these light directions point towards the illuminated object, forming angles of -45, -30, 0, 30 and 45 degrees with the viewing direction, where, as before, a negative angle means the light is located on the left-hand side of the camera and vice versa. We denote these light directions L_1, L_2, L_3, L_4 and L_5, respectively, so as to be consistent with the synthetic dataset.

In Figure 4, we show the rendering results for one of the objects in Figure 3 obtained making use of an incandescent illuminant. In the top row, from left to right, we show the images rendered under illuminant combinations L_3, L_4, $L_1 + L_5$, and $L_2 + L_4$. The bottom row shows the corresponding ground truth images captured under the same light directions. The rendering is successful for the frontal light L_3 and the two-illuminant settings $L_1 + L_5$, and $L_2 + L_4$. In the extreme condition of the oblique light L_4, the rendering result shows the effect of the geometric cue induced by
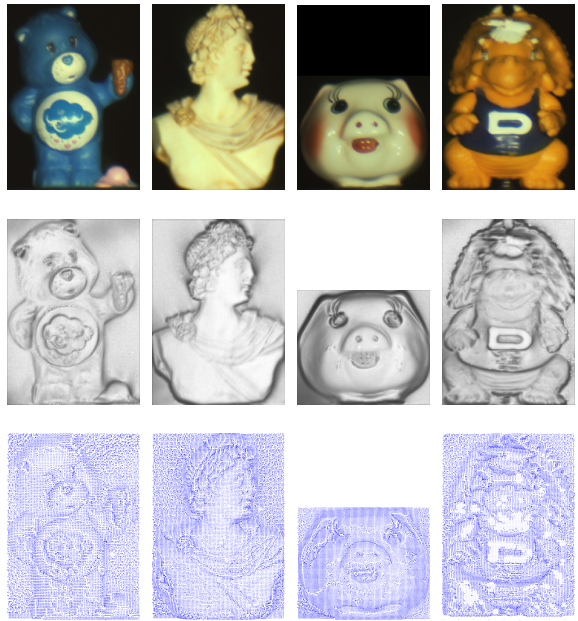

Figure 3. First row: Polarised images of sample real-world objects; Middle row: Recovered shading maps; Bottom row: Needle maps yielded by our method.
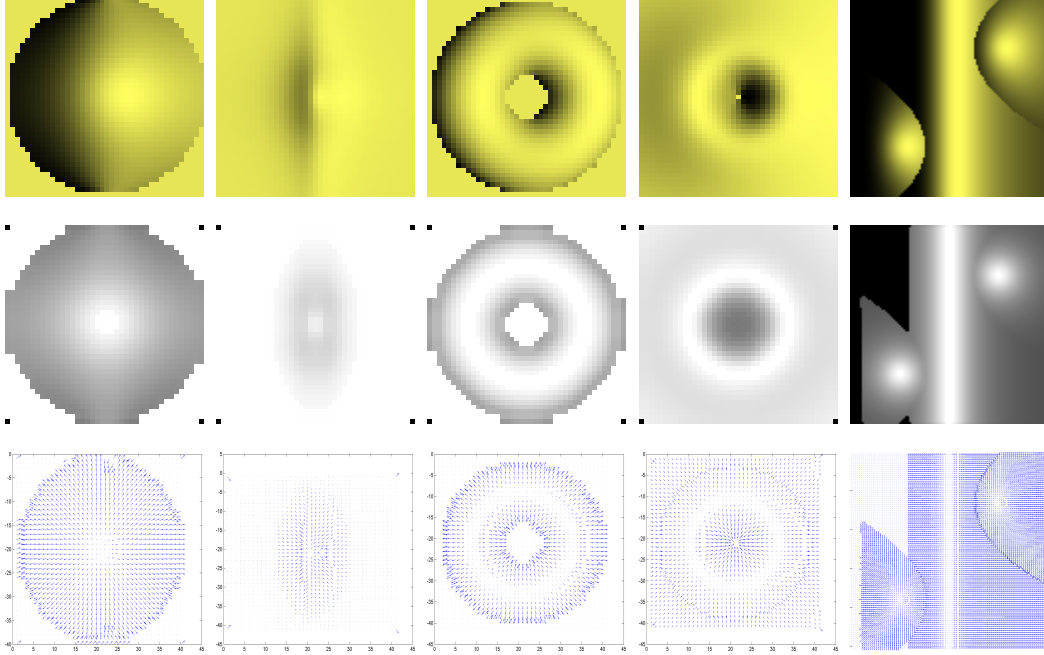
Figure 2. Results on synthetic shapes under an oblique light source direction at an angle of $26.5°$ on the right of the viewing direction. Top row: Input polarisation images; Middle row: Recovered shading maps from the input images. Bottom row: Recovered needle maps.

	Shape error (degrees)					Refractive index error (degrees)				
	L_3	L_4	L_5	$L_2 + L_4$	$L_1 + L_5$	L_3	L_4	L_5	$L_2 + L_4$	$L_1 + L_5$
Dome	7.36	7.38	8.49	7.36	7.35	0.41 ± 0.11	0.41 ± 0.11	0.41 ± 0.11	0.41 ± 0.11	0.41 ± 0.11
Ridge	16.04	15.11	13.12	16.04	16.04	0.28 ± 0.18	0.28 ± 0.18	0.28 ± 0.18	0.28 ± 0.18	0.28 ± 0.18
Torus	10.60	10.92	11.58	10.60	10.60	0.17 ± 0.13	0.17 ± 0.13	0.17 ± 0.13	0.17 ± 0.13	0.17 ± 0.13
Volcano	21.54	21.92	23.94	21.55	21.55	0.19 ± 0.17	0.19 ± 0.17	0.19 ± 0.17	0.19 ± 0.17	0.19 ± 0.17
Test shape	2.81	3.61	4.55	2.91	3.10	1.98 ± 0.17	1.87 ± 0.17	1.63 ± 0.17	1.98 ± 0.17	1.98 ± 0.17

Table 1. The accuracy of the recovered parameters. Columns 2–6 show the error (in degrees) of the needle maps recovered by our method. Columns 7–11 show the error of the estimated refractive indices in terms of Euclidean deviation angle from the ground truth.

material changes in the object as shadows on the resulting image.

For a more quantitative analysis, in Table 2, we show the rendering accuracy under novel lighting directions. The error is measured as the angular difference between the rendered hyperspectral images and their ground-truth on a per-pixel basis. It is worth stressing in passing that, on the panels, so far we have shown pseudocolour images, nonetheless, our error measures are obtained on the spectra and not on the RGB values yielded by the color matching functions. The figures reported in Table 2 are the mean and standard deviation across pixels in each image. The results here are consistent with the qualitative results in Figure 4, in the sense that the rendering quality is higher for the cases of frontal illuminants and two simultaneous illuminants. This is due to the fact that in such conditions, the objects are fully illuminated and the image produced has a smooth shading variation. For oblique illuminant conditions, the lower rendering accuracy is due to non-smooth shading where shadows occur across material boundaries. Nonetheless, the algorithm can produce rendering results that are in good accordance with the ground truth and, being based on polarisation information, it delivers shading maps which conform well to the geometry of the object under study.

4. Conclusion

We have a method for estimating shape and index of refraction simultaneously based upon polarisation in single-view multi-spectral imagery. The method presented here utilises the spectral variation of the phase of polarisation for estimating the azimuth angle of surface normals. We have also drawn upon the Fresnel transmission ratio to recover the zenith angle of the surface normals and the refractive index simultaneously. To make this problem well-posed,

	L_3	L_4	L_5	$L_1 + L_5$	$L_2 + L_4$
Bear	13.36 ± 6.33	14.19 ± 5.62	13.39 ± 4.94	11.07 ± 6.12	11.79 ± 7.15
Statue	12.51 ± 6.82	12.91 ± 7.05	12.24 ± 6.00	10.95 ± 5.68	11.10 ± 6.72
Pig	10.96 ± 6.18	12.03 ± 7.89	11.83 ± 7.43	8.11 ± 5.13	9.69 ± 5.70
Dinosaur	14.43 ± 5.73	16.18 ± 8.70	16.69 ± 9.02	11.10 ± 4.69	11.05 ± 5.46

Table 2. The angular deviation (in degrees) of images rendered under the frontal view direction, from their ground truth, across different light source directions.

we have enforced a wavelength-dependent dispersion equation on the index of refraction. The experimental results on synthetic and real image data demonstrate the merit of our method for purposes of shape and refractive index recovery.

References

[1] G. Atkinson and E. Hancock. Recovery of surface orientation from diffuse polarization. *IEEE Transactions on Image Processing*, 15(6):1653–1664, June 2006.

[2] G. Atkinson and E. R. Hancock. Recovery of surface height using polarization from two views. In *CAIP*, pages 162–170, 2005.

[3] G. A. Atkinson and E. R. Hancock. Multi-view surface reconstruction using polarization. In *IEEE International Conference on Computer Vision (ICCV'05), Volume 1*, pages 309–316, 2005.

[4] G. A. Atkinson and E. R. Hancock. Shape estimation using polarization and shading from two views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11):2001–2017, 2007.

[5] M. Born and E. Wolf. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light (7th Edition)*. Cambridge University Press, 7th edition, 1999.

[6] H. Chen and L. Wolff. Polarization phase based method for material classification in computer vision. *IJCV*, 28(1):73–83, June 1998.

[7] T. F. Coleman and Y. Li. An Interior, Trust Region Approach for Nonlinear Minimization Subject to Bounds. *SIAM Journal on Optimization*, 6(2):418–445, 1996.

[8] Frankot, Robert T. and Chellappa, Rama. A Method for Enforcing Integrability in Shape from Shading Algorithms. *IEEE Trans. Pattern Anal. Mach. Intell.*, 10(4):439–451, 1988.

[9] F. Goudail, P. Terrier, Y. Takakura, L. Bigu, F. Galland, and V. DeVlaminck. Target detection with a liquid-crystal-based passive stokes polarimeter. *Applied Optics*, 43:274–282, 2004.

[10] S. N. Kasarova, N. G. Sultanova, C. D. Ivanov, and I. D. Nikolov. Analysis of the dispersion of optical plastic materials. *Optical Materials*, 29(11):1481 – 1490, 2007.

[11] D. Miyazaki, M. Kagesawa, and K. Ikeuchi. Transparent surface modeling from a pair of polarization images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(1):73–82, 2004.

[12] D. Miyazaki, M. Saito, Y. Sato, and K. Ikeuchi. Determining surface orientations of transparent objects based on polarization degrees in visible and infrared wavelengths. *Journal of Optical Society America A*, 19:687–694, 2002.

[13] D. Miyazaki, R. T. Tan, K. Hara, and K. Ikeuchi. Polarization-based inverse rendering from a single view. *IEEE International Conference on Computer Vision*, 2:982, 2003.

[14] S. K. Nayar, X.-S. Fang, and T. Boult. Separation of reflection components using color and polarization. *Int. J. Comput. Vision*, 21(3):163–186, 1997.

[15] S. Rahmann and N. Canterakis. Reconstruction of specular surfaces using polarization imaging. *IEEE Conference on Computer Vision and Pattern Recognition*, 1:149, 2001.

[16] M. Saito, H. Kashiwagi, Y. Sato, and K. Ikeuchi. Measurement of surface orientations of transparent objects using polarization in highlight. *IEEE Conference on Computer Vision and Pattern Recognition*, 1:1381, 1999.

[17] W. S. Stiles and J. M. Burch. N.P.L. colour-matching investigation: Final report (1958). *Optica Acta*, 6:1–26, 1959.

[18] V. Thilak, D. G. Voelz, and C. D. Creusere. Polarization-based index of refraction and reflection angle estimation for remote sensing applications. *Applied Optics*, 46:7527–7536, 2007.

[19] K. Torrance and E. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57:1105–1114, 1967.

[20] K. E. Torrance, E. M. Sparrow, and R. C. Birkebak. Polarization, Directional Distribution, and Off-Specular Peak Phe-

(e) L_3 (f) L_4 (g) $L_1 + L_5$ (h) $L_2 + L_4$
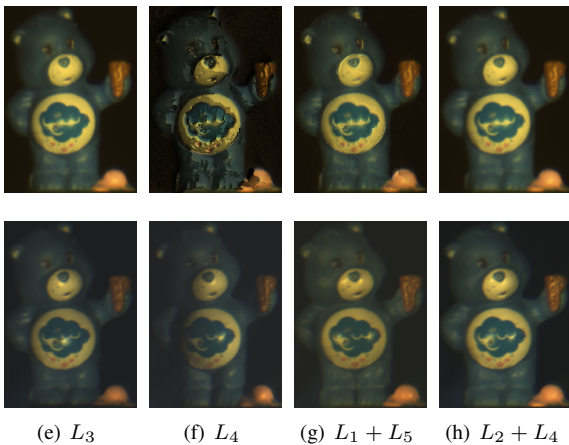
Figure 4. Top row: Rendered images of the bear; Bottom row: Ground truth images.

nomena in Light Reflected from Roughened Surfaces. *Journal of the Optical Society of America (1917-1983)*, 56:916–925, July 1966.

[21] L. Wolff and T. Boult. Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):635–657, 1991.

[22] L. B. Wolff. Polarization-based material classification from specular reflection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(11):1059–1071, 1990.

[23] L. B. Wolff. Shape from polarization images. In *Radiometry*, pages 193–199. Jones and Bartlett Publishers, Inc., USA, 1992.

[24] L. B. Wolff. Diffuse-reflectance model for smooth dielectric surfaces. *Journal of the Optical Society of America*, 11(11):2956–2968, November 1994.

[25] L. B. Wolff. Polarization vision: a new sensory approach to image understanding. *Image Vision Computing*, 15(2):81–93, 1997.

[26] L. B. Wolff and T. E. Boult. Polarization/radiometric based material classification. In *CVPR 89*, pages 387–395, 1989.