

# Shape and Refractive Index from Single-View Spectro-Polarimetric Images

Cong Phuoc Huynh<sup>1,2</sup> \*, Antonio Robles-Kelly<sup>1,2</sup> and Edwin R. Hancock<sup>3</sup>

<sup>1</sup>School of Engineering, Australian National University, Canberra ACT 0200, Australia

<sup>2</sup>National ICT Australia (NICTA) †, Locked Bag 8001, Canberra ACT 2601, Australia

<sup>3</sup>Department of Computer Science, University of York, Heslington, York, YO10 5DD, UK

## Abstract

In this paper, we address the problem of the simultaneous recovery of the shape and refractive index of an object from a spectro-polarimetric image captured from a single view. Here, we focus on the diffuse polarisation process occurring at dielectric surfaces due to subsurface scattering and transmission from the object surface into the air. The diffuse polarisation of the reflection process is modelled by the Fresnel transmission theory. We present a method for estimating the azimuth angle of surface normals from the spectral variation of the phase of polarisation. Moreover, we estimate the zenith angle of surface normals and index of refraction simultaneously in a well-posed optimisation framework. We achieve well-posedness by introducing two additional constraints to the problem, including the surface integrability and the material dispersion equation. This yields an iterative solution which is computationally efficient due to the use of closed-form solutions for both the zenith angle and the refractive index in each iteration. To demonstrate the effectiveness of our approach, we show results of shape recovery and surface rendering for both real-world and synthetic imagery.

**Keywords:** polarisation, shape recovery, refractive index, spectro-polarimetric imagery, multi-spectral imagery, hyperspectral imagery, Fresnel reflection, dispersion equations.

---

\*Corresponding author. E-mail: cong.huynh@nicta.com.au. Tel: +61(2) 6267 6234

†NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program.

# 1 Introduction

Polarisation measures the orientation of the electric field oscillations of light in the plane perpendicular to the direction of propagation. It has been widely utilised to develop powerful measurement and imaging techniques in various branches of physics including astronomy [16], applied optics [7, 23] and crystallography. Although the human vision system is insensitive to polarisation, a number of organisms including the Mantis shrimp, naturally possess a polarisation vision system [24]. In biology, researchers have also observed evidence for biophysical mechanisms of polarisation coding in various species of fish [18]. With recent advances in camera technology, polarisation effects can be captured by devices such as polarimeters and more recently, polarisation cameras [44, 45, 47]. Wolff *et al.* [44, 45, 47] have developed a liquid crystal polarisation video camera. The key to this development is to use electro-optically controlled Twisted Nematic liquid crystals to replace the need for mechanically rotated linear polarisers. The development of such portable, low-cost and fast polarisation camera sensors potentially extends the applications of polarisation imaging to areas such as target detection and segmentation [14, 33] and material property recovery [46].

In image analysis, polarisation provides an important source of information concerning both the shape and the material composition of the object being observed. Object shape and surface material properties such as refractive index, both directly influence the appearance of an object to an observer. They also determine the polarisation properties of the reflected light. In early work, Torrance, Sparrow and Birkebak [40] measured the specular reflectance distribution of rough surfaces for different polarisation orientations. The reflectance model attributes polarisation to specular reflection from a collection of small and randomly disposed mirror-like facets that constitute the surface. Their model includes a specular reflection component based on the Fresnel reflection theory together with a microfacet distribution function that accounts for surface roughness. It is interesting that more recent reflectance models such as the Torrance-Sparrow model [20] and the Wolff model [43] which aim to incorporate more complex surface models, also draw on the Fresnel reflection theory. As a result, they share the common feature of considering the reflected light as a combination of polarisation components parallel to and perpendicular to the plane of reflection. In consequence they can accommodate polarised light sources. One of the common features of these Fresnel-based models, is that physical properties of the surface material and the geometry of the reflection process are expressed in a single equation with multiple degrees of freedom. As a

result, the simultaneous recovery of the photometric parameters and shape information becomes an underconstrained problem.

To remedy the ill-posedness nature of the problem, Miyazaki *et al.* [27] has assumed that the histogram of zenith angles for a given object was similar to that for a sphere. He has used this property to recover a mapping from the degree of polarisation to the zenith angle. Despite being effective, their approach is limited to only surfaces with a uniform distribution of surface normal directions. Moreover, the mapping is not necessarily consistent across different material refractive indices. The related work in [27] and [26] employed the degree of polarisation in the visible and far-infrared regions to resolve ambiguities in determining the surface orientation of transparent objects. The main drawback of the method in [26] is the need for an omni-directional diffuse illumination source, which limits its applicability in real-world settings. Moreover, the method requires measurements of the energy emitted in the far-infrared spectrum, which is susceptible to thermal noise due to heating of the object under study. Using a similar experimental setup involving a spherical optical diffuser, Saito *et al.* [34] were able to recover the shape of transparent objects with known refractive index. Rahmann [30] has presented a method for computing the light source position together with the orientation of rough surfaces through the polarisation information conveyed by specular highlights. However, because the method in [30] relies on the existence of highlights reflected from flat surfaces, it cannot be applied to objects which exhibit purely diffuse reflectance. In a subsequent development, Rahmann [31] employed level sets to reconstruct the surface of an object from a single polarisation image. This was done under the assumption of weak perspective camera projection. Atkinson and Hancock [1] have recovered the surface orientation from the measured diffuse polarisation for smooth dielectric surfaces. However, they assumed a known refractive index in order to estimate the zenith angle of surface normals from the degree of polarisation. The available methods for inferring surface orientation using single viewpoint polarisation images assume either a known refractive index or a known surface orientation distribution, or alternatively require a complex instrumental setup.

Despite this progress, the recovery of object shape from a single view still remains a challenging task due to the presence of photometric artifacts and discontinuities on the object surface. To overcome these difficulties, several attempts have been made to make use of polarisation images captured with varying viewpoint and light source direction. Rahmann and Canterakis [32] developed a polarisation imaging method which recovers the shape of specular surfaces. Their method has made use of the correspondences between the polarisation phases recovered from multiple

viewpoints. They showed that three viewpoints were sufficient for surface reconstruction. Atkinson and Hancock [4] also made use of the link between the phase and the degree of polarisation for shape recovery. Using two viewpoints, they developed a method that disambiguates the two candidate azimuth angles for each surface normal. They also showed how to resolve the ambiguous mapping from the degree of specular polarisation to the zenith angle of the surface normal. However, the method in [4] assumes the homography between the views is known in order to match points across the imagery. This work was later extended in [5], where robust statistics were used to refine the correspondence estimates between the two views of an object. Miyazaki *et al.* [25] disambiguated the two possible zenith angles which yield the same degree of specular polarisation by physically tilting the observed object by a small angle.

As an alternative to the use of multiple viewpoints, polarisation imaging can also be combined with photometric stereo using images of an object from a fixed viewpoint but under different light source directions. For instance, Drbohlav and Šára [11] showed how to disambiguate surface orientations from uncalibrated photometric stereo images by altering the polarisation angles of the incident and emitted light. Their method uses two projections of the object normals onto two planes perpendicular to both the viewing direction and the illumination direction. When combined with the surface integrability constraint, this yields a method that can cope with both the bas-relief [6] and convex/concave ambiguities. Atkinson and Hancock [3] have disambiguated the surface normal directions by combining polarisation measurements with photometric stereo data captured with three known light source positions. Thilak *et al.* [39] presented a nonlinear least-squares estimation algorithm which can be used to extract both the complex index of refraction and the zenith angle of the surface normals from multiple images illuminated by unpolarised light sources. However, their method requires prior knowledge of the light source positions relative to the observer. Furthermore, it employs a polarimetric Bidirectional Reflectance Distribution Function (BRDF) of light scattering, and this is limited to the case where the light source direction, the surface normal direction and the view direction are co-planar.

Polarisation has also proven to be an effective tool in revealing the material properties of surfaces from images. Early work by Wolff and Boult [46] showed how to classify image regions as belonging to metallic or dielectric materials. In this work, the ratio between the two Fresnel reflection components, *i.e.* the so-called Fresnel ratio, was used to characterise the relative electrical conductivity of the surface, and was used as a discriminant feature for classification. Later, the same authors used a Fresnel reflectance model to predict the magnitude of polarisation for an arbitrary

orientation in the image plane [41, 42]. Using this model, the Fresnel ratio can be estimated and used for the classification of dielectrics. More recently, Chen and Wolff [8] employed the phase angle of polarisation for distinguishing conducting metals from dielectric materials. Their approach hinges on the physical observation that upon reflection from a metal surface the phase difference between polarisation components is altered [7], whereas this is not the case for dielectrics.

## 1.1 Contribution

Given the current interest in polarisation vision particularly within the Computer Vision community, it is somewhat surprising that there is little work on the use of polarisation methods for the simultaneous estimation of surface orientation and surface properties such as the index of refraction. In fact, most of the shape recovery methods rely on the assumption that the index of refraction is known [1, 3, 4, 5, 25, 26, 27, 34]. On the other hand, surface material classification methods [8, 46, 42, 41] have often neglected the wavelength-dependence of the index of refraction. This is an important omission since, it is well known in physics that refractive index is wavelength dependent and governed by the material dispersion equations, with concrete examples being provided by Cauchy and Sellmeier's equations [7].

In our work, we utilise the wavelength dependence as an additional constraint to render the recovery problem well-posed. This leads to a novel method that is based on the spectral (wavelength) variation of spectro-polarimetric images captured from a single viewpoint. In contrast to alternative methods reported elsewhere in the literature, our work makes use of three distinct model ingredients, namely a) Fresnel reflection theory, b) the material dispersion equations governing the wavelength dependence of refractive index and c) the surface integrability constraint. We combine these three model ingredients so as to estimate both the shape and refractive index simultaneously from polarisation images acquired from a single viewpoint.

The work presented here assumes minimal knowledge of the illumination conditions, the material under study or the object surface orientation. In fact, our method does not assume either a known illumination direction or a known illuminant power spectrum. This is because the Fresnel transmission ratio can be computed from the maximal and minimal values of the Transmitted Radiance Sinusoid, without any knowledge of the illumination power spectrum or the illumination direction. We can then further relate the Fresnel ratio obtained to the zenith angle and the refractive index. Therefore, our method is different from shape from shading methods that it does not infer

shading information from the angle between the surface normal and the illumination direction.

To the best of our knowledge, this is the first time a method for estimating both shape and refractive index from spectro-polarimetric imagery has been proposed in the Computer Vision literature. Refractive index is an important surface characteristic which can be used to quantify subsurface structure. Our study makes this information available without the need for carefully calibrated optical bench measurements. Moreover, we contribute a shape from polarisation method that requires only information from a single viewpoint. More importantly, the method does not assume any knowledge of the illumination conditions and the index of refraction for the surface under study. This is a major advantage of the method over existing approaches which require either accurate measurements of shape or refractive index, or complicated instrumental setups. Here, we propose an optimisation framework that permits the recovery of shape and refractive index simultaneously. To do this, we render the problem well-posed by drawing on the integrability constraint and the physics of both Fresnel reflection theory and material dispersion. This permits the use of an iterative procedure to find an approximately optimal solution to the above optimisation problem. Further, our iterative approach is computationally efficient due to the use of closed-form solutions for the recovery of the zenith angle of surface normals and the refractive index.

The remainder of the paper is organised as follows. In Section 2, we commence by providing the background on the polarisation of electromagnetic waves. This consists of an introduction to the polarisation theory for dielectric surfaces and the properties of diffuse polarisation due to subsurface scattering. The section also introduces the Fresnel transmission ratio and the phase of polarisation concepts which appear throughout the paper. In Section 3, we present our method for the recovery of shape and refractive index from a set of spectro-polarimetric images captured from a single viewpoint. The section begins with a formulation of the polarisation image radiance as a sinusoidal function. This permits the estimation of the azimuth angle of the surface normals from the polarisation phase across the wavelength domain. Subsequently, we formulate the recovery of the zenith angle and the refractive index as an optimisation problem with an integrability constraint over the surface and a dispersion equation for the material refractive index. In Section 4, we show how this optimisation problem can be tackled in an elegant and efficient manner. In Section 5, we illustrate the utility of the method for the purposes of recovering the shape and refractive index from synthetic and real-world imagery. Finally, in Section 6 we draw conclusions and suggest directions for future work.

## 2 Preliminaries

In this section, we provide the physics background to the theory of the polarisation of electromagnetic waves. Based on the wave theory, light propagation induces harmonic vibrations within the transmission medium. As an electromagnetic wave, light is associated with a magnetic field and an electric field that are mutually orthogonal and vibrate perpendicular to the propagation direction. Such electric and magnetic fields can be represented as field vectors parallel to a plane perpendicular to the propagation direction. According to [7, 19], the polarisation of light characterises the pattern of vibration of the electric and magnetic field vectors as light propagates across space with time.

The section introduces a number of concepts relevant to polarisation, diffuse reflection and the Fresnel transmission theory. For mnemonic purposes, in Figure 1, we provide a list of notation used in the paper. These are defined in order of appearance. Here, we use a right-handed reference coordinate system with the origin located at the view point. The positive  $z$ -axis coincides with the propagation direction of the reflected light as observed in the line of sight. The positive  $x$ -axis points towards the right-hand side of the field of view. Within this right-handed reference system, surface normals can be specified by their azimuth (tilt) and zenith (slant) angles with respect to the positive  $x$  and  $z$  axes.

### 2.1 The Polarisation of Light

Figure 2(a) depicts a circularly polarised illumination flux with its electric field vector  $\vec{E}$  rotating in a helicoidal manner about the propagation direction of the reflected light (as observed from the view-point). In the figure, the direction of propagation is along the  $z$ -axis. As shown in Figure 2(a), the electric field vector can be decomposed into two sinusoidal components  $\vec{E}_x$  and  $\vec{E}_y$  vibrating in the  $x - z$  and  $y - z$  planes. In the figure, the two components have equal magnitudes and  $\vec{E}_y$  is shifted by 1/4-wavelength ( $\frac{\pi}{2}$ ) with respect to  $\vec{E}_x$ . As a result, the wavefront of the electric field has a circular projection on the  $x - y$  plane as it propagates with time. Figure 2(b) illustrates the circular rotation of the electric field, while its  $x$  and  $y$  components vibrate in their respective planes.

In general, the polarisation of a light wave can range from being unpolarised, to being completely polarised. When the phase difference between its orthogonal components varies in a random manner with respect to time, the resultant electric field is incoherent or unpolarised, producing an electric field vector rotating isotropically and randomly in the plane perpendicular to the light

propagation direction. On the other hand, completely polarised light is the result of the coherent superposition of two orthogonal harmonic components with a fixed-phase difference with respect to time. In this case, the light wave is elliptically polarised since the wavefront of the electric field traces out an ellipse whose major and minor axes are determined by the phase difference of the harmonic components.

Notation	Description
$I_{\vartheta}(u, \lambda)$	Transmitted irradiance at a pixel $u$ and a wavelength $\lambda$ through a polariser oriented at an angle $\vartheta$
$\phi(u, \lambda)$	Phase of polarisation.
$I_{max}(u, \lambda)$ and $I_{min}(u, \lambda)$	Maximum and minimum transmitted radiance on the Transmitted Radiance Sinusoid (TRS).
$\vec{N}(u)$	Surface normal at pixel $u$ .
$\alpha(u)$ and $\theta(u)$	Azimuth and zenith angle of the surface normal at pixel $u$ .
$\eta(u, \lambda)$	Refractive index at pixel $u$ and wavelength $\lambda$ .
$\vec{E}_{i\parallel}$ and $\vec{E}_{i\perp}$	Parallel and perpendicular components of the incident light.
$I_{i\parallel}$ and $I_{i\perp}$	Irradiance of the parallel and perpendicular components of the incident light
$\vec{E}_{T\parallel}$ and $\vec{E}_{T\perp}$	Parallel and perpendicular components of the transmitted light.
$I_{T\parallel}$ and $I_{T\perp}$	Radiance of the parallel and perpendicular components of the transmitted light.
$F_{\parallel}$ and $F_{\perp}$	Parallel and perpendicular Fresnel reflection coefficients.
$R(u, \lambda)$	Fresnel transmission ratio function (with respect to the zenith angle $\theta(u)$ and the refractive index $\eta(u, \lambda)$ ).
$C_m(u)$	Dispersion coefficients of Sellmeier's equations.
$B_m(u)$ and $D_m(u)$	Dispersion coefficients of Sellmeier's equations.

Figure 1: Notation used throughout the paper.

## 2.2 Polarisation Upon Diffuse Reflection

We now focus on the polarisation caused by diffuse reflection from a dielectric surface. Diffuse polarisation can be caused by a number of physical scattering processes. For instance, it can be the result of multiple scattering from microfacets [22]. On the other hand, for translucent or multi-layered materials, diffuse reflection is attributed to incident light penetrating the surface, scattering within the dielectric body and then refracting back into the transmission medium. Here, we may assume that after penetrating a surface, light is largely depolarised by the internal scattering process within the dielectric body. Since the refraction of scattered light from the dielectric body through the material-air boundary induces polarisation, the following model of diffuse polarisation, which is subject to Snell's law of refraction and the Fresnel reflection theory, applies to translucent or multi-layered materials with subsurface scattering.

Although diffuse polarisation is weaker, and thus, more difficult to measure than specular polarisation, it is more prevalent in uncontrolled lighting conditions. In fact, specular polarisation only occurs where there is strong specular reflection or inter-reflection. The other advantage of diffuse

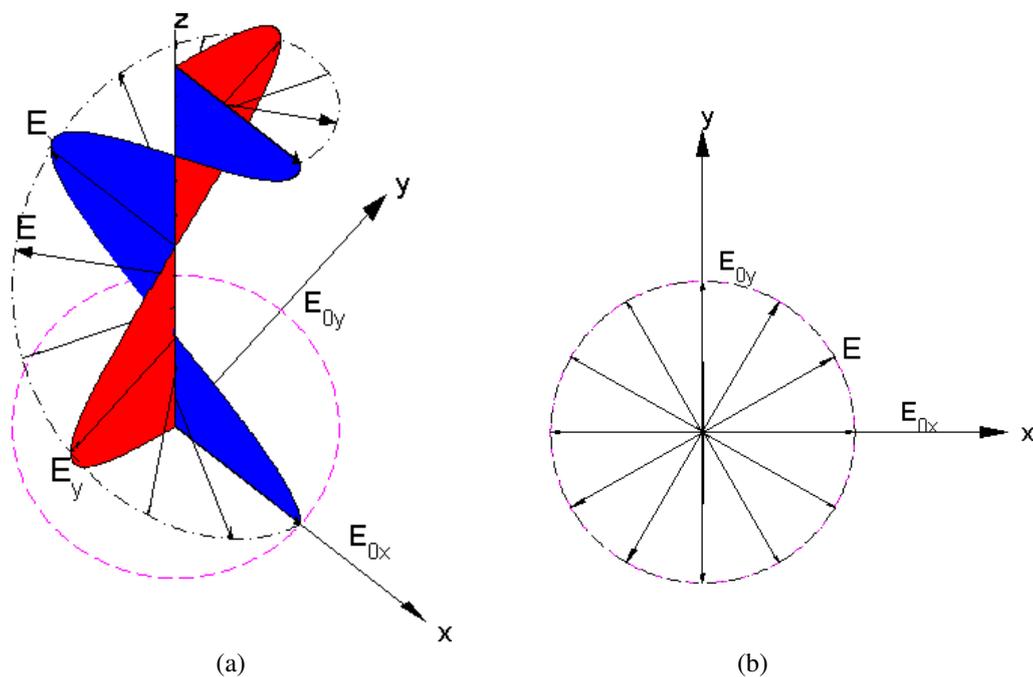


Figure 2: The polarisation of an electromagnetic wave. (a) The polarisation of light is caused by the rotation of its electric field ( $\vec{E}$ ) along its propagation direction  $z$ . The electric field can be decomposed into the two orthogonal harmonic components  $\vec{E}_x$  and  $\vec{E}_y$  with magnitudes  $E_{0x}$  and  $E_{0y}$ . (b) The projection of the trajectory for the front of  $\vec{E}$  as projected onto the  $x - y$  plane perpendicular to the propagation direction.

polarisation is that its degree of polarisation has a one-to-one relationship with the zenith angle, while the degree of specular polarisation corresponds to two candidates of the zenith angle [1]. Our method also applies to incident light with partial or total polarisation since the assumption that light becomes depolarised after penetrating the material surface still remains valid. On the other hand, under polarised incident light, the azimuth angle of the normal cannot be determined from specular polarisation since it alters the plane of polarisation of the reflected light.

Figure 3 shows a diagram of the diffuse reflection process described above. Here, refraction is the result of the change of velocity as the internally scattered ray with propagation direction  $\vec{\mathbf{k}}_i$  travels from the material body with refractive index  $\eta$  to the air whose refractive index is unity. Assume that the incident angle of the ray with the surface boundary is  $\theta_i$ . After emerging from the surface, the transmitted ray  $\vec{\mathbf{k}}_T$  is refracted away from the surface normal  $\vec{\mathbf{N}}$  and forms an emittance (reflection) angle  $\theta$  with respect to the surface normal. The propagation direction of the transmitted ray and the surface normal both lie in the plane of reflection, as shown in Figure 3. Note that  $\theta$  is also the zenith angle of the surface normal with respect to the viewing direction.

Note that the electric field vectors associated with the rays  $\vec{\mathbf{k}}_i$  and  $\vec{\mathbf{k}}_T$  are always perpendicular to their propagation directions. Furthermore, the electric field vectors can be decomposed into two orthogonal harmonic components, *i.e.* one in the plane of reflection while the other one perpendicular to it. In Figure 3, we denote by  $\vec{E}_{i\parallel}$  and  $\vec{E}_{i\perp}$  the parallel and perpendicular components of the electric field associated with  $\vec{\mathbf{k}}_i$ . Similarly,  $\vec{E}_{T\parallel}$  and  $\vec{E}_{T\perp}$  are the parallel and perpendicular components of the transmitted waves travelling in the direction  $\vec{\mathbf{k}}_T$ .

### 2.3 Fresnel Transmission Ratio

We now define two important properties relating to the polarisation of light due to diffuse reflection, namely the Fresnel transmission ratio and the phase of polarisation. To commence, let us denote the radiance of the parallel and perpendicular components emitted at the material surface as  $I_{T\parallel}$  and  $I_{T\perp}$  respectively. The transmission ratios of these components through the material-air interface are related to the Fresnel reflection coefficients as follows

$$\frac{I_{T\perp}}{I_{i\perp}} = 1 - F_{\perp} \quad (1)$$

$$\frac{I_{T\parallel}}{I_{i\parallel}} = 1 - F_{\parallel} \quad (2)$$

where  $I_{i\parallel}$  and  $I_{i\perp}$  are the parallel and perpendicular components of the incident irradiance, respectively and  $F_{\perp}$  and  $F_{\parallel}$  are the Fresnel reflection coefficients for the parallel and perpendicular polarisation components.

Assuming that the light scattered within the dielectric body is unpolarised, we can conclude that its electric field vector oscillates isotropically over all the directions perpendicular to the propagation direction with equal energy and probability. This means that the average magnitudes  $E_{i\parallel}$  and  $E_{i\perp}$  of the perpendicular and parallel components of the internally scattered ray  $\mathbf{k}_i$  are equal, *i.e.*  $\langle E_{i\parallel} \rangle_t = \langle E_{i\perp} \rangle_t$ . As a result,  $I_{i\parallel} = I_{i\perp}$  since the radiance is proportional to the mean square of the

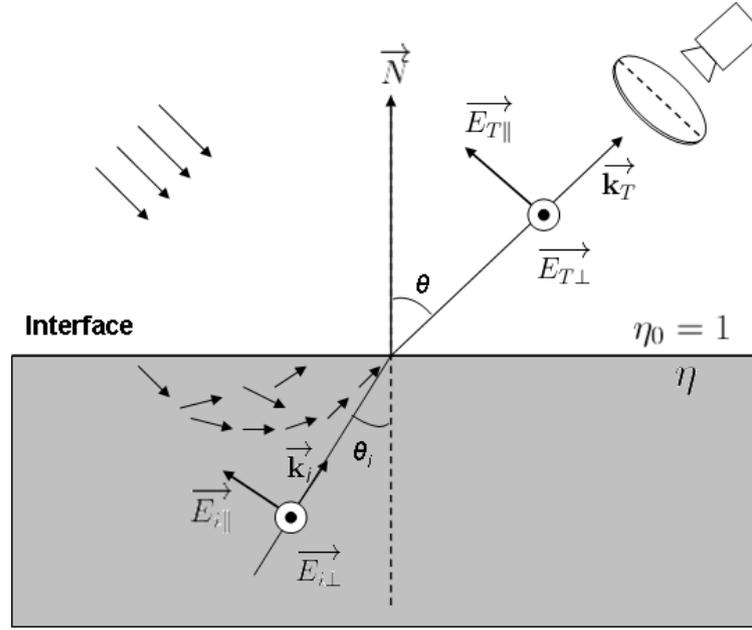


Figure 3: Polarisation resulting from diffuse reflection from a dielectric surface with a normal vector  $\vec{N}$  and refractive index  $\eta$ . The incident light waves penetrate the surface, scattering inside the dielectric body and are finally refracted back into the air through the material-air boundary. The electric fields of the waves incident on and transmitted through the material-air interface are represented by vectors rotating in a plane perpendicular to the propagation direction. These electric field vectors, in turn, can be decomposed into two orthogonal components which are parallel and perpendicular to the plane of emittance, which contains both the surface normal and the emittance direction. The vectors  $\vec{E}_{i\parallel}$  and  $\vec{E}_{i\perp}$  are the components of the internal incident field on the material-air surface boundary before refraction. The vectors  $\vec{E}_{T\parallel}$  and  $\vec{E}_{T\perp}$  are the components of the wave emerging from the surface after refraction.

electric field amplitude, *i.e.*  $I_{i\parallel} = \frac{1}{2}E_{i\parallel}^2$  and  $I_{i\perp} = \frac{1}{2}E_{i\perp}^2$ . Hence

$$\frac{I_{T\perp}}{I_{T\parallel}} = \frac{1 - F_{\perp}}{1 - F_{\parallel}} \quad (3)$$

We refer to the right hand side of Equation 3 as the Fresnel transmission ratio, which is a function of the zenith angle  $\theta$  and the refractive index.

Since the internal light ray  $\vec{\mathbf{k}}_i$  in Figure 3 is scattered randomly (depolarised), its orthogonal components  $E_{i\parallel}^{\vec{}}$  and  $E_{i\perp}^{\vec{}}$  are incoherent, *i.e.* their phase difference varies randomly over time. For dielectric materials, the phase of the polarised light is preserved under refraction [7]. Therefore, the components  $E_{T\parallel}^{\vec{}}$  and  $E_{T\perp}^{\vec{}}$  of the transmitted light are also incoherent after refraction through the material-air boundary.

### 3 Shape and Refractive Index Recovery

Building on the results from the theory of diffuse polarisation presented in Section 2, in this section we develop a method for the recovery of surface orientation and index of refraction from polarisation imagery. The technique discussed here is applicable to convex and continuously twice-differentiable surfaces with material refractive index following the dispersion equation. Further, the method provided assumes an orthographic camera model. In Section 3.1, we commence by extracting the phase angle and the maximal and minimal radiance of the Transmitted Radiance Sinusoid (TRS) from the polarimetric imagery. In Section 3.2.2, we present a method for estimating the azimuth angle of the surface normals from the wavelength-indexed spectrum of phase angles. In addition, we compute the Fresnel transmission ratio from the maximal and minimal radiance, from which to estimate the zenith angle of the surface normal and the refractive index. In Section 3.3, we formulate the estimation as an optimisation problem that takes into account the data error for the Fresnel transmission ratio, the surface integrability and the dispersion of material refractive index across the spectrum. Having formulated the objective function, we present the optimisation process entirely in Section 4.

The steps of shape and refractive index recovery can be summarised as follows.

1. Decomposition of the spectral polarimetric imagery into polarisation components, including the phase, maximal and minimal radiance of the Transmitted Radiance Sinusoid (TRS), as described in Section 3.1.

2. Rough estimation of the azimuth angle from the phase angle for each pixel and wavelength, allowing a 180-degree ambiguity. The relationship between these two angles is described in Section 3.2.
3. Joint estimation of the zenith angle and refractive index from the Fresnel transmission ratio, as presented in Section 3.3.
4. Disambiguation between the two candidates of the azimuth angle using shading information indicated by the estimated zenith angle, as presented in Section 3.2.1.
5. Re-estimation of the azimuth angle as presented in Section 3.2.2.
6. Integration of the surface normal field to reconstruct the surface depth.

It is noted that the 180-degree ambiguity between the candidates of the azimuth angle resulting from the Step 2 does not affect the joint estimation of the zenith angle and refractive index in Step 3. This is because the objective functions involved in this estimation are invariant to the 180-degree shift in the azimuth angle. As we shall show in Section 3.3.3, the relevant objective function consists of the square of the cosine and sine of the azimuth angle. Once the zenith angle has been obtained, we perform disambiguation of the azimuth angle based on the shading information suggested by the estimated zenith angle.

### 3.1 Decomposing Polarisation Images

Using the experimental arrangement shown in Figure 3, the polarisation of light reflected from a surface is measured by mounting a linear polariser in front of the camera optics. By rotating the polariser, we can analyse the polarisation components oriented at different angles on the plane orthogonal to the light propagation direction. The analysis in Section 2.1 can be applied directly to this polarisation imaging system. As a result, the radiance at each point in the scene varies sinusoidally with respect to the angle of rotation of the polariser. This sinusoid is termed the Transmitted Radiance Sinusoid (TRS) and is illustrated in Figure 4.

Let  $\vartheta$  be the polariser orientation angle and  $\phi$  be the phase angle of the sinusoid. Since we use polarimetric imaging spectroscopy data, we let  $\lambda$  denote the wavelength index for the spectra and  $u$  denote the pixel index in the collected images. Moreover, assume that the input comprises  $J$  polarimetric spectral images  $\mathcal{I}_{\vartheta_1}, \mathcal{I}_{\vartheta_2}, \dots, \mathcal{I}_{\vartheta_J}$ . Each of these images is captured at  $J$  discrete

values of the polariser orientation angle  $\vartheta$ , which we denote as  $\vartheta_j, j = \{1, 2, \dots, J\}$ . According to the TRS Sinusoid, the variation of the image intensity at the pixel  $u$ , wavelength  $\lambda$  with polariser orientation  $\vartheta_j$  is given by

$$I_{\vartheta_j}(u, \lambda) = \frac{I_{max} + I_{min}}{2} + \frac{I_{max} - I_{min}}{2} \cos(2\vartheta_j - 2\phi) \quad (4)$$

where  $I_{min}$  and  $I_{max}$  are respectively the minimum and maximum radiance at the pixel  $u$  and  $\phi$  is the corresponding phase angle. Note that here, for the sake of brevity, we have omitted the wavelength, pixel and polarisation angle variables from  $I_{min}$  and  $I_{max}$ .

The recovery of  $I_{min}$ ,  $I_{max}$  and  $\phi$  in Equation 4 may be effected in a number of ways. Note that, by capturing images of 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 [44], where three images are acquired so as to compute the phase, intensity and degree of polarisation. However, the method in [44] is susceptible to noise corruption since it employs only three images. Here, we employ an alternative method that yields a stable estimate of the intensity, phase and degree of polarisation by solving an over-determined linear system of equations. The method is akin to that described by Nayar *et al.* in [28]. We commence by rewriting Equation 4 for each pixel-site  $u$  and

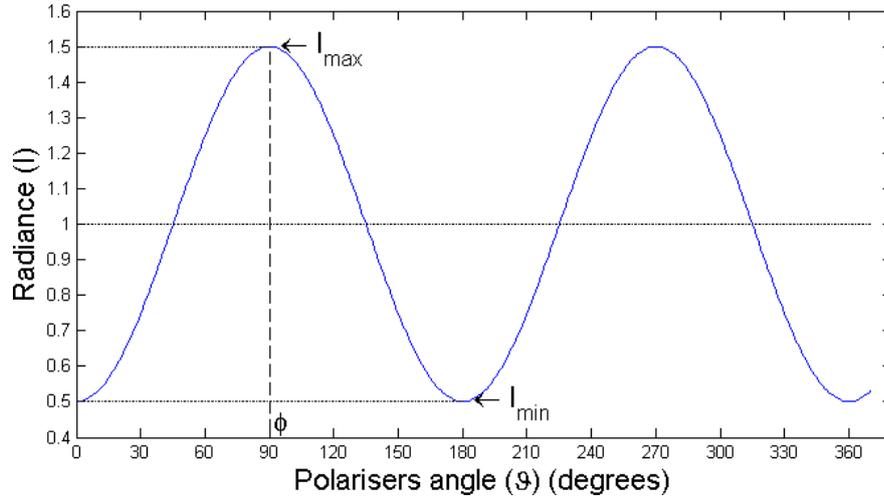


Figure 4: The transmitted radiance through a linear polariser varies sinusoidally with respect to the polariser orientation angle.

wavelength  $\lambda$  in the following vectorial form

$$\begin{aligned} I_{\vartheta_j}(u, \lambda) &= \begin{bmatrix} 1, \cos(2\vartheta_j), \sin(2\vartheta_j) \end{bmatrix} \begin{bmatrix} \frac{I_{max}+I_{min}}{2} \\ \frac{I_{max}-I_{min}}{2} \cos(2\phi) \\ \frac{I_{max}-I_{min}}{2} \sin(2\phi) \end{bmatrix} \\ &= \mathbf{a}_j \mathbf{x} \end{aligned} \quad (5)$$

with  $\mathbf{a}_j = [1, \cos(2\vartheta_j), \sin(2\vartheta_j)]$  and  $\mathbf{x} = \begin{bmatrix} \frac{I_{max}+I_{min}}{2} \\ \frac{I_{max}-I_{min}}{2} \cos(2\phi) \\ \frac{I_{max}-I_{min}}{2} \sin(2\phi) \end{bmatrix}$ .

After collecting  $J \geq 3$  measurements at three or more polariser orientations, we arrive at the following over-determined linear system

$$\mathbb{I} = \mathbb{A} \mathbf{x} \quad (6)$$

where  $\mathbb{I} = [I_{\vartheta_1}(u, \lambda), I_{\vartheta_2}(u, \lambda), \dots, I_{\vartheta_J}(u, \lambda)]^T$  and  $\mathbb{A} = \begin{bmatrix} \mathbf{a}_1 \\ \mathbf{a}_2 \\ \dots \\ \mathbf{a}_J \end{bmatrix}$

Equation 6 is well-constrained when the number of polariser angles is  $J \geq 3$ . Moreover, the coefficient matrix  $\mathbb{A}$  depends solely on the polariser angles and, therefore, allows for an efficient solution of Equation 6 over all the image pixels and wavelengths simultaneously.

Having obtained the solution  $\mathbf{x}^* = [\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3]^T$ , the maximal and minimal radiance on the sinusoid and the phase of polarisation at each pixel  $u$  and wavelength  $\lambda$  can be recovered in a straightforward manner making use of the following relations

$$\begin{aligned} I_{max} &= \mathbf{x}_1 + \sqrt{\mathbf{x}_2^2 + \mathbf{x}_3^2} \\ I_{min} &= \mathbf{x}_1 - \sqrt{\mathbf{x}_2^2 + \mathbf{x}_3^2} \\ \phi &= \frac{1}{2} \arctan \frac{\mathbf{x}_3}{\mathbf{x}_2} \end{aligned}$$

### 3.2 Azimuth Angle and Phase of Diffuse Polarisation

We turn our attention to the estimation of the azimuth angle of the surface normals using the polarisation phase angle  $\phi$ . To determine the relationship between the phase of polarisation  $\phi$  resulting from subsurface scattering and the surface geometry, we consider the geometry of the transmitted electric field observed from the camera's viewpoint.

Figure 5 illustrates a cross-section of the emitted electric field as viewed from the camera's position, where the surface normal  $\vec{N}$  is aligned with the  $y$  axis. In this view,  $\vec{E}_{T\parallel}$  vibrations along the  $y$  axis because it is coplanar with  $\vec{N}$ . The perpendicular component  $\vec{E}_{T\perp}$  lies in the  $x$ -axis as it is orthogonal to  $\vec{E}_{T\parallel}$ . In the Figure, the polariser transmission axis  $\vec{P}$  forms an angle  $\vartheta$  with the horizontal axis. Here, the  $y$ -axis is aligned with the surface normal  $\vec{N}$ . Since the  $\vec{E}_{T\parallel}$  component is in the plane of reflection, it vibrates along the  $y$ -axis. On the other hand, the perpendicular component  $\vec{E}_{T\perp}$  vibrates in the plane spanned by the  $x$ -axis and the line of sight.

Now we consider the polarisation reflectance model developed by Wolff and Boulton [41] with an additional unpolarised component of the light transmitted from the object surface. By including this term, the transmitted irradiance at the polariser angle  $\vartheta$  becomes

$$I_{\vartheta} = \frac{1}{2} (I_{un} + I_{T\perp} \cos^2 \vartheta + I_{T\parallel} \sin^2 \vartheta) \quad (7)$$

where  $I_{T\parallel}$  and  $I_{T\perp}$  are the radiance of the parallel and perpendicular polarised components, and  $I_{un}$  is that of the unpolarised component of the emitted light.

In Equation 7, the radiance of a polarisation component oriented at an arbitrary angle is hence a convex combination of the parallel and perpendicular reflection components. Due to  $F_{\parallel} \leq F_{\perp}$ ,

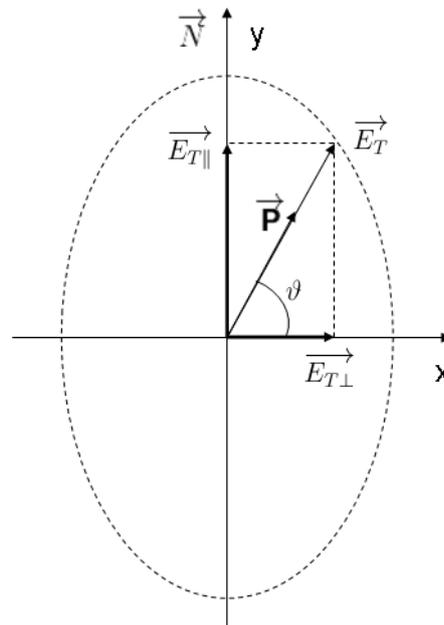


Figure 5: The emitted electric field, where the surface normal  $\vec{N}$  is aligned with the  $y$ -axis.  $\vec{E}_{T\parallel}$  vibrates along the  $y$ -axis because it is in the plane of reflection. The perpendicular component  $\vec{E}_{T\perp}$  lies in the  $x$ -axis and  $\vec{P}$  is the orientation of the polariser transmission axis.

Equation 3 implies that the radiance of the parallel component is always larger than or equal to that of the perpendicular component, *i.e.*  $I_{T\perp} \leq I_{T\parallel}$ . This inequality is derived from the expressions of the Fresnel reflection coefficients as  $F_{\parallel} = F_{\perp} \times \left( \frac{a - \sin \theta_i \tan \theta_i}{a + \sin \theta_i \tan \theta_i} \right)^2$  where  $a = (\eta^2 - \sin^2 \theta_i)^{\frac{1}{2}}$ , with  $\theta_i$  and  $\eta$  being the incident angle and the material refractive index, respectively [43]. Since  $|a - \sin \theta_i \tan \theta_i| \leq |a + \sin \theta_i \tan \theta_i|$ , we obtain  $F_{\parallel} \leq F_{\perp}$ .

As a result, the maximal transmitted radiance occurs when the polariser angle  $\vartheta$  satisfies  $\vartheta = \pm \frac{\pi}{2} (\text{mod } \pi)$ , *i.e.* when the transmission axis lies in the plane of reflection. Moreover, since the plane of reflection contains the surface normal vector, this polariser angle indicates the tilt direction of the surface normal with respect to the viewpoint.

We note the equivalence between this derivation and the Transmitted Radiance Sinusoid in Equation 4. Moreover, the TRS in Figure 4 reaches its maximum at the phase angle. Consequently, the phase angle  $\phi$  must either coincide with the azimuth (tilt) angle  $\alpha$  of the surface normal or differ from it by  $\pi$  radians [1]. In other words, at each pixel  $u$ , either  $\alpha(u) = \phi$  or  $\alpha(u) = \phi \pm \pi$ . This reduces the problem of recovering the azimuth of the surface normal to that of disambiguating the two possible solutions for  $\alpha(u)$ .

### 3.2.1 Disambiguation of the Azimuth Angle

To this end, we assume that the surface under study is convex, *i.e.* the surface normals point in the opposite direction to the gradient of surface shading. This assumption allows us to select the candidate azimuth angle with the closer orientation to the negative gradient direction. The surface shading at a pixel  $u$  is given by  $\cos \theta(u)$ , where  $\theta(u)$  is the zenith angle of the surface normal. Suppose that the zenith angle has been obtained through the process described in Section 4. By sorting the zenith angles available, we are able to detect both the boundary and singular points on the surface and use these to disambiguate the azimuth angles.

The disambiguation process commences by first rotating the azimuth angles by  $180^\circ$  where necessary to satisfy the convexity requirement. In addition, for convex objects, the surface normals at the occluding boundary always point outward from the body of the object. We make use of this observation to initialise the azimuth angles at the occluding boundary. Also, since diffuse polarisation occurs most strongly at the occlusion boundary, the disambiguation of azimuth angle based on polarisation performs most reliably at these positions. On the other hand, weak polarisation is often observed at singular points, *i.e.* where the surface normal is perpendicular to the image

plane. Therefore, abrupt changes in azimuth angle are permitted near these locations. As a result, we start smoothing the azimuth angle at the boundary pixels and propagate the operation toward the singular points of the surface. The smoothing operation proceeds iteratively, aiming to preserve the directional similarity of the azimuth angle within a neighbourhood.

Note that the disambiguation method above only copes with convex surfaces. To handle shapes consisting of a mixture of convex and concave parts, we can adopt an approach similar to that proposed by Zhu and Shi [48]. With the estimated zenith angles, we can detect singular points as those with a zero zenith angle. The method in [48] disambiguates the convexity/concavity by computing the optimal heights of the singular points so as to satisfy a global integrability constraint. To encode this constraint, a configuration graph is constructed with vertices corresponding to the singular points and edges indicating the direction of height differences. With this configuration graph, the above constraint can be stated as a Max-cut problem. Subsequently, the Max-cut problem can be converted into a semi-definite programming problem by relaxing the discrete values of the edges to the real domain. The optimal heights of the singular points can then be found by means of semi-definite programming. Having obtained the height at singular points, the height of the whole surface is propagated from these points via the application of fast marching and patch stitching.

### 3.2.2 Estimation of the Azimuth Angle

In general, the disambiguation process above provides an estimation of the azimuth angle  $\phi^*(u, \lambda)$  at each pixel  $u$  and wavelength  $\lambda$  of the spectral imagery. However, the estimated phase angle for each wavelength may vary widely due to weak polarisation and noise corruption. Hence, we make use of the weighted mean of the spectral phase of polarisation across the spectrum as an alternative to the phase angle at any given wavelength. The weights are 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 using the  $\ell^2$ -norm of the residual  $\epsilon(u, \lambda)$  for the Equation 6, *i.e.*  $\epsilon(u, \lambda) = \|\mathbb{I} - \mathbb{A}\mathbf{x}\|_2$ , where  $\mathbb{I}$ ,  $\mathbb{A}$  and  $\mathbf{x}$  have been computed per pixel and wavelength as per the previous section. The weight  $w(u, \lambda)$  associated with the phase of polarisation at each wavelength is defined via a kernel weighting function. Here we employ the Epanechnikov kernel, which is a popular choice in the

parameter estimation literature. The kernel is given by

$$w(u, \lambda) = \begin{cases} 1 - \frac{\epsilon^2(u, \lambda)}{h} & \text{if } \frac{\epsilon^2(u, \lambda)}{h} < 1 \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

where  $h$  is a bandwidth parameter.

Since the azimuth angle is a directional quantity, instead of averaging the disambiguated azimuth angles  $\phi^*(u, \lambda)$  directly, we estimate the mean of the sines and cosines of these angles for each pixel-site as follows

$$\begin{aligned} \langle \sin(\phi^*(u)) \rangle_\lambda &= \frac{\sum_\lambda \sin(\phi^*(u, \lambda)) w(u, \lambda)}{\sum_\lambda w(u, \lambda)} \\ \langle \cos(\phi^*(u)) \rangle_\lambda &= \frac{\sum_\lambda \cos(\phi^*(u, \lambda)) w(u, \lambda)}{\sum_\lambda w(u, \lambda)} \end{aligned}$$

where  $\langle \cdot \rangle_\lambda$  denotes the mean value across wavelengths. Thus, the estimated azimuth angle at pixel  $u$  becomes

$$\alpha^*(u) = \begin{cases} \arctan\left(\frac{\langle \sin(\phi^*(u)) \rangle_\lambda}{\langle \cos(\phi^*(u)) \rangle_\lambda}\right) & \text{if } \langle \cos(\phi^*(u)) \rangle_\lambda > 0 \\ \arctan\left(\frac{\langle \sin(\phi^*(u)) \rangle_\lambda}{\langle \cos(\phi^*(u)) \rangle_\lambda}\right) + \pi & \text{if } \langle \cos(\phi^*(u)) \rangle_\lambda < 0 \\ \frac{\pi}{2} & \text{if } \langle \cos(\phi^*(u)) \rangle_\lambda = 0 \end{cases} \quad (9)$$

### 3.3 Zenith Angle and Refractive Index

Now we aim to estimate the zenith angle of the surface normals with respect to the viewing direction. Following the previous section, where we showed how the azimuth angle of surface normals can be estimated from the phase of polarisation, we now provide an account of the joint estimation of the zenith angle of surface normals and material refractive index from diffuse polarisation.

Previously, in Section 3.2.2, we concluded that for diffuse polarisation, the TRS curve reaches the minimum  $I_{min}$  when the transmission axis is perpendicular to the plane of reflection. A similar result can be concluded for the maximal radiance  $I_{max}$ . As a result, we have the equalities  $I_{max} = I_{T\parallel}$  and  $I_{min} = I_{T\perp}$  and obtain the ratio of the recovered minimal and maximal radiance on the TRS curve as

$$\begin{aligned} \frac{I_{min}}{I_{max}} &= \frac{I_{T\perp}}{I_{T\parallel}} \\ &= \frac{1 - F_\perp}{1 - F_\parallel} \end{aligned} \quad (10)$$

Following the expressions of the Fresnel reflection coefficients in [43] and Snell's law of refraction, we can relate the ratio on the right-hand side of Equation 10 to the zenith angle and refractive

index through the equation

$$\frac{I_{min}}{I_{max}} = \left( \frac{\cos \theta(u) \sqrt{\eta^2(u, \lambda) - \sin^2 \theta(u)} + \sin^2 \theta(u)}{\eta(u, \lambda)} \right)^2 \quad (11)$$

The right-hand side of Equation 11 is the Fresnel transmission ratio. To simplify computation, we let  $R(u, \lambda) = \sqrt{\frac{I_{min}}{I_{max}}}$ , i.e.

$$R(u, \lambda) \triangleq \frac{\cos \theta(u) \sqrt{\eta^2(u, \lambda) - \sin^2 \theta(u)} + \sin^2 \theta(u)}{\eta(u, \lambda)} \quad (12)$$

In prior literature, the zenith angle of surface normals and the index of refraction can be recovered from the degree of polarisation [1, 2, 5, 25, 26, 39]. Indeed, all of these methods make use of the Fresnel reflection theory to arrive at an equation similar to Equation 11. However, the main limitation to their practical application resides on their reliance upon either known index of refraction [1, 2, 5, 25, 26], imagery captured from multiple viewpoints [25] or under several known light source directions [39]. The need for multiple measurements and instrumental setups makes them impractical for shape and material analysis on real-world data.

The estimation of both the zenith angle and the index of refraction cannot be performed without additional constraints. This is because the Fresnel theory only provides a single equation per wavelength relating the zenith angle and index of refraction to the transmission ratio  $\frac{I_{min}}{I_{max}}$ . As a result, the number of variables to be estimated is one more than the number of equations, rendering the problem ill-posed.

To deal with these limitations, we propose two additional constraints in order to recover the zenith angle and the refractive index in a well-posed manner. To do this, we make use of both the surface integrability constraint on the spatial domain and the material dispersion equations. While the integrability constraint enforces spatial consistency between neighbouring surface locations, the dispersion constraint aims to resolve the ill-posedness of the joint estimation of the zenith angle and refractive index.

### 3.3.1 Integrability Constraint

We commence by formulating the integrability constraint with respect to the zenith and azimuth angles of the surface normal. Assume that the surface under study can be represented by a two-dimensional twice-differentiable function with a continuous second derivative [12]. As a result, its cross derivatives are the same irrespective of the order of the differentiated variable.

Using the reference coordinate system previously defined, let the surface height function at the pixel  $u$  be  $\mathcal{Z}(u)$ . By definition, the normalised surface normal at the pixel  $u$  is given by

$$\vec{N} = \frac{1}{\sqrt{\mathcal{Z}_x^2 + \mathcal{Z}_y^2 + 1}}[-\mathcal{Z}_x, -\mathcal{Z}_y, 1]^T \quad (13)$$

where  $\mathcal{Z}_x$  and  $\mathcal{Z}_y$  are the surface gradients in the  $x$  and  $y$  direction of our coordinate system.

The normalised surface normal direction can also be represented with respect to the azimuth and zenith angles as follows

$$\vec{N} = \begin{bmatrix} \cos \alpha(u) \sin \theta(u) \\ \sin \alpha(u) \sin \theta(u) \\ \cos \theta(u) \end{bmatrix} \quad (14)$$

From Equations 13 and 14, we have

$$\mathcal{Z}_x = -\cos \alpha(u) \tan \theta(u) \quad (15)$$

$$\mathcal{Z}_y = -\sin \alpha(u) \tan \theta(u) \quad (16)$$

Recall that, in Section 3.2, we have obtained an estimation of the azimuth angle  $\alpha(u)$  up to an ambiguity of 180 degree and treat it as a known value in Equations 15 and 16. As a consequence, the cross derivatives can be rewritten as  $\mathcal{Z}_{xy} = -\cos \alpha(u) \frac{\partial \tan \theta(u)}{\partial y}$  and  $\mathcal{Z}_{yx} = -\sin \alpha(u) \frac{\partial \tan \theta(u)}{\partial x}$ . Since the integrability constraint enforces that  $\mathcal{Z}_{xy} = \mathcal{Z}_{yx}$ , we can express it in terms of the zenith and azimuth angles as

$$\cos \alpha(u) \frac{\partial \tan \theta(u)}{\partial y} = \sin \alpha(u) \frac{\partial \tan \theta(u)}{\partial x} \quad (17)$$

According to the chain rule,  $\partial \tan \theta(u) = \frac{\partial \theta(u)}{\cos^2 \theta(u)}$ . Therefore, the integrability constraint in Equation 17 can be rewritten as

$$\cos \alpha(u) \theta_y(u) = \sin \alpha(u) \theta_x(u) \quad (18)$$

with  $\theta_x(u)$  and  $\theta_y(u)$  being the derivatives of  $\theta(u)$  with respect to  $x$  and  $y$ .

### 3.3.2 Material Dispersion

In order to impose further constraints on the variation of the refractive index in the wavelength domain, we note that, for a wide variety of materials, the index of refraction is governed by the material dispersion equation [7]. The material dispersion equation models the wavelength-dependence of the refractive index using a small number of parameters. The number of these parameters is generally lower than the number of spectral bands. Hence, the implication is that the number of

variables to be estimated in the spectral domain can be significantly reduced, especially for spectral images with a high spectral resolution.

Thus, we introduce dispersion equations as a constraint on the refractive index for our optimisation scheme. Among several approximating functions of material dispersion in the physics literature, perhaps Cauchy's and Sellmeier's [7] are the most popular. In the former, Cauchy modelled the empirical relationship between the refractive index of a material and the wavelength of light as follows

$$\eta(u, \lambda) = \sum_{m=1}^M C_m(u) \lambda^{-2(m-1)} \quad (19)$$

where the index of refraction depends solely on the wavelength and the material-specific dispersion coefficients  $C_m(u)$ ,  $m \in \{1, \dots, M\}$ .

In addition, Sellmeier's dispersion equation [36] can handle anomalous dispersive regions by including additional coefficients to represent vacuum wavelengths, *i.e.* where the wave front moves across vacuum, and holds for a wide range of wavelengths, including the ultraviolet, visible and infrared spectrum. Sellmeier's dispersion equation is given by

$$\eta^2(u, \lambda) = 1 + \sum_{m=1}^M \frac{B_m(u) \lambda^2}{\lambda^2 - D_m(u)} \quad (20)$$

where  $B_m(u)$  and  $D_m(u)$  are the material-specific dispersion coefficients.

The dispersion equations above allow a representation of the index of refraction as a linear combination of  $M$  rational functions of wavelength. With these representations, the estimation of refractive index can be treated as that of computing the dispersion coefficients. In practice, an expansion containing up to the sixth term is sufficient to represent a wide range of materials including crystals, liquids, glasses, gases and plastics [21]. For spectral imagery comprising more than seven wavelength-indexed bands, the number of equations relating the Fresnel transmission ratio to the zenith angle and refractive index exceeds the number of dispersion coefficients, rendering the problem solvable. As a result, the problem becomes a well-constrained one that can be formulated in a minimisation setting.

### 3.3.3 Objective Function

Having introduced the integrability and material dispersion constraints in Sections 3.3.1 and 3.3.2, we now focus on the formulation of an objective function for the estimation of the zenith angle and refractive index. The rationale behind our cost function lies in the use of the additional constraints,

including integrability and dispersion equations so as to allow the recovery of the shape and index of refraction to be performed without prior knowledge or predetermined illumination conditions.

The cost function aims at satisfying Equation 11, which equates the square root of the Fresnel transmission ratio defined in Equation 12 to the quantity  $\sqrt{\frac{I_{min}}{I_{max}}}$  while taking into account the integrability of the surface and the material dispersion equation. Our objective function is given by two terms. The first of these accounts for the error of the Fresnel transmission ratio  $R(u, \lambda)$  in Equation 12 with respect to the ratio  $r(u, \lambda) \triangleq \sqrt{\frac{I_{min}}{I_{max}}}$  as computed from the image radiance. The second term measures the error of the integrability constraint described in Equation 18. Thus, the cost function is given by

$$\begin{aligned} \mathcal{E} = & \int_{\mathcal{S}} \int_{\mathcal{W}} (R(u, \lambda) - r(u, \lambda))^2 d\lambda du \\ & + \beta(u) \int_{\mathcal{S}} (\cos \alpha(u)\theta_y(u) - \sin \alpha(u)\theta_x(u))^2 du \end{aligned} \quad (21)$$

subject to the chosen dispersion equation, *i.e.* Equation 19 or 20, where  $\mathcal{S}$  is the image spatial domain and  $\mathcal{W}$  is the wavelength range.

In Equation 21, we assume to have obtained an estimation of the azimuth angle  $\alpha(u)$  up to a 180-degree ambiguity, as described in Section 3.2, and treat it as a constant in the cost function. We note that this cost function is invariant to the 180-degree shift in the azimuth angle, *i.e.*  $(\cos \alpha(u)\theta_y(u) - \sin \alpha(u)\theta_x(u))^2 = (\cos(\alpha(u) + \pi)\theta_y(u) - \sin(\alpha(u) + \pi)\theta_x(u))^2$ . Therefore, we can utilise the rough estimate of the azimuth angle obtained in Section 3.2 without an adverse effect in the estimation of the zenith angle and the refractive index.

In addition, the Fresnel transmission ratio  $R(u, \lambda)$  is related to the zenith angle and refractive through Equation 12. The regularisation parameter  $\beta(u)$  is spatially varying and weighs the relative importance between the data closeness and surface smoothness imposed by the integrability constraint. Here, we use the spatial dependence of  $\beta(u)$  on the surface location so as to reflect the reliability of the azimuth angle  $\alpha(u)$  estimated from polarisation information. To quantify the reliability of the estimate  $\alpha(u)$ , we employ the degree of polarisation as follows

$$\beta(u) = \gamma \left\langle \frac{I_{max}(u, \lambda) - I_{min}(u, \lambda)}{I_{max}(u, \lambda) + I_{min}(u, \lambda)} \right\rangle_{\lambda} \quad (22)$$

where  $\gamma$  is a scaling constant that applies over the whole image and  $\langle \cdot \rangle_{\lambda}$  denotes the mean value over wavelength.

## 4 Optimal Zenith Angle and Refractive Index

The joint estimation of the zenith angle and refractive index amounts to solving a non-linear optimisation problem with a large number of unknowns. Given a spectro-polarimetric image with  $N$  pixels and  $K$  wavelengths, where  $N \geq 100000$ ,  $K = 21$ , we often employ  $M = 5$  dispersion coefficients to model real-world material refractive indices. Consequently, the number of unknown zenith angles and dispersion coefficients is  $N + NM \geq 600000$ . Due to the size of this problem, the use of global optimisation techniques is computationally impractical. On the other hand, direct application of local optimisation methods such as gradient descent, conjugate gradient, Levenberg-Marquardt or the Nelder-Mead simplex method to optimise all the variables at the same time can easily be trapped in a local minimum. Therefore, we iteratively alternate between the estimation of the zenith angle and the refractive index until the estimates are stable. This approach effectively reduces the complexity of each optimisation step in terms of the number of variables. Further, the approach does not require the tuning of optimisation parameters often associated with numerical optimisation methods.

We adopt an iterative approach to the recovery of both the zenith angle and the index of refraction. Algorithm 1 consists of the step sequence of the minimisation strategy. In the first step, we derive a closed-form solution for the zenith angle  $\{\theta(u)|u \in \mathcal{S}\}$  in each iteration while keeping the index of refraction fixed. The second step recovers the index of refraction  $\{\eta(u, \lambda)|u \in \mathcal{S}, \lambda \in \mathcal{W}\}$  using the current estimate of the zenith angle. We iterate between these two steps until the change in the estimates between two successive iterations is below a small threshold.

Our algorithm is an instance of the coordinate search method described by Nocedal and Wright [29]. The general approach of this method is that it partitions the original set of parameters into subsets and performs optimisation over each subset at a time. Generally, this method can iterate infinitely without reaching a critical point. However, some variants of this method have been proven to have a global convergence property. One particular strategy is the “back-and-forth” approach, where the order of variables in which to perform the coordinate search is reversed in successive iterations. With this strategy, the algorithm is proven to converge to an optimum, as discussed in [29]. Since we divide the variables involved in the recovery of the zenith angle and refractive index into two subsets, our strategy of alternating the optimisation between them conforms to the “back-and-forth” approach and would converge to a local optimum.

The algorithm is initialised with a uniform material refractive index  $\eta_0$  across both the spatial

---

**Algorithm 1** Estimation of the zenith angle and refractive index from a polarimetric spectral image with  $K$  wavelength-indexed bands.

---

**Require:** Fresnel transmission ratio  $r(u, \lambda)$  for each pixel  $u$  and band  $\lambda \in \{\lambda_1, \dots, \lambda_K\}$ .

**Ensure:**  $\{\theta(u), \eta(u, \lambda) | u \in \mathcal{S}, \lambda \in \mathcal{W}\}$ , where

$\theta(u)$ : the zenith angle at the pixel  $u$ .

$\eta(u, \lambda)$ : the refractive index at pixel  $u$  and wavelength  $\lambda$ .

1:  $\eta(u, \lambda) \leftarrow \eta_0 \forall u \in \mathcal{S}, \lambda \in \mathcal{W}$

2: **while true do**

3:    $\theta_{old}(u) \leftarrow \theta(u)$

4:    $\eta_{old}(u, \lambda) \leftarrow \eta(u, \lambda)$

5:    $\theta(u) \leftarrow \operatorname{argmin}_{\theta(u)} \mathcal{E}$

6:    $\eta(u, \lambda) \leftarrow \operatorname{argmin}_{\eta(u, \lambda)} \mathcal{E}$

7:   **if**  $|\theta_{old}(u) - \theta(u)| < \tau_\theta$  and  $|\eta_{old}(u, \lambda) - \eta(u, \lambda)| < \tau_\eta, \forall u$  and  $\lambda$  **then**

8:     **break**

9:   **end if**

10: **end while**

11: **return**  $\theta(u), \eta(u, \lambda)$

---

and spectral domains, as indicated in Line 1 of Algorithm 1. It terminates once the estimated parameter values stabilise, *i.e.* the change between the estimates obtained at successive iterations falls below a pre-determined threshold  $\tau_\theta$  for the zenith angle and a threshold  $\tau_\eta$  for the refractive index. For both, the zenith angle and refractive index, the change is measured as the  $\ell^1$ -norm of the difference between successive estimates. In the following two subsections, we elaborate further on the details of the optimisation steps above.

## 4.1 Recovery of the Zenith Angle

Here we derive a solution for the zenith angle while keeping the refractive index fixed in each iteration. We note that the original cost functional in Equation 21 is complex, involving a non-linear Fresnel transmission ratio function  $R(u, \lambda)$  with respect to the zenith angle  $\theta(u)$ . To minimise this cost functional, we need to solve a highly non-linear Euler-Lagrange equation whose analytical solution cannot be derived in a straightforward manner. To this end, we opt for an equivalent, yet simpler formulation of the cost functional, which expresses the data error term using the inverse

function of the Fresnel transmission ratio in Equation 12. Eventually, we reformulate the cost functional with a quadratic function of the zenith angle  $\theta(u)$  in the data error term, whose minimum can be derived analytically.

With the refractive index  $\eta(u, \lambda)$  fixed, we can invert the Fresnel transmission ratio function in Equation 12 with respect to the zenith angle  $\theta(u)$  as follows

$$\sin \theta(u) \equiv \frac{\eta(u, \lambda) \sqrt{1 - R^2(u, \lambda)}}{\sqrt{\eta^2(u, \lambda) - 2R(u, \lambda)\eta(u, \lambda) + 1}} \quad (23)$$

Note that the zenith angle is a geometric variable independent of the wavelength. However, the above equation provides for wavelength-dependent estimates of the zenith angle. In practice, these estimates may not be the same across the spectrum due to measurement error and noise corruption. If the index of refraction is at hand, the value of  $r(u, \lambda)$  computed from the ratio of maximal and minimal image irradiance can be used as an estimate for the function  $R(u, \lambda)$ . This yields a wavelength-dependent estimate  $\varphi(u, \lambda)$  for the zenith angle  $\theta(u)$ , which is given by

$$\varphi(u, \lambda) = \arcsin \left( \frac{\eta(u, \lambda) \sqrt{1 - r^2(u, \lambda)}}{\sqrt{\eta^2(u, \lambda) - 2r(u, \lambda)\eta(u, \lambda) + 1}} \right) \quad (24)$$

Note that in Equation 24, we use the notation  $\varphi(u, \lambda)$  for the wavelength-dependent estimate and distinguish it from the wavelength-independent zenith angle  $\theta(u)$ . We take advantage of this wavelength dependency and, instead of directly minimising the original cost functional in Equation 21, we seek to recover a zenith angle close to the wavelength-dependent estimates in Equation 24 while satisfying the integrability constraint. Thus, we employ the alternative cost functional

$$\mathcal{E}_1 = \int_S \int_{\mathcal{W}} (\theta(u) - \varphi(u, \lambda))^2 d\lambda du + \beta(u) \int_S \left( \cos \alpha(u) \frac{\partial \theta(u)}{\partial y} - \sin \alpha(u) \frac{\partial \theta(u)}{\partial x} \right)^2 du \quad (25)$$

as an alternative to that in Equation 21.

Equation 25 poses the minimisation problem in a simpler setting. The merit of the alternative cost function is that the Fresnel ratio error is quantified as a quadratic form of  $\theta(u)$ . This is important since this quadratic form is more tractable than the original error term, which contains a rational function in the expression for  $R(u, \lambda)$ . Moreover, we can rewrite Equation 25 as follows

$$\mathcal{E}_1 = \int_S f(u, \theta(u), \theta_x(u), \theta_y(u)) du \quad (26)$$

by letting  $f(\cdot)$  be given by

$$f(u, \theta(u), \theta_x(u), \theta_y(u)) \triangleq \int_{\mathcal{W}} (\theta(u) - \varphi(u, \lambda))^2 d\lambda + \beta(u) \left( \cos \alpha(u) \theta_y(u) - \sin \alpha(u) \theta_x(u) \right)^2 \quad (27)$$

The function  $f(\cdot)$  is important since it permits the use of calculus of variations to recover the minimiser of the functional in Equation 26. We do this by noting that the minima must satisfy the following Euler-Lagrange equation

$$\frac{\partial f}{\partial \theta} = \frac{\partial}{\partial x} \left( \frac{\partial f}{\partial \theta_x} \right) + \frac{\partial}{\partial y} \left( \frac{\partial f}{\partial \theta_y} \right) \quad (28)$$

By computing the derivatives of  $f$  so as to satisfy the Euler-Lagrange equation above, we arrive at the following differential equation

$$\theta(u) \int_{\mathcal{W}} d\lambda - \int_{\mathcal{W}} \varphi(u, \lambda) d\lambda = \beta(u) \times \left( \sin^2 \alpha(u) \theta_{xx}(u) - \sin 2\alpha(u) \theta_{xy}(u) + \cos^2 \alpha(u) \theta_{yy}(u) \right) \quad (29)$$

where  $\theta_{xx}(u)$ ,  $\theta_{yy}(u)$  and  $\theta_{xy}(u)$  are the second order and covariant derivatives of  $\theta(u)$  with respect to the  $x$  and  $y$  axes of the coordinate system.

In the discrete case, where the imagery is acquired at  $K$  wavelength-indexed bands  $\lambda \in \{\lambda_1, \dots, \lambda_K\}$ , we have  $\int_{\mathcal{W}} d\lambda = K$ . Therefore,  $\theta(u)$  satisfies the differential equation

$$\theta(u) = \frac{1}{K} \int_{\mathcal{W}} \varphi(u, \lambda) d\lambda + \frac{\beta(u)}{K} \left( \sin^2 \alpha(u) \theta_{xx}(u) - \sin 2\alpha(u) \theta_{xy}(u) + \cos^2 \alpha(u) \theta_{yy}(u) \right) \quad (30)$$

We note that Equation 30 is a second-order partial differential equation with respect to  $\theta(u)$ . We further enforce the continuity and differentiability of the spatial domain by assuming that the function  $\theta(u)$  is continuously twice-differentiable, *i.e.*  $\theta_{xy}(u) = \theta_{yx}(u)$ . This assumption permits the decomposition of  $\theta(u)$  into an orthogonal basis of integrable two-dimensional functions, in a similar manner to that in [12]. Since digital images have a limited band of spatial frequencies, the surface shading can be expressed as a finite linear combinations of the real part of Fourier basis functions, which are cosine functions. This representation allows an analytical solution to the functional minimisation problem above. Moreover, we will show later that this representation also leads to a computationally efficient solution to Equation 30.

Note that the function  $\theta(u)$  can be viewed as a discrete function on a two-dimensional lattice. Let the size of the lattice be  $X \times Y$ , where  $X$  and  $Y$  are the image width and height, respectively. Based on the Nyquist–Shannon sampling theorem [37], the zenith angle  $\theta(u)$  can be reconstructed using frequency components of up to one-half of the sampling frequency of the image. In image processing, these sampling frequencies  $\nu$ , where  $\nu = (\nu_x, \nu_y)$ , are often chosen such that  $\nu_x = \frac{2\pi i}{X}$ , where  $i = 0, 1, \dots, X-1$ , and  $\nu_y = \frac{2\pi j}{Y}$ , where  $j = 0, 1, \dots, Y-1$  [13]. With these two-dimensional frequency components, the function  $\theta(u)$  can be reconstructed with an orthogonal set of Fourier

basis functions  $e^{\mathbf{i}(u^T \mathbf{v})} = e^{\mathbf{i}(u_x x_u + u_y y_u)}$ , where  $\mathbf{i}$  is the imaginary number and  $u = (x_u, y_u)$  is the pixel location. Formally, this is given by

$$\theta(u) = \sum_v \kappa_v e^{\mathbf{i}(u^T \mathbf{v})} \quad (31)$$

Intuitively, Equation 31 means that the shading of the surface can be decomposed into a linear combination of Fourier components with a range of frequencies matching that of the input imagery. In the equation,  $\kappa_v$  is the coefficient (weight) of the Fourier basis function  $e^{\mathbf{i}(u^T \mathbf{v})}$ , which can be computed making use of the equation

$$\kappa_v = \frac{1}{|\mathcal{S}|} \sum_u \theta(u) e^{-\mathbf{i}(u^T \mathbf{v})}$$

where  $|\mathcal{S}|$  represents the number of image pixels.

Similarly, the partial derivatives of  $\theta(u)$  can also be expressed in terms of the Fourier basis, as follows

$$\theta_{xx}(u) = - \sum_v \kappa_v v_x^2 e^{\mathbf{i}(u^T \mathbf{v})} \quad (32)$$

$$\theta_{xy}(u) = - \sum_v \kappa_v v_x v_y e^{\mathbf{i}(u^T \mathbf{v})} \quad (33)$$

$$\theta_{yy}(u) = - \sum_v \kappa_v v_y^2 e^{\mathbf{i}(u^T \mathbf{v})} \quad (34)$$

Let  $h(u) = \frac{1}{K} \int_{\mathcal{W}} \varphi(u, \lambda) d\lambda$ . By substituting Equations 31, 32, 33 and 34 into Equation 30, we obtain

$$h(u) = \sum_v \kappa_v e^{\mathbf{i}(u^T \mathbf{v})} \left( 1 + \frac{\beta(u)}{K} \left( \sin^2 \alpha(u) v_x^2 - \sin 2\alpha(u) v_x v_y + \cos^2 \alpha(u) v_y^2 \right) \right) \quad (35)$$

Note that Equation 35 applies to every image location  $u$  and every spatial frequency  $v$ . By making use of the expressions for  $h(u)$  at every image location and frequency, we can construct a linear system with respect to the unknown vector  $\mathbb{U} = [\kappa_v]^T$ , which is, effectively, the concatenation of all the Fourier coefficients. The recovery of the coefficients  $\kappa_v$  can be then effected by solving the linear system  $\mathbb{L}\mathbb{U} = \mathbb{H}$ , with  $\mathbb{H} = [h(u)]^T$ , which is the vectorial concatenation of the known function values  $h(u)$  for all the image locations  $u$  and  $\mathbb{L}$  is a matrix with rows and columns indexed to the image pixels and spatial frequencies, respectively. In other words, the matrix element  $\mathbb{L}_{u,v}$  corresponding to a given pixel  $u$  and a given frequency  $v$  is  $\mathbb{L}_{u,v} = e^{\mathbf{i}(u^T \mathbf{v})} \left( 1 + \frac{\beta(u)}{K} \left( \sin^2 \alpha(u) v_x^2 - \sin 2\alpha(u) v_x v_y + \cos^2 \alpha(u) v_y^2 \right) \right)$ . With a chosen Fourier basis and the azimuth angle  $\alpha(u)$  obtained as described in Section 3.2.2, the matrix  $\mathbb{L}$  can be computed in a straightforward manner. With the coefficients  $\kappa_v$  at hand, the zenith angle  $\theta(u)$  can be recovered through the application of Equation 31.

## 4.2 Recovery of the Refractive Index

With the zenith angle in hand, we now turn our attention to the estimation of the refractive index at each image location making use of the Fresnel transmission ratio. To derive the refractive index directly from the Fresnel transmission ratio in Equation 12, we are required to solve a quadratic equation involving the index of refraction  $\eta(u, \lambda)$ . This expression is given by

$$(\cos^2 \theta(u) - r^2(u, \lambda)) \times \eta^2(u, \lambda) + 2r(u, \lambda) \sin^2 \theta(u) \times \eta(u, \lambda) - \sin^2 \theta(u) = 0 \quad (36)$$

In general, the quadratic equation above yields no more than two real-valued roots. The choice of refractive index value depends on the physical plausibility of these roots, *i.e.* the refractive index for dielectrics must be a real value greater than one. This choice is straightforward if only one of the roots is physically plausible.

In the case where the two roots are plausible, we adopt an iterative approach which iteratively selects the root closer to the refractive index average at the same wavelength within the local spatial neighbourhood. This approach works under the assumption that there is a single solution to the refractive index at a number of pixels in the image. Initially, we label these pixels as having their refractive index uniquely determined. At each iteration, we assign the refractive index of those pixels with two plausible solutions making use of the regions whose refractive index is already determined. We do this by selecting the root which is in better accordance with the average of the refractive indices within its spatial neighbourhood. A pseudocode of this iterative procedure is illustrated in Algorithm 2.

When neither root is a physically plausible solution to Equation 36, we use an approximation of the Fresnel transmission ratio that provides a single solution to the refractive index. To this end, we consider an approximating function that is a product of two separable factors, one of which depends solely on the refractive index, while the second one depends on the zenith angle. By adopting a formulation similar to Schlick's approximation of the Fresnel reflection coefficient [35], we employ the expression

$$R^*(u, \lambda) = d + c(1 - \cos \theta(u))^b \quad (37)$$

where  $b$ ,  $c$  and  $d$  are constants.

Figure 6(a) shows this graph for values of  $\eta$  ranging between 1.2 and 3, with an increment of 0.3. Empirical observations of this graph give rise to a power function with respect to the zenith angle. At the end-points,  $R = 1$  as  $\theta(u) = 0$  and  $R = \frac{1}{\eta(u, \lambda)}$  as  $\theta(u) = \frac{\pi}{2}$ . To satisfy these conditions,

---

**Algorithm 2** Refractive Index Selection

---

**Require:**  $\eta(u, \lambda)$ : Solutions to the refractive index at pixel  $u$  and wavelength  $\lambda$

**Ensure:**  $\eta^*(u, \lambda)$ : The uniquely determined refractive index at pixel  $u$  and wavelength  $\lambda$

```
1: for all wavelength  $\lambda$  do
2:   for all pixel  $u$  with a single physically plausible refractive index at wavelength  $\lambda$  do
3:      $determined(u, \lambda) \leftarrow \mathbf{true}$ 
4:   end for
5:   while there are pixels with two plausible refractive indices do
6:     for all pixel  $u$  with two solutions  $\eta_1(u, \lambda)$  and  $\eta_2(u, \lambda)$  do
7:        $\mathcal{N}(u) \leftarrow$  the spatial neighbourhood of  $u$ 
8:        $\bar{\eta}(u, \lambda) \leftarrow \text{Average}_{v \in \mathcal{N}(u), determined(v, \lambda) = \mathbf{true}} \eta(v, \lambda)$ 
9:        $\eta^*(u, \lambda) \leftarrow \eta_i(u, \lambda)$  which is closer to  $\bar{\eta}(u, \lambda)$ 
10:       $determined(u, \lambda) \leftarrow \mathbf{true}$ 
11:     end for
12:   end while
13: end for
14: return  $\eta^*(u, \lambda) \forall$  pixel  $u$  and wavelength  $\lambda$ 
```

---

it is necessary that  $c = \frac{1}{\eta(u, \lambda)} - 1$  and  $d = 1$ . Thus, we have

$$R^*(u, \lambda) = 1 + \left( \frac{1}{\eta(u, \lambda)} - 1 \right) (1 - \cos \theta(u))^b \quad (38)$$

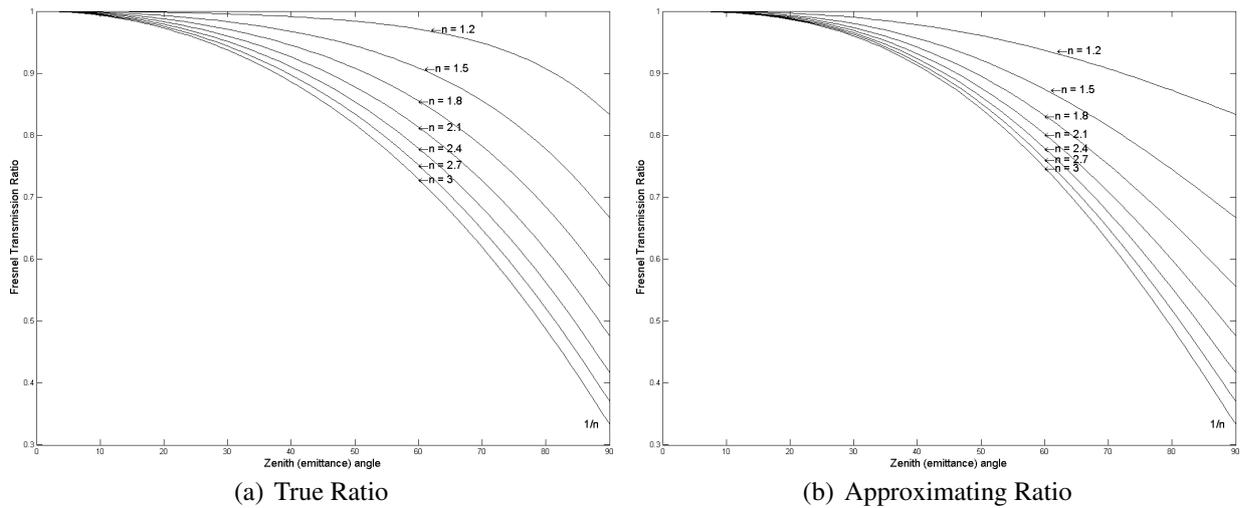


Figure 6: The Fresnel transmission ratio for several refractive indices and zenith angles in the range of  $[0, \frac{\pi}{2}]$ . (a) The true Fresnel transmission ratio function. (b) The approximating function.

In Equation 38,  $b$  is a pre-determined power that provides the best fit with respect to the true Fresnel transmission ratio over a range of parameter values. In this paper, we consider material refractive indices between 1 and 3 and  $\theta(u) \in [0, \frac{\pi}{2}]$ . Using a one-dimensional search for the power  $b$ , we find that  $b = 1.4$  minimises the  $\ell^1$ -error between the approximating and the true Fresnel transmission ratio. In Figure 6(b), we have plotted the approximating Fresnel ratio for  $b = 1.4$  and noted its similarity to the Fresnel ratio function.

To verify the approximation accuracy for the Fresnel transmission ratio, in Figure 7, we plot the error function for  $b = 1.4$ , where it is represented as a surface with respect to both the zenith angle and refractive index. From the Figure, we observe that the absolute error is below 0.04 for all the combinations of zenith angle and refractive index within the considered range. The error is largest near grazing zenith angles ( $70^\circ \leq \theta(u) < 90^\circ$ ) for small refractive indices or at around  $40^\circ$  for refractive indices larger than 2.3. However, these are extreme cases and, generally, the material and the geometry under study do not fall into these cases.

Using the approximated Fresnel ratio function, we arrive at a single approximate solution for the refractive index at each pixel and wavelength. Given the zenith angle  $\theta(u)$  at the current iteration,

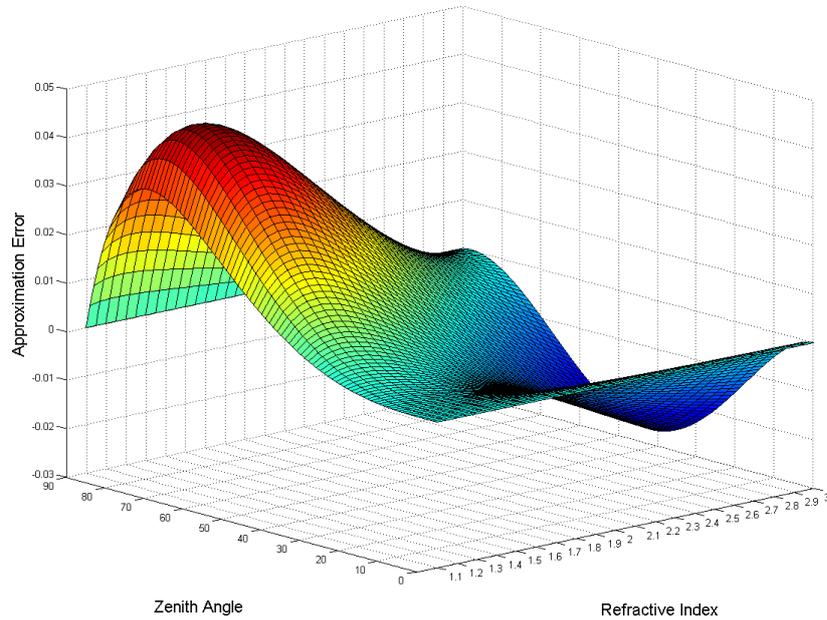


Figure 7: The approximation error for the Fresnel transmission ratio plotted for refractive indices between 1 and 3 and zenith angles in the range of  $[0, \frac{\pi}{2}]$ .

the index of refraction is estimated to be

$$\eta(u, \lambda) = \frac{(1 - \cos \theta(u))^b}{(1 - \cos \theta(u))^b - 1 + r(u, \lambda)} \quad (39)$$

Next, we apply the following strategy to ensure the physical plausibility of the approximating solution in Equation 39. We note that the approximating refractive index is physically plausible, *i.e.*  $\eta(u, \lambda) \geq 1$  if  $(1 - \cos \theta(u))^b - 1 + r(u, \lambda) > 0$  and  $r(u, \lambda) < 1$ . In the case where  $r(u, \lambda) = 1$ , we can conclude that  $\theta(u) = 0$  as can be observed in Figure 6(a). However, in this case, the refractive index value can be arbitrary. Otherwise, when  $r(u, \lambda) < 1$ , *i.e.*  $\theta(u) \neq 0$ , we can guarantee the physical plausibility of the solution in Equation 39 by scaling the zenith angles at all the image pixels such that  $\cos \theta(u) < 1 - (1 - r(u, \lambda))^{\frac{1}{b}}$  for all  $u$  and  $\lambda$ .

Finally, to satisfy the material dispersion constraint, we fit the refractive index spectrum obtained for each surface location to either of the dispersion equations in Equation 19 or 20. For Cauchy’s dispersion equation, coefficients are given by the least-square solution to the following linear system

$$\mathbf{n}(u) = \Lambda \mathbf{C}(u) \quad (40)$$

where  $\mathbf{n}(u) = \begin{bmatrix} \eta(u, \lambda_1) \\ \dots \\ \eta(u, \lambda_K) \end{bmatrix}$ ,  $\mathbf{C}(u) = \begin{bmatrix} C_1(u) \\ \dots \\ C_M(u) \end{bmatrix}$ , and  $\Lambda = \begin{bmatrix} 1 & \lambda_1^{-2} & \dots & \lambda_1^{-2(M-1)} \\ 1 & \lambda_2^{-2} & \dots & \lambda_2^{-2(M-1)} \\ & & \dots & \\ 1 & \lambda_K^{-2} & \dots & \lambda_K^{-2(M-1)} \end{bmatrix}$ . where  $M$  is the

number of dispersion coefficients used.

Since  $\mathbf{n}(u)$  and  $\Lambda$  are known, the system is over-determined if the number of dispersion coefficients is chosen such that  $M \leq K$ , where  $K$  is the number of wavelength indexed bands in the imagery.

For Sellmeier’s equation (Equation 20), the fitting task can be posed as a non-linear least-squares optimisation problem. If the number of dispersion coefficients  $M$  is chosen such that  $M \leq K$  then the non-linear least-squares problem above becomes well-constrained and can be solved numerically by standard line-search or trust-region techniques.

## 5 Experiments

In this section, we perform experiments to illustrate the utility of our method for the purpose of recovering the shape and refractive index of objects in a scene. To this end, we report results

on synthetic and real-world spectro-polarimetric images. The use of synthetic data permits us to effect a quantitative analysis where the accuracy of the estimated parameters can be compared to the ground-truth. On the other hand, the use of real-world images allows us to illustrate the utility of the recovered surface orientation and index of refraction for the purpose of generating views under novel lighting conditions and view points.

For all our experiments, we employed the Cauchy’s dispersion equation as a constraint on the refractive index to estimate the optimal parameters of the cost functional in Equation 21. Here, we consider the eighth-order Cauchy’s dispersion equation consisting of five terms. The fitting of refractive index spectra to this dispersion equation was performed via constrained linear least-squares [9].

## 5.1 Synthetic Data

We commence by performing experiments on a synthetic dataset. Our dataset comprises spectro-polarimetric images of eight 3D surfaces rendered with the Wolff diffuse reflectance model for dielectrics [43]. To generate our dataset, we render these synthetic surfaces with the refractive index of the 19 plastic and liquid materials reported in [21], and the spectral reflectance of Polyvinyl Chloride (PVC) under five illumination directions. These amount to a total of 1520 combinations of shape and photometric parameters. For each combination, five polarisation images were generated corresponding to five polariser orientations at  $0^\circ$ ,  $30^\circ$ ,  $45^\circ$ ,  $60^\circ$  and  $90^\circ$  in the anti-clockwise direction with respect to the horizontal axis of the rendering context, *i.e.* the image coordinate system. All the multispectral images in this dataset are 30 bands in length, spanning the  $430 - 720nm$  range, with a  $10nm$  step between successive bands. In this experiment, we have set the bandwidth  $h$  of the Epanechnikov kernel in Section 3.2.2 to unity and the scaling constant  $\gamma$  of the integrability term in Equation 22 to 0.05.

The different illumination conditions simulated in our dataset are described as follows. Each of the illuminants assumes the power spectrum of sunlight, which we acquired using a StellarNet spectrometer. Each illuminant direction is in the horizontal ground plane that contain the viewing direction and points towards the surface under study. We have used five light source directions, which we denoted  $L_1$ ,  $L_2$ ,  $L_3$ ,  $L_4$  and  $L_5$ . These form angles of  $-26.5^\circ$ ,  $-14^\circ$ ,  $0^\circ$ ,  $14^\circ$  and  $26.5^\circ$  with respect to the viewing direction, where a negative angle means the light direction is on the left-hand side with respect to the view point and a positive angle means otherwise.

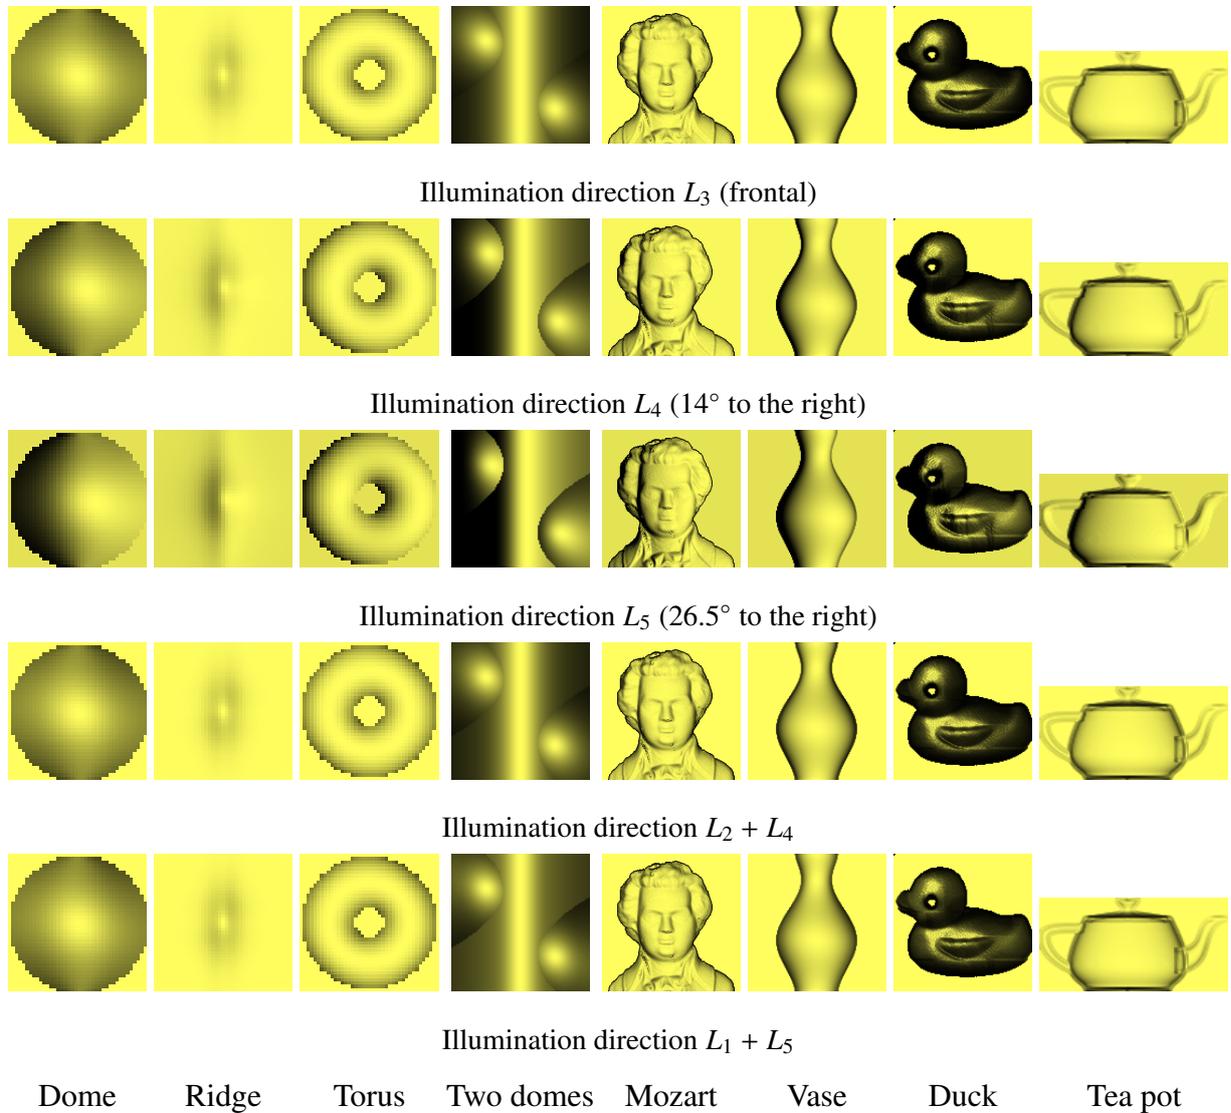


Figure 8: Synthetic images of our test shapes rendered with the refractive index of Polystyrene under five different illumination conditions. The rows, from top to bottom, correspond to the combinations of illumination directions  $L_3$ ,  $L_4$ ,  $L_5$ ,  $L_2 + L_4$  and  $L_1 + L_5$ . The images are rendered in pseudo trichromatic colours synthesized from the multispectral radiance of the polarisation component oriented at  $45^\circ$  with respect to the horizontal axis of the image coordinate system.

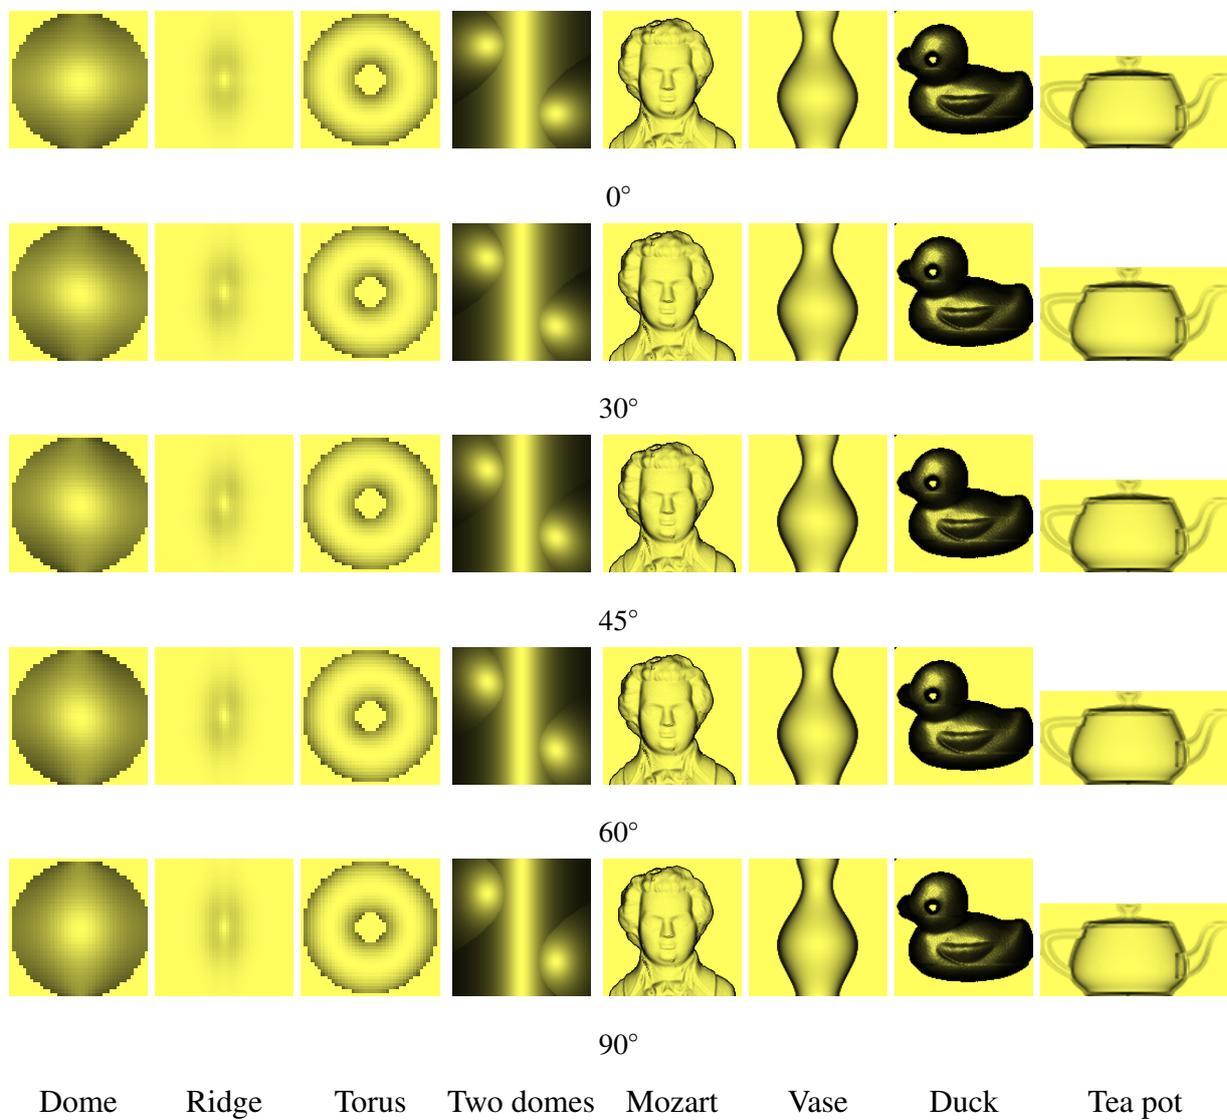


Figure 9: Pseudo trichromatic rendering of synthetic spectral-polarimetric images of eight shapes synthesized with the refractive index of Polystyrene under the frontal light source direction. Each row, from top to bottom, shows the polarisation component oriented at  $0^\circ$ ,  $30^\circ$ ,  $45^\circ$ ,  $60^\circ$ ,  $90^\circ$ .

In Figure 8, we show the variation of shading with respect to the illumination direction. In the Figure, we show the shapes in our datasets. The first three rows include synthetic images under three instances of single illumination directions, which are  $L_3$ ,  $L_4$  and  $L_5$ , and the last two are instances of two simultaneous illuminant directions, which are  $L_2 + L_4$  and  $L_1 + L_5$ . The images are rendered in pseudo trichromatic colours synthesized from the multispectral radiance of the  $45^\circ$  polarisation component, using the Stiles and Burch colour matching functions [38].

Similarly, Figure 9 shows the variation of shading with respect to the angle of polarisation. Each row, from top to bottom, respectively shows the polarisation component oriented at  $0^\circ$ ,  $30^\circ$ ,  $45^\circ$ ,  $60^\circ$ ,  $90^\circ$ . The images of the above components have been synthesized with the refractive index of Polystyrene under the frontal illumination direction. As before, we have employed the Stiles and Burch colour matching functions [38] in order to generate the pseudo trichromatic colours.

In Figure 10, we present the needle maps estimated by our method for the synthetic shapes under the frontal light direction  $L_3$ , and the oblique light directions originating from the right-hand side of the viewing position, at angles of  $14^\circ$  ( $L_4$ ) and  $26.5^\circ$  ( $L_5$ ). The top row shows the ground-truth needle maps corresponding to the surface normals of the synthetic shapes. The second, fourth and sixth rows respectively show the needle-maps recovered from the images rendered under the illuminant directions  $L_3$ ,  $L_4$  and  $L_5$ . In addition, we show the corresponding angular error of the estimated normal fields for the above illumination conditions in the third, fifth and seventh rows. The error maps in these rows are depicted in grayscale values indicating the dot products between the normalised estimated surface normals and the ground truth.

We note that, for the Dome, Ridge, Torus, Two-Domes, Vase and Tea pot, the surface normal orientation has been successfully recovered almost everywhere except for regions with strong shadowing. This observation applies to the shapes above regardless of the illumination direction and shading variations. The low estimation errors for the surface normals of these shapes can be verified visually by inspecting the error maps in the third, fifth and seventh rows. The strategy of smoothing the azimuth angles we described in Section 3.2.1 tends to be effective on convex surfaces, but breaks down at non-convex parts for shapes consisting of a mixture of concave and convex surfaces, such as the Mozart and the Duck. Here, the horizontal symmetry of the recovered surface normals for a number of shapes, such as the Dome, Ridge, Torus, Two-Domes and Vase, indicates that our estimation method relies on polarisation rather than shading to reveal the surface geometry. The results also imply that our method is insensitive to changes in the illuminant power spectrum and direction.

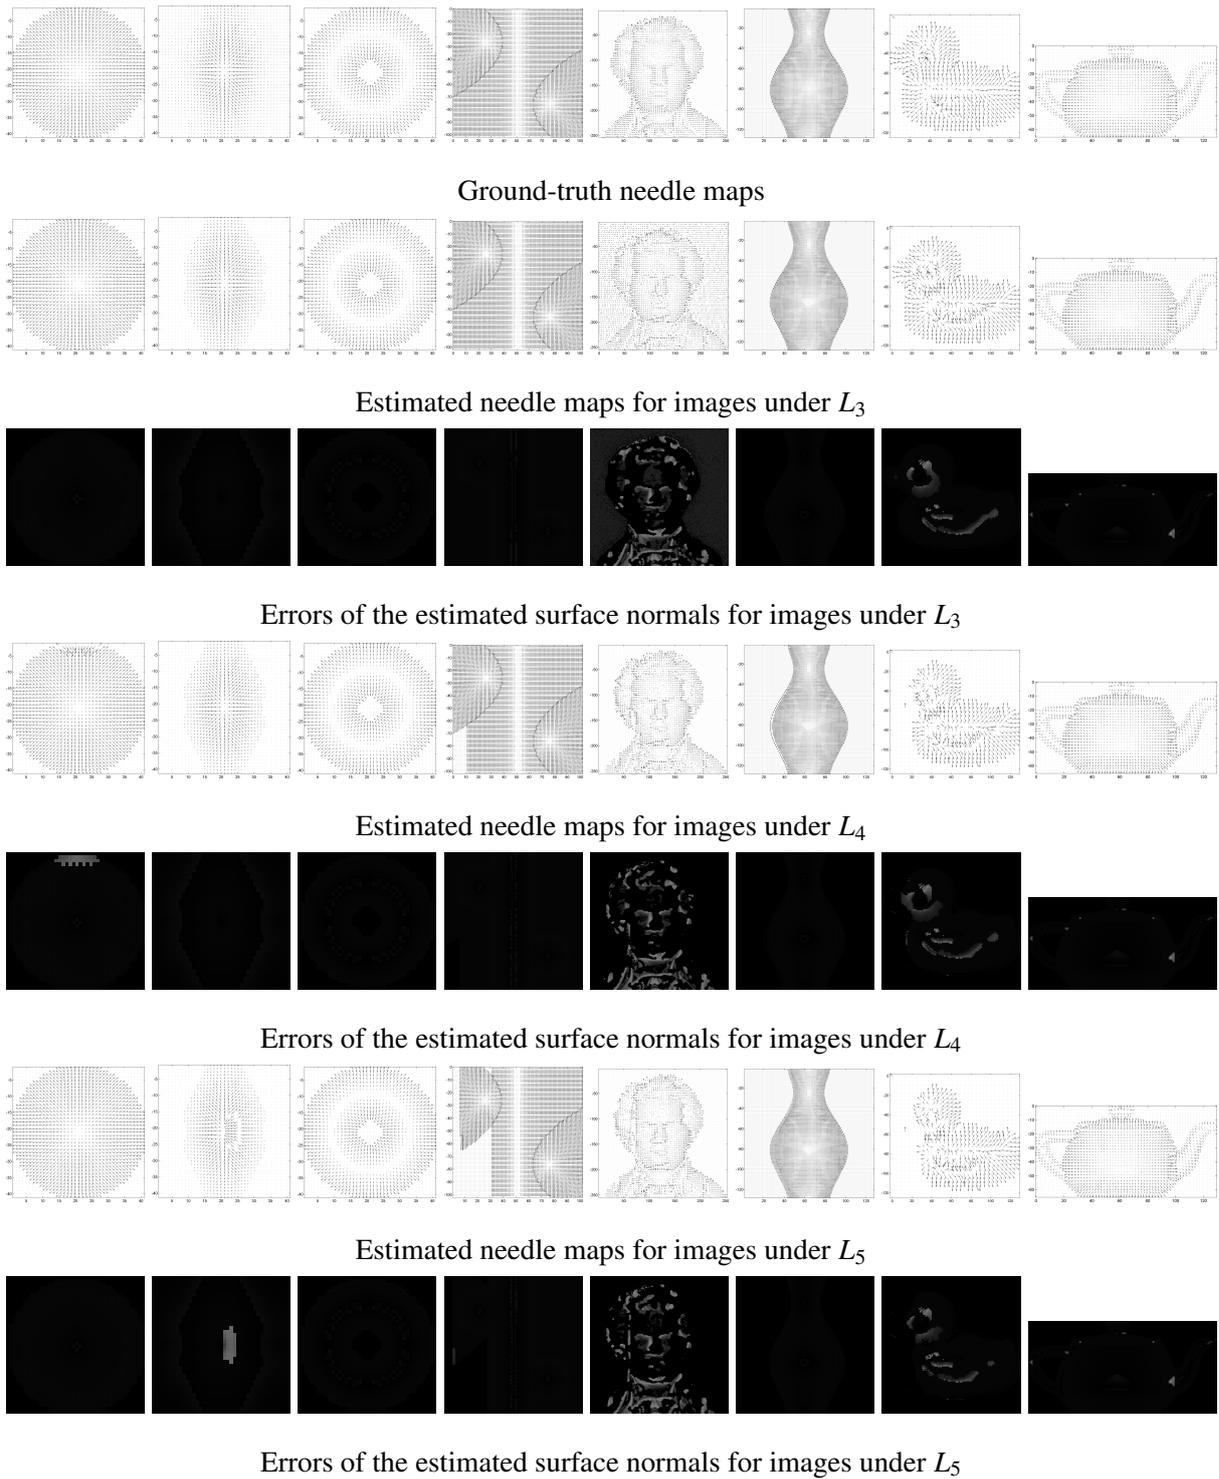


Figure 10: The estimated needle maps of eight different shapes as compared to the corresponding groundtruth. First row: ground truth needle maps. Second, fourth and sixth rows: The needle maps recovered from images illuminated under the illumination directions  $L_3$ ,  $L_4$  and  $L_5$ , respectively. Third, fifth and seventh rows: The error of the estimated needle maps for the illumination directions  $L_3$ ,  $L_4$  and  $L_5$ , respectively.

With surface normal estimates in hand, we reconstruct the surface depth by means of surface normal field integration. To this end, we make use of the surface integration method introduced by Frankot and Chellappa [12] for the purpose of surface reconstruction. In the first row of Figure 11, we present the surface rendering for each of the shapes in the dataset. The subsequent rows show rendering of the recovered surfaces when the input surfaces are illuminated from several different illumination directions, including  $L_3$ ,  $L_4$  and  $L_5$ . Each of the surfaces is rendered under the same novel lighting direction. As shown, the reconstruction of most of the shapes except the Mozart bust and the toy Duck is almost perfect with respect to the groundtruth. This result is consistent with the qualitative results in Figure 10. Here, we also note that strong shadows caused by the oblique illumination directions  $L_4$  and  $L_5$  do not convey any information, and therefore, do not assist surface reconstruction within these regions. The effect of the dark regions on the results of reconstruction can be observed in the rendering of the Ridge, the Two-Domed shape and the toy Duck in the third and fourth rows. However, the reconstruction of the non-shadowed parts still exhibits the curvature of the original shapes.

We now perform a quantitative set of experiments to measure the accuracy of our method. In Table 1, we show the accuracy of the recovered normal map, which is quantified as the Euclidean angular difference, in degrees, between the estimated surface normal direction and the corresponding ground truth, averaged over the pixels of each image studied. In columns 2–6, we report the mean and standard deviation of the angular error for each combination of shape and illuminant condition. Similarly, in Table 2, we report the surface height reconstruction error as the mean absolute difference between the estimated and the ground-truth depth, averaged over the image pixels. For the purpose of comparison, we have normalised both the reconstructed and ground-truth depth maps so that the height varies in the range between 0 and 1.

The numerical results in Table 1 show that the surface normals were recovered at a high level of accuracy, being lower or equal to 3.4 degrees for most of the shapes, except for the Mozart bust and the toy Duck. The error outliers for the Mozart and Duck shapes are understandable since they are more complex than the other shapes, being composed of both convex and concave regions. Overall, the quantitative results are largely consistent with the normal maps and error maps in Figure 10.

Likewise, Table 2 shows negligible surface reconstruction error for most of the shapes except for the Mozart bust and the toy Duck. Of these shapes, the highest error is observed for the Ridge, at *i.e.* approximately 3%–6% of the maximal depth value. As before, the higher depth reconstruction error for the Mozart bust and the toy Duck is due to the complexity introduced by the combination

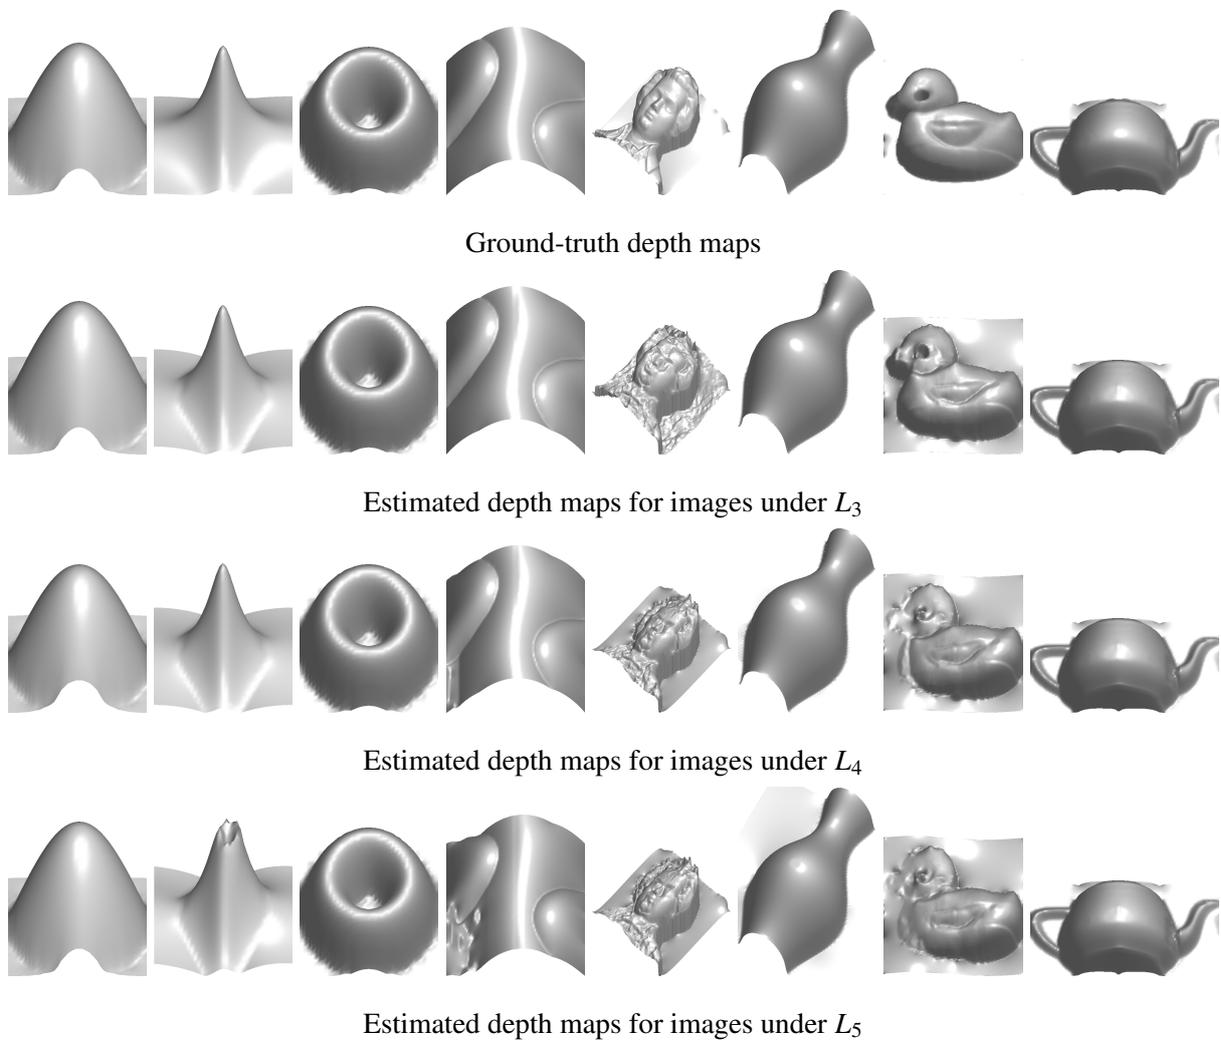


Figure 11: The estimated surface depths of eight different shapes as compared to the corresponding groundtruth. First row: The ground truth depth. From the second to the fourth row: The depth maps recovered from images illuminated under the illumination directions  $L_3$ ,  $L_4$  and  $L_5$ , respectively.

	Surface normal error (degrees)				
	$L_3$	$L_4$	$L_5$	$L_2 + L_4$	$L_1 + L_5$
Dome	2.9311 ± 2.5470	3.0990 ± 2.6147	2.9042 ± 2.5147	2.9311 ± 2.5470	2.9312 ± 2.5470
Ridge	2.3102 ± 0.5731	2.3102 ± 0.5731	3.5714 ± 0.4945	2.3102 ± 0.5731	2.3102 ± 0.5731
Torus	2.4112 ± 1.6977	2.4178 ± 1.7084	2.4369 ± 1.6915	2.4144 ± 1.7056	2.4118 ± 1.7036
Test shape	2.9270 ± 2.3648	3.4837 ± 2.1826	5.9983 ± 1.8518	2.9271 ± 2.3649	2.9271 ± 2.3648
Mozart	14.8964 ± 1.1588	15.1172 ± 1.1909	15.3956 ± 1.1347	14.9276 ± 1.1671	14.9181 ± 1.0614
Vase	2.5177 ± 2.1209	2.7918 ± 2.0990	3.6414 ± 1.9563	2.5177 ± 2.1209	2.5177 ± 2.1209
Duck	6.5573 ± 2.3130	11.2253 ± 1.7184	12.9258 ± 1.7877	6.5547 ± 2.3268	6.1266 ± 2.3304
Teapot	2.6371 ± 1.7283	2.6371 ± 1.7283	2.7051 ± 1.7134	2.6371 ± 1.7283	2.6371 ± 1.7283

Table 1: The accuracy of the recovered surface normals for several illumination directions computed as the average absolute angular difference (in degrees) between the estimated surface normal directions and the ground truth.

	Normalised depth error				
	$L_3$	$L_4$	$L_5$	$L_2 + L_4$	$L_1 + L_5$
Dome	0.0043 ± 0.0056	0.0073 ± 0.0101	0.0038 ± 0.0033	0.0043 ± 0.0056	0.0043 ± 0.0056
Ridge	0.0321 ± 0.0035	0.0321 ± 0.0035	0.0626 ± 0.0093	0.0321 ± 0.0035	0.0321 ± 0.0035
Torus	0.0049 ± 0.0026	0.0048 ± 0.0026	0.0057 ± 0.0028	0.0048 ± 0.0025	0.0049 ± 0.0026
Test shape	0.0074 ± 0.0056	0.0455 ± 0.0173	0.0910 ± 0.0160	0.0074 ± 0.0056	0.0074 ± 0.0056
Mozart	0.0927 ± 0.0399	0.0849 ± 0.0216	0.1005 ± 0.0205	0.0910 ± 0.0402	0.0903 ± 0.0400
Vase	0.0101 ± 0.0078	0.0146 ± 0.0078	0.0413 ± 0.0060	0.0101 ± 0.0078	0.0101 ± 0.0078
Duck	0.0940 ± 0.0349	0.1531 ± 0.0216	0.1509 ± 0.0213	0.0946 ± 0.0348	0.0919 ± 0.0345
Teapot	0.0122 ± 0.0026	0.0122 ± 0.0026	0.0126 ± 0.0031	0.0122 ± 0.0026	0.0122 ± 0.0026

Table 2: The accuracy of surface depth reconstruction for several illumination directions. The depth error is computed as the mean absolute difference between the estimated and the ground-truth depth, with both the depth maps being normalised to the range [0, 1].

	Refractive index error (degrees)				
	$L_3$	$L_4$	$L_5$	$L_2 + L_4$	$L_1 + L_5$
Dome	$0.0443 \pm 0.0583$	$0.0443 \pm 0.0583$	$0.0444 \pm 0.0584$	$0.0443 \pm 0.0583$	$0.0443 \pm 0.0583$
Ridge	$0.0488 \pm 0.0640$	$0.0488 \pm 0.0640$	$0.0488 \pm 0.0640$	$0.0488 \pm 0.0640$	$0.0488 \pm 0.0640$
Torus	$0.0470 \pm 0.0618$	$0.0470 \pm 0.0618$	$0.0470 \pm 0.0618$	$0.0470 \pm 0.0618$	$0.0470 \pm 0.0618$
Test shape	$0.0503 \pm 0.0610$	$0.0494 \pm 0.0590$	$0.0529 \pm 0.0549$	$0.0503 \pm 0.0610$	$0.0503 \pm 0.0610$
Mozart	$0.0458 \pm 0.0593$	$0.0460 \pm 0.0582$	$0.0458 \pm 0.0566$	$0.0458 \pm 0.0593$	$0.0458 \pm 0.0593$
Vase	$0.0474 \pm 0.0601$	$0.0473 \pm 0.0600$	$0.0467 \pm 0.0559$	$0.0474 \pm 0.0601$	$0.0474 \pm 0.0601$
Duck	$0.0528 \pm 0.0475$	$0.0628 \pm 0.0578$	$0.0687 \pm 0.0632$	$0.0528 \pm 0.0475$	$0.0528 \pm 0.0475$
Teapot	$0.0481 \pm 0.0621$	$0.0481 \pm 0.0621$	$0.0479 \pm 0.0612$	$0.0481 \pm 0.0621$	$0.0481 \pm 0.0621$

Table 3: Means and standard deviations of the angular difference (in degrees) between the estimated refractive index spectra and the ground truth, across all the reported materials.

of convex and concave regions. Further, the depth error for the Duck can be partly attributed to shadowed regions in the input images when synthesized under the oblique light source directions ( $L_3$  and  $L_4$ ).

We note that the errors in Tables 1 and 2 often vary only slightly with respect to illuminant direction. In some cases, the errors for images under the frontal and the combined illuminant directions ( $L_2 + L_4$  and  $L_1 + L_5$ ) are lower than those for images under the oblique illuminant directions. The increase in error for the oblique lighting directions is partly due to the shadowed image regions in the input imagery. In fact, the variation of the surface normal error is within one degree and that of the depth error is within 0.03 when the illuminant shifts from the frontal direction to the most oblique direction (at an angle of  $26.5^\circ$  from the camera axis). Moreover, the standard deviations of the surface normal error and depth error in Tables 1 and 2 are negligible, peaking at 1.8689 degrees and 0.0402, respectively. This means that the recovered shape is substantially insensitive to the material refractive indices used in our experiment. These observations, again, support the claim that polarisation is a good cue to surface orientation because it is robust to changes in illumination direction and material refractive index. As such, this differentiates Shape from Polarisation from both Shape from Shading and Photometric Stereo methods, which attribute geometric cues to image shading.

Next, we turn our attention to the accuracy of the recovered refractive index. In Table 3, we

report the Euclidean angular difference, in degrees, between the average estimated refractive index spectra over all the pixels in an image and the ground truth refractive index used for synthesizing that image. The angular error measure is meaningful because our method aims at recovering a scalar multiple of the actual refractive index spectra. Furthermore, for the purpose of recognition, the spectral variation of refractive index is of higher importance than its absolute value. The low angular errors reported in Table 3, peaking at 0.0551 degrees, show that we can recover the refractive index with a high level of accuracy. Moreover, the refractive index estimate is largely consistent across all the illumination directions for each shape. In fact, the change in lighting direction hardly affects the resulting refractive index. This observation is consistent with the shape recovery results, demonstrating that the recovered refractive index is robust against changes in illumination conditions.

Now we turn our attention to evaluating the effectiveness of the surface integrability constraint. In Tables 4, 5 and 6, we present the errors in the surface normals, the normalised surface depth and the refractive index spectra recovered by optimising the cost function in Equation 21 without the integrability constraint. In other words, we set the weight  $\beta(u) = 0$  for all the pixels  $u$  and optimised the cost function per pixel with the refractive index subject to the Cauchy dispersion equation.

As observed in Tables 4, 5 and 6, the mean errors resulted from the alternative optimisation approach are significantly larger than those shown in Tables 1, 2 and 3, which have been produced with the surface integrability taken into account. This trend is consistent across most of the presented shapes and lighting directions. With the integrability constraint added, the improvement in the errors is at least several standard deviations compared to the alternative method. This improvement could be explained by the role of the integrability constraint in enforcing spatial consistency across the image. In the absence of this constraint, the optimisation of the zenith angle and refractive index can only be performed independently for each individual pixel. The comparison above again confirms the importance of the integrability constraint in our framework and justifies the enforcement of this constraint.

Finally, we evaluate the robustness of our method to the initial choice of the refractive index  $\eta_0$  as specified in Algorithm 1. To this end, we have estimated the shape and refractive index from the synthetic dataset with starting values of  $\eta_0$  varying between 1.1 and 1.8 in increments of 0.1. In Figure 12, we show the variation of the surface normal error, the normalised depth error and spectral angle error of the refractive index, averaged over all the synthetic images. From the figure,

	Surface normal error (degrees)				
	$L_3$	$L_4$	$L_5$	$L_2 + L_4$	$L_1 + L_5$
Dome	$4.3168 \pm 1.6151$	$4.3234 \pm 1.5150$	$4.4600 \pm 1.6298$	$4.3197 \pm 1.5134$	$4.3001 \pm 1.5575$
Ridge	$0.6792 \pm 0.4496$	$0.6792 \pm 0.4588$	$1.9701 \pm 0.3950$	$0.6792 \pm 0.4588$	$0.6792 \pm 0.3849$
Torus	$6.0059 \pm 2.7435$	$6.1787 \pm 2.8280$	$8.2153 \pm 3.2483$	$6.0127 \pm 2.7473$	$6.0121 \pm 2.1863$
Test shape	$3.6110 \pm 0.7431$	$3.4253 \pm 0.7302$	$3.1528 \pm 0.6604$	$3.5857 \pm 0.7523$	$3.5757 \pm 0.7709$
Mozart	$17.7184 \pm 1.4330$	$18.2380 \pm 0.9583$	$19.1064 \pm 0.5824$	$17.7863 \pm 1.2256$	$17.6490 \pm 0.6599$
Vase	$2.2294 \pm 0.5856$	$2.2331 \pm 0.5638$	$2.3704 \pm 0.5033$	$2.2296 \pm 0.5826$	$2.2289 \pm 0.4926$
Duck	$27.6443 \pm 2.7751$	$28.0286 \pm 1.9848$	$27.0117 \pm 2.1032$	$27.4962 \pm 2.7953$	$26.1977 \pm 3.1813$
Teapot	$16.1288 \pm 0.6596$	$16.8396 \pm 0.7417$	$19.8454 \pm 0.3913$	$16.0351 \pm 0.8332$	$16.7656 \pm 0.7711$

Table 4: The accuracy of the recovered surface normals for several illumination directions with no integrability constraint enforced. The errors are quantified as the average absolute angular difference (in degrees) between the estimated surface normal directions and the ground truth.

we can observe that the normal and depth errors reach their minima when  $\eta_0$  is between 1.4 and 1.6 and the refractive index error is lowest when  $\eta_0 = 1.5$ . This is also the case for the standard deviations of the error measures above. This is not surprising since the ground truth value for the refractive index of the materials used in our simulation are in the range [1.31, 1.68]. Nonetheless, initial refractive index values outside this range cause the final results to drift away from the true value, the error for the estimated shape and refractive index are low for initial values greater than the ground truth ones. Thus, in practice, although our algorithm may be trapped in local minima, low estimation errors can be achieved by setting  $\eta_0$  to a value corresponding to the upper end of those which are physically plausible for the refractive index.

## 5.2 Real-world Imagery

In this section, we report experimental results on real-world multispectral images acquired in-house. To acquire the imagery, we have used a benchtop hyperspectral camera equipped with an Acousto-Optical Tunable Filter (AOTF) which selects wavelengths in the range of 450 – 650nm at a spectral resolution of 10nm. By applying radio-frequency acoustic waves to compress and relax an optically anisotropic crystal [10, 15, 17], the incoming light is diffracted into an ordinary and an extraordinary component polarised in different directions. In our camera system, the filter only

	Normalised depth error				
	$L_3$	$L_4$	$L_5$	$L_2 + L_4$	$L_1 + L_5$
Dome	$0.0274 \pm 0.0064$	$0.0274 \pm 0.0064$	$0.0300 \pm 0.0077$	$0.0274 \pm 0.0064$	$0.0275 \pm 0.0065$
Ridge	$0.0116 \pm 0.0053$	$0.0116 \pm 0.0052$	$0.0784 \pm 0.0031$	$0.0116 \pm 0.0052$	$0.0116 \pm 0.0053$
Torus	$0.0668 \pm 0.0284$	$0.0654 \pm 0.0277$	$0.1125 \pm 0.0510$	$0.0668 \pm 0.0284$	$0.0668 \pm 0.0286$
Test shape	$0.0356 \pm 0.0005$	$0.0629 \pm 0.0006$	$0.1078 \pm 0.0004$	$0.0352 \pm 0.0005$	$0.0354 \pm 0.0006$
Mozart	$0.1118 \pm 0.0102$	$0.1044 \pm 0.0027$	$0.1151 \pm 0.0045$	$0.1127 \pm 0.0105$	$0.1106 \pm 0.0083$
Vase	$0.0318 \pm 0.0004$	$0.0358 \pm 0.0003$	$0.0450 \pm 0.0004$	$0.0318 \pm 0.0004$	$0.0318 \pm 0.0003$
Duck	$0.1268 \pm 0.0035$	$0.1448 \pm 0.0036$	$0.1492 \pm 0.0040$	$0.1267 \pm 0.0039$	$0.1298 \pm 0.0028$
Teapot	$0.1284 \pm 0.0048$	$0.1324 \pm 0.0084$	$0.1580 \pm 0.0039$	$0.1274 \pm 0.0067$	$0.1329 \pm 0.0074$

Table 5: The accuracy of surface depth reconstruction for several illumination directions with no integrability constraint enforced. The errors are quantified as the mean absolute difference between the estimated and the ground-truth depth, with both the depth maps being normalised to the range between 0 and 1.

	Refractive index error (degrees)				
	$L_3$	$L_4$	$L_5$	$L_2 + L_4$	$L_1 + L_5$
Dome	$0.7962 \pm 0.0788$	$0.7962 \pm 0.0788$	$0.7962 \pm 0.0788$	$0.7962 \pm 0.0788$	$0.7962 \pm 0.0788$
Ridge	$0.1107 \pm 0.0564$	$0.1107 \pm 0.0564$	$0.1107 \pm 0.0564$	$0.1107 \pm 0.0564$	$0.1107 \pm 0.0564$
Torus	$0.2907 \pm 0.0356$	$0.2907 \pm 0.0356$	$0.2907 \pm 0.0356$	$0.2907 \pm 0.0356$	$0.2907 \pm 0.0356$
Test shape	$1.7606 \pm 0.1737$	$1.6583 \pm 0.1644$	$1.4570 \pm 0.1425$	$1.7606 \pm 0.1737$	$1.7606 \pm 0.1737$
Mozart	$0.6979 \pm 0.1290$	$0.5528 \pm 0.0881$	$0.5054 \pm 0.0833$	$0.6979 \pm 0.1290$	$0.6979 \pm 0.1290$
Vase	$0.7782 \pm 0.0657$	$0.7262 \pm 0.0616$	$0.5914 \pm 0.0502$	$0.7782 \pm 0.0657$	$0.7782 \pm 0.0657$
Duck	$4.4762 \pm 0.6824$	$4.1360 \pm 0.5846$	$4.1432 \pm 0.5963$	$4.4762 \pm 0.6824$	$4.4762 \pm 0.6824$
Teapot	$0.4012 \pm 0.0385$	$0.4012 \pm 0.0385$	$0.3904 \pm 0.0380$	$0.4012 \pm 0.0385$	$0.4012 \pm 0.0385$

Table 6: Means and standard deviations of the angular difference (in degrees) between the estimated refractive index spectra and the ground truth, across all the reported materials. These results have been produced with no integrability constraint enforced.

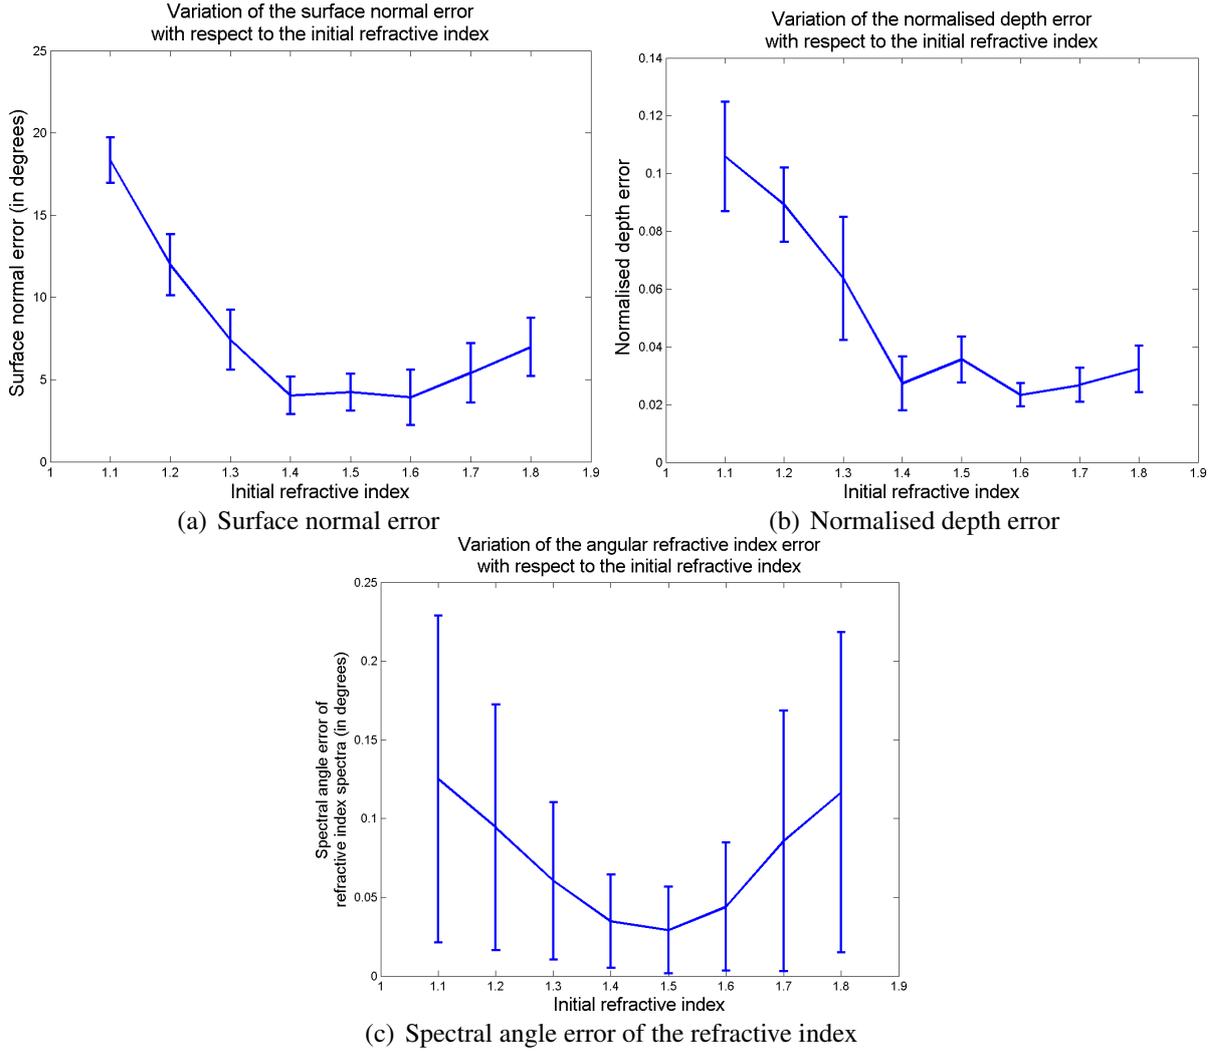


Figure 12: The variation of the mean and standard deviation of the surface normal error, the normalised depth error and spectral angle error of the refractive index with respect to the initial refractive index value when performing our algorithm on the synthetic dataset.

transmits the ordinary ray while blocking the extraordinary ray using total internal reflection. By rotating the camera box about its optical axis, we effectively selected the plane of polarisation of the polarised light at each wavelength.

We acquired multispectral images of five different objects made of matte plastic and porcelain. In our experiment, the object size was significantly smaller than its distance from the camera centre. Therefore, the projection of the object surfaces onto the image plane can be assumed to be orthographic. To measure polarisation, we captured images of these objects when the polariser transmission axis is oriented at each of seven different angles given by  $45^\circ$ ,  $60^\circ$ ,  $75^\circ$ ,  $90^\circ$ ,  $105^\circ$ ,  $120^\circ$  and  $135^\circ$  in the clock-wise direction with respect to the vertical axis of the camera. Theo-

retically, it is sufficient to perform decomposition step in Section 3.1 with only three polarisation angles. However, we have performed experiments on images captured under seven angles to minimise the effect of noise on the shift of the phase angle. The imagery was captured using two unpolarised artificial sunlights as illuminants. These simultaneously illuminated the left and right hand side of the objects under study. Note that the algorithm does not require prior knowledge of the illuminant power spectrum and direction. Here, we used two illuminants to ensure that there were no dark shadows in the captured images. To synthesize the surface, as presented later in the section, we estimated the material reflectance from unpolarised multispectral images of the same objects captured under an illuminant direction aligned to the viewing direction.

Recall that our method delivers the surface orientation and the index of refraction of the surface materials from input polarisation images. To recover the reflectance, we assume that the light penetrating the object sub-surface does not scatter over a significant distance and is therefore emitted at the location of the incident point. Under this assumption, the incidence and reflection angles can be assumed to be identical. This is important since it allows the use of the illuminant power spectrum and the Wolff reflectance model [43] to estimate the material reflectance (albedo), after both the surface orientation and the index of refraction have been recovered by our algorithm.

Figure 13 shows sample input images of the five objects under study, which we name hereafter Bear, Statue, Pig, Dinosaur and Pine Tree. The images in the left-most column show the trichromatic pseudocolour of the input multispectral images, simulated with the Stiles and Burch colour matching functions [38]. These correspond to the images of the objects with the polariser transmission axis at an angle of  $45^\circ$  with respect to the vertical axis of the camera. The second and third columns show the recovered needle-maps and shading maps of the objects in the first column. Note that the needle maps show clear overall surface contours, capturing the surface normal orientation along occlusion boundaries. The shading maps are qualitatively consistent with the shape of the objects, with minor distortions at the boundaries of regions of different material composition. This is due to the fact that our method is based solely on polarisation information. Therefore, the changes at these boundaries have been interpreted by our algorithm as variations in object geometry.

Having obtained the surface orientation and surface material properties of the objects under study, we generate novel views of the original objects under novel illuminants. For this purpose, we again reconstructed the surface height from surface orientation using the integration method introduced by Frankot and Chellappa [12]. For validation, we captured the ground truth images of

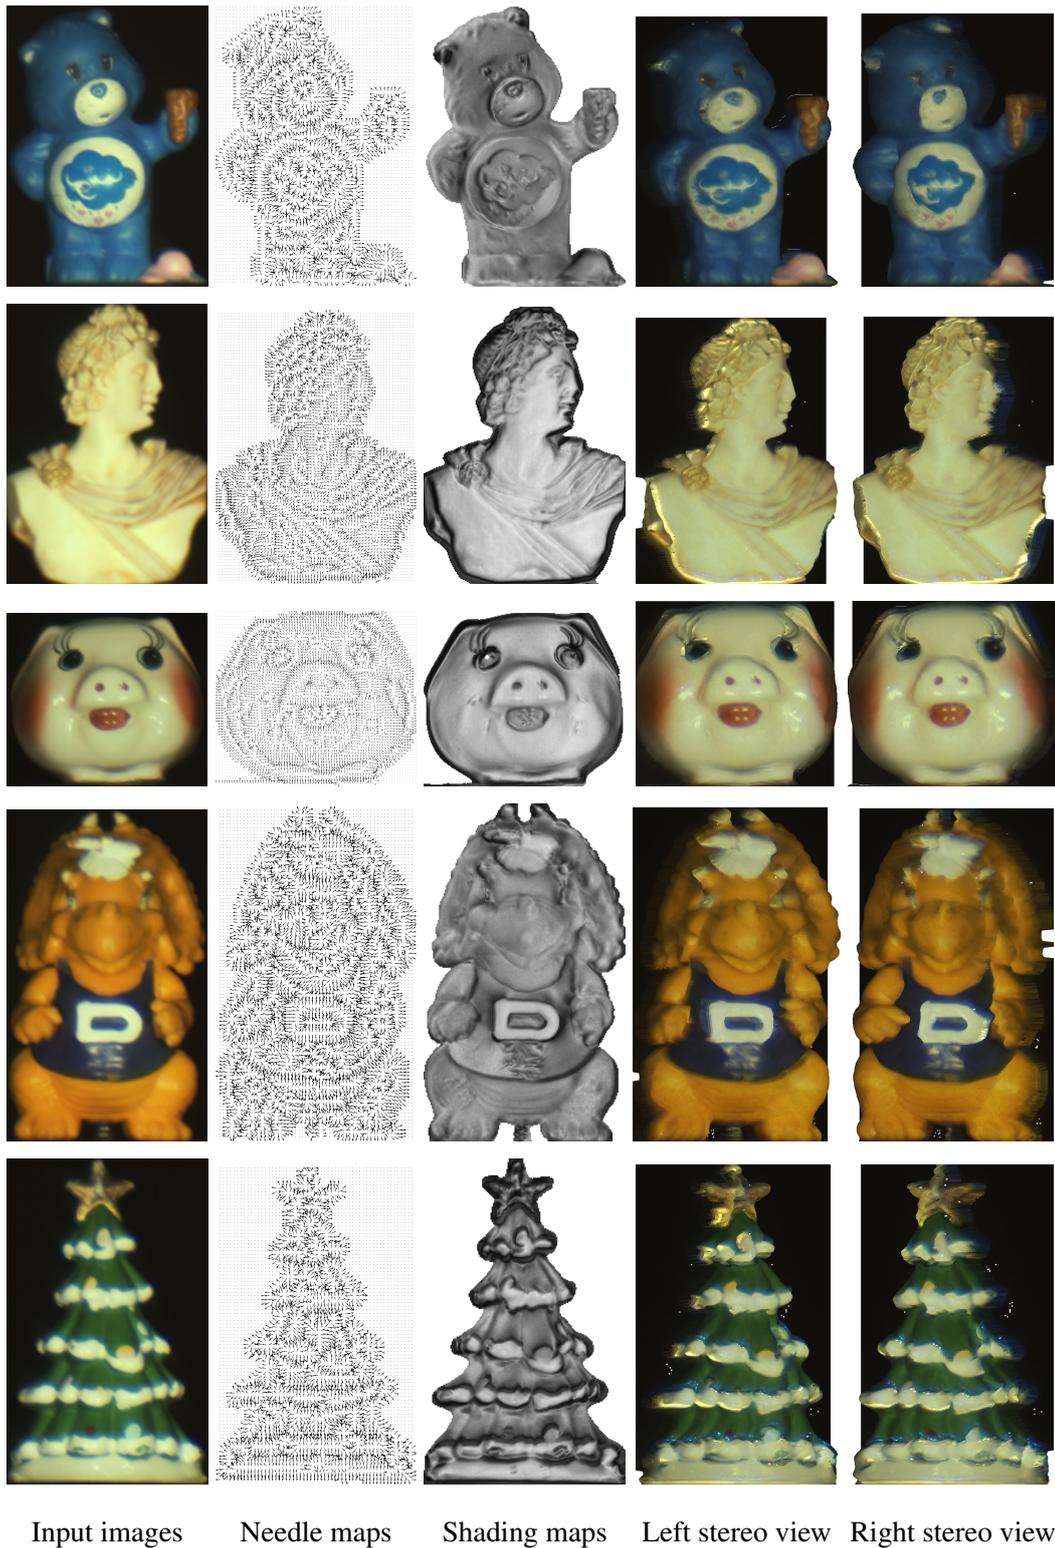


Figure 13: Shape estimation and rendering for real-world images. First column: The  $45^\circ$  polarisation component images rendered in trichromatic pseudo-colours. Second column: Needle maps. Third column: Shading maps. Fourth and fifth columns: Two-view stereo rendering of the input shapes under  $L_3$ . In the fourth row, the view direction is  $10^\circ$  to the left of the original view. In the fifth row, the view direction forms an angle of  $10^\circ$  to the right of the original view.

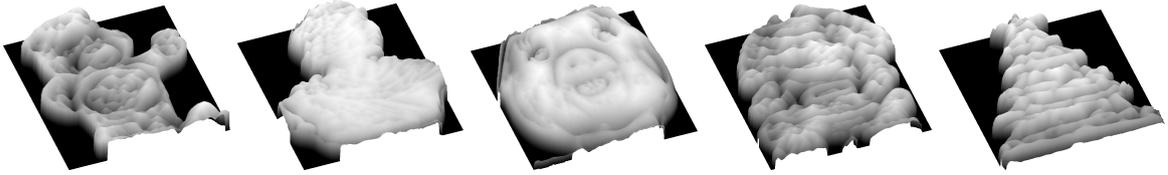


Figure 14: Depth maps of the objects shown in Figure 13.

these objects under incandescent lights with directions coplanar to the horizontal axis of the image plane. Specifically, from left to right, these light source directions point towards the illuminated object, forming angles of  $-45^\circ$ ,  $-30^\circ$ ,  $0^\circ$ ,  $30^\circ$  and  $45^\circ$  with the viewing direction, where a negative angle means the light is located to the left-hand side of the camera and positive to the right. 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 the fourth and fifth columns of Figure 13, we show a pair of stereo images of the original objects rendered under the frontal illumination direction  $L_3$ . The left and right images were generated for two viewing directions displaced  $10^\circ$  to the left and to the right of the original viewing direction. The novel views reveal the object shapes near the left and right occlusion boundaries. This points to potential applications in the three-dimensional visualisation of objects using single-view images.

In Figure 14, we present the depth maps recovered from the real-world imagery. These maps have been produced in such a way that the gray level corresponds to the surface height. We note that the three-dimensional structure of the reconstructed shapes is, in general, perceptually consistent with their original input images and the corresponding stereo pairs in Figure 13. We also note that the distortions across material boundaries are particularly visible in the reconstructed shapes for the bear and the dinosaur. This is because the variation of polarisation across materials has been interpreted as a result of the variation in surface geometry.

To provide a quantitative analysis of the rendering results, we proceed to render the objects under the power spectrum of an incandescent light. In Table 7, we show the rendering accuracy under the novel lighting directions. The error is measured as the angular difference between the rendered image reflectance spectra and their ground-truth on a per-pixel basis. It is worth stressing that, our error measures are obtained from the multi-spectral images rather than the RGB values yielded by the color matching functions. Recall that we estimate the image reflectance spectra from the rendering equation of the Wolff reflectance model [43]. The numerical results reported in Table 7

	$L_3$	$L_4$	$L_5$	$L_1 + L_5$	$L_2 + L_4$
Bear	$11.63 \pm 2.90$	$12.22 \pm 3.67$	$12.94 \pm 5.95$	$12.65 \pm 4.71$	$11.48 \pm 3.17$
Statue	$12.32 \pm 3.43$	$14.11 \pm 3.24$	$14.18 \pm 4.02$	$15.75 \pm 4.03$	$13.46 \pm 2.90$
Pig	$10.70 \pm 3.40$	$11.78 \pm 3.47$	$12.94 \pm 4.14$	$12.87 \pm 4.43$	$10.53 \pm 2.91$
Dinosaur	$10.67 \pm 3.76$	$12.19 \pm 6.94$	$14.01 \pm 8.15$	$9.02 \pm 3.94$	$8.27 \pm 3.60$
Pine Tree	$10.82 \pm 2.69$	$11.33 \pm 3.51$	$14.92 \pm 4.81$	$13.05 \pm 3.94$	$10.35 \pm 3.99$

Table 7: The angular deviation (in degrees) between the spectral reflectance images rendered for the frontal viewing direction and the ground truth images. The mean and standard deviation of these errors across pixels are reported for each image.

are the mean and standard deviation across pixels in each image. The results here are consistent with the qualitative results presented previously in the sense that the rendering quality is better for the cases of the frontal illuminant and the two simultaneous illuminants. This is due to the fact that in such conditions, the objects are fully illuminated and the rendered image has a smooth shading variation. On the other hand, the lower rendering accuracy for the oblique illuminant directions is due to non-smooth shading where cast shadows occur across material boundaries. Nonetheless, our method can produce rendering results that are in good accordance with the ground truth. In addition, the method delivers shading maps which accord well with the geometry of the object under study, based solely on polarisation information. This is an important characteristic of our method since it does not employ shading or chrominance as a cue.

We present qualitative results for refractive index estimated from the real-world imagery. In Figure 15, we show the mean refractive index recovered from the input spectro-polarimetric images, whose  $45^\circ$  polarisation components are shown in pseudocolour in the left-most column. The second, third and fourth columns in the figure correspond to the refractive index corresponding to the wavelengths of 450, 550 and 650nm, respectively. In these, for purposes of visualisation, pure white corresponds to a refractive index of 3, whereas  $\frac{1}{3}$  corresponds to a refractive index of 1. Note the refractive index not only varies with respect to the wavelength but also spatially. Moreover, the spatial variation of the refractive index values at the wavelengths 450 and 650nm reflect the differences between materials in the input images.

Finally, in Figure 16, we plot the mean refractive index as a function of the wavelength for two sample regions in each image. The top row shows the  $45^\circ$  polarisation image component



Figure 15: The variation of the refractive index with respect to the wavelength. First column: The 45° polarisation component of the input images, rendered in trichromatic pseudocolour. The second to fourth columns respectively show the estimated refractive index at the wavelengths of 450, 550 and 650nm. The refractive index value is linearly scaled to the range between 0 and 1, with pure white corresponding to a refractive index of 3 and  $\frac{1}{3}$  corresponding to a refractive index of 1.

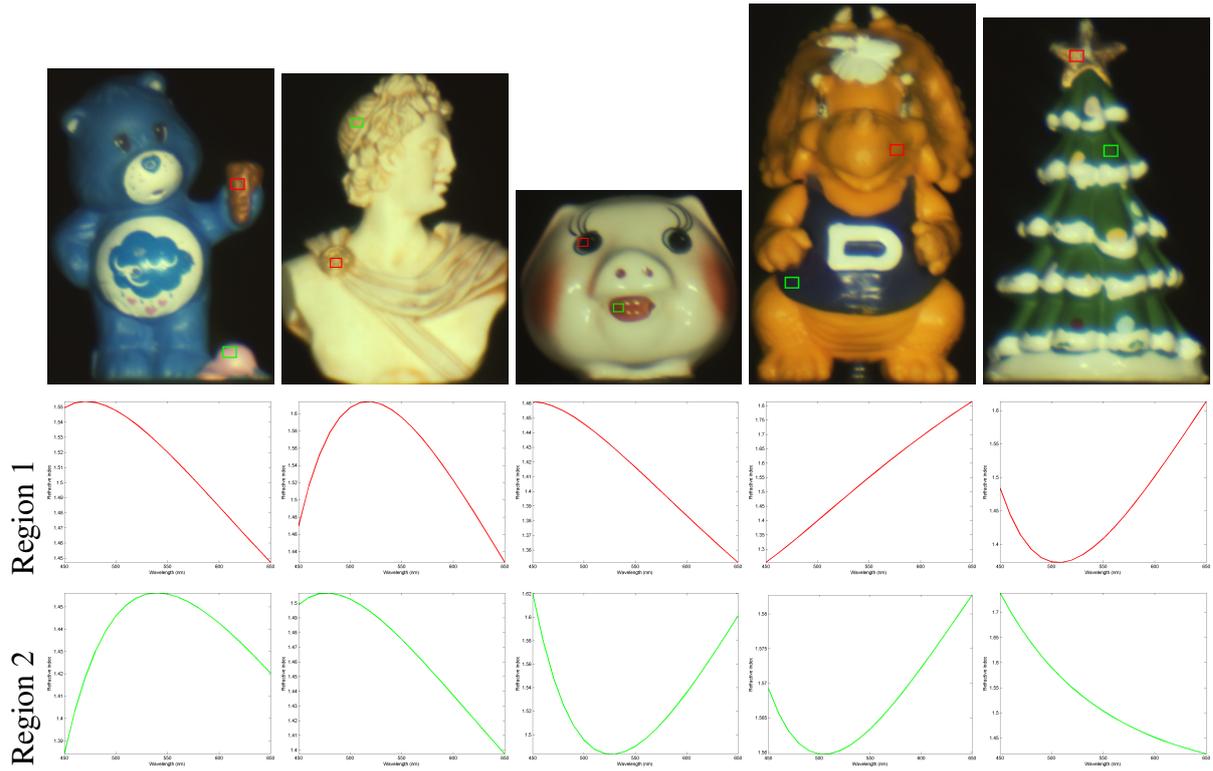


Figure 16: Mean refractive index spectra of a number of selected regions in the input images. First row: The  $45^\circ$  polarisation component images rendered in pseudocolour, with the selected regions indicated by rectangular bounding boxes. Second and third rows: The mean of the estimated refractive index spectra over the pixels in each selected region, fitted to the Cauchy dispersion equation. These spectra are plotted as solid lines with colours matching those of the region boundaries depicted in the first row.

rendered in pseudocolour, with the selected regions indicated by rectangular bounding boxes. The remaining rows show the mean estimated refractive index spectrum for each selected region as a result of fitting to the Cauchy dispersion equation. Note that the refractive index spectra are plotted in colours that match those of the bounding boxes of the selected regions in the input images. In the figure, the regions selected in the image of each object, except for the Statue, are made of different materials. This results in the difference between the estimated refractive index spectra between the two regions in each of these images. Since the Statue is composed of almost the same material across its surface, the dispersion of the refractive index over the wavelength does not vary significantly across both of the selected regions. Furthermore, the refractive index values can provide a hint at the roughness of the surface. For example, the surface of region 1 in the Bear image consists of several ridges and valleys and is rougher than that of region 2. This observation is consistent with the ranges of refractive index values estimated by our method, which are between

1.45–1.55 for region 1 and 1.38–1.46 for region 2.

## 6 Conclusions

This paper has focused on shape and refractive index recovery based on the polarisation information in a multispectral or hyperspectral image acquired from a single-view. Departing from the theory of polarisation in the electromagnetic spectrum, the method hinges on the analysis of polarisation upon diffuse reflection from dielectric surfaces. By capturing the polarisation of light emitted from these surfaces, we observe the variation of the spectral radiance with respect to the polariser angle along a Transmitted Radiance Sinusoidal (TRS) curve. The recovery process commences with the decomposition of the input imagery into the components of the TRS, *i.e.* its phase, minimal and maximal radiance values. We have provided a link between the azimuth angle of surface normal and the phase of polarisation. Moreover, we have shown how to disambiguate the azimuth angle of surface normals from the wavelength-indexed spectrum of the phase angles. We have also drawn upon the Fresnel transmission ratio between the minimal and maximal radiance on the TRS to jointly recover the zenith angle of the surface normal and the refractive index. To make the recovery problem well-posed, we have enforced integrability on the surface and employed the material dispersion as a constraint on the refractive index. We solve the estimation problem using an iterative optimisation approach and derive a closed-form solution for the zenith angle and the refractive index in each iteration. Lastly, we have demonstrated the merit of our method for the purpose of shape and refractive index recovery from synthetic and real-world imagery. The experimental results demonstrate the utility of the method for applications involving non-contact measurement of refractive indices of dielectrics.

There are several strategies to extend the current work. While diffuse polarisation is often due to the penetration, subsurface scattering and back refraction of light at the surface boundary, specular polarisation may occur as light is directly reflected at the air-material interface. In the work presented here, we have dealt with materials that diffusely polarise light. In the future, there is potential research in developing a reflection model that accounts for the air-material interaction. Other possibility is to exploit the utility of polarisation information as a useful cue for the recovery of the object shape near the grazing angle between the viewing direction and the surface normal, where the surface orientation induces a strong degree of polarisation. Future research can benefit from this property by combining shape recovery results near the occluding boundaries of objects

captured from multiple views so as to enhance the quality of shape reconstruction.

## References

- [1] ATKINSON, G., AND HANCOCK, E. Recovery of surface orientation from diffuse polarization. *IEEE Transactions on Image Processing* 15, 6 (June 2006), 1653–1664.
- [2] ATKINSON, G., AND HANCOCK, E. R. Recovery of surface height using polarization from two views. In *CAIP* (2005), pp. 162–170.
- [3] ATKINSON, G., AND HANCOCK, E. R. Surface reconstruction using polarization and photometric stereo. In *CAIP* (2007), pp. 466–473.
- [4] ATKINSON, G. A., AND HANCOCK, E. R. Multi-view surface reconstruction using polarization. In *IEEE International Conference on Computer Vision (ICCV'05), Volume 1* (Washington, DC, USA, 2005), IEEE Computer Society, pp. 309–316.
- [5] ATKINSON, G. A., AND HANCOCK, E. R. Shape estimation using polarization and shading from two views. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 29, 11 (2007), 2001–2017.
- [6] BELHUMEUR, P. N., KRIEGMAN, D. J., AND YUILLE, A. L. The bas-relief ambiguity. *International Journal of Computer Vision* 35, 1 (1999), 33–44.
- [7] BORN, M., AND WOLF, E. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light (7<sup>th</sup> Edition)*, 7th ed. Cambridge University Press, 1999.
- [8] CHEN, H., AND WOLFF, L. Polarization phase based method for material classification in computer vision. *IJCV* 28, 1 (June 1998), 73–83.
- [9] COLEMAN, T. F., AND LI, Y. A Reflective Newton Method for Minimizing a Quadratic Function Subject to Bounds on some of the Variables. *SIAM Journal on Optimization* 6, 4 (April 1996), 1040–1058.
- [10] DENES, L. J. Acousto-optic tunable filters in imaging applications. *Optical Engineering* 37 (1998).
- [11] DRBOHLAV, O., AND ŠÁRA, R. Unambiguous determination of shape from photometric stereo with unknown light sources. In *International Conference on Computer Vision* (Los Alamitos, CA, USA, 2001), vol. 1, IEEE Computer Society, pp. 581–586.
- [12] FRANKOT, R. T., AND CHELLAPPA, R. A Method for Enforcing Integrability in Shape from Shading Algorithms. *IEEE Trans. Pattern Anal. Mach. Intell.* 10, 4 (1988), 439–451.
- [13] GONZALEZ, R. C., AND WOODS, R. E. *Digital Image Processing*, 2nd ed. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1992.
- [14] GOUDAIL, F., TERRIER, P., TAKAKURA, Y., BIGUÉ, L., GALLAND, F., AND DEVLAMINCK, V. Target detection with a liquid-crystal-based passive stokes polarimeter. *Applied Optics* 43, 2 (Jan 2004), 274–282.

- [15] GUPTA, N., DAHMANI, R., AND CHOY, S. Acousto-optic tunable filter based visible- to near-infrared spectropolarimetric imager. *Optical Engineering* 41, 5 (2002), 1033–1038.
- [16] HALL, J. S. Some polarization measurements in astronomy. *Journal of Optical Society America* 41, 12 (Dec 1951), 963–966.
- [17] HARRIS, S. E., AND WALLACE, R. W. Acousto-optic tunable filter. *Journal of Optical Society America* 59, 6 (June 1969), 744–747.
- [18] HAWRYSHYN, C. W. Ultraviolet polarization vision in fishes: Possible mechanisms for coding e-vector. *Philosophical Transactions: Biological Sciences* 355, 1401 (2000), 1187–1190.
- [19] HECHT, E. *Optics*, fourth ed. Addison-Wesley, 2002.
- [20] K. TORRANCE AND E. SPARROW. Theory for Off-Specular Reflection From Roughened Surfaces. *Journal of the Optical Society of America* 57, 9 (Sep 1967), 1105–1112.
- [21] KASAROVA, S. N., SULTANOVA, N. G., IVANOV, C. D., AND NIKOLOV, I. D. Analysis of the dispersion of optical plastic materials. *Optical Materials* 29, 11 (2007), 1481 – 1490.
- [22] M. OREN AND S. K. NAYAR. Generalization of the Lambertian Model and Implications for Machine Vision. *International Journal of Computer Vision* 14, 3 (April 1995), 227–251.
- [23] MANDEL, L., AND WOLF, E. *Optical Coherence and Quantum Optics*. Cambridge University Press, September 1995.
- [24] MARSHALL, N. J., LAND, M. F., KING, C. A., AND CRONIN, T. W. The Compound Eyes of Mantis Shrimps (Crustacea, Hoplocarida, Stomatopoda). I. Compound Eye Structure: The Detection of Polarized Light. *Philosophical Transactions: Biological Sciences* 334, 1269 (October 1991), 33–56.
- [25] MIYAZAKI, D., KAGESAWA, M., AND IKEUCHI, K. Transparent surface modeling from a pair of polarization images. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 26, 1 (2004), 73–82.
- [26] MIYAZAKI, D., SAITO, M., SATO, Y., AND IKEUCHI, K. Determining surface orientations of transparent objects based on polarization degrees in visible and infrared wavelengths. *Journal of Optical Society America A* 19, 4 (2002), 687–694.
- [27] MIYAZAKI, D., TAN, R. T., HARA, K., AND IKEUCHI, K. Polarization-based inverse rendering from a single view. *IEEE International Conference on Computer Vision* 2 (2003), 982.
- [28] NAYAR, S. K., FANG, X.-S., AND BOULT, T. Separation of reflection components using color and polarization. *Int. J. Comput. Vision* 21, 3 (1997), 163–186.
- [29] NOCEDAL, J., AND WRIGHT, S. J. *Numerical Optimization*, 2nd ed. Springer, New York, 2006.
- [30] RAHMANN, S. Inferring 3D Scene Structure from a Single Polarization Image. In *SPIE Proceedings on Polarization and Color Techniques in Industrial Inspection* (1999), vol. 3826, pp. 22–33.
- [31] RAHMANN, S. Polarization Images: A Geometric Interpretation for Shape Analysis. *International Conference on Pattern Recognition* 3 (2000), 538 – 542.

- [32] RAHMANN, S., AND CANTERAKIS, N. Reconstruction of specular surfaces using polarization imaging. *IEEE Conference on Computer Vision and Pattern Recognition 1* (2001), 149 – 155.
- [33] SADJADI, F. A., AND CHUN, C. S. L. Remote sensing using passive infrared stokes parameters. *Optical Engineering* 43 (2004), 2283–2291.
- [34] SAITO, M., KASHIWAGI, H., SATO, Y., AND IKEUCHI, K. Measurement of surface orientations of transparent objects using polarization in highlight. *IEEE Conference on Computer Vision and Pattern Recognition 1* (1999), 1381.
- [35] SCHLICK, C. An inexpensive brdf model for physically-based rendering. *Computer Graphics Forum* 13, 3 (1994), 233–246.
- [36] SELLMEIER. Zur erklärung der abnormen farbenfolge im spectrum einiger substanzen. *Annalen der Physik* 219, 6 (1871), 272–282.
- [37] SHANNON, C. E. Communication in the presence of noise. *Proc. Institute of Radio Engineers* 37, 1 (Jan 1949), 10–21.
- [38] STILES, W. S., AND BURCH, J. M. N.P.L. colour-matching investigation: Final report (1958). *Optica Acta* 6 (1959), 1–26.
- [39] THILAK, V., VOELZ, D. G., AND CREUSERE, C. D. Polarization-based index of refraction and reflection angle estimation for remote sensing applications. *Applied Optics* 46, 30 (Oct 2007), 7527–7536.
- [40] TORRANCE, K. E., SPARROW, E. M., AND BIRKEBAK, R. C. Polarization, Directional Distribution, and Off-Specular Peak Phenomena in Light Reflected from Roughened Surfaces. *Journal of the Optical Society of America* 56, 7 (July 1966), 916–924.
- [41] WOLFF, L., AND BOULT, T. Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 13, 7 (1991), 635–657.
- [42] WOLFF, L. B. Polarization-based material classification from specular reflection. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12, 11 (1990), 1059–1071.
- [43] WOLFF, L. B. Diffuse-reflectance model for smooth dielectric surfaces. *Journal of the Optical Society of America* 11, 11 (November 1994), 2956–2968.
- [44] WOLFF, L. B. Polarization vision: a new sensory approach to image understanding. *Image Vision Computing* 15, 2 (1997), 81–93.
- [45] WOLFF, L. B., AND ANDREOU, A. G. Polarization camera sensors. *Image Vision Computing* 13, 6 (August 1995), 497–510.
- [46] WOLFF, L. B., AND BOULT, T. E. Polarization/Radiometric Based Material Classification. In *Computer Vision and Pattern Recognition* (1989), pp. 387–395.
- [47] WOLFF, L. B., MANCINI, T. A., POULIQUEN, P., AND ANDREOU, A. G. Liquid crystal polarization camera. *IEEE Transactions on Robotics and Automation* 13, 2 (1997), 195–203.

- [48] ZHU, Q., AND SHI, J. Shape from shading: Recognizing the mountains through a global view. In *Computer Vision and Pattern Recognition* (2006), vol. 2, pp. 1839–1846.