

Self-calibrating Polarising Radiometric Calibration

Daniel Teo Guangwei¹ Boxin Shi^{2*} Yinqiang Zheng³ Sai-Kit Yeung¹

¹Information Systems Technology and Design, Singapore University of Technology and Design

²National Engineering Laboratory for Video Technology, School of EECS, Peking University

³Digital Content and Media Sciences Research Division, National Institute of Informatics

Abstract

We present a self-calibrating polarising radiometric calibration method. From a set of images taken from a single viewpoint under different unknown polarising angles, we recover the inverse camera response function and the polarising angles relative to the first angle. The problem is solved in an integrated manner, recovering both of the unknowns simultaneously. The method exploits the fact that the intensity of polarised light should vary sinusoidally as the polarising filter is rotated, provided that the response is linear. It offers the first solution to demonstrate the possibility of radiometric calibration through polarisation. We evaluate the accuracy of our proposed method using synthetic data and real world objects captured using different cameras. The self-calibrated results were found to be comparable with those from multiple exposure sequence.

1. Introduction

Shape from polarisation is a technique that recovers the 3D surface of an object by exploiting the polarisation states of light. A surface normal map can be obtained by taking multiple images from a single viewpoint under different polarising angles. By integrating the surface normal map, the 3D surface of an object can be reconstructed. Like shape from shading or photometric stereo, shape from polarisation provides the surface normal map in a pixel-wise manner, encoding much denser 3D information than multiview stereo. However, shape from polarisation is a passive approach where no active or controlled light sources are placed and calibrated.

As shape from polarisation becomes more established, it provides an avenue for consumers to perform 3D reconstruction simply with a consumer-grade camera and a polariser due to the low cost in comparison to structured light scanners. Mobile phones with add-on polariser clips could enable such a possibility. However, there are two problems

which hinder such a widespread application. The first is that most consumer cameras have a nonlinear mapping between the measured intensity and the irradiance, whose relationship is established by the camera response function (CRF). Such information is usually not provided by the manufacturers, which means that the camera has to be radiometrically calibrated to get polarisation information for 3D reconstruction more accurately. The second problem is due to the fact that a typical shape from polarisation setup requires professional optical equipment to accurately control the angle of polarisation, limiting the technique to the laboratory environment. If these two problems can be solved in a self-calibrating manner, there is great potential to develop a low-cost shape from polarisation 3D camera for daily use.

This work seeks to enable the shape from polarisation technique for mass adoption by formulating a framework to solve for both the inverse camera response function (ICRF) and the polarising angles of the filter. As there have been previous works describing radiometric calibration through various techniques [14, 19, 27, 16], to the best of knowledge none have attempted radiometric calibration using polarisation. Our work is specially designed for the shape from polarisation problem in a self-calibrating and integrated manner. The main contributions are twofold: 1) We solve for the inverse CRF using polarisation information solely from input data in a self-calibrating manner; 2) We integrate the solving of uncalibrated polariser angles with radiometric calibration as a unified optimisation.

The key insight is that the transmitted radiance sinusoid (TRS) under a linear response would be distorted should the CRF be non-linear. By capturing multiple images of the scene under unknown polarising angles, we analyze how the CRF and unknown polarising angles affect the TRS and propose a unified optimisation to undistort the observed TRS, from which the ICRF and polarising angles can be estimated. An overview of the problem this paper addresses is illustrated in figure 1.

The rest of the paper is structured as follows. After reviewing related works in section 2, we provide a background summary of shape from polarisation in section 3 and

*Corresponding author: shiboxin@pku.edu.cn

mathematically formulate the calibration problem in section 4. The experimental details are provided in section 5, and the results for both synthetic and real world data would be discussed and analysed in section 6. Finally, we conclude our discussion in section 7.

2. Related works

Shape from polarisation: Shape from polarisation has been studied extensively but is less widely used compared to other photometric 3D modeling approaches such as photometric stereo, due to some inherent ambiguities. However, it has niche applications such as the reconstruction of transparent surfaces [20, 21] and black specular objects [22, 23]. For shape from polarisation, the surface normal map has an ambiguity in the phase, which has mainly limited its use to be combined with other techniques. Atkinson and Hancock [1] solved the ambiguity problem by assuming a convex object and propagating the surface normals inwards to preserve object smoothness. Followed, they used two views to resolve the ambiguity [2]. Kadambi *et al.* [11, 6] used the depth from Kinect time-of-flight sensors to help resolve the ambiguity. Cui *et al.* [6] used multiview stereo to first build up a point cloud, followed by iso-depth contour tracing which bypasses the π -ambiguity. Smith *et al.* [28] resolved the ambiguity by combining with shape from shading constraints, enabling a passive depth capture for a general outdoor scene. More recently, shape recovery through polarisation has been approached through a differential approach rather than the previous integral approach, circumventing the surface normal ambiguity problem [29, 18].

Radiometric calibration: Radiometric calibration of cameras has been studied extensively for computer vision applications. The earliest approach involved the use of a Macbeth color chart [17], which consists of color patches with reflectances mimicking natural objects. Mann and Picard [14] calibrated the camera using multiple registered images of the same scene, fitting to a gamma function. Mitsunaga and Nayar [19] used different known exposure times as well, fitting the inverse response function to a high-order polynomial. Grossberg and Nayar [9] collected a diverse database of real-world camera response functions (DoRF) and applied PCA to derive a basis which is known as the Empirical Model of Response (EMoR). Other camera calibration methods without using known exposures include using the RGB distribution color edges of a single image [13], using geometric invariants [25], using the noise distribution [16] and using consistent scene reflectance across corresponding pixels for internet images [24].

Self-calibration methods: In terms of radiometric self-calibration aspect, our method shares a similar spirit with

using the linearity of color profiles for photometric stereo images [27], but we explore the regularity of the transmitted radiance sinusoid of shape from polarisation images. In terms of polarising angle self-calibration aspect, our method shares a similar spirit with Schechner's method [26], but has the essential difference in that we do not assume a linear camera, and we use a unified optimisation method.

3. Background

3.1. Degree of polarisation

Shape from polarisation works by assuming that the light source is unpolarised, which is generally true for most scenes. When unpolarised light bounces off a smooth surface, it becomes partially polarised as shown by Fresnel reflection theory [10, 3]. When a polarising filter is rotated at angle ϕ_{pol} and placed in front of a camera, the irradiance varies given by the transmitted radiance sinusoid (TRS):

$$I(\phi_{pol}) = \frac{I_{max} + I_{min}}{2} + \frac{I_{max} - I_{min}}{2} \cos(2(\phi_{pol} - \psi)). \quad (1)$$

By capturing multiple photographs at different polarising angles, the quantities I_{max} , I_{min} and ψ can be found by fitting a sine curve. The surface normal at a particular point is controlled by 2 parameters, the azimuth angle ψ and the zenith angle θ . Due to the factor of 2 in the argument of the TRS, the periodicity is halved leading to an ambiguity in the azimuth angle. Note that ψ and $\psi + \pi$ both give the same value for equation 1. To find the zenith angle, it is required to first calculate a related quantity known as the degree of polarisation (DoP), ρ as

$$\rho = \frac{I_{max} - I_{min}}{I_{max} + I_{min}}, \quad (2)$$

which is simply interpreted as the ratio of the amplitude to the offset of the given sinusoid. Using the Fresnel equations [10] and equation 2, it is possible to obtain an analytic expression for ρ in terms of θ . Inverting the expression would then give a value for the zenith angle. The azimuth and zenith angles produce a surface normal map, which can be integrated through an integrator, such as the Frankot-Chellapa algorithm [8], to produce a 3D surface.

3.2. Radiometric calibration

The image of a camera is formed when photons from the scene strike the camera sensor array. In a modern camera, the sensor array is typically a CMOS or CCD, which records an intensity value proportional to the amount of photons hitting the array. However, for modern cameras, a nonlinear effect is added between the measured intensity, M , and the irradiance, I . Consumer cameras typically have image processing algorithms, which may be scene-dependent, to maximise visual aesthetics [4]. The nonlinear effect can come

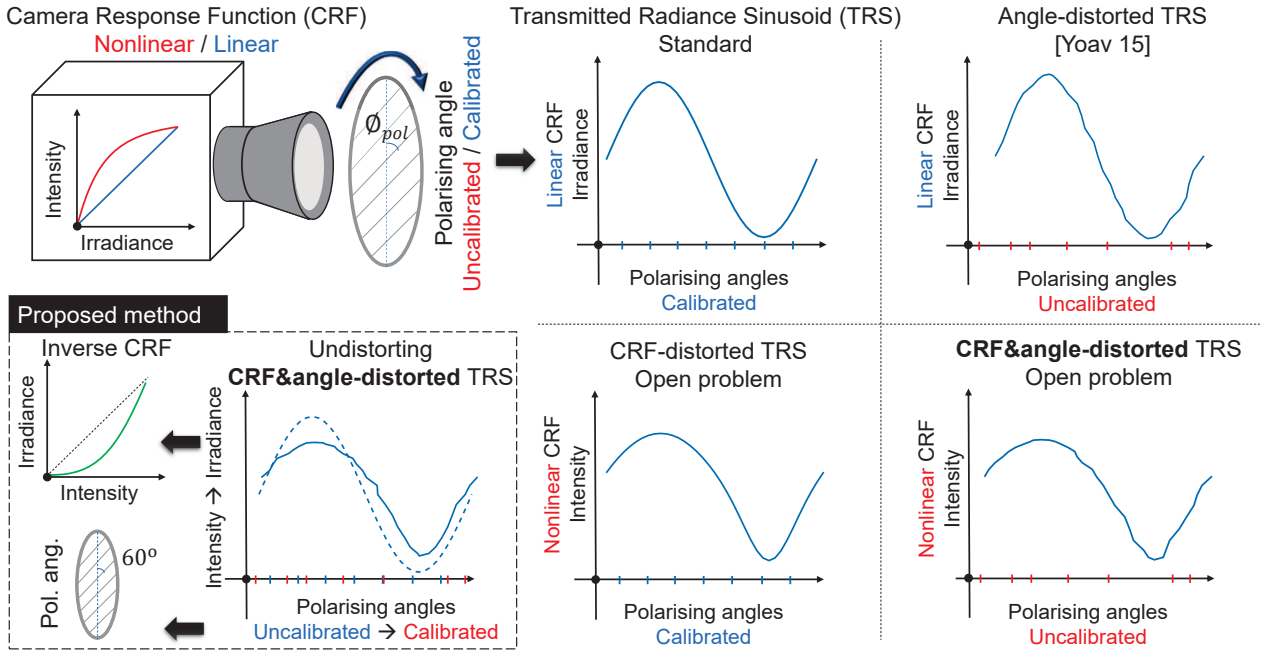


Figure 1: Overview for the proposed problems in polarising radiometric self-calibration. Angle-distorted TRS has been solved, but CRF-distorted TRS and CRF&angle-distorted is still an open problem. Our goal is to solve for the ICRF and the polarising angle by undistorting the CRF&angle-distorted TRS.

from several sources, such as demosaicking, JPEG compression, image enhancement and gamut mapping which cause a distortion in the TRS. Our proposed method takes into account of the aggregated nonlinear effect of all the TRS distorting factors, although we cannot tell how much distortion is brought by each component. The main goal of radiometric calibration would be to map the relationship between the measured intensity and the irradiance, which is also known as the inverse camera response function, g . We chose to model g with a polynomial function [7], although there are other camera response models such as the standard gamma function, generalised gamma curve model [5, 25] and the EMoR [9].

$$I = g(M) = c_1 M + c_2 M^2 + \dots + c_5 M^5 = \sum_{j=1}^5 c_j M^j, \quad (3)$$

where M is the measured intensity and c_j are the coefficients to be fitted. We only consider up to the fifth power which is sufficient to reproduce the inverse camera response function accurately. c_0 is not used because we define $g(0) = 0$. We assume, in this paper, that the inverse camera response function is either concave or convex, ie: $\frac{d^2 g}{dM^2} \geq 0$ or $\frac{d^2 g}{dM^2} \leq 0$, which holds true for most modern day cameras.

4. Proposed Method

4.1. Matrix representation

We can rewrite equation 1 for every pixel, and taking into account the ICRF, we obtain

$$g(M_{k,p}) = [t_p + a_p \cos 2(\phi_k - \psi_p)], \quad (4)$$

where k indexes the polarising angle, and p indexes the pixels of the image. In the variables where only a single subscript is used, it means the values are invariant with respect to the change in the other index. Expanding the cosine term we obtain,

$$g(M_{k,p}) = \begin{bmatrix} 1 & \cos 2\phi_k & \sin 2\phi_k \end{bmatrix} \begin{bmatrix} t_p \\ u_p \\ v_p \end{bmatrix}, \quad (5)$$

where $u_p = a_p \cos 2\psi_k$ and $v_p = a_p \sin 2\psi_k$. Representing equation 5 in the form of a matrix, we obtain

$$D = PO, \quad (6)$$

where

$$D = \begin{bmatrix} g(M_{1,1}) & g(M_{1,2}) & \dots & g(M_{1,p_N}) \\ g(M_{2,1}) & g(M_{2,2}) & \dots & g(M_{2,p_N}) \\ \vdots & \vdots & & \vdots \\ g(M_{L,1}) & g(M_{L,2}) & \dots & g(M_{L,p_N}) \end{bmatrix}, \quad (7)$$

and

$$\mathbf{P} = \begin{bmatrix} 1 & \cos 2(\phi_1) & \sin 2(\phi_1) \\ 1 & \cos 2(\phi_2) & \sin 2(\phi_2) \\ \vdots & \vdots & \vdots \\ 1 & \cos 2(\phi_L) & \sin 2(\phi_L) \end{bmatrix}, \quad (8)$$

and

$$\mathbf{O} = \begin{bmatrix} t_1 & t_2 & \dots & t_{p_N} \\ u_1 & u_2 & \dots & u_{p_N} \\ v_1 & v_2 & \dots & v_{p_N} \end{bmatrix}. \quad (9)$$

4.2. Optimisation procedure

We solve equation 6 by minimising

$$\|\mathbf{D} - \mathbf{P}\mathbf{O}\|^2, \quad (10)$$

using an iterative optimisation procedure described in further detail in the following subsections.

4.2.1 Initialisation

As \mathbf{D} , \mathbf{P} and \mathbf{O} all have unknown parameters, they are initialised prior to the optimisation. The angles in \mathbf{P} are initialised as the actual experimental angle added with a random angular error of up to $\pm 15^\circ$. The values of \mathbf{D} and \mathbf{O} were not explicitly coded, hence they were initialised according to Matlab's internal method [15].

4.2.2 Finding c_j and \mathbf{O}

At each step of the iteration, we find c_j and \mathbf{O} . Using equation 10, we rewrite it as

$$\begin{aligned} \{c_j^*, \mathbf{O}^*\} &= \underset{c_j, \mathbf{O}}{\operatorname{argmin}} \mathbf{y}^T \mathbf{J}^T \mathbf{J} \mathbf{y}, \\ \text{s.t., } \sum_{j=1}^5 c_j &= 1, g(M_{int} + 0.1) \geq g(M_{int}), \end{aligned} \quad (11)$$

$$\frac{d^2 g}{dM^2} \big|_{M=M_{int}} \geq 0 (\text{or } \leq 0),$$

where $M_{int} = 0.1, 0.2, \dots, 0.9$, and \mathbf{y} is the unknown vector such that

$$\mathbf{y} = [c_1 \quad \dots \quad c_5 \quad t_1 \quad u_1 \quad v_1 \quad \dots \quad t_{p_N} \quad u_{p_N} \quad v_{p_N}]^T, \quad (12)$$

and \mathbf{J} can be constructed by

$$\mathbf{J} = \begin{bmatrix} M_{1,1} & \dots & M_{1,1}^5 & -\mathbf{P} & \mathbf{0} \\ \dots & \dots & \dots & \ddots & \dots \\ M_{L,p_N} & \dots & M_{L,p_N}^5 & \mathbf{0} & -\mathbf{P} \end{bmatrix}. \quad (13)$$

The above convex quadratic program can be solved using quadprog in Matlab.

4.2.3 Finding angles ϕ_k

Given \mathbf{D} and \mathbf{O} , we are able to calculate the angle $\phi_k, k = 1, 2, \dots, L$, independently. Let D_k denote the k -th row of \mathbf{D} , R the first row of \mathbf{O} , \mathbf{S} the last two rows of \mathbf{O} , the constrained optimization reads

$$\begin{aligned} \{\phi_k^*\} &= \underset{\phi_k}{\operatorname{argmin}} \left\| \mathbf{S}^T \begin{bmatrix} \cos(2\phi_k) \\ \sin(2\phi_k) \end{bmatrix} - (D_k - R)^T \right\|_2^2, \\ &= \underset{\phi_k}{\operatorname{argmin}} \begin{bmatrix} \cos(2\phi_k) \\ \sin(2\phi_k) \end{bmatrix}^T \mathbf{A} \begin{bmatrix} \cos(2\phi_k) \\ \sin(2\phi_k) \end{bmatrix} - 2b_k^T \begin{bmatrix} \cos(2\phi_k) \\ \sin(2\phi_k) \end{bmatrix}, \\ &\quad \text{s.t., } \cos^2(2\phi_k) + \sin^2(2\phi_k) = 1, \end{aligned} \quad (14)$$

where $\mathbf{A} = \mathbf{S}\mathbf{S}^T$ and $b_k = \mathbf{S}(D_k - R)^T$. To find its global minimum, we calculate its first-order optimality condition:

$$\begin{aligned} (\mathbf{A} - \lambda \mathbf{I}) \begin{bmatrix} \cos(2\phi_k) \\ \sin(2\phi_k) \end{bmatrix} &= b_k, \\ \text{s.t., } \cos^2(2\phi_k) + \sin^2(2\phi_k) &= 1, \end{aligned} \quad (15)$$

where λ is the Lagrange multiplier and \mathbf{I} is the 2×2 identity matrix. By letting $x = (\mathbf{A} - \lambda \mathbf{I})^{-1} b_k$, we have $x = (\mathbf{A} - \lambda \mathbf{I})^{-1} [(\mathbf{A} - \lambda \mathbf{I})^{-1} b_k] = (\mathbf{A} - \lambda \mathbf{I})^{-1} \begin{bmatrix} \cos(2\phi_k) \\ \sin(2\phi_k) \end{bmatrix}$, which leads to $\begin{bmatrix} \cos(2\phi_k) \\ \sin(2\phi_k) \end{bmatrix} = (\mathbf{A} - \lambda \mathbf{I})x$. Since $\cos^2(2\phi_k) + \sin^2(2\phi_k) = 1$, $x^T (\mathbf{A} - \lambda \mathbf{I})^2 x = x^T b_k = 1$, implying $b_k b_k^T x = b_k$, we expand $(\mathbf{A} - \lambda \mathbf{I})^2 x = b_k$, such that $(\lambda^2 \mathbf{I} - 2\lambda \mathbf{A} + \mathbf{A}^T \mathbf{A})x = b_k$. By substituting b_k by $b_k b_k^T x$, we obtain

$$(\lambda^2 \mathbf{I} - \lambda(2\mathbf{A}) + (\mathbf{A}^T \mathbf{A} - b_k b_k^T))x = \mathbf{0}. \quad (16)$$

The above quadratic eigenvalue problem can be solved using polyeig in Matlab. As the objective value in equation 14 is proportional to the Lagrange multiplier λ , we get the global minimum using the minimal real/positive eigenvalue λ^* . x^* is its corresponding eigenvector and the angle ϕ_k is determined from:

$$\begin{bmatrix} \cos(2\phi_k) & \sin(2\phi_k) \end{bmatrix}^T = (\mathbf{A} - \lambda^* \mathbf{I})x^*. \quad (17)$$

4.2.4 Iterations

The steps described in sections 4.2.2 and 4.2.3 were repeated alternatively for 300 times. Towards the end of the optimisation process, we find that the cost function does not change significantly.

4.3. Additional priors

We fully utilise the properties of the shape from polarisation to develop the following two techniques.

Pixel selection. The naive way is to use all pixels in the captured image. However, the degree of polarisation provides information as to which pixels are more reliable. For pixels with a low degree of polarisation, noise from the camera sensor affects the overall shape of the curve. Ideally, the pixels selected should cover a wide range of intensities, have high values of ρ and low noise so that the optimisation procedure would yield accurate results.

We approximate the values of ρ for each pixel using the maximum and minimum of the measured intensity across the different polarising angles with equation 2. We adopt a threshold filter to randomly select 10 pixels with $0.3 < \rho < 1$ and intensity values > 0.2 to filter out the background dark pixels. The noise level of the pixels was determined by fitting the measured intensities across different polarising angles with a 10 order polynomial and measuring the residue. However, we note that the order of the polynomial is a parameter that should be tuned to the data size to prevent over-fitting. A threshold value of the residue < 0.02 was used to filter away noisy pixels.

Convexity prior. By plotting the intensity of a pixel across the different polarising angles as illustrated in figure 1, we obtain the CRF&angle-distorted TRS. The curvature of the TRS tells an important prior, to improve the robustness of the optimisation method. We note that the curvature of the actual camera response function affects the distortion as illustrated in figure 2. The figure was generated from two of the curves in the DoRF [9]. For instance, if the camera response function is convex, the width of the trough will tend to be wider than the width of the peak. The vertical distance between the maximum value and the mean is also greater than the vertical distance between the mean and the minimum value. The properties would be the opposite if the CRF is concave.

To determine the curvature of the CRF for real data, we assume that the polariser is rotated in one direction only, and the distorted sine curve is sufficiently sampled such that the separation between consecutive angles do not vary too widely. We determine the curvature of all the selected pixels by measuring the vertical distances between the peak and the mean, and the mean and the trough. A count on the number of pixels which indicate either concave or convex was maintained. The curvature with the most counts would be selected as the constraint for the optimisation procedure.

4.4. Well-posedness and reliability of our method

Although there are unknown parameters on both sides of equation 10, our results in section 6 show that it is indeed solvable. Starting from rough initialization of ϕ_k , our algorithm mostly converges to a solution that is consistent with the ground truth. There is only an arbitrary global shift in ϕ_k , which can be resolved by using $\phi_k - \phi_1$ to find the rel-




Figure 2: Plot of the distorted sine of a convex CRF. The red line is the mean of the values.

ative angles. Although it should be noted that the angle initialisation cannot deviate significantly from the ground truth ($< 90^\circ$) as it may at times converge to a wrong result. The necessary condition is that $L \times p_N \geq L + 3p_N + n - 2$. For example, given $L = 4$ angles and $n = 5$ terms in ICRF, the number of independent polarised pixels should be $p_N \geq 7$.

5. Experimental Procedure

Synthetic data. Ten sets of sinusoidal curves with different amplitude, D.C. offset, and phase were generated and sampled at 17 angles of different spacing in a monotonic manner. White gaussian noise with Signal-to-Noise Ratio (SNR) of 99 was added to simulate real data. We employed the DoRF [9] to distort the set of synthetically generated sine curves and ran it through the optimisation process.

Real-world data. Images were taken with consumer grade cameras, the Nikon D800E and the Canon EOS M2. The polarising filter was mounted onto a rotating mount from Thorlabs to measure the polarising angle. The polariser used was the Hoya HD2 CIR-PL filter with an arrow indicating the direction of polarisation. The setup is shown in figure 3. Multiple images at the same polarising angle were taken and averaged to reduce the noise.

Images were taken for two different objects. The first object is that of a ceramic orange tiger, while the second object is a ceramic green and white lion as shown in figure 5. The objects were placed in a diffuse light tent with a black background and illuminated with an LED desk lamp. We chose ceramic coloured objects for calibration because the degree of polarisation is generally high for ceramic objects, and the different colours would provide a greater range of intensi-



Figure 3: Image of the polariser and camera setup

Table 1: Table of results from synthetic data using the DoRF

	Average	Standard Deviation
RMSE	0.0299	0.0413
Disparity	0.0541	0.0799
Angular RMSE	0.7124	2.50
Angular Disparity	1.27	4.01




Figure 4: Boxplot for the error in ICRF and angular estimations for the DoRF. The red line is the median, the green dot is the mean, the blue box marking the first and third quartile, the black lines are the maximum and minimum which are not outliers, and the red crosses are the outliers.

ties sampled. Images were taken at every 5° in the interval $[0, 175]$ with the camera set to manual mode saved in .TIFF format. Care was taken not to move the camera, object or light source during the capture process. The captured images were normalised ranging from $[0, 1]$.




Figure 5: Ceramic tiger (left) and lion (right) captured at 0° polarising angle. The blue points are those selected for optimisation for the Nikon D800E while the red points are those for the Canon EOS M2.

6. Results and Discussion

6.1. Synthetic data

Of the 201 curves in the DoRF, our method has identified 9 concave ICRF. The other 192 curves either have a convex ICRF or a change in curvature. Doing a manual count on the DoRF, we have identified 10 concave ICRF, which is similar to our result. The statistics for the ICRF and the angular estimations over the entire set of DoRF are presented in table 1 and figure 4. The disparity is calculated by finding the largest difference between the recovered ICRF and the ground truth. The RMSE is calculated across all sampled points in the ICRF. Our method can be seen to be accurate, as the average disparity and RMSE are reasonably low, and the angular errors are approximately within $\pm 1^\circ$. The estimated angles are all measured relative to the first angle. From figure 4, it is observed that the median lies towards the lower end of the blue bounding box for both the ICRF and the angular estimation. This would imply that the errors are in general small for most of the CRF, but the errors are skewed higher due to a few CRF which give larger errors. This is expected because our proposed method assumes either a concave or convex curvature in the CRF, whereas there exists a few CRFs with varying curvature in the DoRF.

6.2. Real-world data

The images of the objects are shown in figure 5. The dots in the figures represent the points selected for use in the optimisation process for each camera. The recovered ICRF and estimated angles for the Nikon D800E are shown in figure 6 for both the tiger and lion objects. The ground truth, which is assumed to be correct, was found by using a multi-exposure rank minimization method as described in [12]. The camera was found to have a convex ICRF for




Figure 6: Recovered ICRF and angle estimation for the tiger (left) and the lion (right) captured using the Nikon D800E




Figure 7: Recovered ICRF and angle estimation for the tiger (left) and the lion (right) captured using the Canon EOS M2

both methods, justifying the validity of our ICRF convexity prior. The recovered ICRF follows closely with that from the rank minimization method with low values of disparity and RMSE, and is comparable to the errors reflected by the




Figure 8: Top-left: CRF&angle-distorted TRS for points of high ρ (Green) and low ρ (Red). Top-right: CRF&angle-distorted TRS for points of low ρ (Red). Bottom-left: CRF&angle-distorted TRS for points of high ρ . Bottom-right: Corresponding undistorted TRS for points of high ρ .

synthetic data. The angle estimation error varies between the different objects and cameras, with the disparity as low as $\pm 1^\circ$ to $\pm 4^\circ$. The ground truth angles here refer to the reading of the scale on the rotating mount marked every 2° . Hence, it should be noted that the ground truth would also be susceptible to experimental human error of $\pm 1^\circ$. The increase in angular error compared to the synthetic data could be attributed to noise from the camera. However, the recovered ICRF from the two objects have small error values and are consistent with each other, demonstrating the robustness of our method against different objects.

The points selected for the Canon EOS M2 are different from that of the Nikon D800E due to the random pixel selection process. The recovered ICRF and estimated angles of the Canon EOS M2 for both objects are shown in figure 7. The RMSE and disparity of the ICRF are both small and comparable to the Nikon D800E and the synthetic results.

Figure 8 shows the real data taken from the lion object captured by the Nikon D800E. From the given CRF&angle-distorted TRS, the values of ρ are approximated. Points with sufficiently high ρ are selected, as shown by the green curves, while those with low ρ , depicted by the red curves, are discarded. We undistort the CRF&angle-distorted TRS in the bottom-left figure to obtain the undistorted TRS shown in the bottom right figure. It can be seen that the CRF&angle-distorted TRS curves are more noisy than the undistorted TRS. It is also observed that after the undistortion, the widths of the peak and the trough appear more




Figure 9: Average ICRF with standard deviation and the angle estimates for the lion captured by the Nikon D800E.

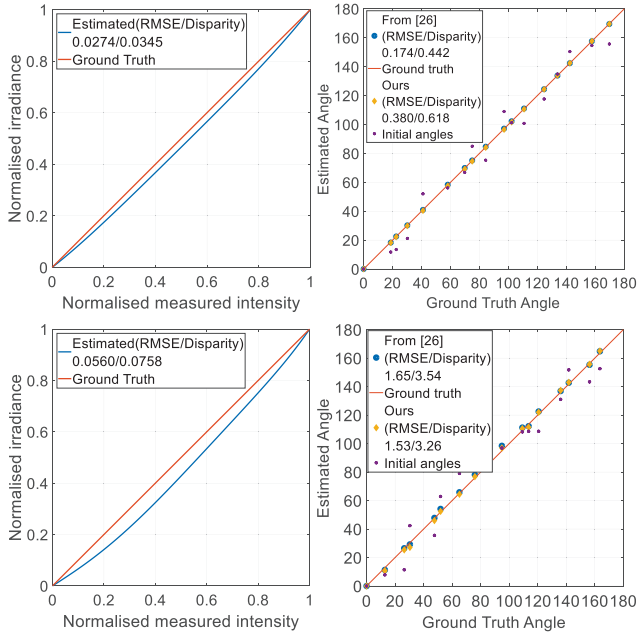


Figure 10: Angular comparison using synthetic data with white gaussian noise. Top: SNR = 50, Bottom: SNR = 30

symmetric and resembles closer to an actual sinusoid. We note that the TRS of different pixels have different phases due to having different surface normals.

Robustness of pixel selection. The robustness of the method was tested against different sets of pixels selected for the Nikon D800E lion dataset. We ran the random pixel selection and the optimisation for 10 times and plotted the results in figure 9. The error bars represent the standard deviation of the 10 sets of data. The average of the ICRF lies close to the ground truth with low values of disparity and RMSE. The RMSE and disparity of the average angular estimation values compared to the ground truth are also small.

Table 2: Comparing [26] and our method

		Using [26]	Ours
Lion	Angular Disparity	1.3	1.02
Nikon D800E	Angular RMSE	0.471	0.446
Tiger	Angular Disparity	3.9	3.85
Nikon D800E	Angular RMSE	2.28	2.11

Experiments have been done using more pixels within the optimisation, but with limited improvements.

6.3. Comparison of angular estimation with [26]

Synthetic data described in section 5 was generated and ran through our method and [26]. We used a linear CRF to enable a fair comparison for both methods as shown in figure 10 and table 2. Our method shows a slightly better performance for lower SNR. As for real-world data, we linearised the images using the ICRF obtained from our method before passing it through the code by [26]. Our method is seen to be comparable to that of [26].

7. Conclusion

We have recovered the ICRF for two scenes which are close to the ground truth, demonstrating the robustness of our proposed method against different scenes and different cameras. The angular estimations are seen to converge towards the ground truth with disparity values reaching as low as $\pm 1^\circ$, validating our integrated optimisation approach. One drawback of the method is that multiple images of a static scene must be captured, which prevents real-time applications. Should shape from polarisation become widely adopted, our approach would aid in obtaining a more accurate shape. We hope that our work would enable the usage of polarisation in more computer vision applications.

Acknowledgements

Daniel Teo and Sai-Kit Yeung are supported by the SUTD Digital Manufacturing and Design Centre which is supported by the Singapore National Research Foundation (NRF). Boxin Shi is supported by the Recruitment Program of Global Experts (Youth Program) in China (a.k.a. 1000 Youth Talents). Yinqiang Zheng is supported in part by JSPS KAKENHI Grant Number JP15H05918 and Microsoft Research Asia through the 2017 Collaborative Research (Core13). Sai-Kit Yeung is also supported by Singapore MOE Academic Research Fund MOE2016-T2-2-154, Heritage Research Grant of the National Heritage Board, Singapore, and Singapore NRF under its IDM Futures Funding Initiative and Virtual Singapore Award No. NRF2015VSGAA3DCM001-014.

References

- [1] G. A. Atkinson and E. R. Hancock. Recovery of surface orientation from diffuse polarization. *IEEE Transactions on Image Processing*, 15(6):1653–1664, 2006. 2
- [2] G. A. Atkinson and E. R. Hancock. Shape estimation using polarization and shading from two views. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(11):2001–2017, 2007. 2
- [3] M. Born and E. Wolf. *Principles of optics: Electromagnetic theory of propagation, interference and diffraction of light*. Elsevier, 2013. 2
- [4] A. Chakrabarti, D. Scharstein, and T. E. Zickler. An empirical camera model for internet color vision. In *Proc. of British Machine Vision Conference*, 2009. 2
- [5] X. Chen, F. Li, J. Yang, and J. Yu. A theoretical analysis of camera response functions in image deblurring. In *Proc. of European Conference on Computer Vision*, 2012. 3
- [6] Z. Cui, J. Gu, B. Shi, P. Tan, and J. Kautz. Polarimetric multi-view stereo. In *Proc. of Computer Vision and Pattern Recognition*, 2017. 2
- [7] P. E. Debevec and J. Malik. Recovering high dynamic range radiance maps from photographs. In *Proc. of ACM SIGGRAPH*, 2008. 3
- [8] R. T. Frankot and R. Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):439–451, 1988. 2
- [9] M. D. Grossberg and S. K. Nayar. What is the space of camera response functions? In *Proc. of Computer Vision and Pattern Recognition*, 2003. 2, 3, 5
- [10] E. Hecht. Optics, 4th. International edition, Addison-Wesley, San Francisco, 3, 2002. 2
- [11] A. Kadambi, V. Taamazyan, B. Shi, and R. Raskar. Polarized 3d: High-quality depth sensing with polarization cues. In *Proc. of International Conference on Computer Vision*, 2015. 2
- [12] J.-Y. Lee, Y. Matsushita, B. Shi, I. S. Kweon, and K. Ikeuchi. Radiometric calibration by rank minimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(1):144–156, 2013. 6
- [13] S. Lin, J. Gu, S. Yamazaki, and H.-Y. Shum. Radiometric calibration from a single image. In *Proc. of Computer Vision and Pattern Recognition*, 2004. 2
- [14] S. Mann and R. Picard. *Being ‘undigital’ with Digital Cameras*. MIT Media Lab Perceptual, 1994. 1, 2
- [15] MathWorks. Quadratic programming algorithms. <https://www.mathworks.com/help/optim/ug/quadratic-programming-algorithms.html>. 4
- [16] Y. Matsushita and S. Lin. Radiometric calibration from noise distributions. In *Proc. of Computer Vision and Pattern Recognition*, 2007. 1, 2
- [17] C. S. McCamy, H. Marcus, J. Davidson, et al. A color-rendition chart. *Journal of Applied Photographic Engineering*, 2(3):95–99, 1976. 2
- [18] C. Mecca, Logothetis. A differential approach to shape from polarization. In *Proc. of British Machine Vision Conference*, 2017. 2
- [19] T. Mitsunaga and S. K. Nayar. Radiometric self calibration. In *Proc. of Computer Vision and Pattern Recognition*, 1999. 1, 2
- [20] D. Miyazaki, M. Kagesawa, and K. Ikeuchi. Polarization-based transparent surface modeling from two views. In *Proc. of International Conference on Computer Vision*, 2003. 2
- [21] D. Miyazaki, M. Saito, Y. Sato, and K. Ikeuchi. Determining surface orientations of transparent objects based on polarization degrees in visible and infrared wavelengths. *Journal of the Optical Society of America A*, 19(4):687–694, 2002. 2
- [22] D. Miyazaki, T. Shigetomi, M. Baba, R. Furukawa, S. Hiura, and N. Asada. Polarization-based surface normal estimation of black specular objects from multiple viewpoints. In *Proc. of International Conference on 3D Imaging, Modeling, Processing, Visualization and Transmission*, 2012. 2
- [23] D. Miyazaki, T. Shigetomi, M. Baba, R. Furukawa, S. Hiura, and N. Asada. Surface normal estimation of black specular objects from multiview polarization images. *Optical Engineering*, 56(4):041303–041303, 2017. 2
- [24] Z. Mo, B. Shi, S.-K. Yeung, and Y. Matsushita. Radiometric calibration for internet photo collections. In *Proc. of Computer Vision and Pattern Recognition*, 2017. 2
- [25] T.-T. Ng, S.-F. Chang, and M.-P. Tsui. Using geometry invariants for camera response function estimation. In *Proc. of Computer Vision and Pattern Recognition*, 2007. 2, 3
- [26] Y. Y. Schechner. Self-calibrating imaging polarimetry. In *Proc. of International Conference on Computational Photography*, 2015. 2, 8
- [27] B. Shi, Y. Matsushita, Y. Wei, C. Xu, and P. Tan. Self-calibrating photometric stereo. In *Proc. of Computer Vision and Pattern Recognition*, 2010. 1, 2
- [28] W. A. Smith, R. Ramamoorthi, and S. Tozza. Linear depth estimation from an uncalibrated, monocular polarisation image. In *Proc. of European Conference on Computer Vision*, 2016. 2
- [29] S. Tozza, W. A. P. Smith, D. Zhu, R. Ramamoorthi, and E. R. Hancock. Linear differential constraints for photopolarimetric height estimation. In *Proc. of International Conference on Computer Vision*, 2017. 2