

# Robust Shape from Polarisation and Shading

Cong Phuoc Huynh<sup>1</sup> Antonio Robles-Kelly<sup>1,2</sup> Edwin 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 present an approach to robust estimation of shape from single-view multi-spectral polarisation images. The developed technique tackles the problem of recovering the azimuth angle of surface normals robust to image noise and a low degree of polarisation. We note that the linear least-squares estimation results in a considerable phase shift from the ground truth in the presence of noise and weak polarisation in multispectral and hyperspectral imaging. This paper discusses the utility of robust statistics to discount the large error attributed to outliers and noise. Combining this approach with Shape from Shading, we fully recover the surface shape. We demonstrate the effectiveness of the robust estimator compared to the linear least-squares estimator through shape recovery experiments on both synthetic and real images.*

## 1. Introduction

Shape recovery is a common subject of interest to which much of the research effort in the Computer Vision and Graphics communities has been devoted. Shape from Shading focuses on the recovery of shapes from shading information in single images. One of the earliest work in Shape from Shading [3] was based on the compliance with the image irradiance equation and local surface smoothness, under the Lambertian reflectance assumption. However, this treatment leads to the need to resolve the azimuth angle of the object surface normals as it is not constrained by the irradiance equation. As a result, regularisation techniques have been employed to enforce surface curvature consistency [11] or to impose the Lambertian reflectance constraint in a hard manner [10].

Polarisation, on the other hand, adds another dimension to the image analysis at hand. Although the human vision system is oblivious to polarisation, its effects can be captured by devices such as polarimeters and more recently, polarisation cameras [9]. There has also been research aimed at exploiting polarisation information for shape recovery. Miyazaki *et al.* [7] mapped the degree of polarisation to zenith angle, under the assumption that the histogram of zenith angles of a given object is similar to that of a sphere. More recently, Atkinson and Hancock [1] published their work concerning the recovery of surface orientation from diffuse polarisation applied to smooth dielectric surfaces.

This paper focuses on providing a robust estimation of polarisation components from multi-spectral images captured from a single view. Further, it shows how to utilise both the polarisation and shading information for the task of surface orientation recovery. In this paper, we study the diffuse polarisation of light reflection from smooth surfaces made of unknown materials. The recovery method is formulated making use of robust statistics so as to achieve robustness against noise effects and weak degrees of polarisation. As a result, the resulting polarisation phase angles enable the estimation of the azimuth angle, while the shading map gives rise to an estimation of the zenith angle.

## 2. Robust Shape Recovery

In this section, we present a two-step process to recover surface orientation from a set of polarisation images captured from a single view. The first of these involves robust statistics to reliably estimate the surface normal's azimuth angle. In the latter step, we estimate the zenith angle using a technique akin to Shape from Shading.

Light, in general, exhibits polarisation upon reflection from a smooth surface. Let us consider a linear polarisation component propagating in the line-of-sight, which oscillates in a plane oriented at an angle  $\vartheta$  with respect to a reference direction in the image plane. The intensity captured varies sinusoidally with respect to the

---

\*Edwin Hancock is supported by a Royal Society Wolfson Research Merit Award.

†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.

orientation of the plane of polarisation. At image pixel  $u$  and wavelength  $\lambda$ , this variation conforms to a Transmitted Radiance Sinusoid (TRS) as follows

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

where  $\phi$  is the phase angle of the polarised reflected light, and  $I_{min}$  and  $I_{max}$  are, respectively, the minimum and maximum intensities on the sinusoidal curve.

## 2.1. Robust Recovery of the Azimuth Angle

We now note that there exists a relationship between the azimuth angle of surface normals and the phase angle  $\phi$  described in Equation 1. In this paper, we consider a diffuse reflection process from dielectric surfaces. Recall that diffuse polarisation is caused by the refraction of subsurface scattered light through the material-air boundary. Further, the linear polarisation component with the maximal amplitude oscillates in a plane parallel to the plane of reflection according to Snell's law of refraction and Fresnel's reflection theory [1]. In addition, Equation 1 shows that the TRS curve peaks when the angle of polarisation  $\vartheta$  coincides with the phase angle  $\phi$ . Thus, we have the relation  $\alpha(u) = \phi$  or  $\alpha(u) = \phi \pm \pi$  [1]. To disambiguate the azimuth angle between these two possibilities, we can assume convexity on the surface under observation.

Now we turn our attention to estimating the parameters  $I_{min}$ ,  $I_{max}$  and  $\phi$  in Equation 1 from a set of single-view polarisation images captured at  $N$  polariser's orientations. Here, we denote  $\vartheta_i$  as the  $i^{th}$  step-angle of the polariser's transmission axis with respect to a reference direction. We index these images to the wavelength  $\lambda \in \{\lambda_1, \dots, \lambda_K\}$  by denoting  $I_i(u, \lambda) \triangleq I(u, \lambda, \vartheta_i)$  as the measured intensity at the pixel-site  $u$  and wavelength  $\lambda$  on the polarised image corresponding to the  $i^{th}$  polariser angle  $\vartheta_i$ .

We commence by rewriting Equation 1 in the following vector form, using the shorthands  $x = \left[ \frac{I_{max} + I_{min}}{2}, \frac{I_{max} - I_{min}}{2} \cos(2\phi), \frac{I_{max} - I_{min}}{2} \sin(2\phi) \right]^T$  and  $f_i = [1, \cos(2\vartheta_i), \sin(2\vartheta_i)]^T$ .

$$I_i(u, \lambda) = f_i^T x, \forall i = 1, \dots, N \quad (2)$$

Since we have measurements at  $N \geq 3$  polariser's orientations, the linear system in Equation 2 becomes over-determined and a least-squares solution can be found by standard techniques. However, least-squares solutions are prone to outliers induced by image noise and more importantly, low degrees of polarisation, which may render the TRS curve nearly "flat". This may ultimately cause bogus results due to significant shifts of the estimated phase angle from the actual

one. To remedy this problem, we resort to a robust M-estimator which, instead, minimises the following objective function, in which  $g_\sigma(\cdot)$  is a robust kernel function with a kernel width  $\sigma$ .

$$E(u, \lambda) = \sum_{i=1}^N g_\sigma(f_i^T x - I_i(u, \lambda)) \quad (3)$$

Formally speaking, a robust kernel function  $g_\sigma(\eta)$  quantifies an energy associated with both the residual  $\eta$  and its influence [4]. Each residual is assigned a weight as defined by an influence function  $h(\eta)$ . Thus the energy is related to the first-moment of the influence function as  $\frac{\partial g_\sigma(\eta)}{\partial \eta} = \eta h(\eta)$ . Some well-known kernel functions include Tukey's bi-weight [2] and Huber's M-estimators [4].

To minimise the function in Equation 3, it is necessary that its derivative vanishes at the solution, *i.e.*  $\frac{\partial E(u, \lambda)}{\partial x} = \sum_{i=1}^N (\eta_i h_\sigma(\eta_i)|_{(u, \lambda)}) f_i = \mathbf{0}$ , where  $\eta_i = f_i^T x - I_i(u, \lambda)$ . In other words

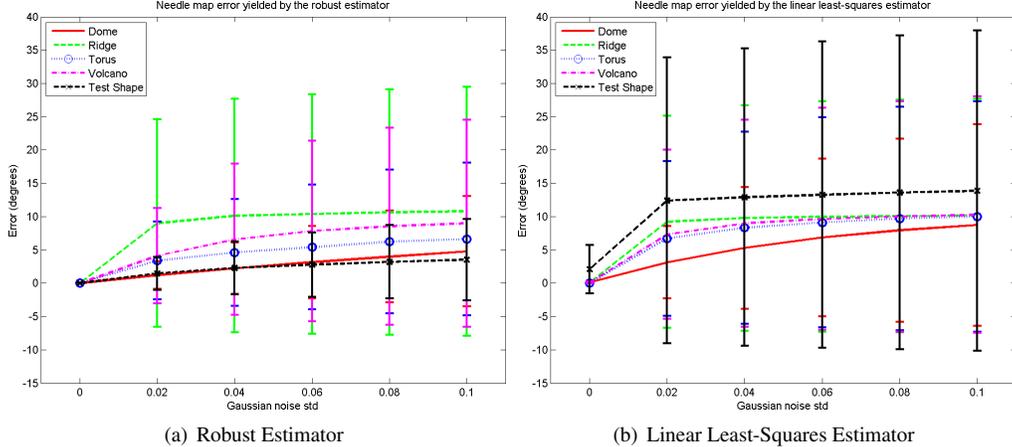
$$\sum_{i=1}^N (h_\sigma(\eta_i) f_i^T x) f_i = \sum_{i=1}^N h_\sigma(\eta_i) I_i(u, \lambda) f_i \quad (4)$$

Since  $h_\sigma(\eta_i)$  is normally a non-linear function of  $x$ , it is not always possible to find a closed-form solution to Equation 4. Nonetheless, we can adopt an iterative approach to find a fixed-point solution to Equation 4 by fixing its right hand side as a constant given the current value of the iterate  $x_k$ , where  $k$  is the iteration number. As a result, we obtain the update equation for the iterate  $x_{k+1}$  given by the following expression

$$x_{k+1} = \begin{bmatrix} \sum_{i=1}^N h_\sigma(\eta_i)|_{x_k} f_i^T \\ \sum_{i=1}^N h_\sigma(\eta_i)|_{x_k} \cos(2\vartheta_i) f_i^T \\ \sum_{i=1}^N h_\sigma(\eta_i)|_{x_k} \sin(2\vartheta_i) f_i^T \end{bmatrix} x_{k+1} = \begin{bmatrix} \sum_{i=1}^N h_\sigma(\eta_i)|_{x_k} I_i(u, \lambda) \\ \sum_{i=1}^N h_\sigma(\eta_i)|_{x_k} \cos(2\vartheta_i) I_i(u, \lambda) \\ \sum_{i=1}^N h_\sigma(\eta_i)|_{x_k} \sin(2\vartheta_i) I_i(u, \lambda) \end{bmatrix} \quad (5)$$

Note that the objective function in Equation 3 is a composition of  $g_\sigma(\cdot)$  with an affine function of  $x$ , therefore being convex in  $x$  if  $g_\sigma(\cdot)$  is convex in its corresponding composed domain. This assumption holds for the robust kernel functions mentioned above. Therefore, the sequence of solutions  $x_{k+1}$  to the system in Equation 5 eventually converges to the global minimum of the objective function in Equation 3. Note, however, that the global solution may be dependent on the choice of the kernel width  $\sigma$ .

To mitigate this issue, we employ a deterministic annealing process [6] which starts with a large enough kernel width to guarantee the convexity of the cost function in a large interval of  $x$ . In subsequent stages, we track



**Figure 1.** (a) Error measurements, in degrees, of the needle maps recovered for several shapes versus the Gaussian noise standard deviation. Left: errors yielded by the robust estimator. Right: errors yielded by the linear least-squares estimator.

the optimal solution of the cost function as the kernel width is reduced. As  $\sigma \rightarrow 0^+$ , we approximate the global minimum  $x^*$ . Having obtained the globally optimal value  $x^* = [x_1, x_2, x_3]^T$ , the phase of polarisation at pixel  $u$  and wavelength  $\lambda$  is recovered through the equation  $\phi(u, \lambda) = \frac{1}{2} \arctan \frac{x_3}{x_2}$ .

Finally, we note that, at each pixel, the phase angles of polarisation across various wavelengths are deemed to approximate the tilt of surface orientation. Here, we treat the problem of recovering the azimuth angle as an estimation from the set of spectral phase angles at the same image location. Under this setting, the azimuth angle at pixel-site  $u$  can be estimated using the following equation  $\tan \alpha(u) = \frac{\sum_{\lambda} \tan \phi(u, \lambda) h_{\gamma}(\delta(u, \lambda))}{\sum_{\lambda} h_{\gamma}(\delta(u, \lambda))}$ , where  $\delta(u, \lambda) \triangleq \phi(u, \lambda) - \alpha(u)$  and  $h_{\gamma}(\cdot)$  is an influence func-

tion with a chosen kernel width  $\gamma$ .

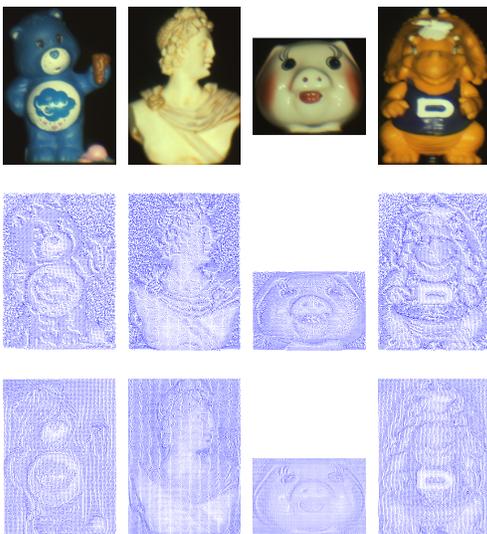
## 2.2. Recovery of the Zenith Angle via Shape from Shading

We now focus on the recovery of the zenith angle from the shading image given by  $I_{un}(u, \lambda) = \frac{1}{2}(I_{max}(u, \lambda) + I_{min}(u, \lambda)) = x_1(u, \lambda)$ . Suppose that the illuminant direction coincides with the viewing direction and its power  $L(\lambda)$  is known at each wavelength  $\lambda$ . Moreover, the shading image complies with Lambert's law, *i.e.*  $I_{un}(u, \lambda) = L(\lambda)G(u)\rho_{\mathcal{M}}(\lambda)$ , where  $\rho_{\mathcal{M}}(\lambda)$  is the spectral reflectance of the material  $\mathcal{M}$  of pixel  $u$ . As a result, the geometric shading factor  $G(u)$  is proportional to  $\cos \theta(u)$ , where  $\theta(u)$  is the normal's zenith angle with respect to the view direction.

Therefore, the wavelength-indexed reflectance vector  $\rho_{\mathcal{M}} \triangleq [\rho_{\mathcal{M}}(\lambda_1), \dots, \rho_{\mathcal{M}}(\lambda_K)]^T$  can be recovered at each pixel up to a multiple scalar. By quantifying the similarity between pixel reflectance vectors making use of their spectral angle, we can cluster pixels into materials using a standard clustering approach. Having obtained the pixel cluster belonging to a material  $\mathcal{M}$  and its normalised spectral reflectance  $\hat{\rho}_{\mathcal{M}}(\lambda)$ , the cosine of the zenith angle at pixel  $u$  for  $u \in \mathcal{M}$  is determined up to scale as  $\cos \theta(u) = s(\mathcal{M}) \frac{I_{un}(u, \lambda)}{L(\lambda)\hat{\rho}_{\mathcal{M}}(\lambda)}$  where  $s(\mathcal{M})$  is a per-material scaling constant. Here,  $s(\mathcal{M})$  is sought as to minimise the difference between the shading factors along the boundaries of adjacent materials.

## 3. Experiments

In this section, we present results of shape recovery on synthetic and real-world polarisation images. Specifically, we compare the performance of a robust estimator based on the Tukey's bi-weight function [2] with the linear least-squares estimator which directly solves the linear system in Equation 2.



**Figure 2.** Top row: Pseudo RGB colour images of real-world objects; Middle row: Needle-maps recovered by the robust estimator; Bottom row: Needle-maps yielded by the linear least-squares estimator.

### 3.1. Synthetic Images

We first perform experiments on a synthetic dataset generated from the 3D surfaces of a dome, a ridge, a torus, a volcano and a two-dome test shape. We render the images of these synthetic shapes using the refractive index of plastic and liquid materials reported in [5], under a frontal light source direction. The image dataset is synthesized using the Wolff diffuse reflectance model [8] with the total diffuse albedo being that of Polyvinyl Chloride (PVC).

For each combination of shape and material, five polarisation images are generated corresponding to five polariser orientations at  $0^\circ$ ,  $30^\circ$ ,  $45^\circ$ ,  $60^\circ$  and  $90^\circ$  in the clockwise direction with respect to the vertical orientation. The spectral dimension of the images spans the  $430 - 720nm$  range, with a resolution of  $10nm$ . To simulate noise effects, we perturb the clean images with zero-mean Gaussian noise with increasing local standard deviation  $\nu$  from 0.02 to 0.1 times the image intensity, in increments of 0.02.

In Figure 1, we plot the mean and standard deviation of the shape error yielded by both the robust and linear least-squares estimators. The shape error is quantified as the Euclidean angle, in degrees, between the estimated surface normal direction and the corresponding ground truth, averaged per pixel. Overall, the robust estimator significantly outperforms the alternative, in terms of accuracy and stability, for most shapes. In fact, the disparity in their performance becomes more distinct as the noise increases. When the noise is most extreme ( $\nu = 0.1$ ), the robust estimator achieves a mean angular error of about 5 degrees lower than that yielded by the alternative method for the dome, torus and two-dome test shape.

### 3.2. Real Images

We now perform shape recovery using the above two estimators on a real-world dataset consisting of multispectral images of four different objects made of plastic and porcelain. These have been acquired in house using a hyperspectral camera with an acousto-optic tunable filter, which we have tuned to the wavelength range of  $400 - 650nm$ , with a spectral resolution of  $5nm$ . To measure polarisation, we align the polariser's transmission axis to each of the seven angles of 45, 60, 75, 90, 105, 120 and 135 degrees in the clock-wise direction with respect to the vertical direction.

In Figure 2, we present the images of the objects polarised at an angle of 45 degrees in the first row and their needle-maps recovered by the robust and linear least-squares estimators in the subsequent rows. Note that the real-world images are subject to sensor noise and a

low degree of polarisation that can occur at surface locations where the normal vector is nearly aligned with the camera's optical axis. Here, we notice that the robust estimator produces needle-maps which clearly indicate the curvature and locations of the folds on the shapes, whereas the alternative estimator yields a flatter appearance of these surfaces.

## 4. Conclusion

In this paper, we have demonstrated that polarisation and shading information can complement each other for the task of shape recovery from single-view polarisation images. We have presented a method based upon robust statistics aimed at recovering the azimuth of the object surface normals. The estimator is robust to noise and a low degree of polarisation. We have shown the effectiveness of the method for purposes of shape recovery on both synthetic and real-world multispectral images.

## References

- [1] G. Atkinson and E. Hancock. Recovery of surface orientation from diffuse polarization. *IEEE Transactions on Image Processing*, 15(6):1653–1664, June 2006.
- [2] D. C. Hoaglin, F. Mosteller, and J. W. Tukey. *Understanding Robust and Exploratory Data Analysis*. Wiley-Interscience, 2000.
- [3] B. K. P. Horn and M. J. Brooks. The variational approach to shape from shading. *Comput. Vision Graph. Image Process.*, 33(2):174–208, 1986.
- [4] P. J. Huber. *Robust Statistics*. Wiley-Interscience, 1981.
- [5] S. N. Kasarova, N. G. Sultanova, C. D. Ivanov, and I. D. Nikolov. Analysis of the dispersion of optical plastic materials. *Optical Materials*, 29(11):1481 – 1490, 2007.
- [6] S. Z. Li. Robustizing robust M-estimation using deterministic annealing. *Pattern Recognition*, 29(1):159–166, 1996.
- [7] D. Miyazaki, R. T. Tan, K. Hara, and K. Ikeuchi. Polarization-based inverse rendering from a single view. *IEEE International Conference on Computer Vision*, 2:982, 2003.
- [8] L. B. Wolff. Diffuse-reflectance model for smooth dielectric surfaces. *Journal of the Optical Society of America*, 11(11):2956–2968, November 1994.
- [9] L. B. Wolff. Polarization vision: a new sensory approach to image understanding. *Image Vision Computing*, 15(2):81–93, 1997.
- [10] P. L. Worthington and E. R. Hancock. New constraints on data-closeness and needle map consistency for shape-from-shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(12):1250–1267, 1999.
- [11] Q. Zheng and R. Chellapa. Estimation of illuminant direction, albedo, and shape from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):680–702, 1991.