



Recovering Material Reflectance from Polarization and Simulated annealing

Gary Atkinson, Edwin R. Hancock

► To cite this version:

Gary Atkinson, Edwin R. Hancock. Recovering Material Reflectance from Polarization and Simulated annealing. Proceedings of the First International Workshop on Photometric Analysis For Computer Vision - PACV 2007, Oct 2007, Rio de Janeiro, Brazil. 8 p. inria-00265255

HAL Id: inria-00265255

<https://inria.hal.science/inria-00265255>

Submitted on 18 Mar 2008

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Recovering Material Reflectance from Polarization and Simulated Annealing

Gary A. Atkinson*

University of the West of England
Bristol, BS16 1QY UK

gary.atkinson@uwe.ac.uk

Edwin R. Hancock

University of York
Heslington, York, YO10 5DD UK

erh@cs.york.ac.uk

Abstract

This paper presents a novel image-based method for estimating BRDF information. The goal is to determine the dependence of the reflected radiance on the surface orientation for any given illumination direction. The proposed method consists of two stages. In the first stage, the polarization state of light reflected from an object of unknown shape is analysed using a linear polarizer and digital camera. This is interpreted using Fresnel theory to estimate a field of surface normals. This stage of the algorithm is largely based on existing ideas. In the second stage, a 3D histogram of surface normals and measured pixel brightnesses is constructed. The BRDF is recovered by fitting an arbitrary surface to this histogram data using simulated annealing and elastica. The experiments presented in this paper show that realistic image rendering is possible based on the recovered BRDFs of smooth dielectric materials.

1. Introduction

Detailed material reflectance determination is key to many tasks in computer graphics and computer vision. In graphics, reflectance properties are essential for realistic image rendering. For vision, they are used for surface analysis. Usually, reflectance properties are characterised using the *bidirectional reflectance distribution function*, or BRDF [4]. The BRDF accounts for the reflected radiance distribution from a surface for any given light source direction.

The BRDF has four angular degrees of freedom, but for some applications it is not necessary for it to be recovered in full. For example, in some surface reconstruction algorithms, it is only necessary to know the relationship between the surface orientation and the reflected radiance for the current illumination distribution [1]. The estimation of this relationship is the main aim of this paper. The method presented uses polarization to fit a 2D surface, representing a portion of the BRDF, to measured image data.

Previous efforts to estimate BRDFs can be broadly divided into two categories based on whether the approach is theoretical or experimental. Theoretical methods attempt to predict the BRDF using either ray or wave optics, or by using a semi-empirical methods that give phenomenological agreement with measured data. For example, the Phong model [13] has been widely used in graphics for several decades. The model has no physical basis, but has been shown to allow visually impressive image rendering. Torrance and Sparrow [16] attempted to incorporate more realistic ideas by modelling rough surfaces as arrays of V-shaped grooves that can be parameterized by a small number of roughness coefficients. Ray optics were then used to predict the distribution of the reflected light. Wolff [17] developed a model restricted to smooth surfaces entirely based on real physics. He showed how internal light scattering affects the reflected light distribution by using results from Fresnel theory and radiative transfer theory.

Experimental BRDF estimation methods can be either image-based or photometer-based. The most direct methods use gonireflectometers [7]. These devices use photometers to measure the reflected radiance from a planar sample of the material as the light source direction and/or the sample orientation are adjusted. Although these methods are the most accurate and reliable, they are also very time-consuming and not always practical. Marschner et al. [8] greatly reduced experimental BRDF acquisition time by replacing the photometer with a camera and the planar sample with a spherical one. This effectively allowed many measurements to be made simultaneously.

Robles-Kelly and Hancock [14] proposed a purely image-based method for reflectance function estimation from a single view. Their idea was to determine a mapping from a single image onto a Gauss sphere using the cumulative distribution of intensity gradients. Atkinson and Hancock [1] also proposed an image based approach. Their method used polarization to obtain the mapping. Unfortunately, both of these methods are restricted to the case where the light source and camera directions are identical.

The method that is presented in this paper consists of two

*This work was carried out while the author was at the University of York

main stages. In the first stage, a linear polarizer and a digital camera are used to ambiguously estimate a field of surface normals from Fresnel theory. These estimates are disambiguated using a simple method based on ideas from photometric stereo [19]. This first stage of the proposed technique is mainly based on existing methods and is described in Section 2. The second stage of the technique, which is entirely novel, aims to construct a BRDF from the estimated surface normals and the measured pixel brightnesses. This is done by fitting a surface (a representation of the BRDF) to a 3D histogram of surface orientations and pixel brightnesses. Simulated annealing is used for this task, where the objective function to be minimized consists of the histogram frequencies and the “elastic energy” of the surface (which ensures a smooth result). Section 3 describes this stage of the technique in detail. With a suitable choice of cooling parameters, simulated annealing gives the *globally* optimal solution. Another advantage of simulated annealing is that the associated cost function is based on a directly interpretable quantity, namely histogram frequencies.

Other options for surface fitting include thin-plate splines, Markov random fields, or simply fitting parametric BRDFs such as the Wolff [17] model or the Phong model [13]. It is also worth noting that polarization is not the only method of surface normal estimation that could be used. For example, four-source photometric stereo could be used, which would help make the method more general to surface type. These alternative methods will be studied further in future work.

Using polarization analysis for BRDF estimation has several advantages over most existing methods. Most notably, because polarization allows for the direct calculation of surface orientations, we do not need ground truth information on the geometry of the material sample object. Another benefit of polarization analysis is that it allows entirely image-based data processing, so there is no need for structured light sources. Finally, we only make relatively unrestrictive assumptions about the form of the reflectance function being estimated. We do however, assume that the surface is dielectric and is either smooth or only slightly rough.

2. Estimating Surface Normals

Consider the reflection of a ray of light from a smooth surface. Fresnel theory [5] provides a means to calculate the ratio of the incident light intensity to the reflected light intensity for a given angle of incidence. Further to this, if the incident light is unpolarized, the theory predicts that the reflected ray will become partially linearly polarized, again depending on the angle of incidence.

In this paper, we study diffuse reflection, where light penetrates the surface and is scattered internally before being re-emitted (in contrast to specular reflection, where the




Figure 1. Relationship between the zenith angle and the degree of polarization.

light is reflected from the surface directly). For diffuse reflection, Fresnel theory can be applied to light as it is re-emitted from the medium into air [18]. This provides a relation between the polarization state of the reflected light and the angle of the reflection.

The standard approach to measure the polarization state of reflected light is to take a succession of images of the reflecting surface with a polarizer placed in front of a camera rotated to different angles. The measured intensity at each pixel varies sinusoidally with the polarizer angle. By performing (for example) a least squares fit on the measured pixel brightnesses as a function of the polarizer angle, the minimum and maximum intensities on the sinusoid, I_{\min} and I_{\max} , can easily be determined. The angle of polarization, or *phase* angle, ϕ , can also be determined [10].

The polarization state can then be expressed using the *degree of polarization*:

$$\rho = (I_{\max} - I_{\min}) / (I_{\max} + I_{\min}) \quad (1)$$

Fresnel theory predicts that the degree of polarization is related to the reflectance angle, θ , and the refractive index, n , by [10]

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

We assume that the refractive index is 1.4 throughout this work, which is a typical value. If a surface point is observed such that the angle between its normal and the viewing direction (the zenith angle) is θ , then the observed degree of polarization is given by (2). Fig. 1 shows the relationship between the degree of polarization and the zenith angle.

Fresnel theory has also been used to provide constraints on the azimuth angles of the surface normals, α (the angles of the projections of the normals onto the image plane) [10, 1]. These constraints are provided by the phase angle of the reflected light, ϕ . It turns out that, for diffuse reflection, the azimuth angle can be determined only ambiguously and is either equal to the phase angle or differs by 180° .

Fig. 2a shows a greyscale image of a porcelain bear model. The degree of polarization and phase images are




Figure 2. (a) Greyscale image of a porcelain bear. (b) Degree of polarization (dark areas indicate higher values). (c) phase (light areas indicate larger angles).




Figure 3. Experimental arrangement.

shown in Figs. 2b and c respectively. These were calculated using the standard method with the rotating linear polarizer mentioned above. The zenith angles can then be estimated at each pixel by numerically solving (2). The azimuth angle is less straightforward to estimate due to the two-way ambiguity of 180° . For this paper, we disambiguate the angles using three light source directions, as shown in Fig. 3. This simple method was inspired by standard photometric stereo [19]. It is easy to show that the correct choice of azimuth angle for pixel k is given by [2]

$$\begin{aligned} \text{if } \phi_k < 45^\circ & \quad \text{then } \alpha_k = \begin{cases} \phi_k & \text{if } I_k^{(2)} > I_k^{(1)} \\ \phi_k + 180^\circ & \text{otherwise} \end{cases} \\ \text{if } 45^\circ \leq \phi_k < 135^\circ & \quad \text{then } \alpha_k = \begin{cases} \phi_k & \text{if } I_k^{(3)} > I_k^{(1)} \\ \phi_k + 180^\circ & \text{otherwise} \end{cases} \\ \text{if } 135^\circ \leq \phi_k & \quad \text{then } \alpha_k = \begin{cases} \phi_k & \text{if } I_k^{(3)} > I_k^{(2)} \\ \phi_k + 180^\circ & \text{otherwise} \end{cases} \end{aligned} \quad (3)$$

where $I_k^{(1)}$, $I_k^{(2)}$ and $I_k^{(3)}$ refer to the pixel brightnesses measured under light sources 1, 2 and 3 respectively.

3. Two-dimensional BRDF Estimation

The aim of this section of the paper is to use simulated annealing optimization to fit a 2D surface (a representation of the BRDF) to the 3D histogram of surface normals and pixel intensities. The starting point is the field of zenith and azimuth angles estimated using the technique above, and the measured pixel brightnesses, I . We will be referring to this data collectively as a set of Cartesian triples:

$$\mathcal{D} = \{(\theta_d, \alpha_d, I_d) ; d = 1, 2, \dots, |\mathcal{D}|\} \quad (4)$$

The BRDF, $f(\theta_i, \theta_r, \alpha_i, \alpha_r)$, for a particular material is the ratio of reflected radiance to incident irradiance for any illumination and viewing directions. It is measured per unit solid angle per unit foreshortened area and is usually given by [4]

$$f(\theta_i, \theta_r, \alpha_i, \alpha_r) = \frac{L_r(\theta_r, \alpha_r)}{L_i(\theta_i, \alpha_i) \cos \theta_i d\omega} \quad (5)$$

where the subscripts i and r denote incidence and reflectance, and $d\omega$ is the differential solid angle subtended by the light source.

We do not attempt to estimate the reflectance function directly, but instead estimate the *radiance function*, $L_r(\theta_i, \theta_r, \alpha_i, \alpha_r)$. The radiance function incorporates the $\cos \theta$ term in (5), so that $L_r(\theta_i, \theta_r, \alpha_i, \alpha_r) \propto f(\theta_i, \theta_r, \alpha_i, \alpha_r) \cos \theta$. For the remainder of this paper, we therefore concentrate on the estimation of the radiance function. The incident radiance is unknown since we work with an uncalibrated light source. Consequently, we estimate the reflected radiance up to an unknown constant. Throughout this paper, all surfaces are assumed to exhibit isotropic reflectance properties.

3.1. Simulated Annealing Algorithm Paradigm

Simulated annealing can be stated generally as follows [15]. Given a configuration space, a system state, s , and an objective function, $\varepsilon(s)$, simulated annealing finds the optimum state s_{opt} that minimizes $\varepsilon(s)$. That is, the algorithm searches for

$$s_{\text{opt}} = \underset{s}{\operatorname{argmin}} \varepsilon(s) \quad (6)$$

In our algorithm, s_{opt} is to represent the BRDF, or strictly speaking, the radiance function.

Let the current state of the system be s and the “neighbour state” (i.e. the state after a perturbation) be s' . The probability of accepting this new state is usually given by the Boltzmann probability:

$$p_B = \begin{cases} 1 & \Delta\varepsilon \leq 0 \\ \exp(-\Delta\varepsilon/T) & \Delta\varepsilon > 0 \end{cases} \quad (7)$$

where $\Delta\varepsilon = \varepsilon(s') - \varepsilon(s)$ and T is the “temperature”. The process of forming perturbations and accepting the new

state with probability p_B continues as T is reduced (typically logarithmically) until the state converges.

3.2. Potential field

The technique must fit a surface, representing a 2D slice of the BRDF, to 3D histogram data. The raw data for the histogram comes from the set of Cartesian triples, \mathcal{D} (4), acquired using the method described in Section 2. Simulated annealing is used by repeatedly perturbing the surface inside a “potential field” derived from the histogram. Clearly, regions of the histogram with high frequencies should have an associated low energy and vice-versa.

We now detail the steps necessary to acquire this potential field. Firstly, it is possible to increase the data density in the histogram by noting that for isotropic surfaces

$$f(\theta, \alpha_L - \Delta\alpha) = f(\theta, \alpha_L + \Delta\alpha) \quad (8)$$

where $\Delta\alpha \in [0, 180^\circ]$ is the absolute difference between the light source azimuth angle, α_L , and the surface azimuth angle, α . With this in mind, we replace our α values in \mathcal{D} with the reduced azimuth angles given by [2]

$$\Delta\alpha = \min(|\alpha - \alpha_L|, 360^\circ - |\alpha - \alpha_L|) \quad (9)$$

Define the histogram contents as $HC_{a,b,c} = |\mathcal{R}_{a,b,c}|$, where

$$\mathcal{R}_{a,b,c} = \left\{ \left(\theta_d, \Delta\alpha_d, I_d \right); \begin{array}{l} \theta_d \in \left[\theta_a^{(\text{bin})} - \frac{90^\circ}{2k_{\max}}, \theta_a^{(\text{bin})} + \frac{90^\circ}{2k_{\max}} \right) \wedge \\ \Delta\alpha_d \in \left[\Delta\alpha_b^{(\text{bin})} - \frac{180^\circ}{2k_{\max}}, \Delta\alpha_b^{(\text{bin})} + \frac{180^\circ}{2k_{\max}} \right) \wedge \\ I_d \in \left[I_c^{(\text{bin})} - \frac{255}{2k_{\max}}, I_c^{(\text{bin})} + \frac{255}{2k_{\max}} \right) \end{array} \right\} \quad (10)$$

Here, the sets $\{\theta_a^{(\text{bin})}; a=1, 2, \dots, k_{\max}\}$, $\{\Delta\alpha_b^{(\text{bin})}; b=1, 2, \dots, k_{\max}\}$ and $\{I_c^{(\text{bin})}; c=1, 2, \dots, k_{\max}\}$ are uniformly spaced bin centres for zenith angles, reduced azimuth angles and intensities respectively. We arbitrarily use $k_{\max} = 32$ for this paper.

We wish to convert the histogram into a potential field, $\Phi(HC)$, using as simple a means as possible in order to preserve a direct relation between the histogram frequency and the energy. This is achieved by first negating the bin contents, HC , so that greater bin contents correspond to lower energies. The result is then convolved with a 3D Gaussian smoothing kernel, Ξ . The original histogram frequencies are then subtracted from this smoothed version. This step ensures that the potential field corresponding to non-zero histogram contents have significantly lower energy than surrounding areas. The potential field is, therefore, given by

$$\Phi(HC) = -HC * \Xi - HC \quad (11)$$

where $*$ denotes the convolution operation. Fig. 4 shows a potential field generated in this way.

3.3. Representation of the BRDF

The BRDF is to be represented in the simulated annealing paradigm by a discretely sampled set of points in Cartesian space, $\hat{s} = \{(\theta_k, \Delta\alpha_k, I_k), k=1, 2, \dots, k_{\max}^2\}$. This corresponds to the state s being optimized by (6). To fit the optimum surface, we require an initial state for the surface and a means to perturb the surface. It was found that, for the completed algorithm, there was virtually no dependence on the initial state. We therefore simply initialize s as a plane with corners at the following co-ordinates: $(\theta, \Delta\alpha, I) = (0, 0, 255), (90^\circ, 0, 0), (0, 180^\circ, 255)$ and $(90^\circ, 180^\circ, 0)$. The plane is sampled at k_{\max} equally spaced points for a given $\Delta\alpha$ and at k_{\max} equally spaced values of $\Delta\alpha$. This means that the number of points in the initial state is $k_{\max}^2 = 1024$. This number is fixed throughout the annealing process.

Each perturbation is performed as follows. Let the points in the perturbed state be \hat{s}' . We wish to disturb the surface such that a random point is chosen as the epicentre of the perturbation. Surrounding points are then disturbed by an




Figure 4. A selection of slices of the potential field for a porcelain object using $\theta_L = 45^\circ$ (lighter areas have higher energy).

amount that decreases with the distance to the epicentre. The index of the random point is an integer chosen from a uniform distribution in the interval $[1, k_{\max}^2]$.

We choose the perturbation to be Gaussian-like. That is, the point chosen as the epicentre, $(\theta_{\text{epi}}, \Delta\alpha_{\text{epi}}, I_{\text{epi}})$, moves the greatest distance, Υ_{epi} , and surrounding points move by an amount that decays as a Gaussian. Consider the motion of the epicentral point. To preserve the stochastic nature required for simulated annealing, we let the magnitude of the perturbation, Υ_{epi} , also be chosen from a Gaussian distribution. We restrict the motion of any point, k , (at the epicentre or otherwise) to be along the straight line that connects the points $(\theta_k, \Delta\alpha_k, I_k)$ and $(0, \Delta\alpha_k, 0)$. In other words, points on the surface can only move *directly* towards the $\theta = I = 0$ line or directly away from it. This ensures that points remain evenly spread across the surface.

The Gaussian PDF used to select Υ_{epi} has zero mean. When $\Upsilon_{\text{epi}} < 0$, the perturbation is towards the $\theta = I = 0$ line, while $\Upsilon_{\text{epi}} > 0$ indicates the opposite. The distance moved by point k is therefore

$$\Upsilon_k = \Upsilon_{\text{epi}} \exp\left(-\frac{(\theta_{\text{epi}} - \theta_k)^2 + (\Delta\alpha_{\text{epi}} - \Delta\alpha_k)^2 + (I_{\text{epi}} - I_k)^2}{2\sigma^2}\right) \quad (12)$$

where σ is the Gaussian standard deviation.

3.4. Energy of State

We now have the configuration space and move class required for simulated annealing. The next ingredient is a means to calculate the energy of a particular state. The total energy of state that we use, $\varepsilon_{\text{tot}}(\hat{s})$, has two terms: the “bin contents energy,” $\varepsilon_{\text{bc}}(\hat{s})$, derived from the histogram, and the “elastic energy,” $\varepsilon_{\text{el}}(\hat{s})$, used to maintain a smooth surface.

The bin contents energy is defined as the sum of the potentials at each point on the current state. This energy is therefore given by

$$\varepsilon_{\text{bc}}(\hat{s}) = \sum_{k=1}^{k_{\max}^2} \Phi(\mathbf{s}_k) \quad (13)$$

where \mathbf{s}_k is a vector from the origin to the point \hat{s}_k in the current solution and $\Phi(\mathbf{s}_k)$ is the potential at location \mathbf{s}_k . The potential field, Φ , is fixed and given by (11). Using (13) as the total energy of state, without the introduction of the elastic term, we found that some areas of the fitted surfaces become invariably and unrealistically irregular. We therefore introduce the elastic energy term, which we describe below.

In general, the elastic energy of a surface [3] quantifies its bending energy and is based on curvature. Larger curvatures result in higher elastic energies, thus promoting smooth solutions to the problem being solved. For our case,




Figure 5. A barycell (large shaded area) surrounding point \mathbf{s}_k . The eight neighbouring points are $\{\mathbf{s}_l\}$, which are here bonded to \mathbf{s}_k by dark broken lines (the light broken lines are used to bond other points). The barycentres, B_{kl}^- and B_{kl}^+ , chevron area, A_{kl} , and angles, ϑ_{kl}^- and ϑ_{kl}^+ , associated with point \mathbf{s}_l are indicated.

the problem is to find a surface, S , embedded in \mathbb{R}^3 , that minimizes the integral [6]

$$\varepsilon_{\text{el}}(S) = \int_S (a\kappa^2 + b) dA \quad (14)$$

where κ is the mean curvature and dA is the differential unit of surface area. The second term in (14) results in a multiple of the area. We experimented with different values of b for our work but found that there was negligible overall impact. We therefore set $a = 1$ and $b = 0$, so our elastic energy becomes

$$\varepsilon_{\text{el}}(S) = \int_S \kappa^2 dA \quad (15)$$

To calculate the integral in (15), a method is needed to approximate the mean curvature at any surface point. Given the discrete nature of our problem, the curvature, κ_k , at point k , is assumed to be constant over the area, A_k , of the barycell that contains the point. A barycell, as shown in Fig. 5, is defined as to connect the midpoints of the “bonds” between \mathbf{s}_k with its eight neighbours and the barycentres (centres of mass) of the triangles formed by these neighbours. Consider the point at location $\mathbf{s}_k = (\theta_k, \Delta\alpha_k, I_k)^T$. Let the point be bonded to eight neighbours, $\{\mathbf{s}_l(k); l \in \text{Nei}(k)\}$. The neighbours, $\text{Nei}(k)$, of a particular point, k , are fixed throughout the annealing process, and are defined as the eight closest points at the *initial state*. Fixing the bonds in this manner means that it is unlikely for any two points that are initially close together, to drift far apart at low temperatures.

It can be shown [9] that the mean curvature of a surface at point k can be calculated by

$$\kappa_k = \frac{1}{4A_k^{(\text{tot})}} \sum_{l \in \text{Nei}(k)} (\cot \vartheta_{kl}^- + \cot \vartheta_{kl}^+) |\mathbf{s}_k - \mathbf{s}_l| \quad (16)$$

where $A_k^{(\text{tot})}$ is the barycell area for point k and the angles ϑ_{kl}^- and ϑ_{kl}^+ are defined in Fig. 5.

The barycentric cell area is calculated as follows. For point s_k , we have the set of neighbours, $\{s_l; l \in \text{Nei}(k)\}$. The barycell can be divided into a set of chevron shapes for each l , one of which is highlighted in Fig. 5. If the area of each chevron is A_{kl} , then the total cell area is clearly

$$A_k^{(\text{tot})} = \sum_{l \in \text{Nei}(k)} A_{kl}$$

The location of the barycentres of the triangles associated with point k are

$$\mathbf{B}_{kl}^+ = \frac{\mathbf{s}_k + \mathbf{s}_l + \mathbf{s}_l^+}{3} \quad \mathbf{B}_{kl}^- = \frac{\mathbf{s}_k + \mathbf{s}_l + \mathbf{s}_l^-}{3} \quad (17)$$

where \mathbf{s}_l^+ refers to the neighbouring point that appears next in a clockwise direction and \mathbf{s}_l^- is in the anti-clockwise direction around \mathbf{s}_k . Using the vector product, we can write

$$A_k^{(\text{tot})} = \frac{1}{4} \sum_{l \in \text{Nei}(k)} (|(\mathbf{s}_l - \mathbf{s}_k) \times (\mathbf{B}_{kl}^- - \mathbf{s}_k)| + |(\mathbf{s}_l - \mathbf{s}_k) \times (\mathbf{B}_{kl}^+ - \mathbf{s}_l)|) \quad (18)$$

Collectively, (16) to (18) allow for the calculation of a set of mean curvatures for each point in \hat{s} . The elastic energy of a state is then found using (15) where the differential area is the barycell.

We now need to combine the bin contents energy, $\varepsilon_{\text{bc}}(\hat{s})$, given by (13), with the elastic energy $\varepsilon_{\text{el}}(\hat{s})$, given by (15), to give the total energy of state $\varepsilon_{\text{tot}}(\hat{s})$. We choose to construct the total energy by using $\varepsilon_{\text{bc}}(\hat{s})$ as our starting point. The elastic energy can then act as a multiplier so that the energy increases for rougher surfaces. Since the elastic energy $\varepsilon_{\text{el}}(\hat{s})$ has a minimum of zero and we would like the total energy to be equal to the bin contents energy for the smoothest possible surface (a plane), the following form is adopted:

$$\varepsilon_{\text{tot}}(\hat{s}) = \varepsilon_{\text{bc}}(\hat{s}) (\varepsilon_{\text{el}}(\hat{s}) + 1) \quad (19)$$

3.5. Fine Tuning

This section concludes with a discussion of a few details to help the algorithm obtain best results. These include extra constraints on the BRDF and the choice of parameters for the cooling schedule.

Shadow Boundary Constraint

It is possible to provide an extra constraint on the BRDF by consideration of the shadow boundary. Under the distant light source assumption, the unit vector in the direction of the light source, at any point, is

$$\mathbf{n}_L = \begin{pmatrix} \cos \alpha_L \sin \theta_L \\ \sin \alpha_L \sin \theta_L \\ \cos \theta_L \end{pmatrix} \quad (20)$$

Taking the scalar product of this and the normal vector, and noting that in shadowed areas the angle between these two vectors exceeds 90° , we arrive at the following condition:

$$\cot \theta + \tan \theta_L \cos(\Delta\alpha - \alpha_L) \geq 0 \quad (21)$$

If this condition is not met, then that part of the surface must be in shadow.

The condition in (21) is applied by detecting points in the current state, \hat{s} , that do not satisfy this condition and mapping them onto the nearest point in the potential field where the condition is satisfied. Note that this does not overcome the problem of cast shadows. These occur where one part of the surface casts a shadow over another part that would otherwise be illuminated.

Monotonicity Constraint

Previous studies [11] have shown that most BRDFs are monotonic in $\Delta\alpha$ for a given zenith angle. That is, the following inequality can be enforced:

$$\frac{\partial f(\theta, \Delta\alpha)}{\partial(\Delta\alpha)} \leq 0 \quad \forall \Delta\alpha, \theta \quad (22)$$

This is used as a final constraint, which is appended to the end of the algorithm. It turned out that the simulated annealing algorithm returned BRDFs that largely satisfy this condition anyway, so we enforce the constraint using as simple a method as possible. The following transformation was applied to achieve this task:

$$I_i^{(\theta)} \leftarrow \max \{I_j^{(\theta)}; j = i, i+1, \dots, i_{\text{max}}\} \quad (23)$$

where $\{I_i^{(\theta)}; i = 1, 2, \dots, i_{\text{max}}\}$ is the set of intensities for points in \hat{s} with zenith angle θ (actually within a small range of θ). The indices i in $\{I_i^{(\theta)}\}$ run such that $\Delta\alpha_{i+1} > \Delta\alpha_i \forall i$. The algorithm also smooths the data by applying moving averages to the set $\{I_i^{(\theta)}\}$ to remove any sudden gradient variations.

Temperature Control

There are three aspects of temperature control that must be considered in simulated annealing algorithms: the initial temperature, the cooling schedule and the convergence criterion. We currently use the following simple rule as a convergence criterion: when the fraction of transition probabilities equal to zero (< 0.01) exceeds 75% as the temperature drops by an order of magnitude, then the surface has converged. To *guarantee* convergence to the global minimum, the initial temperature must be set high and the rate of temperature reduction should be very low. All our results were obtained using an initial temperature of 10 and a cooling factor of 0.999.

In future work, the initial temperature could be set using a few random perturbations to estimate a small number of




Figure 6. Annealing progression from an initial state at $T = 10$ to a final state at $T = 0.01$. Points not connected by lines are at locations with zero bin contents energy.




Figure 7. Spheres rendered using the BRDFs estimated from the bear model with (a) $\theta_L = 45^\circ$ and (b) $\theta_L = 60^\circ$. The white lines indicate the calculated shadow boundary.

$\Delta\epsilon$ values. These values could then be used with (7), allowing to set an initial temperature that would give an even distribution of probabilities. The choice of the cooling factor is more difficult. Its selection should ideally be determined by the specific application. If efficiency is important, then a lower value should be adopted compared to when this is not the case.

4. Experiments

The results presented in this section use the illumination configuration shown in Fig. 3. The laboratory walls were matte black so there were no environmental inter-reflections. Fig. 6 shows the state, \hat{s} , of the surface, at various temperatures during the annealing process when applied to the porcelain bear model from Fig. 2. The large indentation to the upper-right of the initial state is due to the enforcement of the shadow boundary constraint (21). The algorithm took approximately three minutes to run.

Fig. 7 shows two spheres that were rendered using the radiance functions calculated from the porcelain bear. Here, we used both $\theta_L = 45^\circ$ and $\theta_L = 60^\circ$. Fig. 8 shows a comparison of the intensities predicted by the estimated radiance functions and ground truth (estimating an object of known shape). Figs. 7 and 8 together, show that the method has recovered the BRDF well enough for visually realistic rendering, but that more work is needed before it can be used for high-precision computer vision applications.

We have also conducted experiments with colour im-




Figure 8. Intensities as a function of the distance from the rendered sphere centre in (a) the vertical direction (in relation to Fig. 7) and (b) the horizontal direction. The thick lines indicate ground truth while the thin lines are the estimates. Solid lines have $\theta_L = 45^\circ$ and broken lines have $\theta_L = 60^\circ$.




Figure 9. (a) Image of a slightly rough yellow plastic duck. (b) Sphere rendered using the recovered BRDF.

ages. These involved applying the algorithm independently to each of the red, green and blue colour channels. One problem encountered with the *independent* estimations for each colour channel was that the colour balance present in the raw images did not always match that in the recovered BRDFs. A simple method we adopted to counter this issue was to linearly rescale the colour components of the BRDFs. This was done in such a way that the total colour ratios in the BRDFs matched the ratios in the raw image.

The colour BRDFs have been estimated for a yellow plastic duck and an orange under conditions where $\theta_L = 45^\circ$. Fig. 9 shows the estimated BRDF of the plastic duck, the surface of which is slightly rough. The rendered image looks realistic and appears to be made of the same material as the object. Since the object is slightly translucent, a small part of object at the base is visible where it would otherwise have been in shadow.

Fig. 10 shows results for an orange. On this occasion the




Figure 10. (a) Image of an orange. (b) Sphere rendered using the recovered BRDF.




Figure 11. Intensities of the rendered spheres of the duck (thin solid line) and the orange (thin broken line). (a) relates to vertical cross-sections of the spheres while (b) is for horizontal cross-sections. The Lambertian curve is also shown (thick broken line).

texture and roughness have caused greater problems for the algorithm and the results are less impressive. In particular, the colour balance is still incorrect for a few small regions. This is mainly due to the sparsity of data in certain regions of the histogram. Despite this, the results show good potential for the technique, even for these complicated surfaces.

The greyscale intensities of cross-sections of the plastic and orange rendered spheres are plotted in Fig. 11. The figure also shows the intensities predicted by the Lambertian model. Note that the greyscale versions of the BRDFs are approximately the same for the horizontal cross-section. Note also that they take a form which is more complicated than one would expect. A potential method to overcome this, is to fit the data to a theoretical model, such as the Oren-Nayar BRDF [12].

5. Conclusion

This paper has presented a new method for reflectance function estimation, which relies on polarization information. The method first estimates the surface zenith and azimuth angles based on the degree of polarization and the phase angle of the reflected light. The proposed algorithm then constructs the 3D histogram of surface angles and pixel brightnesses. A surface is then fitted to this data using simulated annealing and elastica. The result is a 2D function which can predict the reflected radiance for any zenith and azimuth angle combinations for the given illumination.

The results are promising and allowed for good image rendering for porcelain and plastic. In future work, we will use a greater range of light source positions to increase both

the density of the histogram data and the portion of the BRDF estimated (e.g. to incorporate the incident light direction into the recovered radiance function).

References

- [1] G. A. Atkinson and E. R. Hancock. Shape estimation using polarization and shading from two views. *IEEE Trans. Patt. Anal. Mach. Intell.*, 29:2001–2017, 2007.
- [2] G. A. Atkinson and E. R. Hancock. Surface reconstruction using polarization and photometric stereo. In *Proc. CAIP*, pages 466–473, 2007.
- [3] A. M. Bruckstein, R. J. Holt, and A. N. Netravali. Discrete elastica. *Applicable Analysis*, 78:453–485, 2001.
- [4] D. A. Forsyth and J. Ponce. *Computer Vision, A Modern Approach*. Prentice Hall, Upper Saddle River, NJ, 2003.
- [5] E. Hecht. *Optics*. Addison Wesley Longman, third edition, 1998.
- [6] L. Hsu, R. Kusner, and J. Sullivan. Minimising the squared mean curvature integral for surfaces in space forms. *Experimental Mathematics*, 1:191–207, 1992.
- [7] H. Li, S. C. Foo, K. E. Torrance, and S. H. Westin. Automated three-axis gonioreflectometer for computer graphics applications. *Optical Engineering*, 45:043605, 2006.
- [8] S. R. Marschner, S. H. Westin, E. P. F. LaFortune, and K. E. Torrance. Image-based bidirectional reflectance distribution function measurement. *Applied Optics*, 39:2592–2600, 2000.
- [9] M. Meyer, M. Desbrun, P. Schröder, and A. H. Barr. Discrete differential-geometry operators for triangulated 2-manifolds. In *Proc. VisMath*, pages 35–57, 2002.
- [10] D. Miyazaki, R. T. Tan, K. Hara, and K. Ikeuchi. Polarization-based inverse rendering from a single view. In *Proc. ICCV*, volume 2, pages 982–987, 2003.
- [11] J. A. Ogilvy. *Theory of Wave Scattering from Random Rough Surfaces*. Adam Hilger, Bristol, 1991.
- [12] M. Oren and S. K. Nayar. Generalization of the Lambertian model and implications for machine vision. *Intl. J. Comp. Vis.*, 14:227–251, 1995.
- [13] B. T. Phong. Illumination for computer generated pictures. *Comm. ACM*, 18:449–455, 1975.
- [14] A. Robles-Kelly and E. R. Hancock. Estimating the surface radiance function from single images. *Graphical Models*, 67:518–548, 2005.
- [15] P. Salamon, P. Sibani, and R. Frost. *Facts, Conjectures and Improvements for Simulated Annealing*. Society for Industrial and Applied Mathematics, 2002.
- [16] K. Torrance and M. Sparrow. Theory for off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.*, 57:1105–1114, 1967.
- [17] L. B. Wolff. Diffuse-reflectance model for smooth dielectric surfaces. *J. Opt. Soc. Am. A*, 11:2956–2968, 1994.
- [18] L. B. Wolff and T. E. Boult. Constraining object features using a polarisation reflectance model. *IEEE Trans. Patt. Anal. Mach. Intell.*, 13:635–657, 1991.
- [19] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19:139–144, 1980.