

OPTICAL ESTIMATION OF THE 3D SHAPE OF A SOLAR ILLUMINATED, REFLECTING SATELLITE SURFACE

Julian Antolin, Zhixian Yu, and Sudhakar Prasad

Department of Physics and Astronomy, University of New Mexico, Albuquerque, NM 87131

The spatial distribution of the polarized component of the power reflected by a macroscopically smooth but microscopically roughened curved surface under highly directional illumination, as characterized by an appropriate bi-directional reflectance distribution function (BRDF), carries information about the three-dimensional (3D) shape of the surface. This information can be exploited to recover the surface shape locally under rather general conditions whenever power reflectance data for at least two different illumination or observation directions can be obtained. We present here two different parametric approaches for surface reconstruction, amounting to the recovery of the surface parameters that are either the global parameters of the family to which the surface is known *a priori* to belong or the coefficients of a low-order polynomial that can be employed to characterize a smoothly varying surface locally over the observed patch.

1 Introduction

A smooth surface reflects light incident on it in an angular distribution that contains information about its local 3D shape around its specular-reflection, or glint, points. For a fixed illumination direction, as the surface is roughened microscopically, the angular distribution, still centered about the glint points, broadens spatially around them in a manner that corresponds almost uniquely to the varying surface shape over the visible reflectance footprint. The spatial distribution of specularly reflected light, largely corresponding to single geometrical reflections from random but smoothly oriented facets on the roughened surface, carries sensitive information about the 3D shape and the degree of roughness of the surface.

The machine vision, artificial intelligence, and psychophysics communities have a longstanding and active interest in the problem of shape recovery, and have adopted two approaches, namely shape-from-shading (SFS) and photometric stereo (PS) [1–4], that exploit the spatial distribution of largely Lambertian, diffuse reflections from reflecting surfaces. The non-Lambertian character of diffuse reflections from rough surfaces has also been carefully modeled [6, 7] by allowing isotropically reflecting, but randomly oriented microfacets on such surfaces, which presents an additional complication that must be carefully accounted for in any SFS and PS approaches.

Here we present a more rigorous, physics-based approach for local shape recovery, one which in fact exploits the spatial distribution of specularly reflected light, rather than that of the diffuse albedo, that results from two different illuminations of a single 3D shape. The specular component of reflected light carries sensitive information about the 3D shape and roughness characteristics of the surface, while the diffuse component carries only gross information about surface shape. By a polarization-differencing strategy applied to reflectance data, the essentially unpolarized nature of the latter component, in sharp contrast with the highly polarized character of the former, allows it to be removed thus enabling the possibility of sensitive surface shape recovery over all glint spots.

Under the condition that illumination and observation are along fixed directions, which we assume here, since the specular BRDFs in the s and p linear polarization channels are simply proportional to each other via the ratio of the power Fresnel reflection coefficients for the two channels, their difference, defining the Stokes parameter s_1 , is also proportional to the specular BRDF in either channel, both proportionality constants being independent of the location of the observed surface point. Since the polarization character of the specular BRDF is spatially uniform under this condition, we drop any reference to the polarimetric properties of the specular BRDF, with the understanding that polarization differencing has been done to extract the specular reflection fraction of the image data, focusing simply on its spatial character that constrains the local 3D shape around glint points. A previous work [8] reported on the use of polarization differencing, combined with a compressive spectral-polarimetric acquisition protocol, to demonstrate the recovery of global shape and roughness parameters from simulated data that included noise but no instrumental or turbulence-induced blur. We extend this work in two different ways - first by including a telescope blur and second, by including two different parameterizations of the surface, specifically the superquadric one [9] that was already covered in Ref. [8] and a local polynomial one. Future work will consider degradations of our surface estimation from partially compensated, turbulence-induced random blur as well as more general roughness models.

The second characterization is the more general of the two, requiring *no* prior knowledge of the surface other than that it admits a smooth patch approximation locally around a brightly reflecting point, *i.e.*, a glint spot. By

contrast, the first relies on the prior knowledge of the surface in question as belonging to a known parametrized family. Nevertheless, our first characterization is expected to apply to other low-dimensional global parameterizations of the surface as well. In both descriptions, we must estimate surface parameters, namely the SQ parameters and the orientation Euler angles of the surface in the first instance and the coefficients of a modest-degree bivariate polynomial governing the surface shape locally in the second, as well as the degree of surface roughness. The generality of the second description is further extended by noting that shape estimation can be enlarged over the entire observable area of the surface by stitching together individually estimated local patches contiguously.

In this work, the surface roughness, which extends the footprint of specular reflection around what would otherwise be isolated glint points carrying no local shape information, is described by a single parameter which is also recovered in all three estimation techniques. We use a roughness model based on the simplifying assumption that the surface consists of V shaped microscopic indentations with planar surfaces, or microfacets, of a characteristic linear dimension, say l , that are randomly and isotropically oriented with respect to the underlying smooth mean surface [10]. The random orientation of a microfacet may be characterized statistically [11] by means of a standard deviation σ_h of the height differential across the facet above the mean surface. We shall take, as commonly done, the distribution of height differentials to be a zero-mean, circular Gaussian probability density function (PDF) with standard deviation σ_h , and the ratio $\chi \equiv \sigma_h/l$ to be the characteristic surface roughness parameter. We shall assume further that l is large compared to the illumination wavelength so that each facet scatters light essentially by means of specular reflection. Other roughness models based on different geometrical primitives from the planar microfacets assumed here, such as dual-scale microfacets [12], curved microfacets [13], hemispherical pits [14], etc., and even a statistical admixture of them may be treated analogously.

The main focus of this work is on the recovery of 3D shape of solar-illuminated space objects such as man-made satellites illuminated in the terminator mode and observed from a ground station. While satisfying our assumption of fixed illumination and observation directions, it has the complication that we must include image blur and noise in the specular reflectance data. We shall assume here that we know the blur function perfectly, as might be approximately true for either a perfect adaptive-optics (AO) based wavefront-corrected telescope or a space-based telescope. Future work will address the case of random turbulence-induced blurs that would require a multi-frame blind deconvolution (MFBD) approach [15] to tease out uncorrupted BRDF data.

2 Problem Setup

We use an expression for the BRDF derived by Sun [11], which was extended by Hyde, *et al.* [16] to capture the polarimetric dependence of the specular BRDF, under the microfacet-based roughness model,

$$F(\theta_i, \theta_r, \phi, \beta, \alpha) = \frac{P(\alpha)G(\theta_i, \theta_r, \phi, \beta)}{4 \cos \theta_i \cos \theta_r \cos \alpha} \quad (1)$$

where β is half the angle between the illumination and observation directions, θ_i and θ_r are the angles that the illumination and observation directions make with the mean surface normal, ϕ is the difference of the azimuthal angles of the two directions with respect to the mean surface normal, and α is the angle between the mean surface normal and the (fixed) bisector of the angle between the two directions. By Snell's law, this bisector is also the direction of the normal of those micro-facets that cause observable specular reflections in the micro-facet based roughness model of interest here. The function $P(\alpha)$ is the PDF that represents the probability, per unit solid angle, for the normal of the microfacet at a surface point to make an angle α with the mean surface normal locally at any point, which gives expression (1) a local character. A good working model for the distribution of facet-normal orientations is based on a Gaussian height distribution [7, 11]

$$P(\alpha) = \frac{1}{4\pi\chi^2 \cos^3 \alpha} e^{-\tan^2 \alpha / 4\chi^2}, \quad (2)$$

where χ is the characteristic roughness parameter defined earlier. The symbol G denotes the shadowing-masking function that represents the fraction of incident light eliminated from reaching the observer via single specular reflections due to either one tilted facet partially blocking the illumination of another (shadowing) or the specular reflection from a facet being partially masked by another (masking) because of oblique incidence and reflection directions. These two effects are modeled approximately by a single simplified function [17]:

$$G(\theta_i, \theta_r, \phi, \beta) = \min \left(1, \frac{2 \cos \alpha \cos \theta_i}{\cos \beta}, \frac{2 \cos \alpha \cos \theta_r}{\cos \beta} \right). \quad (3)$$

From (1), one can calculate the s_1 Stokes component of the specular BRDF [16] by multiplying the right-hand side (RHS) of this expression by a fixed factor, $0.5(|r_s|^2 - |r_p|^2)$, where r_s and r_p are the Fresnel amplitude reflectivities for single reflection off planar facets for unpolarized sunlight incidence along a fixed direction and reflection along the fixed observer direction. This fact is independent of the local 3D shape of the mean surface.

Minimum Number of BRDF Datasets for Recovery of Surface Normals Determining the unit normal, \hat{n} , at a point, \vec{r} , of the mean surface amounts to evaluating its components n_x, n_y, n_z in a fixed Cartesian coordinate system. We conveniently choose this coordinate system to have its z axis along the observation direction, $\hat{z} = \hat{e}_r$, and two orthogonal directions in the sensor plane, say along the pixel edges, as its x and y axes. Since the angle between the known solar-angle bisector, given by $\hat{s} = (1/2)(\hat{e}_i + \hat{e}_r)$, and the unknown mean-surface normal \hat{n} is α , on which the BRDF (1) depends most sensitively via $P(\alpha)$, we have

$$\hat{n} \cdot \hat{s} = \cos \alpha. \quad (4)$$

If α is small, as would be the case near a glint point, then the shadowing-masking factor G may be approximately set equal to 1 and the angles θ_i and θ_r to β in expression (1) for the BRDF, which would allow us to estimate from this expression $\cos \alpha$ that occurs on the RHS of relation (4). In view of this and since \hat{s} is known in advance, we need only two relations of the type (4) to solve for the two unknown components of the unit vector, \hat{n} , along with the “outwardness” condition, $n_z > 0$, required for surface points that are observable along the camera direction, which we denote as \hat{z} . With surface normals determined everywhere on the surface areas that are solar-illuminated and observed by the camera, we may then by integrating the normal, related to the gradient of the local shape function, over the surface recover the surface shape over these areas.

In the presence of noise, more than two datasets may be needed, but we avoid that requirement because of our low-dimensional parameterizations of the surface for which the number of parameters being estimated is small and two datasets with far more numerous pixels of non-redundant intensity values will more than suffice. The different BRDF frames of data may be generated by either changing the illumination direction \hat{e}_i itself or by changing the direction of observation \hat{e}_r . The latter alternative may be more practical for an object like a sun-illuminated satellite in LEO in the terminator mode if either two geographically separated observatories could make simultaneous observations or a single observatory could take a sequence of two (or more) observations as such objects tend to orbit rapidly and thus their elevation angle in the sky changes significantly even over a short time period. For definiteness, we shall assume in this paper two different illumination directions in the xz plane, but a fixed observation direction along \hat{z} .

3 Superquadric and Polynomial Families of Shapes

We consider here two different parameterizations of the 3D shape of an opaque, reflecting surface. The first employs a global surface function that is a generalization of the quadric surface, the so-called superquadric (SQ), while the second is based on the use of low-order bivariate polynomials, which assumes an analytic, differentiable surface, and is a local characterization about some fiducial point.

The SQ family of surfaces is defined in its canonical coordinates, (ξ, η, ζ) , defined to be along its symmetry axes as

$$\left[\left(\frac{|\xi|}{a_1} \right)^{\frac{2}{\epsilon_2}} + \left(\frac{|\eta|}{a_2} \right)^{\frac{2}{\epsilon_2}} \right]^{\frac{\epsilon_2}{\epsilon_1}} + \left(\frac{|\zeta|}{a_3} \right)^{\frac{2}{\epsilon_1}} = 1, \quad (5)$$

within the ranges, $|\xi| \leq a_1$, $|\eta| \leq a_2$, $|\zeta| \leq a_3$, $\epsilon_1 > 0$, $\epsilon_2 > 0$. The surface can be rotated and displaced with respect to the observation coordinate system, (x, y, z) , which can be specified in terms of three Euler angles, $(\alpha_0, \beta_0, \gamma_0)$, and coordinates of its center, (x_0, y_0, z_0) . Roughening the surface according to the microfacet model adds the roughness parameter, χ , introduced in (2) for a total of 12 parameters to be estimated. We shall, for simplicity, however take the surface to be centered in both coordinate systems at the origin, $(x_0, y_0, z_0) = (0, 0, 0)$, which leaves only 9 parameters to be estimated. The two coordinate systems are then linearly related to each other simply by the matrix of Euler rotation. The choice of shape parameters, $\epsilon_1 = \epsilon_2 = 1$, yields the ellipsoid family, while for values of the shape parameters $\epsilon_{1,2} < 1$ the superquadric is box-shaped and for values $\epsilon_{1,2} \geq 2$, the surface develops sharp edges and cusps. We shall see two examples of such surfaces in the next section.

On the other hand, the polynomial class of surfaces of N th order is defined, with respect to the (x, y, z) coordinates set by the Sun and the observer, by the function

$$z(x, y) = \sum_{n=0}^N \sum_{k=0}^n c_{k,n-k} (x - x_0)^k (y - y_0)^{n-k} \quad (6)$$

or, equivalently by

$$z(x, y) = \sum_{\substack{k+\ell \leq n \\ k, \ell \geq 0}} c_{k\ell} (x - x_0)^k (y - y_0)^\ell, \quad (7)$$

where x_0 and y_0 are the x and y coordinates of a point near the center of the glinting regions of the illuminated surface. Ignoring the lowest-order coefficient, c_{00} , which being a constant does not affect the the surface shape and the associated BRDF, leaves $N(N+3)/2$ polynomial coefficients at monomial orders $1, 2, \dots, N$ to be estimated. Along with χ , this means estimating 15-28 parameters for N between 4 and 6, which will suffice for our purposes.

4 Global Optimization via Cost-Function Alternation between Low-Order Constraints and Residuals

We introduce here a novel optimization approach for surface recovery that is based on the imposition of low-order physical constraints that are derived from the data themselves and requires no regularization. We start by noting that global optimization to recover highly embedded parameters, which in general is a highly non-linear and non-convex inverse problem, and is attended by stagnation in local minima. A number of approaches have been suggested in the literature, but they all depend on modifying the the problem to improve its convex relaxation from local minima. An example of such an approach is provided by [18] in which a branch-and-reduce approach is taken to improve the tightness of the vectors of lower and upper bounds on the parameters that yield the least residual. The problem is greatly exacerbated by the presence of noise and only the best solution that, via a fit-to-data least squares cost function, fits the data within noise (χ^2) levels is sought.

The approach we adopt here is based on an alternating minimization between a cost function that requires a fit of the estimate to the data only coarsely and a second cost function consisting of fine, pixel-wise data residuals, both cost functions having a least-squares form. The first cost function only imposes low-order constraints all derived from the data, specifically the condition, $\alpha = 0$, at each glint point, and the coordinates of the center-of-mass (CM) and the eigenvalues and eigenvectors of the matrix of second-order spatial moments of pixel intensities, on the estimates of the two datasets. By requiring a fit to these low-order constraints, the first cost function relaxes a stagnated local minimum solution obtained via the second cost function, forcing the solution vector to explore other potential locations of the global minimum where the estimate (model) must fit the data pixel-wise to a tolerance determined by the noise variance at each pixel. We found this approach of alternation between the two cost functions to work particularly well to determine the 10-30 parameters of our problem when we combine it with a pre-optimization step which we describe next. No regularization is needed in either cost function, and the constraints are derived from the BRDF data alone, so no prior bias is introduced into the inverse problem.

Necessity of a Pre-Optimization Step An important pre-requisite for the motion of any optimizer to lower the value of the cost function is at least a partial spatial overlap of the starting value of the model and the data. Without such overlap, the force responsible for the motion of the optimizer toward a minimum, which can expressed as the gradient of the residuals with respect to the parameter vector, will vanish. We guarantee this overlap for the present problem by requiring that the locations of the glint points be matched on the starting surface for each of the datasets. The starting surface was chosen to belong to a lower-dimensional subspace of the full family of surfaces, specifically an ellipsoid, determined by only six parameters, for the full nine-parameter SQ family and a quadratic surface with only five parameters as the starting surface for a polynomial surface of 4th-6th orders with 14-27 parameters. The four glint conditions for the two datasets, namely that the surface normal at each glint point match the known solar bisector direction corresponding to each dataset, are easily implemented with these quadratic surfaces that each have more than four free parameters, all of which we could quickly determine by a steepest-descent Matlab algorithm like *fmincon* via a least-squares fitting step. Such a pre-optimization step, as we found in our simulations, is essential for the convergence of the alternating runs of the final global optimization step to the global minimum.

5 Results of Simulation

We now describe the results of simulation for the recovery of the two different families of surfaces we considered in this work. The two entail different global optimizers - the SQ recovery fared poorly with local solvers like *fmincon*, which stagnated immediately at one of numerous local minima, but worked rather well with a global optimizer like *Pattern Search* (PS), but the polynomial recovery worked more efficiently with a local solver. This difference presumably has to do with the fact that the parameters of the SQ are fairly comparable in magnitude to one another, all being of the order 1 in our simulations, while the polynomial coefficients were quite disparate in their magnitude across the different monomial orders of the polynomial. Local gradient-descent solvers can accommodate these disparities well, but PS imposes a single step size for the change of the value of every parameter, say P in all, one at a time both forward and backward, in order to lower the value of the cost function along any of these $2P$ directions. If it succeeds then it doubles the step size for the next update, but if it cannot, then it typically halves the step size. Although PS is typically slower as it is exhaustive and does not utilize any information about local descent directions, it handles local minima better, particularly since constraints are alternately applied in the minimization of the data residuals in our approach.

5.1 Results for a SQ Surface

We simulated BRDF data for a large variety of SQ surfaces, *e.g.*, with $\underline{p} = (a_1, a_2, a_3, \epsilon_1, \epsilon_2, \epsilon_3, \alpha_0, \beta_0, \gamma_0, \chi)$, the vector of parameters to be estimated, taking its value to be $(4, 1, 0.5, 1.5, 1.5, 2, 1, 1, 4)$ for a needle-shaped surface and $(4, 2.5, 2, 1.5, 0.5, 2, 1, 1, 6)$ for a rounder surface, as shown in Fig. 1, for two different values of the solar half angle, β , equal to 0.2 and 0.6 *radian*. The difference-polarization Stokes component, s_1 , of the corresponding specular BRDF data for Gaussian image blur width, $w = 0.4$, and the peak mean photon number, $K_{max} = 1000$, is plotted in Fig. 2 for the dataset corresponding to $\beta = 0.6$ *radian* for the two surfaces (top row). The bottom row of this plot displays the estimates formed from our recovered SQ surfaces for the same value of β . Note that the BRDFs are quite extended on the two surfaces because of large surface roughness (small χ). The set of figures corresponding to the other value of β , not shown here for lack of space, has a very similar character, the estimates being visually indistinguishable from the truths.




Figure 1: (a) A needle-shaped SQ surface; (b) a rounder SQ surface

While the reconstructed SQ surfaces themselves are not displayed, the errors in its nine parameters are shown via a scatter plot for ten frames of Poisson data for each surface at two different peak mean photon numbers in the s_1 data for each view. We also show in Fig. 3 how the alternation between the fit-to-data cost function (*i.e.*, the residual) and fit-to-constraints cost function between successive runs of the PS solver drives the former (and latter) to lower values for the rounder surface shown in Fig. 1 (b). The data residual was normalized by dividing it by its expected value for a perfect fit, which can be easily calculated both for Poisson and Gaussian additive noise models.

5.2 Results for a Polynomial Surface Patch

We employed a SQ surface specified by the SQ parameter vector, $\underline{p} = (2, 2, 3, 1.2, 1.2, 0, 0, 0, 12)$, to give us the surface patch over which we simulated BRDF datasets for two closely spaced values of the solar half angle, namely $\beta = \pi/5$ and $\pi/5 + \pi/36$, for which the datasets had overlapping brightness over an appropriate 8×8 or 12×12 patch




Figure 2: The true noisy (a) and estimated (c) blurred difference-polarimetric BRDF datasets for $\beta = 0.6$ radian for the needle-shaped SQ surface; (b) and (d) refer to the corresponding BRDFs for the rounder SQ surface.

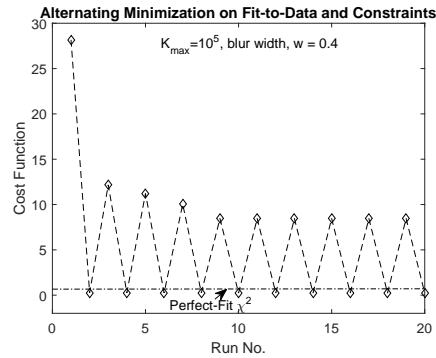


Figure 3: The alternation of cost-function minimization runs from fit-to-data residuals to fit-to-low-order-constraints

of the full 64×64 array on which the surface was constructed in the sensor plane. The Gaussian image-blur function had width equal to 5 pixels, or $5/16$ in the length units in which the sensor array had side length 4. The patch shape was parameterized by a 4th or 6th order polynomial, respectively, for the smaller or larger patch. The datasets shown in Fig. 4 refer to the smaller patch for the two solar half angles contain both image blur and Poisson shot noise with a peak SNR of value 10. The true values of the 14 (27) coefficients for the 4th (6th) order polynomial were calculated by minimizing a cost function equal to the squared 2-norm of the difference between the fitted polynomial and the true SQ over the patch.

The pre-optimized starting quadratic surface with five polynomial coefficients was determined by means of the solver *fminunc* ensuring that the two glint points in the two datasets were correctly located on the starting surface by fitting the glint condition, namely that at each such point the normal to the surface be along the solar-angle bisector. Both the starting surface and the true polynomial patch are displayed in Fig. 5 along side the surface estimated by the




Figure 4: Two sets of BRDF simulations for the two illumination directions with the Gaussian blur and additive Gaussian noise specified in the text for the 8×8 patch. (a) and (b) are the two sets of BRDF on the entire object, while (c) and (d) present a magnified view of those BRDFs on the chosen patch.

alternating optimization approach proposed in this paper. The only difference from the problem of recovering the SQ surface was in the use of *fminunc*, rather than PS, as the optimizer. Note the excellent shape approximation provided by the starting surface itself for both polynomial orders, which meant rapid convergence to the final surface estimate in our alternating-minimization approach.




Figure 5: The true polynomial surface, the pre-optimized quadratic starting surface, and the estimated polynomial surface, shifted relative to one another along z for ease of visualization, are shown for (a) 4th order and (b) 6th order polynomial parameterizations. The patch of overlap between the two datasets in the second case was larger, at 12×12 , due to larger roughness, requiring a higher-order polynomial for accurate representation.

In Fig. 6, we present a scatter plot of the squared error in the x and y components of the unit normals averaged over all $P \times P$ pixels of the patch, namely

$$E_{P^2} = \frac{1}{P^2} \sum_{P \times P} [(\hat{n}_x - n_x)^2 + (\hat{n}_y - n_y)^2], \quad (8)$$

where \mathbf{n} and $\hat{\mathbf{n}}$ are, respectively, the true and the finally estimated unit normal vectors at the patch pixels, and P is equal to 8 (12) for the 4th (6th) order polynomial surface patch. At all but the lowest PSNR values, the errors are quite quite tightly bunched over all 50 frames of noise realizations.

The inverse problem of recovering a set of polynomial shape coefficients from the distribution of normals which the two BRDF datasets determine, in effect, as we saw earlier, is likely to be approximately convex at least over a small enough patch containing the glint points for which the angle α between the normals and the solar-angle bisector to which the BRDF (1) is most sensitive near a glint point remains small. It is presumably this approximate convexity of the problem that yields good convergence of a local solver to the correct values when it is started from a good guess

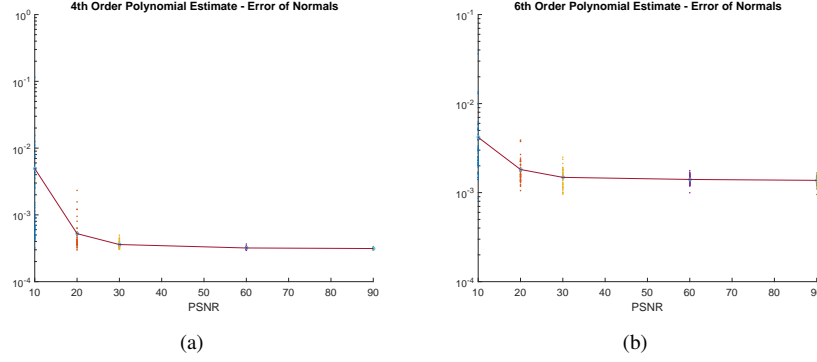


Figure 6: The squared error of normals averaged over the patch for the (a) 4th order 8×8 patch and (b) 6th order 12×12 patch for Poisson data.

of the surface determined by the pre-optimization step. For this reason it is unclear that given a good starting guess if, unlike the problem of SQ surface recovery, the alternation between fit-to-constraints and fit-to-data cost function minimizations is even necessary for the success of the problem of polynomial surface recovery.

A Different Starting Guess for the Polynomial Parameters We also devised a different approach to arrive at an alternate starting guess of the surface polynomial parameters. It was based on requiring an approximate fit of the distribution of candidate polynomial-surface normals to an estimate of the true surface normals determined as a first step by fitting the BRDF model with the BRDF data via the following minimization:

$$\hat{\mathbf{n}}^* = \underset{\hat{\mathbf{n}}}{\operatorname{argmin}} \sum_{k=1}^K \varphi(F_k - \hat{F}_k(\hat{\mathbf{n}})) \quad (9)$$

where F_k is the k th BRDF dataset, $\hat{F}_k(\hat{\mathbf{n}})$ is the corresponding BRDF estimate as a function of the normals and φ is the squared 2-norm (unweighted least-squares) function. Here $K = 2$. The roughness parameter, χ , was directly estimated from the largest value of one of the two BRDF datasets by the relation, $\chi = (16\pi F_{\max} \cos^2 \beta)^{1/2}$, which is obtained simply by setting $\alpha = 0$ in Eqs. (1) and (2). For BRDF data with blur and noise, such direct evaluation of χ is not accurate, but we can still use that value as a rough estimate, which we refine in the final optimization step, to estimate an approximate normal vector field.

We followed this first step by fitting a polynomial surface of the order being sought but only over the patch of interest. This is equivalent to the following least-squares minimization:

$$\hat{\mathbf{p}}_0 = \underset{\mathbf{p}_0}{\operatorname{argmin}} \|\hat{\mathbf{n}}^* - \hat{\mathbf{n}}(\mathbf{p}_0)\|_2^2, \quad (10)$$

where \mathbf{p}_0 is the $N(N+3)/2$ -dimensional vector of polynomial coefficients and the double vertical bars enclosing a vector represent its 2-norm. This second step generated a good starting guess for the polynomial coefficients with which we made a final run of the local solver *fminunc* to refine their estimates and that of the roughness parameter, χ . In this approach, we did not need any alternating minimization approach like the one used for SQ surface recovery.

The results of the recovery of the polynomial coefficients using this alternate approach are displayed below in Fig. 7 for the single pair of noisy, blurred BRDF datasets noise shown in Fig. 4. Visually, little differences can be discerned in shape between the two, but the small differences are characterized quantitatively and rather exhaustively in Figs. 8, which displays the squared error (8) in the unit normals averaged over a $P \times P$ sub-patch of the full 8×8 patch. We see from the figure that the inner 6×6 sub-patch normals are better determined, on average, compared to the outer square-annular patch with the remaining 28 pixels, for which the error, in a slight abuse of notation, is denoted as E_{28} . The low-order polynomial approximation of the surface is evidently more accurate, the closer a pixel on average is to the center of the patch. The upper pair of plots in this figure are scatter plots of the surface-normal error, E_{64} , over the patch for 50 realizations of Gaussian and Poisson data at different values of peak SNR in the BRDF datasets, while the lower pair of figures plot the mean values of E_{36} and E_{28} over these noise realizations for the two noise

models. We also note that while the error performance is similar to that of our previous approach, as a comparison of Figs. 8 and 6 shows, the current approach is considerably slower.




Figure 7: Shape comparison of the estimated polynomial patch to the true polynomial patch, the two shifted to each other along z for ease of visualization




Figure 8: Semi-log plots of E_{64} , E_{36} and E_{28} vs PSNR, with (a) and (b) being scatter plots of E_{64} with 50 realizations for Gaussian additive and Poisson shot noise, respectively, and (c) and (d) plots of the corresponding mean values of E_{36} and E_{28} over the 50 realizations.

6 Conclusions

We have presented two types of parametric surface recovery: a global estimation based on a superquadric (SQ) parameterization and a local polynomial patch estimation. We reported excellent, noise and blur robust recovery of surface parameters for both the SQ and polynomial representations of the surface in the vicinity of glint points in the solar-reflected light. This work has successfully advanced a novel global optimization protocol that by alternating between a fit-to-data based cost function and a fit-to-low-order constraints based cost function rapidly converges to the global minimum even for the case of 28 parameters that we estimated for a 6th-order polynomial. The use of a pre-optimization step to start with a surface in the sub-space of these two families on which the two BRDF datasets have guaranteed overlap was shown to be essential for this alternating global minimizer to reach the global minimum in an efficiently systematic manner. This new global optimizer should work with other low-dimensional parametric surface models just as well, and in fact may even be a method of choice more generally for a variety of different parametric problems in science and engineering requiring a global optimizer to navigate the minefield of local minima.

Our shape recovery methods are quite robust to known image blur. It is of much practical importance, however, to investigate how including a random component in the overall image blur, e.g., that introduced by any uncompensated atmospheric turbulence induced optical phase errors, would degrade the recovery of mean surface shape. Future work will address this problem.

References

- [1] H.K.P Horn, "Obtaining shape from shading information," in *The Psychology of Vision*, pp. 115-155 McGraw-Hill Book Co. (1975)
- [2] R.J. Woodham, "Analysing images of curved surfaces," *Artificial Intelligence*, vol. 17, pp. 117-140 (1981).
- [3] A. Hertzmann, S. M. Seitz, "Example-based photometric stereo: Shape reconstruction with general, varying BRDFs," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 27, pp. 1254-1264 (2005).
- [4] R. Zhang, P.-S. Tsai, J.E. Cryer, and M. Shah, "Shape from Shading: A Survey," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 21, pp. 690-706, 1999
- [5] J.T. Barron, J. Malik, "Shape, Illumination, and Reflectance from Shading," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, vol. 37, pp. 1670-1687, 2015
- [6] M. Oren and S. Nayar, "Generalizations of the Lambertian model and implications for machine vision," *Intern. J. Comp. Vision*, vol. 14, pp. 227-251 (1995).
- [7] B. van Ginneken, M. Stavridi, and J. Koenderink, "Diffuse and specular reflectance from rough surfaces," *Appl. Opt.*, vol. 37, pp. 130-139 (1998).
- [8] S. Prasad, Q. Zhang, and R. Plemmons, "Shape and pose recovery of solar illuminated surfaces from compressive spectral polarimetric image data," *Proc. AMOS Tech. Conf.* (2013), available at <http://www.amostech.com/TechnicalPapers/2013/POSTER/PRASAD.pdf>
- [9] A. Barr, "Rigid Physically Based Superquadrics," in *Graphic Gems III*, pp. 137-159, Academic Press (1992).
- [10] K. Torrance and E. Sparrow, "Theory for off-specular reflection from roughened surfaces," *J. Opt. Soc. Am.*, vol. 57, pp. 1105-1114 (1967).
- [11] Y. Sun, "Statistical ray method for deriving reflection models of rough surfaces," *J. Opt. Soc. Am. A*, Vol. 24, pp. 724-744 (2007).
- [12] J. Leader, "Analysis and prediction of laser scattering from rough-surface materials, *J. Opt. Soc. Am.*, vol. 69, pp. 610-628 (1979).
- [13] C. Sung and W. Eberhardt, "Explanation of the experimental results of light backscattered from a very rough surface," *J. Opt. Soc. Am.*, vol. 68, pp. 323-328 (1978).
- [14] S. Pont and J. Koenderink, "Bidirectional reflectance distribution function of specular surfaces with hemispherical pits," *J. Opt. Soc. Am. A*, vol. 19, pp. 2456-2466 (2002).
- [15] C. Matson, K. Borelli, S. Jefferies, C. Beckner, K. Hege, and M. Lloyd-Hart, "Fast and optimal multiframe blind deconvolution algorithm for high-resolution ground-based imaging of space objects," *Appl. Opt.*, vol. 48, pp. A75-A92 (2009).
- [16] M. Hyde, J. Schmidt, and M. Havrilla, "Geometrical optics polarimetric BRDF for dielectric and metallic surfaces," *Opt. Exp.*, vol. 17, pp. 22138-22153 (2009).
- [17] J. Blinn, "Models of light reflection for computer synthesized pictures," in *Proc. SIGGRAPH*, vol. 11, pp. 192-198 (1977).
- [18] S. Amaran and N. Sahinidis, "Global optimization of nonlinear least-squares problems by branch-and-bound and optimality constraints," *Top*, vol. 20, pp. 154-172 (2012).