

Radiance Function Estimation for Object Classification

Antonio Robles-Kelly and Edwin R. Hancock

Department of Computer Science,
University of York, York YO1 5DD, UK
{arobkell, erh}@cs.york.ac.uk

Abstract. This paper describes a simple method for estimating the surface radiance function from single images of smooth surfaces made of materials whose reflectance function is isotropic and monotonic. The method makes use of an implicit mapping of the Gauss map between the surface and a unit sphere. By assuming the material brightness is monotonic with respect to the angle between the illuminant direction and the surface normal, we show how the radiance function can be represented by a polar function on the unit sphere. Under conditions in which the light source direction and the viewer direction are identical, we show how the recovery of the radiance function may be posed as that of estimating a tabular representation of this polar function. A simple differential geometry analysis shows how the tabular representation of the radiance function can be obtained using the cumulative distribution of image gradients. We illustrate the utility of the tabular representation of the radiance function for purposes of material classification.

1 Introduction

The modeling of surface reflectance is a topic that is of pivotal importance, and has hence attracted considerable effort in both, computer vision and computer graphics communities. Broadly speaking, the methods used to model or approximate the bidirectional reflectance distribution function (BRDF) can be divided into those that are physics-based, semi-empirical or empirical in nature. Although the literature from physics is vast, it is perhaps the work of Beckmann on smooth and rough surface reflectance that is the best known in the vision and graphics communities [1]. Despite being based upon physically meaningful surface parameters, the Beckmann theory is intractable for analysis problems since it relies on the evaluation of the Kirchhoff wave scattering integral. Further, it breaks down when either the surface roughness or the scattering angle are large. However, recently, Vernold and Harvey [2] have overcome this latter problem by developing a model which accounts for self shadowing on rough surfaces. By contrast, in the graphics community it is the development of computationally efficient tools for the purposes of realistic surface rendering that is of primary interest, and hence it is empirical models that have been the focus of activity [3, 4]. One of the most popular models is that developed by Phong [4]. However,

neither the models developed in physics nor the computational models developed in graphics are well suited for surface analysis tasks in computer vision. It is for this reason that Wolff [5] and Nayar and Oren [6] have developed phenomenological, or semi-empirical, models that account for departures from Lambertian reflectance. Despite these efforts, the physical and phenomenological modeling of the BRDF remains an elusive task. An alternative is to empirically estimate or to learn the BRDF under controlled lighting and viewing conditions of rough and specular objects [7–11]. There have also been attempts to model the reflectance properties of human skin from real-world imagery [12, 13]. Hertzmann and Seitz [14] have shown how the BRDF can be recovered making use of a reference object and multiple views of the scene.

The main problem with existing approaches is that the BRDF has four degrees of freedom that correspond to the zenith and azimuth angles for the light source and the viewer relative to the surface normal direction. As a result, the tabulation of empirical BRDF's can be slow and labour intensive. Furthermore, extensive lighting control and prior knowledge of the surface geometry is often required for the BRDF estimation process.

In this paper, we focus our attention in estimating the radiance function from single images without the use of expensive cameras and complex calibration procedures. Hence, we are interested in a computationally cheap alternative to the complicated setups employed by measurement-intensive approaches. Of course, acquiring a BRDF for purposes of photorealistic rendering from a single image is hard due to technical issues. Here, we aim at recovering a qualitatively good estimate that can be used for purposes of object classification or material library indexing.

We present an essentially non-parametric method for estimating the reflectance function from image data that avoids using basis functions or a predetermined parameterisation of the BRDF to characterise the specular spike and limb. Our method makes implicit use of the Gauss map, i.e. the projections of the surface normals onto a unit sphere. We map implicitly the brightness values for a single image onto locations on the unit sphere which have identical surface normal directions. Under conditions in which the light source and the viewer directions are identical, we show how the reflectance function can be represented by a polar function on the unit sphere. We pose the problem of recovering the reflectance function as that of estimating a tabular representation of this polar function. To overcome the problem that we do not have the field of surface normals at hand, under conditions of isotropic reflectance, we show how to estimate the zenith angles on the unit sphere using image brightness gradients. A simple analysis shows how the tabular representation of the reflectance function can be obtained using the cumulative distribution of image gradients.

2 Preliminaries

In this section, we provide the background for our method. Our overall aim is to make an estimate of the reflectance distribution function from a single image of a

piecewise smooth surface. We limit our study to those surfaces whose reflectance function is isotropic and monotonic. Surfaces of this kind are those of shiny and moderately rough objects, terse objects and mate materials. Examples of these are porcelain, smooth terracotta, plastics, etc.

We simplify the problem of estimating the radiance function by using the differential properties of the Gauss map of the surface under study onto a unit sphere. For an orientable surface $S \in \mathbb{R}^3$, the Gauss map $G : S \mapsto \hat{S}$ maps points on the surface S onto locations on the unit sphere \hat{S} which have identical surface normal directions. Our aim is to use correspondences between surface normal directions to map brightness values from the image onto the unit sphere. The polar distribution of brightness values on the unit sphere \hat{S} is the radiance function for the surface. To avoid ambiguities, we assume that points on the surface with identical surface normal directions have identical brightness values.

Of course, when only a single image is at hand, the mapping of the brightness values from the image onto the unit-sphere is not straightforward. In fact, the task of estimating surface normal directions from measured brightness values is an underconstrained one, which has preoccupied researchers in the field of shape-from-shading for several decades. Even for surfaces which exhibit simple Lambertian reflectance, the problem is not tractable in closed-form. Furthermore, for non-Lambertian reflectance the situation is more complex. In the case of non-Lambertian reflectance, provided that the reflectance properties of the surface are isotropic and homogeneous, the problem is simplified considerably if the viewer and light source directions are identical. The isotropy assumption will ensure that circles of latitude on the unit sphere will still have constant brightness. The problem of recovering the distribution of the brightness with respect to the latitude becomes that of estimating the zenith angle from the distribution of brightness values.

Hence, we limit our discussion to the case where the image plane Π is chosen so that the viewer direction vector \mathbf{V} and the light-source direction vector \mathbf{L} are coincident, i.e. $\mathbf{L} = \mathbf{V}$. Suppose that the point p on the unit sphere has zenith angle θ and azimuth angle α . Under the Gauss map, the brightness value associated with this point is denoted by the polar radiance function $f_O(\theta, \alpha) = I$, where I is the measured brightness at the point s in the image of the surface S . Thus, when the viewer and light source directions are identical, then provided that the reflectance process is isotropic, the distribution of radiance across the unit sphere can be represented by a function $g(\theta)$ of the zenith angle alone. As a result, the observed brightness values mapped onto the unit sphere by the Gauss map G can be generated by revolving the function $g(\theta) = f_O(\theta, 0)$ in α about the axis defined by the viewer and light source directions. The problem of describing the brightness distribution over the Gauss sphere hence reduces itself to that of approximating the function $g(\theta)$ and computing its trace of revolution.

To develop our analysis, and to show how to map brightness values onto the unit sphere, it will prove useful to consider the image of the unit sphere under orthographic projection onto the plane $\hat{\Pi}$ that is perpendicular to the viewer direction. The Cartesian image can be represented using the polar coordinates

of the unit sphere and is given by $I_{\hat{\Pi}}(\sin(\theta) \cos(\alpha), \sin(\theta) \sin(\alpha)) = f_O(\theta, \alpha)$. In fact, when the light source and viewer directions are identical, then the image is circularly symmetric and we can write $I_{\hat{\Pi}}(\sin(\theta) \cos(\alpha), \sin(\theta) \sin(\alpha)) = g(\theta)$.

3 Radiance Function Estimation

When the viewer and light source directions are identical, then the task of estimating the radiance function reduces to that of estimating the distribution of brightness values with respect to the zenith angle on the unit sphere, i. e. to estimate $g(\theta)$. We show how this can be performed by using the differential structure of the observed brightness on the image plane $\hat{\Pi}$. Hence, we commence by rewriting $g(\theta)$ as the integral of the partial derivative of the observed brightness with respect to the angular variable θ . To do this, we assume the radiance function $f_O(\theta, \alpha)$ to be monotonically decreasing for $\theta \in [0, \frac{\pi}{2}]$ and write

$$g(\theta) = \frac{1}{2\pi} \int_0^{2\pi} \left(f_O(0, \alpha) + \int_0^\theta \frac{\partial f_O(\theta, \alpha)}{\partial \theta} d\theta \right) d\alpha \quad (1)$$

In other words, the generating function $g(\cdot)$ on the unit sphere can be expressed in terms of the cumulative distribution of the derivatives of the radiance function or, alternatively, the derivatives of the image brightness.

We now turn our attention to the image of the unit sphere on the plane $\hat{\Pi}$. Suppose that $F(r, \theta)$ is a parametric polar function that represents the distribution of radiance values over the image of the unit sphere. The radial coordinate of the function is the Euclidean distance between the point p and the center-point of the unit sphere \hat{S} on the viewer plane $\hat{\Pi}$, i.e. $r = \sin(\theta) = \sqrt{(\sin(\theta) \cos(\alpha))^2 + (\sin(\theta) \sin(\alpha))^2}$. Hence

$$F(r, \theta) = \begin{bmatrix} r \\ g(\theta) \end{bmatrix} = \begin{bmatrix} \sin(\theta) \\ \frac{1}{2\pi} \int_0^{2\pi} \left(f_O(0, \alpha) + \int_0^\theta \frac{\partial f_O(\theta, \alpha)}{\partial \theta} d\theta \right) d\alpha \end{bmatrix} \quad (2)$$

As noted earlier, since the surface normals are not at hand, the correspondences between locations on the surface and the unit sphere are not available. Hence, the quantity θ is unknown. In other words, the function $F(r, \theta)$ only allows the surface S to be mapped onto the unit sphere \hat{S} in an implicit manner.

To overcome this lack of correspondance information, we commence by showing the relation between the image brightness gradient and the function $g(\theta)$. Let the magnitude of the brightness gradient be given by

$$|\nabla I| = \sqrt{\left(\frac{\partial I}{\partial x}\right)^2 + \left(\frac{\partial I}{\partial y}\right)^2}$$

Since the image is circularly symmetric, the image gradient can be rotated about the z-axis without any loss of generality. We align the image gradient with the azimuth direction using the rotation matrix R_z , which rotates the vector ∇I by an angle α in a clockwise direction about the z-axis. The image brightness

derivatives may be related to those of the function $g(\theta)$ using the inverse Jacobian J^{-1} via the following matrix equation

$$R_z \begin{bmatrix} \frac{\partial I}{\partial x} \\ \frac{\partial I}{\partial y} \end{bmatrix} = J^{-1} \begin{bmatrix} \frac{\partial g(\theta)}{\partial \theta} \\ \frac{\partial g(\theta)}{\partial \alpha} \end{bmatrix} \quad (3)$$

Using the rotation on the image plane and the coordinate transformation between the image plane and the unit sphere, we find that

$$|\nabla I| = \frac{1}{\cos(\theta)} \frac{\partial g(\theta)}{\partial \theta} = \frac{\partial g(\theta)}{\partial \sin(\theta)} \quad (4)$$

In this way, we can relate the image gradient to the derivative of the function $g(\theta)$ with respect to the zenith angle θ . In terms of finite differences, the relationship between the magnitude of the image gradient and the changes $\Delta g(\theta)$ in $g(\theta)$ and $\Delta \sin(\theta)$ in θ is the gradient of the function $F(r, \theta)$, i.e. $|\nabla I| = \frac{\Delta g(\theta)}{\Delta \sin(\theta)}$.

The image gradient ∇I can be computed using the formula

$$\nabla I = \frac{1}{\delta} \left[\begin{array}{c} I(j+1, k) - I(j-1, k) \\ I(j, k+1) - I(j, k-1) \end{array} \right] \quad (5)$$

where δ is the spacing of sites on the pixel lattice. Furthermore, on the unit sphere \hat{S} , it is always possible to choose points to be sampled so that the difference in brightness is a constant τ . As a result, we can write

$$\Delta \sin(\theta) = \frac{\tau}{|\nabla I|} \quad (6)$$

To recover θ from the expression above we perform numerical integration. To do this, we sort the image gradients according to the associated image brightness values. Accordingly, let ∇I_l be the image gradient associated with the brightness value l . The numerical estimate of $\sin(\theta)$ is then given by

$$\sin(\theta) = \int_{l=0}^{l=m} \frac{\tau}{|\nabla I_l|} dI_l + \kappa \approx \sum_{l=0}^m \frac{\tau}{|\nabla I_l|} + \kappa \quad (7)$$

where κ is the integration constant and m is the maximum brightness value for the surface under study. Hence, we can use the cumulative distribution of inverse gradients to index the zenith angle on the unit sphere. This indexation property means that we can approximate the function $F(r, \theta)$, or equivalently $g(\theta)$, by tabulation.

To pursue this idea, in principle, we only require a single image gradient corresponding to each of the distinct brightness levels in the image. In practice, we make use of the cumulative distribution of image gradients in order to minimise the approximation error by averaging. Let $Q_l = \{s \mid I = l\}$ be the set of pixels with brightness value l . For the brightness value $l = g(\theta)$, the average gradient is given by

$$h(l) = \frac{\sum_{s \in Q_l} |\nabla I|}{|Q_l|} \quad (8)$$

The distribution of average gradients is then stored as a vector h . Zero entries of the vector, which correspond to brightness values that are not sampled in the image, can cause divide-by-zero errors when the radiance function is computed. To overcome this problem, we smooth the components of the vector by performing piecewise linear interpolation of the adjacent non-zero elements. The resulting vector is denoted by \hat{h} . With the average image gradient at hand, we define the tabular approximation \hat{F} to $F(r, \theta)$ as the set of Cartesian pairs

$$\hat{F} = \left\{ \left(\left(\tau \sum_{i=0}^l \hat{h}(i)^{-1} + \kappa \right), l \right); l = 0, 1, 2, \dots, n_{max} \right\} \quad (9)$$

All that remains is to compute the constants τ and κ . We do this by making use of the maximum and minimum values of $\sin(\theta)$ for $\theta \in [0, \frac{\pi}{2}]$. Since the maximum and minimum values of $\sin(\theta)$ are unity and zero when $\theta = \frac{\pi}{2}$ and $\theta = 0$, we can set κ to unity and evaluate the numerical integral for $l = m$, which yields

$$\tau = - \left(\sum_{i=0}^m \hat{h}(i)^{-1} \right)^{-1} \quad (10)$$

4 Experiments

In this section, we illustrate the utility of the method for purposes of classification of shiny and rough materials. To this end, we have computed a set of pairwise distances for eight terracotta and porcelain objects captured using a simple setup which comprises only a Olympus



Fig. 1. Images used in our experiments.

E10 digital camera and a collimated 200W white tungsten light source. With the images at hand, we compute a similarity measure between pair of objects making use of the the reflectance map on the sphere computed from the reflectance function delivered by our method. This similarity measure is then used for purposes of separating images of shiny objects from those that correspond to rough objects.

The process described above is as follows. We compute a distance matrix D whose entry $d(a, b)$ is given by the error between the reflectance-map spheres corresponding to the objects indexed a and b . Hence, for a pair of reflectance-map spheres with a set of sample points Ω projected onto the pixel lattice, the

pairwise distance is given by $d(a, b) = \frac{1}{|\Omega|} \sum_{q \in \Omega} |I_a(u) - I_b(u)|$, where $I_a(u)$ and $I_b(u)$ are the measured brightness values at the point q on the pixel lattice whose coordinates are $u = (i, j)$. With the matrix of pairwise distances at hand, we compute an affinity matrix W . Ideally, the smaller the distance, the stronger the weight, and hence the mutual affinity to a cluster. The affinity weights are required to be in the interval $[0, 1]$. Hence, for the pair of reflectance spheres indexed a and b the affinity weight is taken to be $W(a, b) = \exp(-k \frac{d(a, b)}{\max(D)})$ where k is a constant.

For visualisation purposes, we have performed multidimensional scaling (MDS) [15] on the pairwise distance matrices. We have done this in order to embed the reflectance-map spheres in an eigenspace. Broadly speaking, the method can be viewed as embedding the objects in a pattern space using a measure of their pairwise similarity to one another. It is interesting to note that when the distance measure used is the L2 norm, then MDS is equivalent to principal components analysis.

In the top row of Figure 2, we show, from left-to-right, the distance matrix, the affinity matrix and the MDS plot for the eight porcelain and terracotta objects. The indexing of the objects corresponds to that shown in shown in Figure 1. Based on the visualisation provided by MDS, it is clear that the reflectance function delivered by our method may be suitable for the purposes of separating shiny and rough objects. It is important to stress that, since the distance has been computed using the reflectance map on the sphere, computed from the input images, and not the images themselves, we are effectively capturing the differences in the reflective properties of the objects and not only the differences in color.




Fig. 2. Left-hand column: Distance matrices; Middle column: affinity matrices; Right-hand column: MDS plots.

To take our analysis further, we have performed experiments on four sandpaper rolls whose grades are 100, 150, 180 and 240. In the bottom row of Figure 2, we show the distance matrix, the affinity matrix and the MDS plot for the four sandpaper rolls. From the MDS plot it is clear that the eigenspace embeddings describe a scattering that is in close accordance with a straight line. Further, the distribution of the distances describes a trace whose arrangement suggests an ordering from coarse to fine in the eigenspace, and hence, a way of classifying by grade the sandpaper rolls.

5 Conclusions

In this paper, we have presented a novel approach for approximating the radiance functions and hence the BRDF of objects whose reflectance is isotropic and monotonic from a single image. Although the new method is applicable only when the light source and viewer directions are approximately equal, it can be used as a computationally cheap alternative to other methods which use measurement-intensive approaches. Thus, the reflectance function estimated using the method may be used for tasks which require a computationally cheap estimate of the reflectance properties of the object, such as material classification.

References

1. P. Beckmann and A. Spizzochino. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon, New York, 1963.
2. C. L. Vernold and J. E. Harvey. A modified beckmann-kirchoff scattering theory for non-paraxial angles. In *Scattering and Surface Roughness*, number 3426 in Proc. of the SPIE, pages 51–56, 1998.
3. R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Trans. on Graphics*, 1(1):7–24, 1982.
4. B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.
5. L. B. Wolff. On the relative brightness of specular and diffuse reflection. In *Int. Conf. on Comp. Vision and Patt. Recognition*, pages 369–376, 1994.
6. S. K. Nayar and M. Oren. Visual appearance of matte surfaces. *SCIENCE*, 267:1153–1156, 1995.
7. S. Westin, J. Arvo, and K. Torrance. Predicting reflectance functions from complex surfaces. In *SIGGRAPH 92 Conference Proceedings*, pages 255–264, 1992.
8. X. He, P. Heynen, R. Phillips, K. Torrance, D. Salesin, and D. Greenberg. A fast and accurate light reflection model. In *Siggraph 92 Conference Proceedings*, volume 26, pages 253–254, 1992.
9. G. J. Ward. Measuring and modeling anisotropic reflection. *Computer Graphics*, 26(2):265–272, 1992.
10. E. P.F. Lafortune, Sing-Choong Foo, K. E. Torrance, and D. P. Greenberg. Non-linear approximation of reflectance functions. In *SIGGRAPH 97 Conference Proceedings*, pages 117–126, 1997.

11. K. J. Dana and S. K. Nayar. Correlation model for 3d texture. In *Int. Conf. on Comp. Vision*, pages 1061–1066, 1999.
12. S. R. Marschner, S. H. Westin, E. P. F. Lafortune, K. E. Torrance, and D. P. Greenberg. Image-based brdf measurement including human skin. In *10th Eurographics Rendering Workshop*, 1999.
13. P. Debevec, T. Hawkins, C. Tchou, H.-P. Duiker, W. Sarokin, and M. Sagar. Acquiring the reflectance field of a human face. In *SIGGRAPH 2000*, pages 145–156, 2000.
14. A. Hertzmann and S. M. Seitz. Shape and materials by example: A photometric stereo approach. In *Int. Conf. on Comp. Vision and Patt. Recognition*, pages 533–540, 2003.
15. I. Borg and P. Groenen. *Modern Multidimensional Scaling, Theory and Applications*. Springer Series in Statistics. Springer, 1997.