Recovery of Surface Orientation From Diffuse Polarization

Gary A. Atkinson, *Student Member, IEEE,* and Edwin R. Hancock

Abstract—When unpolarized light is reflected from a smooth dielectric surface, it becomes partially polarized. This is due to the orientation of dipoles induced in the reflecting medium and applies to both specular and diffuse reflection. This paper is concerned with exploiting polarization by surface reflection, using images of smooth dielectric objects, to recover surface normals and, hence, height. This paper presents the underlying physics of polarization by reflection, starting with the Fresnel equations. These equations are used to interpret images taken with a linear polarizer and digital camera, revealing the shape of the objects. Experimental results are presented that illustrate that the technique is accurate near object limbs, as the theory predicts, with less precise, but still useful, results elsewhere. A detailed analysis of the accuracy of the technique for a variety of materials is presented. A method for estimating refractive indices using a laser and linear polarizer is also given.

Index Terms—Diffuse polarization, refractive index, sensitivity study, surface recovery.

I. INTRODUCTION

POLARIZATION has proven to be a useful source of information in the analysis of light scattering from surfaces in computer vision. There are a number of ways in which polarization arises and can be used in surface analysis. One familiar example is when the incident light is polarized and the polarization of the scattered light is analyzed [1]. When the scattering process is coherent, then the polarization of the incident light is preserved, whereas polarization is destroyed when the scattering process is incoherent. This property can be used to remove specularities from surfaces. However, there are subtler polarization effects that can be exploited. For instance, under certain conditions, initially unpolarized light becomes polarized as a result of the scattering process [2]. This applies to both specular reflection (which we refer to as specular polarization) and diffuse reflection (diffuse polarization) and is due to the directionality of the molecular electron charge density interacting with the electromagnetic field of the incident light [3].

There is a considerable amount of literature on the use of polarization for surface analysis. Most research aimed at extracting and interpreting information from polarization data involves placing a linear polarizer in front of a camera and taking images of an object or a scene with the polarizer oriented at different angles [2], [4]. Recently, however, Wolff and others have used liquid crystal technology to rapidly acquire polarization images at a rate such that several complete sets of polarization data (i.e., degree of polarization and orientation of polarized light at each point) can be obtained per second [5]. The rapid acquisition of polarization data made possible by this technology extends the range of potential applications to include slow moving object analysis.

Much of this work relies on specular polarization and very specific lighting conditions so that specular reflection occurs across the whole object. In previous work by Ikeuchi *et al.* [6], [7], this global specularity was obtained by placing the object under investigation inside a spherical diffuser, with several light sources outside and a hole for the camera. With this setup, light impinges the surface from all directions and, since specular reflection is generally much stronger than diffuse reflection, the latter component can be ignored. Other possibilities include separating diffuse and specular reflection components using color [8], [9], a probabilistic framework [1], or polarization [2], [9], [10]. The relevant theory can then be applied to the specular or diffuse images. Existing work has demonstrated the usefulness of polarization in surface height recovery [7], [11], [12]; overcoming the surface orientation ambiguity associated with photometric stereo [13], [14]; image segmentation [2]; recognition and separation of reflection components [2], [10]; and distinguishing true laser stripes from interreflections for triangulation-based laser scanning [15].

An important contribution to shape from *diffuse* polarization was made by Miyazaki *et al.* [4], who use a similar method to that described in this paper and also estimate the illumination distribution from specularities. However, interreflections, which we show here to be of significance, are ignored.

Drbohlav and Šára [16] also use diffuse polarization, but only a single opaque sphere is used for testing and so the range of applicability is not made clear. In this paper, we present a theoretical account of how diffuse polarization can be used to estimate surface normal directions and present results for a greater range of objects. One of our main aims here is to explore the extent to which surface orientation can be estimated from diffuse polarization alone. We clearly identify the limitations of the method and point toward possible ways of overcoming them.

The work in this paper is based on the Fresnel theory for the reflection of electromagnetic waves from interfaces between media of different refractive indices. The Fresnel equations, which are the principal results of this theory, give the ratio of reflected to incident light wave amplitudes for light that is linearly polarized either parallel to the plane of incidence (the plane that contains the incident and reflected light rays) or perpendicular to the plane of incidence. The equations provide

The authors are with the Department of Computer Science, University of York, York YO10 5DD, U.K. (e-mail: atkinson@cs.york.ac.uk; erh@cs.york.ac.uk).

a means by which the local surface normals of an object can be estimated. The zenith angle (i.e., the angle between the surface normal and the viewing direction) can be determined by the degree of polarization and may be recovered by solving a nonlinear equation. The azimuth angle of the surface normal (the angle of the projection of the normal onto the image plane relative to a reference) is determined by the orientation of the polarization. Once a field of surface normals for the object under study is to hand, then the surface height function may be recovered using standard surface integration techniques. To determine the required polarization, at least three images are needed.

An analysis of the accuracy and reliability of the method for a variety of shapes and materials has also been performed. We believe this is important since, like most vision techniques, the bounds of applicability should be clearly identified. Our final contribution is a method for refractive index estimation that uses a planar sample of the material under study, placed on a rotatable table, a laser and a linear polarizer. This helps to demonstrate when the theory is adequate and when further complications are present.

Of course, the topic of recovering local surface orientation and, hence, height is that of shape-from-shading [17]–[19]. Here, the aim is to use the physics of light reflectance to estimate the zenith and azimuth angles of the surface normal. The process is an underconstrained one since the two degrees of freedom of each surface normal cannot be recovered from a single measured brightness value. Hence, assumptions must be made. The surface is usually assumed to be matte, of constant albedo, and illuminated by a single point light source placed at infinity. Shape-from-shading algorithms frequently assume the surface to be Lambertian, i.e., a surface that appears equally bright from all viewing directions, although better but more computationally intensive models do exist [20], [21]. To overcome the problem that the recovery of surface normals is an underconstrained problem, the surface under study is assumed to be smooth and constraints provided by the direction of the surface normals at occluding boundaries are used. However, an alternative way of overcoming the underconstrained nature of the problem is to use multiple images. For instance, in photometric stereo [22]–[24], the object under study is kept static with respect to the camera, and the direction of the light source is varied. In Helmholtz stereopsis [25], on the other hand, unique surface normal recovery is possible if the light source and camera are interchanged.

When compared with single-view shape-from-shading, our technique offers a number of advantages. First, although it is confined to smooth dielectric surfaces, the reflectance does not need to be Lambertian. Second, the light source direction does not need to be known. Finally, the surface normals are fully constrained except for a single $180°$ ambiguity in azimuth angle.

The outline of the paper is as follows. Section II outlines the physics of polarization by reflection and explains how this theory has been used to estimate surface orientation using information concerning the polarization state of the reflected light. The theory is provided for both specular and diffuse reflection. Section III describes the experimental setup and the algorithm that we used to recover surface shape. A representative set of experimental results that illustrate both the strengths and limitations of the technique are given Section IV. We also include an error analysis and describe a method for measuring the refractive indices for the surface under study. Finally, Section V presents conclusions and suggests directions for future research.

II. POLARIZATION AND REFLECTION

In this section, some of the key physics of reflection from smooth surfaces is discussed, with particular emphasis on how measurements of the polarization state of reflected light can be used in computer vision. It is assumed that the incident light is unpolarized. As we will show, the Fresnel reflectance theory can be used to predict the angle of polarization of the reflected light and the extent of the partial polarization for a given material at a given orientation. This prediction is then applied to diffusely reflected light across the entire visible surface of objects under investigation to obtain surface orientation.

The electric field of an electromagnetic wave incident on a surface causes the electrons of the reflecting medium near the surface to vibrate, forming dipoles [3], [26]. These vibrating electrons reradiate, generating the reflected rays. For electric fields perpendicular to the plane of incidence the electrons also vibrate perpendicular to the plane and, thus, so does the electric field of the reflected light. For light polarized parallel to the plane of incidence, on the other hand, the electrons do not vibrate perpendicularly to the reflected ray, as Fig. 1 shows, resulting in a more attenuated wave. The effect is particularly marked for smooth surfaces, as significant roughness tends to depolarize the light. As Fig. 1 suggests, the degree of polarization depends on the angle of incidence. Later, we show how this description can be extended, by considering internal scattering, to account for partial polarization of *diffusely* reflected light.

A. Review of Fresnel Coefficients for Dielectrics

The Fresnel equations give the ratios of the reflected wave amplitude to the incident wave amplitude for incident light that is linearly polarized perpendicular to, or parallel to, the plane of specular incidence [27]. These ratios depend upon the angle of incidence and the refractive index of the reflecting medium. Since the incident light can always be resolved into components perpendicular to, and parallel to, the plane of incidence, the Fresnel equations are applicable to all incident polarization states.

For the geometry of Fig. 1, where $E_{0i\perp}$ and $E_{0r\perp}$ are the amplitudes of the incident and the reflected waves, respectively, the amplitude reflection coefficient for light polarized perpendicular to the plane of incidence $r_\perp$ at a boundary between two media is given by

$$r_\perp (n_i, n_t, \mu_i, \mu_t, \theta_i) \equiv \frac{E_{0r\perp}}{E_{0i\perp}} = \frac{\frac{n_i}{\mu_i} \cos \theta_i - \frac{n_t}{\mu_t} \cos \theta_t}{\frac{n_i}{\mu_i} \cos \theta_i + \frac{n_t}{\mu_t} \cos \theta_t}. \quad (1)$$

In this equation, n_i and n_t are the refractive indices of the first and second media, and μ_i and μ_t are the magnetic permeabilities of the two media. The angles θ_i and θ_t for the incident and transmitted light are defined in Fig. 1. When the first medium is air, then $n_i \approx 1$. In computer vision, we are not usually interested
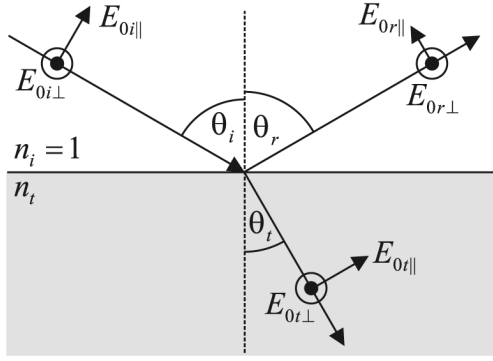
Fig. 1. Reflection of an electromagnetic wave, with electric field directions and relative amplitudes indicated. Within the medium, the electrons vibrate parallel to the electric field. Since this direction is not perpendicular to the reflected wave for the component parallel to the plane of incidence, only a component of the vibrations cause a reflected ray.

in ferromagnetic materials, so $\mu_i = \mu_t = \mu_0$, the permeability of free space. Therefore, (1) reduces to

$$r_\perp (n_i, n_t, \theta_i) \equiv \frac{E_{0r\perp}}{E_{0i\perp}} = \frac{n_i \cos \theta_i - n_t \cos \theta_t}{n_i \cos \theta_i + n_t \cos \theta_t}. \quad (2)$$

For light polarized parallel to the plane of incidence

$$r_\| (n_i, n_t, \theta_i) \equiv \frac{E_{0r\|}}{E_{0i\|}} = \frac{n_t \cos \theta_i - n_i \cos \theta_t}{n_t \cos \theta_i + n_i \cos \theta_t}. \quad (3)$$

The angle θ_t can be obtained from the well-known Snell's Law

$$n_i \sin \theta_i = n_t \sin \theta_t. \quad (4)$$

Generally, it is not the amplitude of the wave that is measured by the detector, but the intensity, which is proportional to the square of the amplitude [3]. With this in mind, it is possible to show that the *intensity coefficients*, which relate the reflected power to the incident power, are $R_\perp = r_\perp^2$ and $R_\| = r_\|^2$. We are assuming that the refractive index is wavelength independent. In fact, there is some residual wavelength dependence because the wavelength of reflected light more closely matches that of the incident light near grazing angles [28], but the equations above provide accurate results for most situations.

Fig. 2(a) shows the Fresnel intensity coefficients for a typical dielectric as a function of the angle of the incident light. Both reflection and transmission coefficients are shown, where the latter refers to the ratio of transmitted to incident power.[1] At about 60° (typical for most dielectrics), the component of the reflected light polarized parallel to the plane of incidence is completely extinguished. This angle is known as the *Brewster* or *polarizing* angle. The figure shows that there is some information contained in the polarization state of the reflected light. For specular reflection, a certain fraction of the light is reflected for each polarization component (i.e., the component parallel to the plane of incidence and the component perpendicular to it). This fraction is greater for the component polarized perpendicular to the plane of incidence. Therefore, the reflected light is *partially linearly polarized*, i.e., consists of an unpolarized component, and a completely polarized component.

[1]The transmission coefficients are simply $T_\perp = 1 - R_\perp$ and $T_\| = 1 - R_\|$.

As one would expect, most of the light incident on a *metallic* surface is reflected, as Fig. 2(b) shows. The problem of overall shape recovery from metallic specular reflection is simplified since no diffuse reflection occurs in metals, i.e., there is no partial reduction in the degree of polarization due to a diffuse reflection component. This is particularly important near occluding boundaries where diffuse reflection has a greater influence on the polarization state of the reflected light for nonmetallic materials. Since this paper focuses on *diffuse* reflection, metals are not studied further here.

B. Polarization Image

We now consider how the above theory can be used in computer vision. As a polarizer placed in front of a camera is rotated, the measured pixel brightness varies according to the *transmitted radiance sinusoid*

$$I (\theta_{\text{pol}}, \phi) = \frac{I_{\max} + I_{\min}}{2} + \frac{I_{\max} - I_{\min}}{2} \cos (2\theta_{\text{pol}} - 2\phi). \quad (5)$$

Here, $I_{\max}$ and $I_{\min}$ are the maximum and minimum observed pixel brightnesses as the polarizer is rotated, θ_{pol} is the angle which the polarizer makes with the arbitrary reference direction (here vertically upwards) and ϕ is the *phase angle*, or angle of polarization of the reflected light. The maximum pixel brightness is, therefore, observed when the polarizer is aligned with the angle of polarization, i.e., $I (\theta_{\text{pol}} = \phi) = I_{\max}$, and the minimum is observed when the polarizer is rotated by a further 90° in either direction, i.e., $I (\theta_{\text{pol}} = \phi \pm 90°) = I_{\min}$. Fig. 3 illustrates this relationship.

Fig. 2 shows that the maximum and minimum intensities detected for a particular surface orientation are

$$I_{\max} = \frac{R_\perp}{R_\perp + R_\|} I_s ; \quad I_{\min} = \frac{R_\|}{R_\perp + R_\|} I_s \quad (6)$$

where I_s is the magnitude of the specular component of reflection (assume for now that there is no diffuse reflection). The *degree of polarization* or *partial polarization*, which is frequently used in computer vision, is defined to be

$$\rho = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}}. \quad (7)$$

Each polarization image, i.e., the full set of polarization data for a given object or scene, is comprised of three separate components. The first of these is the *intensity image*, which is simply the image that would be obtained using a normal camera.

Second, there is the *phase image* which encodes the orientation of the polarizer corresponding to maximum transmission through it. The phase is, therefore, directly related to the angle of the linearly polarized component of the reflected light and can be defined as the angle of maximum (as in Fig. 3) or minimum transmission. Note that polarizers cannot distinguish between two angles separated by 180°, so the range of initially acquired phase measurements is $[0, 180°)$. There is, therefore, a 180° ambiguity, since two maxima in pixel brightness are found as the polarizer is rotated through 360°. Possible methods for dealing with this problem include initially directing the surface normal estimates away from the object at the occluding boundary [4], and propagating into the object, rotating the vectors by 180°
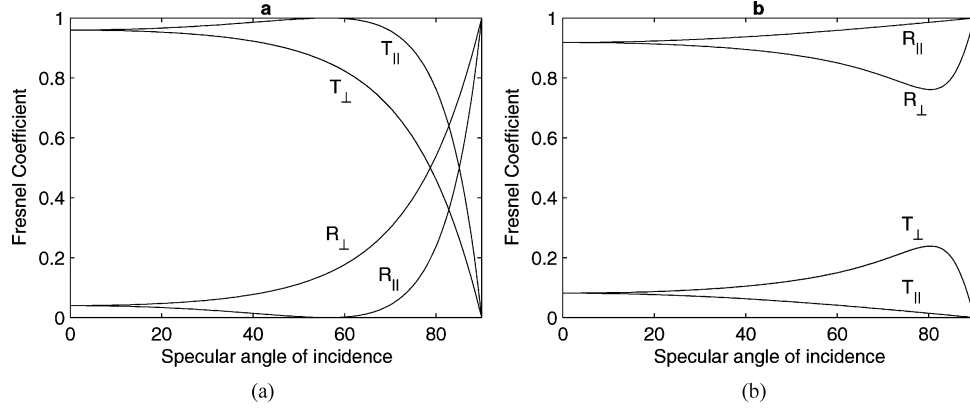
Fig. 2.   (a) Reflection and transmission coefficients for a dielectric ($n = 1.7$). For comparison, (b) shows the coefficients for a metal ($n = 0.8$, $k = 6$, where k is the extinction coefficient [27]. Note that n can be less than 1 if $k > 0$. For dielectrics, $k = 0$ and the equations presented in this paper assume that this is the case).
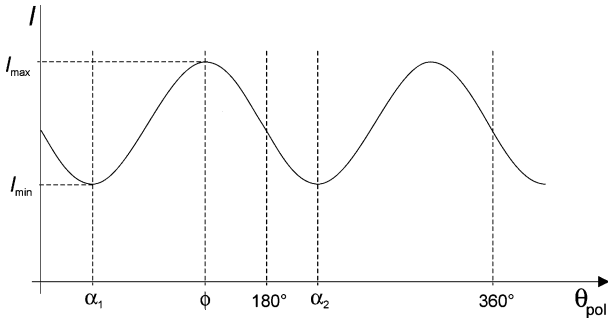

Fig. 3.   Transmitted radiance sinusoid. α_1 and α_2 are the two possible surface azimuths for a given phase angle ϕ. The two candidates for azimuth angle shown here are for specular reflection. For diffuse reflection, they correspond to θ_{pol} where $I_{\max}$ is observed.

if this produces better local alignments of surface normals (see Section III), tracing level curves [29] or applying some form of optimization algorithm possibly involving smoothing. If the light source direction and viewing direction are different, then the consistent viewpoint constraint [14] can be used to recover information from specularities. Finally, the well-known integrability constraint given below can be enforced, which ensures a smooth surface [30]. The constraint is given by

$$\frac{\partial}{\partial y}\frac{\partial z}{\partial x} = \frac{\partial}{\partial x}\frac{\partial z}{\partial y} \qquad (8)$$

where $z = f(x, y)$ is the surface height.

The final component of the polarization image is the *degree of polarization image*, as determined by (7). The phase and degree of polarization components of the polarization image are usually found by taking images with the polarizer at three or more angles and fitting the data to (5).

Wolff [5] suggests taking three images I_0, I_{45}, and I_{90} (as used by liquid crystal polarization cameras) corresponding to polarizer orientations of $0°$, $45°$, and $90°$, respectively, and using the following equations to determine the phase, intensity, and degree of polarization:

Phase

$$\phi = \frac{1}{2}\arctan\left(\frac{I_0 + I_{90} - 2I_{45}}{I_{90} - I_0}\right) + 90°$$

if $(I_{90} < I_0)$ [if $(I_{45} < I_0)$ $\phi = \phi + 90°$ else $\phi = \phi - 90°$]. (9)

Intensity

$$I = I_0 + I_{90}. \qquad (10)$$

Degree of polarization

$$\rho = \frac{I_{90} - I_0}{(I_{90} + I_0)\cos 2\phi}. \qquad (11)$$

Results presented in this paper utilize these equations to estimate the phase and the degree of polarization. We have we also applied the Levenberg–Marquardt nonlinear curve fitting algorithm to recover these quantities from images taken with the polarizer oriented at $10°$ intervals. This was obviously more time consuming but much less affected by noise.

C. Shape From Specular Polarization

Following the dichromatic reflectance model [31], the reflected light is a superposition of specular and diffuse components. Specular reflection is a result of direct surface reflection [21]. As Fig. 2(a) shows, the reflected light is attenuated to a greater extent if it is polarized parallel to the plane of incidence. Thus, greatest transmission through the polarizer occurs when the polarizer is oriented at an angle $90°$ from the azimuth angle of the surface α. Throughout this work, we assume that the image is formed by orthographic projection.

The zenith angle can be computed by considering the degree of polarization. Substituting (6) into (7) gives the degree of specular polarization in terms of the Fresnel coefficients

$$\rho_s = \frac{R_\perp(n, \theta_i) - R_\parallel(n, \theta_i)}{R_\perp(n, \theta_i) + R_\parallel(n, \theta_i)}. \qquad (12)$$

Using (12) with the Fresnel equations (2) and (3) gives ρ_s in terms of n and the zenith angle θ

$$\rho_s = \frac{2\sin^2\theta\cos\theta\sqrt{n^2 - \sin^2\theta}}{n^2 - \sin^2\theta - n^2\sin^2\theta + 2\sin^4\theta}. \qquad (13)$$

This equation has two real solutions for θ as shown in Fig. 4(a). The existence of two solutions means that there is another ambiguity that must be solved. Miyazaki *et al.* [7] solve this problem using two views of the object under study.
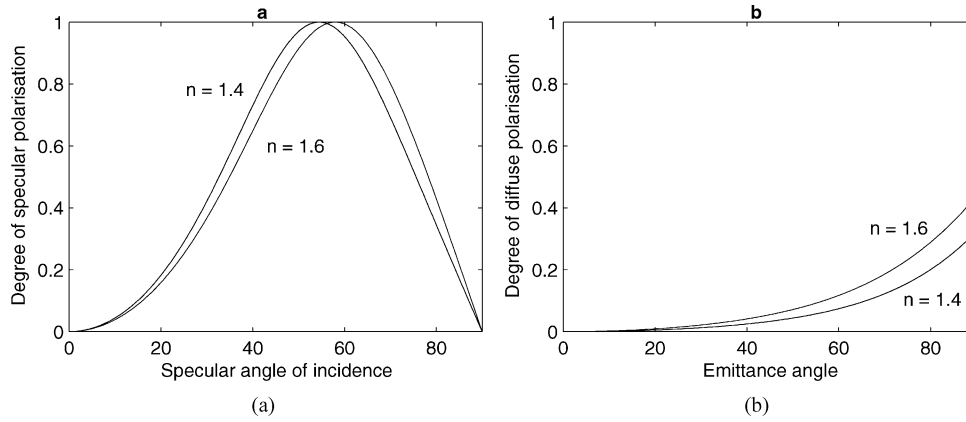
Fig. 4. Degree of polarization for (a) specular and (b) diffuse reflection for two different refractive indices. Many opaque dielectrics have refractive indices between these two values. Note that at the Brewster angle for specular reflection $\rho_s = 1$ since the light is totally polarized. In practice, ρ_s will be a little less than 1 due to a small, but finite, diffuse component of reflection.

Unfortunately, the refractive index is not generally known, but, for most dielectrics, it falls between 1.3 and 1.6, and the dependence of ρ_s on n is weak. Thus, with a known value or estimate of n, and with ρ_s measured using (7), the zenith angle θ can be determined up to the ambiguity. This equation is only applicable to specular reflection since, as explained below, the process that causes diffuse polarization is different. The components of the surface normal vectors can be obtained using

$$\begin{pmatrix} p_x \\ p_y \\ p_z \end{pmatrix} = \begin{pmatrix} \cos\alpha\sin\theta \\ \sin\alpha\sin\theta \\ \cos\theta \end{pmatrix} \qquad (14)$$

where α takes values of either α_1 or α_2 (see Fig. 3).

D. Shape From Diffuse Polarization

Diffuse polarization is a result of the following process [2], [21]: A portion of the incident light penetrates the surface, is partially polarized in the process, as predicted by the Fresnel equations, and is refracted. Due to the random nature of internal scattering, the light becomes depolarized. Some of the light is then refracted back into the air and is, once again, refracted and partially polarized.

When light approaches the surface-air interface from within the medium (after penetration and internal scattering), as shown in Fig. 5, a similar process to that discussed earlier takes place but with the relative index of refraction being $1/n$ instead of n (assuming refractive index of air $= 1$). If the internal angle of incidence is above a critical angle ($\arcsin 1/n$), then total internal reflection occurs. Otherwise, Snell's Law (4) can be used to find the angle of emittance for any given angle of internal incidence. The Fresnel transmission coefficient can then be calculated for a given *emittance* angle. Fig. 6 shows the result of this calculation for a typical dielectric with an additional factor of $1/n$ introduced due to a difference in wave impedance. It should be pointed out that diffuse reflection also results from multiple scattering from microfacets, to which this theory clearly cannot be applied, although this extra contribution is small for smooth surfaces.
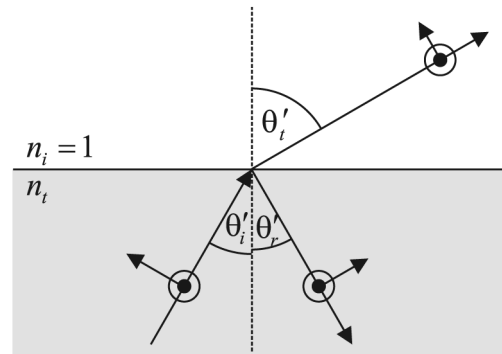

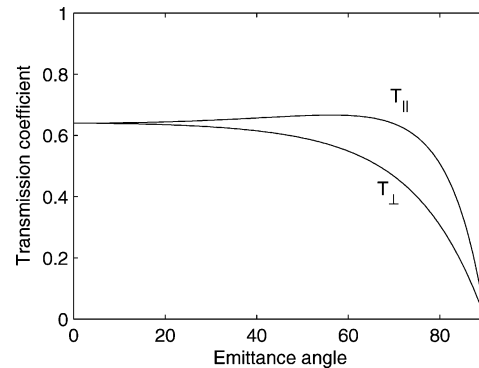Fig. 5. Transmission of internally scattered light back into air.

Fig. 6. Fresnel coefficients for Fig. 5 ($n = 1.7$).

Using (7), the degree of diffuse polarization is

$$\rho_d = \frac{T_\parallel\left(1/n, \theta_i'\right) - T_\perp\left(1/n, \theta_i'\right)}{T_\parallel\left(1/n, \theta_i'\right) + T_\perp\left(1/n, \theta_i'\right)} \qquad (15)$$

$$= \frac{R_\perp\left(1/n, \theta_i'\right) - R_\parallel\left(1/n, \theta_i'\right)}{2 - R_\perp\left(1/n, \theta_i'\right) - R_\parallel\left(1/n, \theta_i'\right)}. \qquad (16)$$

Snell's Law (4) can be used to interchange between the internal angle of incidence θ_i' and the more useful angle of emittance θ_t'. When the surface is viewed from this angle of emittance, θ_t' is the zenith angle, which, from here on, shall be referred to simply as θ. The relevant Fresnel equations (2) and (3) can be substituted into (16) to obtain ρ_d in terms of the refractive
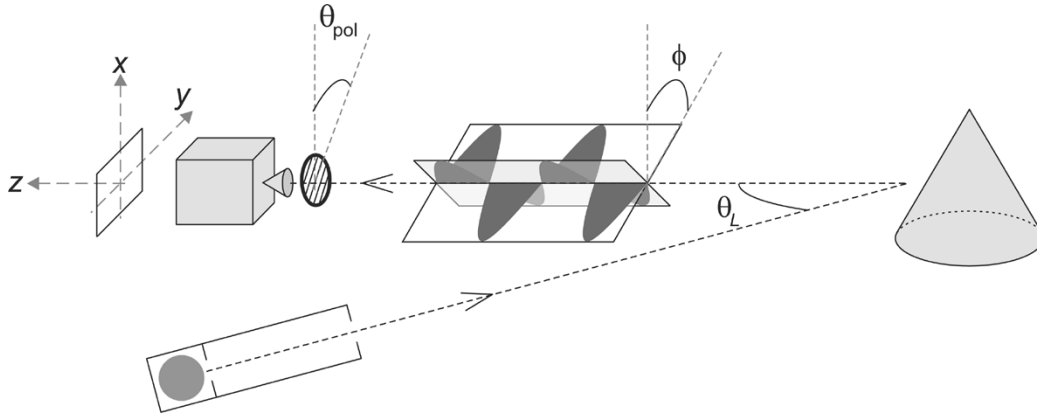
Fig. 7. Experimental setup. The light source was a 200-W tungsten lamp collimated so that reflections from the environment onto the object were minimized. The sine wave of greater amplitude illustrates the component of the electromagnetic wave in the direction of polarization, i.e., is at an angle ϕ to the vertical. The other sine wave is the component 90° to this. The axes on the left define the orthographic image plane.

index and the surface zenith angle. The resulting equation for the degree of diffuse polarization is

$$\rho_d = \frac{(n - 1/n)^2 \sin^2 \theta}{2 + 2n^2 - (n + 1/n)^2 \sin^2 \theta + 4 \cos \theta \sqrt{n^2 - \sin^2 \theta}}. \quad (17)$$

The dependence of the diffuse polarization ρ_d on the zenith angle θ is shown in Fig. 4(b).

Using the experimental setup described in the Section III, the zenith angle of the normal can be calculated using (17), which is the central equation for this paper. The azimuth angle of the normal can be determined using the same method as that used for specular reflection, except that a phase shift of 90° is necessary. The need for a phase shift is illustrated by Fig. 6, which shows that light polarized parallel to the plane of incidence has the highest transmission coefficient and so greater intensity is observed with a polarizer at that orientation. This is in contrast to specular reflection (Fig. 2). The surface normals can then be calculated using (14), where α is either ϕ or $\phi \pm 180°$.

Comparing the plots of diffuse and specular polarization in Fig. 4, it is clear that there is slightly stronger dependence of the degree of polarization on refractive index for diffuse reflection than for specular reflection. Moreover, in the diffuse case, the degree of polarization is lower and, thus, more difficult to measure. On the other hand, the graphs show no zenith angle ambiguity for diffuse reflection. Another advantage of using diffuse polarization for shape recovery is the fact that less controlled lighting conditions are required than when specular polarization is used where a global specularity is needed. One final important advantage is that, since in shape from diffuse polarization, we can assume complete depolarization of the incident light after the light penetrates the surface; the technique remains valid if the incident light is completely or partially polarized. This is not the case for shape from specular polarization, where any polarized component of incident light will severely distort results.

III. ADOPTED METHOD

The process we used to recover surface height can be summarized as follows.

1) Image acquisition using the arrangement shown in Fig. 7.

2) Fitting the data to the transmitted radiance sinusoid (5) to obtain the phase and the degree of polarization at each point.
3) Calculation of the initial surface normal vector estimates using (14) and (17).
4) Disambiguation of the azimuth vectors.
5) Integration of the surface normal vectors to recover surface height.

As (9)–(11) show, three images, obtained with the polarizer oriented at 0°, 45°, and 90° are sufficient to obtain the phase and degree of polarization of a scene. We, therefore, took images of smooth porcelain objects, with a standard linear polarizer placed in front of a Nikon D70 digital SLR camera at these angles. Fig. 7 shows the experimental setup used and defines the coordinate axes.

In principle, the results should not be critically dependent upon lighting conditions, as it is only the orientation of the surface and its index of refraction that determine the polarization state of the reflected light. To simplify matters, however, only one light source was used, which was placed close to the camera ($\theta_L \approx 0$) and collimated, as Fig. 7 shows. This means that specularities occurred only at points on the surface where the zenith angle of the surface was zero. Any nonuniformity in illumination does not deteriorate the obtained results since pixels are (initially) point processed. The camera's aperture was fixed at $f5.6$, with exposures of 0.25 s. The model [2] assumes that light that has penetrated the surface becomes completely depolarized so that a small polarized component in incident light should not distort the results.

Although Fig. 7 shows the two components of the reflected light to be in phase, this is not generally true. However, since only the amplitude is needed for shape recovery, this does not complicate matters. The walls of the room and the table on which the objects lay were black so that pixels of the images having a brightness below a certain threshold could be treated as background.

The degree of polarization is converted to zenith angle by numerically solving (17). It is assumed throughout this work that the refractive index of the reflecting medium is 1.6, except in the uncertainty analysis. This is actually an overestimate, but the quality of the results demonstrate the fact that a precise estimate of n is not required. Next, (14) was used to calculate the Cartesian components of the surface normal vectors at this stage, as-

suming that the azimuth angle α, corresponds to ϕ (Fig. 3). After normalizing the vectors so that they were of unit length, the following disambiguation routine was applied. The routine was designed with efficiency in mind so that the general shape of the object can be rapidly recovered. More accurate methods will be the focus of future work.

In the first stage of disambiguation, the pixels of the image are ordered according to the zenith angle with the greatest angles listed first. We know that the vectors at the occluding contours must be pointing away from the body of the object, so we rotate the azimuth vectors by 180° where necessary in order to meet this requirement. Next, using the rank order of the pixels previously mentioned, we propagate into the object rotating azimuth angles by 180° where this helps to preserve object smoothness. Abrupt changes in azimuth angle are, of course, permitted where the zenith angle is small. As we will see later, this method gave good results for simple objects, but variable results far from the limbs of complex objects due, in part, to insoluble convex/concave ambiguities. Because shape from diffuse polarization is very reliable near object boundaries, a good indication of overall object shape is still obtainable, despite the obvious shortcomings of this disambiguation method.

Throughout this work, the camera was setup such that the diffusely reflected light causes pixels to have a brightness that varies across the whole dynamic range of the camera with specularities saturated.[2] In fact, for $\theta_L \approx 0$, this causes the algorithm to return zero polarization and, hence, $\theta = 0$ for such areas, which is correct.

The noise reduction technique that we used involved processing pixels by taking the median brightness over local neighborhoods, with central pixels being counted more than once, giving them an extra "weighting." This was improved by smoothing over larger neighborhoods for areas with a lower degree of polarization. We later implemented an adaptive smoothing algorithm [32], which was used to obtain the results shown in Section IV. We found that, without adaptive smoothing, it was often necessary to take several images at each polarizer angle and use the average image at each angle. When eighteen images were used and the transmitted radiance sinusoid fitted using the Levenberg–Marquardt algorithm, smoothing was not important.

The final step in processing was to recover surface depth from the surface normals. This was performed using the surface integration algorithm reported in [33]. This method uses the changes in surface normal directions to estimate sectional curvature on the surface. An eigenvector, or graph-spectral, method is used to locate a curvature minimizing path through the field of surface normals. By traversing this path and using the estimated surface orientation, simple trigonometry is used to compute the height offset from a reference level.

IV. EXPERIMENTAL RESULTS

This section describes a set of experiments conducted with the aim of illustrating both the possibilities and limitations of the method described above. First, we present the results of applying the technique to a variety of objects with complex geometries. Second, a discussion of some of the difficulties with shape recovery is given. We then focus on cylindrical objects, with the aim to assess the accuracy of the measurements for various materials. Finally, experiments were performed to measure refractive indices. These are described in Section IV-D, and a comparison of these results to those expected is made.

A. Shape Recovery

A representative set of experimental results are shown in Fig. 8. The top row shows the greyscale images of the objects studied. The first three objects being shiny white porcelain and the fourth object being slightly rough yellow plastic. The second row shows the phase angle of the reflected light. Note that this has been successfully recovered for the plastic duck. This is due to the fact that the phase is not affected by isotropic roughness (except that noise is greater for rougher surfaces). The measured degree of polarization is shown in the third row. Note that the effects of roughness here are more noticeable as the degree of polarization appears slightly lower for the rough plastic at limbs. The recovered fields of surface normals, or needle maps, are shown in the bottom row. Fig. 9 shows the result of applying the needle map integration algorithm to the field of surface normals. It is clear that the basic shapes have been recovered, but with noise being problematic at small zenith angles.

B. Limitations

There are three main limitations of the technique described in this paper. The first of these is a common problem in single view computer vision, namely the convex/concave ambiguity. This can be solved either by knowing information about the illumination conditions and combining the method with shape-from-shading or by using multiview algorithms.

The second limitation is a result of surface roughness. This can be explained by assuming that the surface is composed of planar microfacets of random orientation as described by Torrance and Sparrow [34]. Although this is the only model we use here to account for the effect of roughness, other models have equivalent consequences. Unlike the ideal case shown in Fig. 5, some of the light transmitted from the medium into air toward the camera would undergo reflections from the microfacets that constitute the roughness. Assuming that individual surface microfacets are not resolved by the camera, the result of a rough surface is a reduction in the measured degree of polarization. This, in turn, leads to underestimates of the zenith angles. However, assuming that the distribution of surface slopes is such that more microfacets are parallel to the mean surface than any other orientation (for example, the distribution of microfacet slopes about the local average may be Gaussian with zero mean), then the phase angle of the light should still match the surface azimuth angle up to the ambiguity. A final consequence of roughness is that specularities are broadened by the microfacet angular distribution.

Unlike the three other objects in Fig. 8, the surface of plastic duck [Fig. 8(d)] is slightly rough. This is important since it shows that the phase angle can indeed be accurately measured

---

[2]It is straightforward to set a camera so that single-reflection specularities from a shiny material are saturated. This is not true for rougher surfaces and interreflections are not generally saturated.
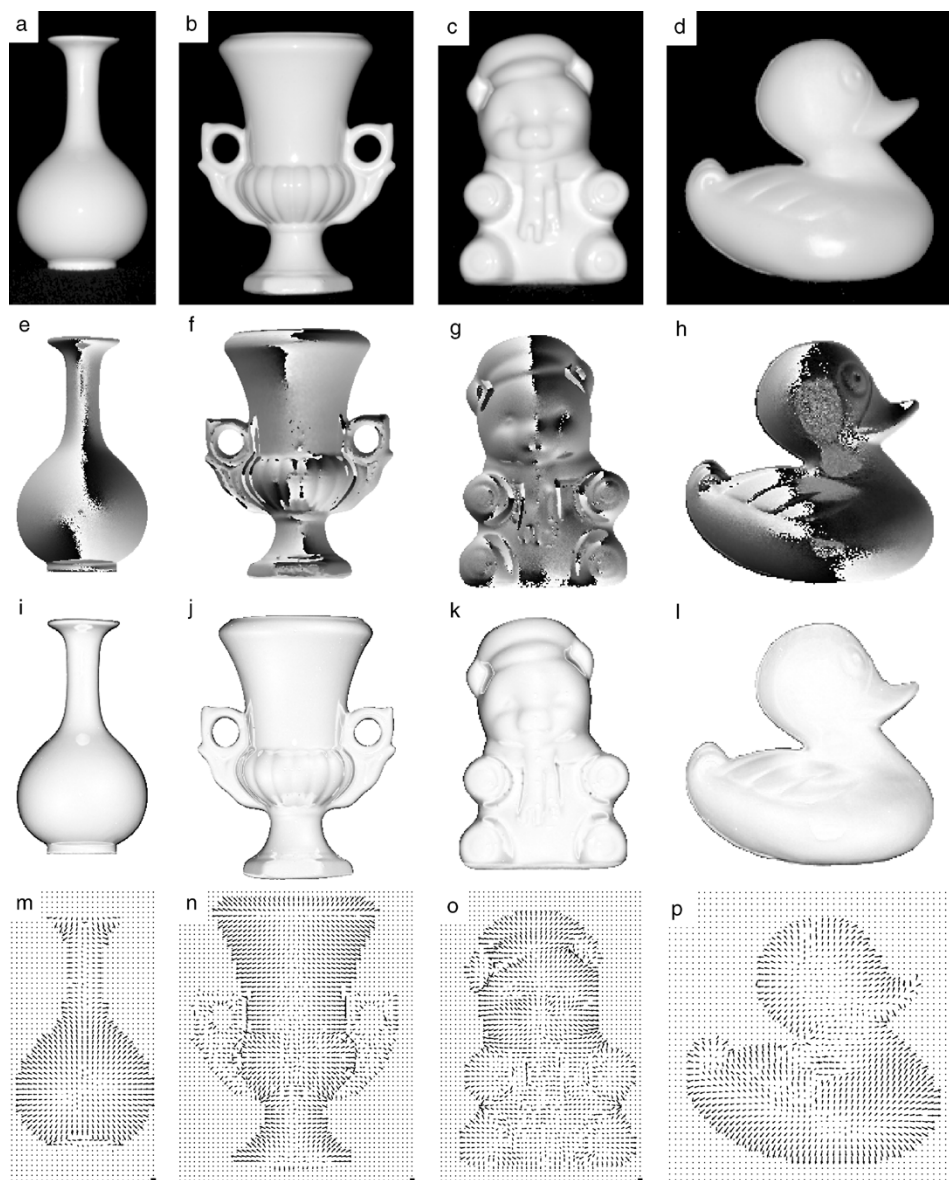
Fig. 8. Greyscale images of smooth porcelain (a) vase, (b) urn, (c) bear, and (d) slightly rough plastic duck. The images are approximately 350 pixels in height. (e)–(h) Phase images showing angle of polarization. (i)–(l) Degree of polarization, dark areas have highest polarization. (m)–(p) Normal vectors (reduced resolution).

in such circumstances. The recovered depth in Fig. 9(d) shows that the zenith angle is at least approximately correct.

The third and final limitation of the method is that caused by specularities (either direct reflections from a light source or interreflections). It may be possible to isolate specularities using abrupt $90°$ phase shifts and the fact that if $\rho > \rho_d\,(\theta = 90°)$, where ρ is the measured degree of polarization, then the reflection must be specular (Fig. 4). Note, however, that the underlying diffuse reflection would reduce the polarization from that expected for pure specular reflection. The intensity image also provides clues as to where interreflections are since the intensity is less affected by polarization.

In many situations, there are only a small number of different light sources and each source subtends only a small angle from the object. In such situations, which include those under which the experiments described here were conducted, direct specularities may be ignored, or if they cause saturation, interpolated across.

Fig. 10(a) shows a closeup of the handle of the urn, and Fig. 10(b) shows a closeup of the bear's hand [center right of Fig. 8(o)]. Both images demonstrate success at recovering some details near object limbs. They also illustrate a problem with shape from diffuse polarization caused by interreflections such as those shown in Fig. 11. To the right of the urn handle and to the left of the bear's hand, the estimated surface azimuth angles are $90°$ from the true azimuth angles and the zenith angles are overestimated. This is explained by the different processes involved in specular and diffuse reflection described earlier.

C. Evaluation of Accuracy

To assess the accuracy of the method, a set of vertically oriented cylinders of various materials were used. The geometry of a cylinder is convenient for three reasons. First, the structure is simple enough for shape recovery to be performed exactly from
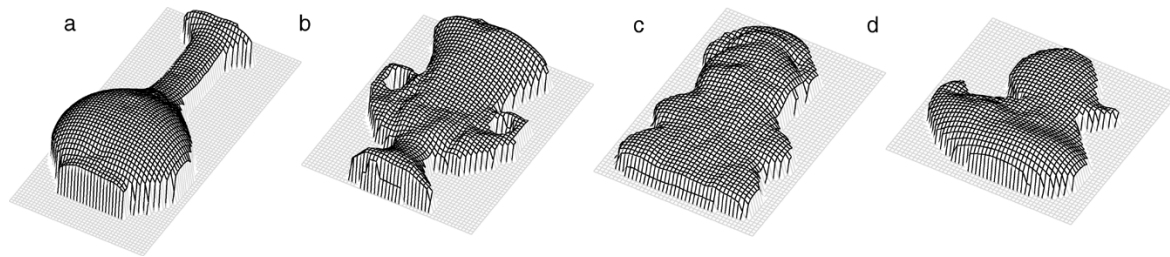
Fig. 9. Depth maps of porcelain vase, urn, bear, and plastic duck computed by applying adaptive smoothing and the algorithm of [33] to the needle maps and rotating the objects to different angles.
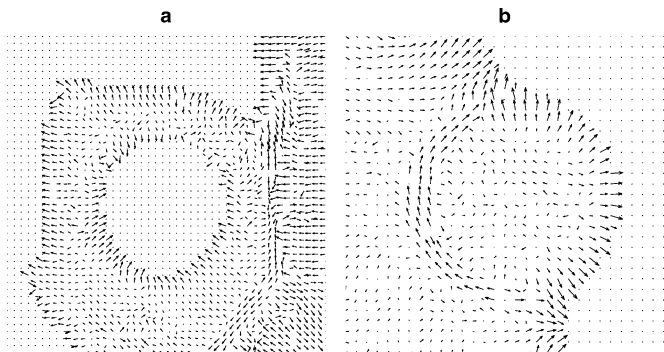

Fig. 10. Closeups of Fig. 8(n) and (o).

a single image.[3] Second, the analysis can easily be performed for all possible zenith angles. Third, noise can be reduced by taking the average image intensity for each column of pixels. To ensure orthogonality between the optical and cylinder axes, spirit levels were employed. The camera was set to maximum resolution and placed far away from the object. When the top and bottom of the cylinder appeared horizontal in the images (i.e., not curved), then the ground truth was accurate.

Fitting (5) to the polarization images as before, we obtained a set of graphs showing the measured and theoretical zenith angle against position across the cylinder for different materials. Since the azimuth angle of the cylinder is constant, we can also see how the accuracy of azimuth angle estimates vary with zenith angle, if at all. A sample of these results for porcelain, blank photographic paper, and photographic paper coated with cling film and normal paper are shown in Fig. 12. Here, the solid lines show the measurements and the broken lines show the ground truth. The photographic paper is much smoother than normal paper, as it is coated to produce glossy photographs. The normal paper, on the other hand, is matte and much rougher. Several other material samples were also analyzed, including different paper types, plastics, wax, terracotta, and papers coated with inks. The first point to note about the figures is that, even for normal paper, which, at the fine scale, is very rough, the azimuth angles have been recovered to a very high degree of accuracy. However, more noise is associated with the rougher surfaces.

There was more variation in the accuracy of the zenith angle estimates. For Fig. 12, the refractive index used was simply the value that produced greatest similarity between theory and experiment for the material in question. The shiny white porcelain object [the same type of porcelain can be seen in Fig. 8(a)–(c)] produced excellent agreement with theory down

to very small zenith angles. The similarity between the exact and experimental curves supports the theory, within the limits of its assumptions.

The remaining graphs in Fig. 12 demonstrate the complications that can cause the measured zenith angles to deviate from the expected values. The result for blank white photographic paper, for example, is very accurate for large zenith angles but an increasing discrepancy is present as the zenith angle approaches zero. When the paper is coated in cling film, the discrepancy is less marked. Clearly, this suggests that there is a process occurring that is not accounted for by the theory and may be due to microfacet interreflections or imperfect depolarization of incident light upon surface penetration. Although it is useful to be aware of this, it is not considered useful to investigate this phenomenon further for two reasons. First, the effect is only present for regions of small zenith angle, where the technique is of limited use anyway. Second, for a camera with 256 grey levels, the intensity may vary by just a few grey levels in such regions. Therefore, intensity quantization errors prevent extraction of useful data. The results for paper, which, of course, is a rough matte surface, also show the phenomenon of finite polarization at low zenith angles, as well as the depolarizing effects of roughness nearer to the limbs.

For comparison, Fig. 13 shows the zenith angle prediction using the Lambertian reflectance model. Clearly, on this occasion, the polarization analysis gave much better results. Of course, for materials such as paper (a classic example of a Lambertian surface), the Lambertian prediction is better due to roughness.

An analysis of the dyed photographic paper shown in Fig. 14 was also conducted. Fig. 15 shows the zenith angles recovered from a cylinder of photographic paper coated in ink of different colors (red, yellow, blue, grey, and green). It is obvious from the graphs where the transitions between colors occur. It appears that darker inks cause greatest variation, suggesting that good shape recovery would be possible if inks of similar albedo were present in an image.

D. Refractive Index Measurements

As Fig. 4 shows, there is greater dependence of degree of polarization on refractive index for diffuse polarization than specular polarization. However, as results in this paper show, it is possible to obtain reasonable height reconstruction simply by using a fixed value of the refractive index, taken to be about 1.4–1.6. Clearly, of course, this is inadequate for precision applications. This method is also inadequate for assessing the accuracy of the theory for different materials as in this work.
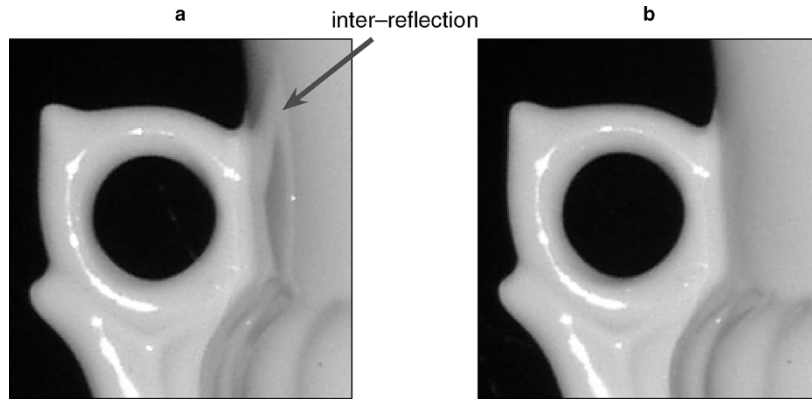
---

[3]This is simply done by isolating the cylinder from the background and placing semicircles that arch from one side of the object to the other.

Fig. 11. Handle of porcelain urn for (a) $\theta_{\text{pol}} = 0°$ and (b) $\theta_{\text{pol}} = 90°$. An interreflection of the handle is clearly visible in (a) as, here, the polarizer is oriented parallel to the direction of polarization of the interreflected light for that part of the urn. In (b), the polarizer is at 90° to this angle so that the interreflected light is not transmitted through the polarizer.
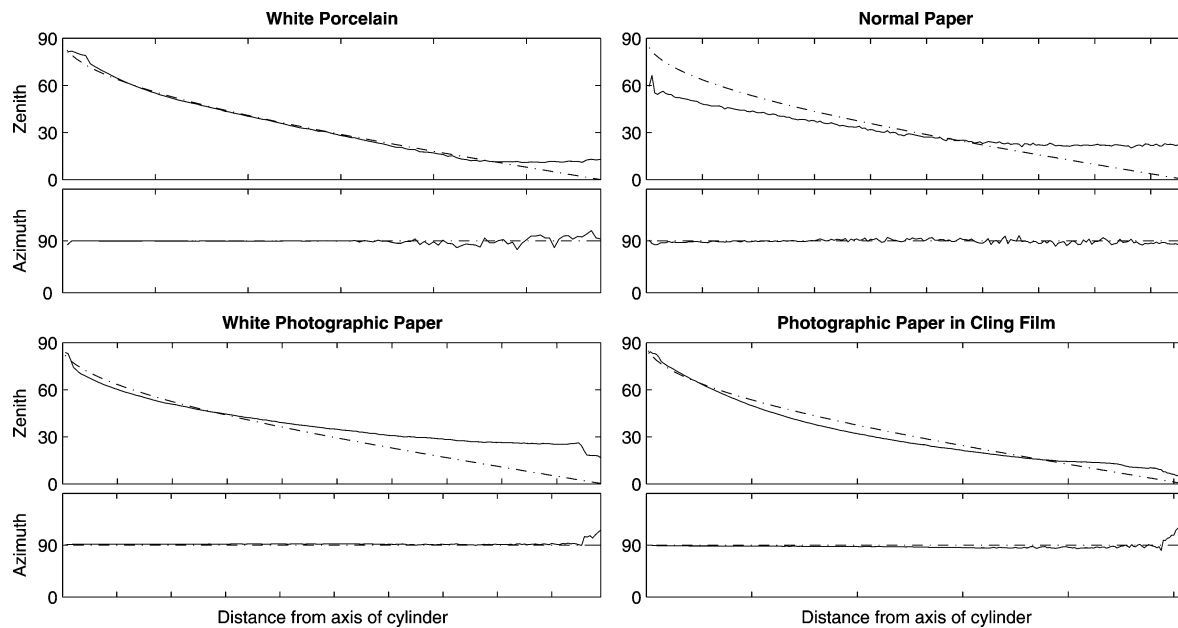

Fig. 12. (Solid lines) Plots of measured zenith and azimuth angles across the surfaces of cylinders of different materials. The exact values are indicated by the broken curves.
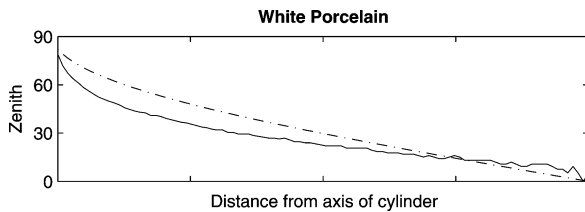

Fig. 13. Estimation of zenith angles using the (solid line) Lambertian reflectance model compared to ground truth.

Fig. 15. (Solid line) Measured and (broken line) exact zenith angles across the surface of a cylinder coated in different colors of ink (shown in Fig. 14). It should be stressed that using the correct refractive index for each ink provides much better agreement.
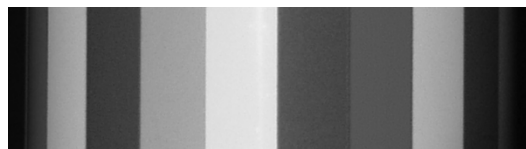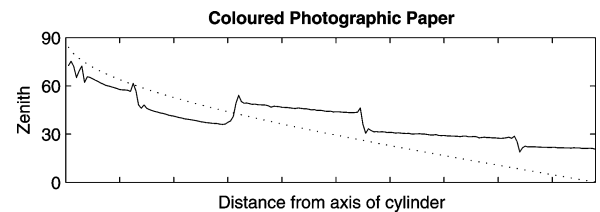

Fig. 14. Dyed photographic paper wrapped around a cylinder used to test shape recovery of different colored inks. The left-hand half of this image was used to obtain Fig. 15.

In order to measure the index of refraction, the device shown in Fig. 16 was constructed. The device works by placing a sample of the material under study on a table that can be rotated by small angles (measured to within 1° if the accuracy of the

refractive index is to be within 0.1 of the true value). Laser light is then directed at a plane, vertical surface of the sample, which then reflects the light through a polarizer with horizontal transmission axis onto a white screen. In fact, the laser light was partially polarized and so was oriented such that the angle of polarization was horizontal (i.e., so that the smaller component was absorbed by the horizontal linear polarizer). The sample is then rotated manually until the spot of light on the screen formed by the reflected laser light is extinguished. Conducting the experiment in a dark room means that the extinguishment
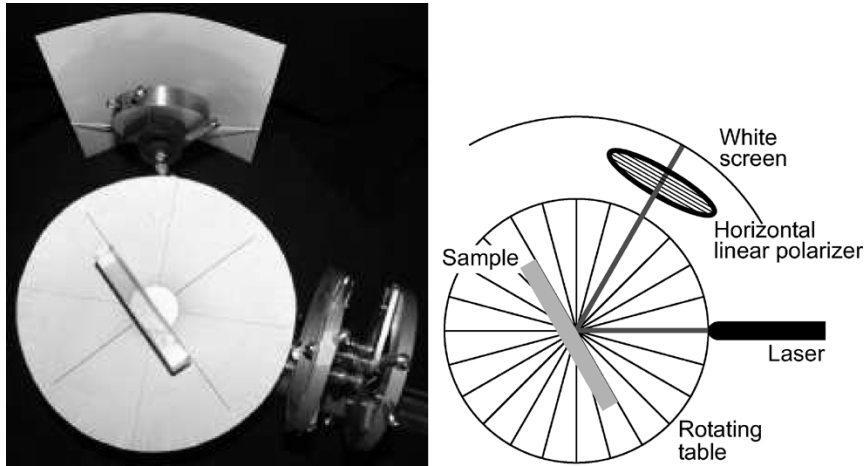
Fig. 16. Refractometer. As the sample is rotated, the intensity of the spot of light on the screen changes according to (3). The refractive index can be found by orienting the sample such that the spot disappears.

TABLE I
MEASURED REFRACTIVE INDICES

Material	n_{ref}	n_{exp}	Difference
Photographic paper	1.30	1.30	0.00
Photographic paper in cling film	1.54	1.70	-0.16
Photographic paper in thicker clear film	1.54	1.70	-0.16
Smooth white porcelain	1.28	1.30	-0.02
Smooth white plastic	1.28	1.28	0.00
Slightly rough orange plastic	1.54	1.25	0.29
Slightly rough wax	1.48	1.23	0.25
Cello tape	1.43	1.30	0.13
Grey ink	1.35	1.35	0.00
Blue ink	1.37	1.45	-0.08
Yellow ink	1.30	1.32	-0.02
Red ink	1.37	1.40	-0.03
Green ink	1.40	1.45	-0.05

can easily be detected by the naked eye. This occurs at the Brewster angle where there is no horizontal component of the electric field of the reflected ray, and the vertical component is absorbed by the polarizer. The Brewster angle θ_B is simply related to the index of refraction by $n = \tan \theta_B$. Calculations show that, for smooth surfaces of average refractive index, the uncertainty in refractive index measurements is 4%. This is increased for rougher surfaces, since the laser light is made diffuse by the reflection. Variation in refractive index by this small amount has little effect on the relationship between zenith angle and degree of polarization.

Table I gives a comparison between measured refractive indices and the values that produced the best agreement with experiment. n_{ref} denotes measurements taken from the refractometer. n_{exp} refers to the values that give best agreement with the experimental curves. Some of the roughest materials dispersed the laser light too much to allow reliable measurements to be made, and so are not included. Results for all the smooth, opaque materials, and all but one of the inks, fell within experimental error. As one might expect, given the discussion in the previous section, rougher materials, translucent materials, and materials coated in transparent media caused greater differences between the two values.

It should be noted that the refractive index should not be measured using infrared or any other nonvisible light, since it is a function of wavelength (we assume this dependence is negligible within the range of visible wavelengths). A comparison of refractive indices measured with the refractometer and those estimated from the simple method described in Section IV-C shows that, for smooth, opaque materials, agreement is good, with discrepancies of less than 0.05. Unfortunately, for items causing more complicated reflections, such as rough, translucent, and coated materials, the estimates from the two methods were less well matched. Nevertheless, as results in Section IV-A demonstrate, the *exact* value of refractive index is not a requirement for shape recovery. Values measured using the refractometer for smooth materials should be sufficiently accurate, therefore, for reasonable shape recovery.

V. CONCLUSION

This study of shape from diffuse polarization has resulted in several accomplishments. First, it has shown that there is potential for using diffuse polarization in the recovery of surface shape, especially near occluding contours, despite the common assumption that the polarization is negligible for diffuse reflection. This is achieved with no knowledge of illumination direction and the only assumption being that all reflections are diffuse. Second, however, it illustrates the fact that, away from the occluding contours of complex objects, diffuse polarization is limited in applicability. Diffuse polarization is, perhaps, best used, therefore, in conjunction with other techniques such as shape from shading, shape from specular polarization, or multiple view vision.

The difference in accuracy of the method between regions of high and low zenith angles is clearly illustrated by Fig. 12 in the error analysis, which also provides a detailed picture of the effects of roughness. Importantly, we see that the surface normal *azimuth* angles can be *accurately* determined even for moderately rough surfaces. The refractive index measurements provide a link between the underlying theory and the experimental results, helping to validate the theory.

There is potential for future work with diffuse polarization from rough surfaces. The zenith angles are more difficult to obtain since the random microfacet orientation has a depolarizing effect. This means that either more information on the microscopic surface structure would be needed, or an empirical lookup table generated from images of an object of known shape. Due to the lower degree of polarization, noise would be more problematic, especially

away from occluding contours. For metals [2] and transparent materials [7], the diffuse component of reflection is almost zero so specular reflection must be used.

Another avenue for future research involves multiview methods. Two of the main problems of conventional stereo are the self occlusion at limbs, where a portion of the object is only visible in one of the two views, and difficulties associated with finding correspondences. Shape from diffuse polarization is most reliable near object limbs. Hence, by taking multiple views as the object is rotated in front of the camera, a full surface can be reconstructed by zippering the boundary regions together.

REFERENCES

[1] H. Ragheb and E. R. Hancock, "A probabilistic framework for specular shape from shading," *Pattern Recognit.*, vol. 36, pp. 407–427, 2003.

[2] L. B. Wolff and T. E. Boult, "Constraining object features using a polarization reflectance model," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 7, pp. 635–657, Jul. 1991.

[3] E. Hecht, *Optics*, 3rd ed. Reading, MA: Addison Wesley, 1998.

[4] D. Miyazaki, R. T. Tan, K. Hara, and K. Ikeuchi, "Polarization-based inverse rendering from a single view," in *Proc. Int. Conf. Computer Vision*, vol. 2, 2003, pp. 982–987.

[5] L. B. Wolff, "Polarization vision: A new sensory approach to image understanding," *Image Vis. Comput.*, vol. 15, pp. 81–93, 1997.

[6] M. Saito, Y. Sato, K. Ikeuchi, and H. Kashiwagi, "Measurement of surface orientations of transparent objects using polarization in highlight," in *Proc. Computer Vision Pattern Recognition*, vol. 1, 1999, pp. 381–386.

[7] D. Miyazaki, M. Kagesawa, and K. Ikeuchi, "Transparent surface modeling from a pair of polarization images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 1, pp. 73–82, Jan. 2004.

[8] R. T. Tan and K. Ikeuchi, "Separating reflection components of textured surfaces using a single image," in *Proc. Int. Conf. Computer Vision*, vol. 2, 2003, pp. 870–877.

[9] S. K. Nayar, X. Fang, and T. Boult, "Separation of reflectance components using color and polarization," *Int. J. Comput. Vis.*, vol. 21, pp. 163–186, 1997.

[10] S. Umeyama, "Separation of diffuse and specular components of surface reflection by use of polarization and statistical analysis of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 26, no. 5, pp. 639–647, May 2004.

[11] D. Miyazaki, M. Saito, Y. Sato, and K. Ikeuchi, "Determining surface orientations of transparent objects based on polarization degrees in visible and infrared wavelengths," *J. Opt. Soc. Amer. A*, vol. 19, pp. 687–694, 2002.

[12] S. Rahmann and N. Canterakis, "Reconstruction of specular surfaces using polarization imaging," in *Proc. Computer Vision Pattern Recognition*, 2001, pp. 149–155.

[13] O. Drbohlav and R. Šára, "Unambiguous determination of shape from photometric stereo with unknown light sources," in *Proc. Int. Conf. Computer Vision*, 2001, pp. 581–586.

[14] ——, "Specularities reduce ambiguity of uncalibrated photometric stereo," in *Proc. ECCV*, vol. 2, 2002, pp. 46–62.

[15] J. Clark, E. Trucco, and L. Wolff, "Using light polarization in laser scanning," *Image Vis. Comput.*, vol. 15, pp. 107–117, 1997.

[16] O. Drbohlav and R. Šára, "Using polarization to determine intrinsic surface properties," in *Proc. SPIE Conf. Polarization and Color Techniques in Industrial Inspection*, 1999, pp. 253–263.

[17] B. K. P. Horn and M. J. Brooks, *Shape from Shading*. Cambridge, MA: MIT Press, 1989.

[18] R. Zhang, P. Tsai, J. Cryer, and M. Shah, "Shape from shading: A survey," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 21, no. 6, pp. 690–706, Jun. 1999.

[19] J. Oliensis and P. Dupius, "A global algorithm for shape from shading," in *Proc. Int. Conf. Computer Vision*, 1993, pp. 690–706.

[20] L. B. Wolff, "Diffuse-reflectance model for smooth dielectric surfaces," *J. Opt. Soc. Amer. A*, vol. 11, pp. 2956–2968, 1994.

[21] L. B. Wolff, S. K. Nayar, and M. Oren, "Improved diffuse reflection models for computer vision," *Int. J. Comput. Vis.*, vol. 30, pp. 55–71, 1998.

[22] R. J. Woodham, "Photometric method for determining surface orientation from multiple images," *Opt. Eng.*, vol. 19, pp. 139–144, 1980.

[23] A. L. Yuille, D. Snow, R. Epstein, and P. Belhumeur, "Determining generative models for objects under varying illumination: Shape and albedo from multiple images using SVD and integrability," *Int. J. Comput. Vis.*, vol. 35, pp. 203–222, 1999.

[24] P. Belhumeur and D. Kriegman, "What is the set of images of an object under all possible lighting conditions?," *Int. J. Comput. Vis.*, vol. 28, pp. 245–260, 1998.

[25] T. Zickler, P. N. Belhumeur, and D. J. Kriegman, "Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction," *Int. J. Comput. Vis.*, vol. 49, pp. 215–227, 2002.

[26] W. T. Doyle, "Scattering approach to Fresnel's equations and Brewster's law," *Amer. J. Phys.*, vol. 53, pp. 463–468, 1985.

[27] M. Born and E. Wolf, *Principles of Optics. Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. London, U.K.: Pergamon, 1959.

[28] R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics," *ACM Trans. Comput. Graph.*, vol. 1, pp. 7–24, 1982.

[29] S. Rahmann, "Polarization images: A geometric interpretation for shape analysis," in *Proc. ICPR*, vol. 3, 2000, pp. 542–546.

[30] R. T. Frankot and R. Chellappa, "A method for enforcing integrability in shape from shading algorithms," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 10, no. 4, pp. 439–451, Jul. 1988.

[31] S. Shafer, "Using color to separate reflection components," *Color Res. Appl.*, vol. 10, pp. 210–218, 1985.

[32] C. Kervrann, "An adaptive window approach for image smoothing and structures preserving," in *Proc. ECCV*, vol. 3, 2004, pp. 132–144.

[33] A. Robles-Kelly and E. R. Hancock, "A graph-spectral method for surface height recovery from needle-maps," in *Proc. Computer Vision Pattern Recognition*, 2001, pp. 141–148.

[34] K. Torrance and M. Sparrow, "Theory for off-specular reflection from roughened surfaces," *J. Opt. Soc. Amer.*, vol. 57, pp. 1105–1114, 1967.

Gary A. Atkinson (S'05) received the M.Sci. degree in physics from the University of Nottingham, Nottingham, U.K., in 2003. His final-year project was on noise reduction in audio signals. He is currently pursuing the Ph.D. degree in the Department of Computer Science, University of York, York, U.K., under the supervision of Prof. E. R. Hancock.

His research is concerned with improving shape recovery algorithms for computer vision, with a particular emphasis on physical optics. Primarily, he has investigated the polarization of light and its exploitation in surface analysis algorithms.

Mr. Atkinson is a member of the British Machine Vision Association and the Institute of Physics.

Edwin R. Hancock received the degree in physics and the Ph.D. degree in high-energy physics from the University of Durham, Durham, U.K., in 1977 and 1981, respectively.

For ten years, he was a Researcher in the fields of high-energy nuclear physics and pattern recognition at the Rutherford-Appleton Laboratory (now the Central Research Laboratory of the Research Councils), Cheshire, U.K. During this period, he also held adjunct teaching posts at the University of Surrey, Surrey, U.K., and the Open University, Milton Keynes, U.K. In 1991, he joined the University of York, York, U.K., as a Lecturer in the Department of Computer Science. He was promoted to Senior Lecturer in 1997 and to Reader in 1998. In 1998, he was appointed to a Chair in Computer Vision. He now leads a group of some 15 faculty, research staff, and Ph.D. degree students working in the areas of computer vision and pattern recognition. He has been a Guest Editor for special issues of the journals *Image and Vision Computing* and *Pattern Recognition*. He has been a member of the Editorial Board for *Pattern Recognition*. He has published over 80 journal papers and 300 refereed conference publications. He has been on the program committees for numerous national and international meetings. In 1997, with Marcello Pelillo, he established a new series of international meetings on energy minimization methods in computer vision and pattern recognition. His main research interests are in the use of optimization and probabilistic methods for high- and intermediate-level vision. He is also interested in the methodology of structural and statistical pattern recognition. He is currently working on graph matching, shape-from-X, image databases, and statistical learning theory. His work has found applications in areas such as radar terrain analysis, seismic section analysis, remote sensing, and medical imaging.

Prof. Hancock became a Fellow of the International Association for Pattern Recognition in 1998. He has been a member of the Editorial Board of the IEEE TRANSACTIONS ON PATTERN ANALYSIS AND MACHINE INTELLIGENCE. He was awarded the Pattern Recognition Society medal in 1991 and an Outstanding Paper Award in 1997 by the journal *Pattern Recognition*.