

Using Frontier Points to Recover Shape, Reflectance and Illumination

George Vogiatzis[†] Paolo Favaro^{‡*} Roberto Cipolla[†]
gv215@eng.cam.ac.uk paolo.favaro@siemens.com cipolla@eng.cam.ac.uk

[†] Dept. of Engineering, University of Cambridge, Cambridge, CB2 1PZ, UK

[‡] Integrated Data Systems Department, Siemens Corporate Research, Princeton, NJ 08540

Abstract

We describe a method to recover the surface reflectance and the 3-d shape of a non-Lambertian object as well as illumination, from a collection of images. It is based on the so-called frontier points, which are extracted from the outlines of an object. Frontier points provide 3-d locations on the object surface where the surface normal is known. This information is exploited to infer the surface reflectance of the object and the light distribution of the scene both under varying illumination and fixed vantage point, and under varying vantage point and fixed illumination. We also show how to apply frontier points for shape recovery in photometric stereo. The effectiveness of frontier points for recovering reflectance, illumination and shape is confirmed by a number of experiments on both real and synthetic data.




Figure 1. Examples of non-Lambertian surfaces and reconstructed material. Using frontier points to sample the surface of a target object, the scene’s illumination and material properties of the object can be obtained. This information can then be used to render synthetic ‘example’ objects of the same material under the same illumination conditions.

1 Introduction

In this paper we consider the problem of recovering reflectance and 3-d shape of a non-Lambertian object from a collection of images. The main challenge is how to establish correspondence between image regions that are projections of the same 3-d point in space. While for Lambertian objects one can solve correspondence by direct matching of image regions¹, in the non-Lambertian case such matching is not possible as the measured intensity may vary dramatically between images, due, for example, to the unknown specular component of the reflectance. Hence, in general, correspondence for non-Lambertian surfaces is established while recovering the reflectance and the surface of an object, together with the illumination distribution, which is a highly ill-posed and computationally challenging problem [10, 21, 6].

In our solution, rather than venturing into such a cumbersome optimization problem, we propose to establish corre-

spondence by exploiting occlusions, and, more precisely, *frontier points*². We notice that occlusions are more resilient than other visual cues to changes in illumination or reflectance properties, which makes them suitable for non-Lambertian objects. Rather than matching image intensities, one can automatically determine correspondence by extracting the occluding boundary from each image region and then by searching along the boundary for the 2D points with tangents lying on the epipolar plane. This procedure allows not only to recover the 3-d location of a point on the surface (the frontier point), but also its corresponding normal vector. Hence, by working at frontier points one can solve a much easier problem where shape is locally given, and one is only left with recovering reflectance and illumination (see Figure 1), which we pose as a *blind deconvolution* problem in section 4. Furthermore, if an object is made of the same material, the reflectance and illumination distribution recovered at frontier points can then be used to infer its full 3-d shape. In this paper, we apply this scheme to the

*This work was carried out while Paolo Favaro was with Cambridge University, Dept. of Engineering.

¹The intensity measured at points on a Lambertian surface is invariant to the viewing direction.

²Given two images each obtained from a different camera view, a frontier point is the 3-d location where the epipolar plane (defined by the two cameras) is tangential to the surface of the object. In addition, the normal to the surface at a frontier point coincides with the normal to the epipolar plane.

case of uniform albedo where very little work has been done in the general scenario where we operate (see next section).

In our current implementation the pose of each camera is assumed to be known, although such an assumption can be relaxed by using, for instance, the methods described in [19, 20].

1.1 Contributions and related work

It has been shown in [14] that given the geometry of an object, it is possible to infer its reflectance and the illumination of the scene. Similarly, in our approach, we use frontier points to recover shape (locally) and infer the remaining unknowns as a blind deconvolution problem [1]. This problem falls within the field of *inverse rendering* in the community of computer graphics [5] and involves the study of the reflectance of objects [13, 17, 8, 12, 2].

Once reflectance and illumination are reconstructed, then one can use them together with frontier points to initialize and/or constrain a global optimization problem [10, 21, 6, 15], and recover the full shape of the observed object. In this paper, however, we are interested in presenting the potential of frontier points in the context of photometry, and hence we revisit a number of previously studied problems and show how easily our method can be adapted. For example, we present a novel solution to *photometric stereo*, i.e. the problem of recovering 3d shape given multiple images captured from the same viewing point, but under different unknown illumination conditions. As in [9], we impose that points in the scene that are subject to the same variations in their reflectance due to the same changes in viewing position, must share the same normal, the so-called *orientation-consistency* cue. This cue allows one to reconstruct first the normal map of the object and then to recover the depth map of the object. Notice that the method in [9] is based on the insertion of an object with known geometry, an *example*, of the same material of the object of interest in the scene.

In some situations, the insertion of such an example may not be possible. Rather, in our method, we do not need to insert any additional object in the scene, as frontier points provide information that can be used to reconstruct a *virtual example*, i.e. a virtual object with known geometry and the same reflectance as the object of interest (see Figure 1). Similarly, our work also relates to [6]; however, while our method works indistinctly in the case of both Lambertian and non-Lambertian objects, [6] is restricted to non-Lambertian objects and suffers from the convexity/concavity ambiguity.

The paper is organized as follows: in section 2 we introduce our reflectance model of non-Lambertian objects and the general problem of shape, reflectance and illumination recovery; then, we show how to use frontier points together with the reflectance model to recover the properties of the material and the illumination (section 4), and the shape of

the object in the scene (section 5).

2 Recovering shape, reflectance and illumination

In this section, we introduce the general problem of shape, reflectance and illumination recovery of a non-Lambertian object. To do so, we need first to introduce the reflectance model of non-Lambertian objects.

2.1 BRDF of non-Lambertian objects

The reflectance of a large class of objects is well approximated by the so-called bidirectional reflectance distribution function (BRDF). This is defined³ at each point P of a surface as a function $\beta(\theta_i, \phi_i; \theta_o, \phi_o)$ mapping the cartesian product between the hemisphere of incoming light directions (θ_i, ϕ_i) and the hemisphere of outgoing light directions (θ_o, ϕ_o) to nonnegative values. The BRDF predicts how much light will be reflected at a point on a surface along a certain direction, due to incoming light.

In the simplest instance of a Lambertian object, the BRDF is a constant, i.e. light is reflected equally in all directions. In the case of non-Lambertian objects the BRDF is much more involved. In this case, a number of models have been proposed for the BRDF, which can be divided into physics-based models and empirical-based models [17, 8, 13]. We adopt the Ward model [18], since it is a good tradeoff between accuracy of approximation and computational complexity. The Ward model is defined as:

$$\beta(\theta_i, \phi_i; \theta_o, \phi_o) = \frac{\rho_d}{\pi} + \frac{\rho_s e^{-\tan^2 \delta (\cos^2 \gamma / \alpha_x^2 + \sin^2 \gamma / \alpha_y^2)}}{4\pi \alpha_x \alpha_y \sqrt{\cos \theta_i \cos \theta_o}} \quad (1)$$

where ρ_d is the diffuse reflectance coefficient and ρ_s is the specular reflectance coefficient; α_x and α_y are the standard deviations of the surface slope at the microscopic level (surface roughness). For simplicity, in our approach we will assume that $\alpha_x = \alpha_y = \alpha$, i.e. that roughness is isotropic. Let h be the bisector of the vectors (θ_i, ϕ_i) and (θ_o, ϕ_o) ; δ is the angle between h and N , where N is the normal to the surface at P . γ is the phase angle between h and the x-axis on the tangent plane at P . Then, the irradiance observed at a pixel p , the projection of P on the image plane, is given by

$$I(p) \doteq I(\theta_o, \phi_o) = \int_0^{2\pi} \int_0^{\pi/2} \beta(\theta_i, \phi_i; \theta_o, \phi_o) L(R_P(\theta_i, \phi_i)) \cos \theta_i \sin \theta_i d\theta_i d\phi_i \quad (2)$$

³In this paper, we define the BRDF in *local coordinates*, i.e. we define a reference system at each point P on the surface of the object and set the z-axis parallel to the normal to the surface, while the x and y-axis lie on the tangent plane.

where the pixel p defines the local direction (θ_o, ϕ_o) . L is the light distribution and since it is defined in global coordinates, we need to introduce the rotation R_P that transforms local coordinates at P to global coordinates.

Notice that eq. (2) depends on the shape of the object via the normal field N , on the BRDF β at each point P and on the global illumination L , which are, in general, all unknown. For simplicity, here we focus on objects made of the same material, i.e. we assume that the BRDF β is the same at each point on the surface. In the next section, we will pose the problem of recovering these unknowns by matching the model in eq. (2) to measured images.

2.2 Problem statement

Suppose we are given a number of images $I_{1,1}, \dots, I_{K,M}$ obtained from K different vantage points and M different illumination conditions, then, as mentioned in the previous section, one may be interested in recovering the shape S of the object in the scene, which we identify with its normal field N and a 3-d point, its BRDF β and the light distribution L . This problem can be posed as the following minimization

$$\hat{S}, \hat{\beta}, \hat{L}_1, \dots, \hat{L}_M = \arg \min_{S, \beta, L_1, \dots, L_M} \sum_{k=1}^K \sum_{m=1}^M \Phi(I_{k,m}, I_k(S, \beta, L_m)) \quad (3)$$

where $I_k(S, \beta, L_m)$ is a short-hand notation for model (2), and its dependency on the unknowns has been made explicit. Φ is a function that accounts for the discrepancy between $I_{k,m}$ and $I_k(S, \beta, L_m)$. We require Φ to be zero if and only if $I_{k,m} = I_k(S, \beta, L_m)$, and to be strictly positive otherwise. In the next sections, we will choose Φ to be either the L_2 norm or the extended Kullback-Leibler distance [16].

Notice that, given one of the unknowns, for instance, the shape S of the object, the minimization task (3) is dramatically simplified. Consider the case of $M = 1$, i.e. fixed illumination conditions, then, one can easily show that the minimization (3) can be cast as a classic *blind deconvolution* problem [1] where β is the convolving kernel and L is the input signal. Similarly, if β and L were given, then recovery of the shape S would be greatly simplified.

In the next section, we will show how to exploit *frontier points* to this purpose. We will show that they provide some partial shape information that can be used to solve for β and L , which, in turn, can then be used to infer the complete shape S . Before presenting our solution we need to briefly introduce frontier points and how they are automatically extracted from images.




Figure 2. Frontier points. A frontier point is a 3-d point on the surface of the object where the plane passing through it and the two camera centers is tangent to the object. It can be retrieved by searching for the pair of epipolar lines that are both tangent to the two outlines of the object.

3 Sampling the surface via frontier points

Frontier points have been introduced in [7] in the context of structure from motion. In this paper, we will show that such points can also be exploited to infer photometric quantities.

The study of frontier points requires the introduction of notions of *contour generators* and *epipolar geometry* [11], which we cannot include here for lack of space. The interested reader is referred to [3] for an extensive analysis of frontier points. In this paper, we will only give a sketch of how frontier points are characterized, and how they can be obtained. Suppose we are given two images of the same object from two different vantage points. A *frontier point* is defined as (one of) the point(s) given by the intersection of the object and the *epipolar plane* T tangent to the object (see Figure 2). Notice that in this way we simultaneously define a point P on the surface and the normal N to the surface at that point.

An alternative and practical way to obtain frontier points is to look at the object’s outlines. An outline is defined as the projection on the image plane of a 3-d curve lying on the object, such that any line connecting this 3-d curve with the camera center is tangent to the object. Given two outlines, a frontier point P can be defined as the location in space that simultaneously satisfies the following properties:

- the projection p of P on each camera lies on both the outlines
- the tangent vector of each outline at the projection of the frontier point must lie on the *same* epipolar plane.

In our algorithm, we find frontier points by defining a cost functional that is minimized only when the two properties above are satisfied. The overall scheme for the auto-

matic extraction of frontier points on the surface of an object is as follows: (a) Obtain a number of images of the object (calibrated for pose and internal parameters) (b) Extract the object's outlines in those images (c) Compute a number of frontier points lying on the extracted outlines and satisfying the properties above.

4 Recovering illumination and BRDF

As mentioned above, frontier points not only define a point P on the object, but also a vector N normal to the object at P . Now, suppose that one is given a set of such pairs (P, N) that have been recovered by following the procedure in the previous section. If we collect intensities at the selected frontier points into a single vector I , then we can think of solving problem (3) in the unknown light distribution L and BRDF β , since shape is given. We will consider this problem in two separate settings, namely when $M = 1$ and $K > 1$, i.e. in the case of fixed illumination and varying vantage point (section 4.1), and when $M > 1$ and $K = 1$, i.e. in the case of varying illumination and fixed vantage point (section 4.2). Empirical evaluation of these schemes is presented in section 6.1.

4.1 Case I: Fixed illumination

This arises in inverse rendering problems [5] and can be useful in settings such as augmented reality. As mentioned in section 2.2, in this case the minimization problem amounts to solving a blind deconvolution. We choose the extended Kullback-Leibler pseudo-distance as our discrepancy measure, i.e. we set

$$\Phi(I_k, J_k) \doteq I_k \log \frac{I_k}{J_k} - I_k + J_k \quad (4)$$

so that the optimization problem 3 becomes

$$\hat{\beta}, \hat{L} = \arg \min_{\beta, L} \sum_{k=1}^K \Phi(I_k, J_k) \quad (5)$$

where we did not explicitly write the second summation in eq. (3) and the respective index since $M = 1$. $\{I_k\}_{k=1, \dots, K}$ are the measured intensities; J_k is directly derived from eq. (2) as

$$J_k = \int_0^{2\pi} \int_0^{\pi/2} \beta(\theta_i, \phi_i; \theta_k, \phi_k) L(R_{P_k}(\theta_i, \phi_i)) \cos \theta_i \sin \theta_i d\theta_i d\phi_i. \quad (6)$$

To adapt our problem to blind deconvolution, we do a change of coordinates on (θ_i, ϕ_i) so that $(\theta'_i, \phi'_i) = R_{P_k}(\theta_i, \phi_i)$, and then set $h(\theta'_i, \phi'_i; \theta_k, \phi_k) \doteq \beta(\theta'_i, \phi'_i; \theta_k, \phi_k) \cos \theta'_i \sin \theta'_i \Delta_k$, where Δ_k is the Jacobian of the change of coordinates. As a result, we obtain

$$J_k = \iint h(\theta'_i, \phi'_i; \theta_k, \phi_k) L(\theta'_i, \phi'_i) d\theta'_i d\phi'_i. \quad (7)$$

We choose to estimate L and the parameters of h (i.e. the parameters of the BRDF β) by running the alternating minimization scheme employed in [4], which is provably minimizing the chosen Φ while preserving the nonnegativity of L . The scheme consists of the following iterations:

1. Fix the parameters of the BRDF and recover L by using the Lucy-Richardson iteration [16, 4]
2. Fix the light distribution L and recover the parameters of the BRDF in h by gradient descent.

Results obtained using the method presented above are shown in Figure 3.

4.2 Case II: Varying illumination

In this section, we assume that $K = 1$ and $M > 1$. This typically arises in photometric stereo where the camera vantage point is kept fixed while the lighting changes between image captures. The optimization problem 3 becomes

$$\hat{\beta}, \hat{L}_1, \dots, \hat{L}_m = \arg \min_{\beta, L_1, \dots, L_m} \sum_{m=1}^M \Phi(I_m, J_m). \quad (8)$$

To reduce the dimensionality of the problem, we restrict our representation of the light distribution to a single moving point light source, i.e. we assume that $L(R_{P_1}(\theta_i, \phi_i)) = \lambda \delta(\theta_i - \theta_L) \delta(\phi_i - \phi_L)$ where δ is the Dirac delta. Then, we immediately obtain that

$$J_m = \lambda_m \beta(\theta_L, \phi_L; \theta_m, \phi_m) \cos \theta_L \sin \theta_L. \quad (9)$$

Since in this case the nonnegativity of L is automatically guaranteed, we do not need to resort to the Kullback-Leibler pseudo-distance and the corresponding alternating minimization scheme. For simplicity, we choose Φ to be simply the L_2 norm of the difference of the measured intensities and the ones predicted by the model (9), i.e.

$$\Phi(I_m, J_m) \doteq (I_m - J_m)^2. \quad (10)$$

To solve problem (8) one can simply run a gradient descent or a standard nonlinear optimization method, as the space of the unknowns is very small (in practice, 4 parameters for the BRDF β and $3M$ parameters for the illumination). We chose to use the standard implementation of *lsqnonlin* in Matlab[®]. The accuracy of this estimation with respect to the number of frontier points is evaluated in the experiment shown in Figure 4.

5 Recovering 3-d shape in photometric stereo

Once light distribution and BRDF parameters have been estimated by collecting the intensities at the frontier points, one can think of recovering the full shape of the object

again by solving (3). For simplicity, we consider the case of $M > 1$ and $K = 1$ also known as *photometric stereo*, but we would like to point out that our method is by no means limited to such a case, other possibilities being $M = 1$ and $K > 1$ (multi-view Shape from Shading) and $M > 1$ and $K > 1$ [22].

In the case of photometric stereo, the reconstruction process is particularly simple and straightforward. We assume that a number of frontier points have been extracted using the procedure described in section 3. Then, we collect images of the scene from the same vantage point ($K = 1$) and for different illumination conditions ($M > 1$) by moving a single point light source. We collect the intensities I_m , $m = 1, \dots, M$ at the frontier points and use them to recover the light direction and intensity for each illumination setting together with the parameters of the BRDF β as described in section 4.1. Once we have estimated L_1, \dots, L_M and the parameters of β , we generate a *synthetic* example in the spirit of [9] (see Figure 7 (c) for example). This virtual example is then employed in the recovery of the shape exactly as it is prescribed in Seitz et al. [9]. Furthermore, to improve our estimates, we also enforce both the depth map and the normal map to match the depth and normals of the selected frontier points.

6 Experiments

6.1 Light recovery evaluation

In this section we describe two experiments performed in order to empirically evaluate the recovery of illumination information from a scene, using frontier points on an object.

In the first experiment, we explore the case of fixed illumination and moving vantage point (section 4.1). 441 frontier points were defined by extracting silhouettes on a semi-diffuse plastic sphere. 35 images of the sphere are obtained from 35 different viewpoints. The goal here is to capture a general light distribution using the frontier points by solving the optimization problem of equation 5. Figure 3 shows the recovered light field mapped on the sphere of all directions. To qualitatively evaluate the recovered light field, a view of the light source is included which shows that the two-peaked light source has correctly produced a two-peaked light intensity field. As a further assessment we perform a synthetic rendering of the sphere specularity, using the estimated light field, as seen from two images not in the input set. The comparison between rendered and real images in Figure. 3 (c) shows that the complex shape of the specularity has been correctly captured. A byproduct of the optimization problem of equation 5 is the roughness of the surface (the α parameter of the Ward model) which is 0.24.

In the second experiment we evaluate the scheme laid out in section 4.2. The same plastic sphere is illuminated by a complex fixed light source. A set of 9 images were

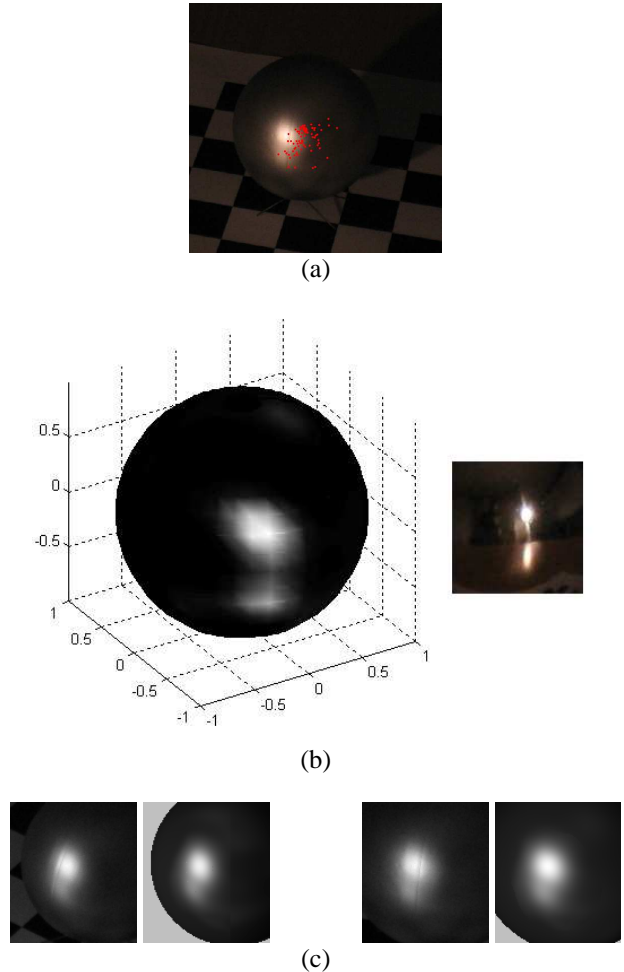


Figure 3. General light distribution recovery with frontier points. (a) the plastic sphere with 80 frontier points defined on it. (b) Left: the recovered light field for the scene mapped onto a sphere. Right: a view of the light source (two bright spots) - note how the two light distribution peaks are preserved in the estimated field. (c) Two close-up views of the specularity on the sphere. For each view, the specularity in the real image is followed by a synthetic rendering of the same view using the estimated light-field and BRDF. The structure of the specularity has been captured correctly.

obtained from the same viewpoint, each with a single directional light source illuminating the scene from a different but unknown direction. The goal of the experiment is to estimate these light directions using the frontier points on the object. To provide ‘ground truth’ light directions, a mirror ball is also placed on the scene (fig. 4). The estimation, described in equation 8, is repeatedly performed with randomly selected subsets of a number of frontier points. The error of each estimation is measured by the maximum angle between true and estimated directions. The graph in Figure 4 shows the decrease in estimation error with an increasing number of frontier points. It demonstrates the feasibility of accurate recovery of light direction using as few as 100 frontier points. The surface roughness parameter obtained from the solution of equation 8 is found to be equal to 0.27 with a standard deviation of 0.03, matching the estimate obtained in the previous experiment.




Figure 4. Light direction recovery with frontier points. Left: The setup with the diffuse and mirror balls for the evaluation experiment. Right: The maximum angle between estimated light directions and true directions vs the number of selected FPs. The error bars show the ± 3 std. dev. intervals




Figure 5. Quantitative evaluation of reconstruction method. Input images and resulting 3-d surfaces. Mean depth error is 4.6% and 10.1%.

6.2 Photometric stereo evaluation

In this section we describe a set of experiments that explore the feasibility of using frontier points to perform photometric stereo on non-Lambertian objects.

The first experiment is performed on synthetic 256x256 images of the ‘Mozart’ and ‘Penny’ scenes for which



(a)



(b)

Figure 6. Bottle reconstruction. (a) Two sets of images of the bottle (right), real example ball (bottom left) and synthetic ball (top left) generated from 100 sample points. (b) Views of the 3-d reconstruction of the bottle

ground truth is provided [23], to quantitatively measure the accuracy of the reconstruction method described in section 5, in isolation from the accuracy of frontier point placement. The frontier points here are simulated by randomly picking points on the true surface. Input images and results are shown in Figure 5. Using 80 random sample points (with their surface normals) on the true surface, the mean depth error for ‘Penny’ is 4.6% and 10.1% for ‘Mozart’ where the percentages are taken with respect to the size of the depth range ($Z_{max} - Z_{min}$) for each scene.

The second experiment illustrates the feasibility of replacing the ‘example’ object, required by the method of [9] by a sparse sampling of the reflectance function and by fitting an appropriate parametric model. The setup consists of 8 images of a bottle from the same viewpoint under 8 different single light source illuminations. Additionally we are given 8 images of an ‘example’ sphere of the same material as the bottle, under the same light conditions. In [9], the entire example sphere images are used to recover a normal direction for every pixel location on the bottle. Figure 6 demonstrates that similar results can be obtained with just a sparse sample of 100 points randomly selected on the example sphere. The figure shows the bottle and example sphere images, synthetically rendered example spheres using the

100 sample points as well as views of the 3-d reconstruction using these synthetic example spheres.

The third experiment uses frontier points to perform a photometric stereo reconstruction of a real, highly specular porcelain figurine. Contours of the figurine were extracted in a set of 20 images two of which are shown in Figure 7 (b). From these contours 100 frontier points were defined on a small convex region on the figurine. Subsequently, 14 images of the figurine were obtained from a single view-point under 14 different single light sources. The directions and intensities of the lights as well as BRDF parameters of the porcelain surface were estimated by solving equation 9. Figure 7 (c) shows the synthetic example porcelain balls next to the porcelain figurine. Each example ball is rendered with the same BRDF parameters and under the same light conditions as the corresponding real image. The photometric information thus obtained was then used to reconstruct the figurine and images of the reconstruction are presented in Figure 7 (d).

The final experiment is a photometric stereo reconstruction of a shiny stone statue. As before, 20 images were used for contours which resulted in 100 frontier points, which are shown in Figure 8 (a). Example balls made of the same material are shown in 8 (c) with images of the 3-d reconstruction shown in Figure 8 (d).

7 Discussion

This paper advocates the use of frontier points for the extraction of photometric information from images. We have presented a practical, robust and efficient solution for the recovery of illumination, surface reflectance and 3-d shape of a non-Lambertian uniform object from a number of images. The accuracy of the method has been evaluated empirically on synthetic scenes and the feasibility of shape and reflectance reconstruction using frontier points has been demonstrated on challenging real objects.

Frontier points have so far mainly been used for structure and motion recovery. However, the high degree of robustness with which they can be extracted as well as their relative independence from scene lighting and reflectance makes them ideal for extracting photometric information. While this paper has focused on a few test cases, the overall aim is to highlight the potential of using frontier points for photometry, either on their own, as a starting point, or as extra constraints to most photometric calculations.

Acknowledgements This work is supported by the Gates Cambridge Trust and Toyota Corporation.

References

- [1] Various Authors. Special issue on blind system identification and estimation. *Proceedings of the IEEE*, 86(10), 1998.




Figure 7. Reconstruction of porcelain figurine using frontier points. (a) Left to right: Images of the figurine with contours extracted. The frontier points defined on a small convex region of the figurine. (b) 3 input images with varying light source (c) Two pairs of images of the real figurine next to a synthetically rendered example sphere of the same material under the same illumination. (d) Three images from front, side and top of the 3-d reconstruction of the porcelain object.

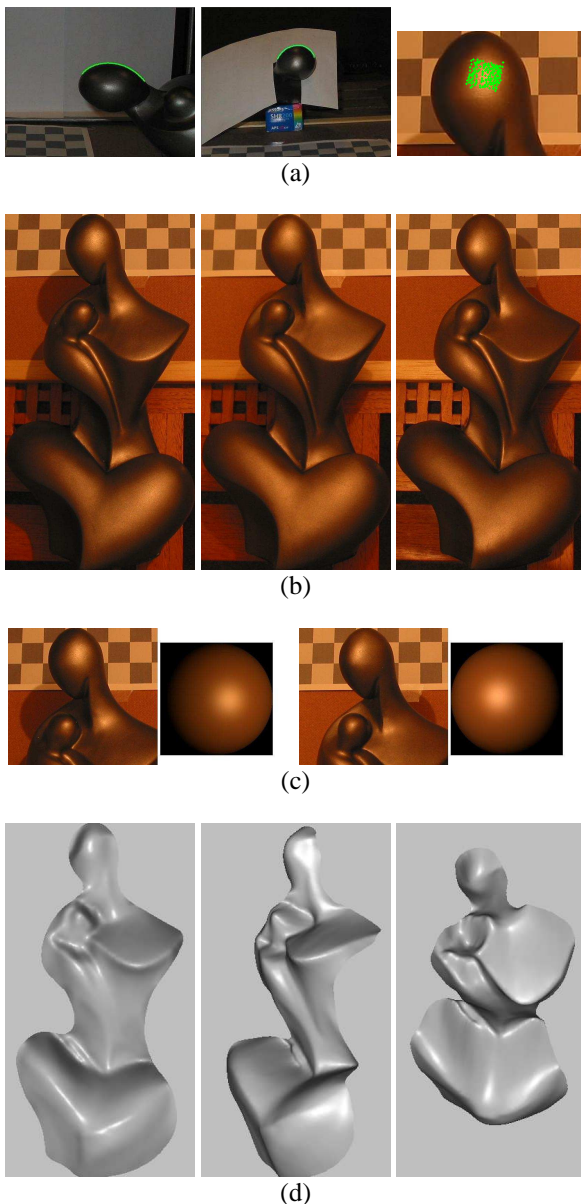


Figure 8. Reconstruction of shiny stone sculpture using frontier points. (a) Left to right: Images of the sculpture with contours extracted. The statue appears with different color due to camera flash, used to maximize contrast. The frontier points are defined on the head of the statue. (b) 3 input images with changing light source (c) Two pairs of images of the real statue next to a synthetically rendered example sphere of the same material and same illumination. (d) Three images from front, side and top of the 3-d reconstruction of the statue.

- [2] W. C. Chen, J.-Y. Bouguet, M. H. Chu, and R. Grzeszczuk. Light field mapping: efficient representation and hardware rendering of surface light fields. In *Proc. of the ACM SIGGRAPH*, 2002.
- [3] R. Cipolla and P. J. Giblin. *Visual Motion of curves and surfaces*. Cambridge University Press, 1999.
- [4] P. Favaro and S. Soatto. Shape and reflectance estimation from the information divergence of blurred images. In *European Conference on Computer Vision*, pages 755–768, June 2000.
- [5] Patow G. and Pueyo X. A survey of inverse rendering problems. In *Computer Graphics Forum*, volume 22, pages 663–87, 2003.
- [6] A.S. Georghiadis. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *ICCV03*, pages 816–823, 2003.
- [7] P. Giblin, F. Pollick, and J. Rycroft. Recovery of an unknown axis of rotation from the profiles of a rotating surface. *J. Optical Soc. America*, 11A:1976–1984, 1994.
- [8] Xiao D. He, Kenneth E. Torrance, Francois X. Sillion, and Donald P. Greenberg. A comprehensive physical model for light reflection. In *SIGGRAPH '91*, pages 175–186. ACM Press, 1991.
- [9] A. Hertzmann and S.M. Seitz. Shape and materials by example: a photometric stereo approach. In *CVPR03*, pages I: 533–540, 2003.
- [10] H. Jin, S. Soatto, and A. J. Yezzi. Multi-view stereo beyond lambert. In *Proc. IEEE Conf. on Comp. Vision and Pattern Recogn.*, pages 1–171–178, June 2003.
- [11] D. Marr. Analysis of occluding contour. In *Proc. Royal Soc. London B*, volume 197, page 441475, 1977.
- [12] Shree K. Nayar, Katsushi Ikeuchi, and Takeo Kanade. Surface reflection: Physical and geometrical perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 13(7):611–634, 1991.
- [13] B. T. Phong. Illumination for computer generated pictures. In *Communications of the ACM*, volume 18(6), pages 311–317, 1975.
- [14] Ravi Ramamoorthi and Pat Hanrahan. A signal-processing framework for inverse rendering. In *SIGGRAPH '01*, pages 117–128. ACM Press, 2001.
- [15] Tood Zickler Sebastian Magda, David J. Kriegman and Peter N. Belhumeur. Beyond lambert: Reconstructing surfaces with arbitrary brdfs.
- [16] D. Snyder, T. Schulz, and J. O'Sullivan. Deblurring subject to nonnegativity constraints. *IEEE Trans. on Signal Processing*, 40(5):1143–1150, 1992.
- [17] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughed surfaces. *J. of the Opt. Soc. of Am.*, 57(9):1105–1114, 1967.
- [18] G. Ward. Measuring and modeling anisotropic reflection. In *SIGGRAPH*, pages 265–272, 1992.
- [19] K.-Y. K. Wong, P. R. S. Mendonça, and R. Cipolla. Camera calibration from surfaces of revolution. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(2):147–161, 2003.
- [20] Jean Ponce Yasutaka Furukawa, Amit Sethi and David Kriegman. Structure and motion from images of smooth textureless objects. In *European Conference on Computer Vision*, 2004.
- [21] T. Yu, N. Xu, and N. Ahuja. Recovering shape and reflectance model of non-lambertian objects from multiple views. In *CVPR04*, pages II: 226–233, 2004.
- [22] L. Zhang, B. Curless, A. Hertzmann, and S. Seitz. Shape and motion under varying illumination: Unifying structure from motion, photometric stereo, and multi-view stereo. In *Proc. 9th Int. Conf. on Computer Vision*, 2003.
- [23] R. Zhang, P. S. Tsai, J Cryer, and M. Shah. Shape from shading: A survey. *PAMI*, 21(8), 1999.