Photometric Stereo With Non-Parametric and Spatially-Varying Reflectance

Neil Alldrin[†]

nalldrin@cs.ucsd.edu

[†]University of California, San Diego

Todd Zickler[††]

zickler@eecs.harvard.edu

[††]Harvard University

David Kriegman[†]

kriegman@cs.ucsd.edu

Abstract

We present a method for simultaneously recovering shape and spatially varying reflectance of a surface from photometric stereo images. The distinguishing feature of our approach is its generality; it does not rely on a specific parametric reflectance model and is therefore purely "data-driven". This is achieved by employing novel bi-variate approximations of isotropic reflectance functions. By combining this new approximation with recent developments in photometric stereo, we are able to simultaneously estimate an independent surface normal at each point, a global set of non-parametric "basis material" BRDFs, and per-point material weights. Our experimental results validate the approach and demonstrate the utility of bi-variate reflectance functions for general non-parametric appearance capture.

Figure 1: (Left) One of 102 single-viewpoint input images. (Right) Rendering from a novel viewpoint, using shape and reflectance acquired by our algorithm.

1. Introduction

Capturing the "appearance" of objects from images has become increasingly important in recent years, especially as computer graphics applications demand a level of photo-realism unattainable by hand modeling. By "appearance", we mean a model that is able to predict images of the object under all possible view and illumination conditions. To adequately sample the appearance of an object which truly varies arbitrarily in both shape and reflectance would require images from every combination of view and illuminant, which is impractical in most situations. Fortunately, objects in the real world typically exhibit regularity that can be exploited to drastically reduce the number of images required. Choosing constraints that are valid (or close to valid), yet powerful enough to be useful in practical systems is thus essential to appearance capture.

In this paper, we consider the special case of photometric stereo – recovering an *explicit* appearance model (i.e., separate shape and reflectance models) from images taken at a single viewpoint under varying, known illumination. This is an important special case of appearance capture, since it

relies solely on photometric cues, avoids solving the correspondence problem, and is relatively simple to extend to multiple views if necessary. It is also important because explicit appearance models of this kind have been shown to be useful for visual tasks such as face recognition [3]. Simultaneously recovering surface normals and reflectance from such input data remains a challenging problem, however. Typically, the form of the reflectance function is restricted by either assuming a parametric model or by the existence of a set of homogeneous reference objects in the scene. The obvious downside to these methods is reduced generality – if the materials in the object being measured differ from the assumed reflectance models, the accuracy of the recovered appearance model will be poor.

Our technique differs from most previous approaches in that we do not impose a parametric model on the reflectance function. Rather, we restrict the form of the reflectance function to satisfy empirically observed physical properties shared by many materials. These physical properties allow us to reduce the domain of the bi-directional reflectance distribution function (BRDF) from a function of four variables to a function of two variables without a significant loss in accuracy. This approximation has both theoretical and empirical motivation. Theoretically, Stark et al. [23] have shown that many of the parametric BRDF

models commonly used in computer graphics are in fact bi-variate functions, which suggests that bi-variate approximations can have at least some level of accuracy. They also show impressive empirical results for (an albeit small) set of measured isotropic BRDFs. We provide additional analysis in this paper, both theoretical and empirical, to further support the validity of bi-variate BRDF approximations.

Our main contributions are (1) to present a technique capable of simultaneously recovering shape and non-parametric reflectance from photometric stereo, and (2) to introduce bi-variate representations of reflectance as a useful tool for vision applications.

2. Background and Related Work

Photometric stereo has long been an active area of research in computer vision. Early work, such as that by Woodham [27] and Silver [22] made strong assumptions on the reflectance function across the surface, typically requiring either explicit knowledge of the BRDF or simple parametric models.

Much of the emphasis in subsequent research has been to weaken constraints on the reflectance function, thus enabling photometric stereo to work on broader classes of objects. For example, it has been observed that the reflectance of many materials is well approximated by the sum of a specular and a diffuse lobe, which has motivated an entire line of research [4, 2, 11, 15]. Many of these approaches assume a Lambertian diffuse lobe, while not imposing a parametric form on the specular lobe. Examples include Coleman and Jain [4] and Barsky and Petrou [2] who treat specular pixels as outliers as well as Schlüns and Wittig [20], Sato and Ikeuchi [19], and Mallick et al. [14] who assume the color of the specular lobe differs from the color of the diffuse lobe, allowing separation of the specular and diffuse components.

A different approach is to place reference objects in the scene that have similar reflectance to the test object. This method was used in early photometric stereo research [22] and was later reexamined by Hertzmann and Seitz [8, 9]. The basic idea is that the reference objects provide a direct measurement of the BRDFs in the scene, which is then matched to points on the test object. This works for arbitrary BRDFs, but requires reference objects of the same material as the test object. Spatially varying BRDFs can also be handled by assuming that the BRDF at each point on the test object is a linear combination of the "basis" BRDFs defined by the set of reference objects. This approach to spatially varying BRDFs is similar in spirit to work by Lensch et al. [13], although their method uses parametric (Lafortune) BRDFs and assumes known surface shape.

Building upon the idea of considering the reflectance at each surface point to be a linear combination of a small set of BRDFs, Goldman et al. [7] removed the need for reference objects by iteratively estimating the basis BRDFs and surface normals. Their method assumes an isotropic Ward model for each basis material, whose parameters are estimated at each iteration. While it requires the solution of a difficult optimization problem, their approach is still one of a very few capable of recovering surface normals and relatively flexible parametric BRDFs in tandem.

While parametric models are very good at reducing the complexity of BRDFs, they are usually only valid for a limited class of materials [16, 23]. An alternative is to exploit physical properties common to large classes of BRDFs. For example, it is well known that all real-world BRDFs satisfy energy conservation, non-negativity, and Helmholtz reciprocity. Utilizing these properties, while not as simple as utilizing parametric models, is nonetheless possible. Helmholtz stereopsis, introduced by Zickler et al. [28, 29], is one such technique, exploiting reciprocity to obtain (multi-view) surface reconstruction with no dependence the BRDF. Isotropy is another physical property which holds for materials without "grain". While isotropy is implicitly assumed in almost all parametric models used in computer vision, only recently has it been explicitly utilized for photometric stereo. Tan et al. [24] use both symmetry and reciprocity present in isotropic BRDFs to resolve the generalized bas-relief ambiguity. More relevant to this paper is work by Alldrin and Kriegman [1], who show that isotropy, with no further assumptions on surface shape or BRDF, can be utilized to recover the surface normal at each surface point up to a plane. In particular, no parametric model is used and the BRDF is allowed to vary arbitrarily across the surface.

Another recent development in non-parametric BRDF acquisition is the concept of factorizing sampled BRDF values into the product of a material weight matrix and a BRDF matrix. The most prominent of these approaches is work by Lawrence et al. [12] who solve the factorization problem using alternating constrained least squares. Their algorithm is again based on the assumption that spatially varying reflectance can be represented as a weighted sum of a small set of materials. Although their technique is primarily focused on BRDF acquisition, they also show limited examples of surface normal estimation.

In this paper, we build upon and improve three recent advances. First, we exploit isotropy, as in [1], to constrain surface normals to a single degree of freedom. Second, we utilize a non-parametric bi-variate approximation of the BRDF. Finally, we assume that surfaces are composed of a small number of "basis" materials and solve a factorization problem similar to that of Lawrence et al. [12], but tailored to our differing setup (single viewpoint, recovery of surface geometry, fewer image measurements).

3. Imaging Setup and Assumptions

Consider a photometric stereo setup with fixed object, fixed orthographic camera, and m images taken under distant point source illumination, with known source positions scattered about the sphere of incident directions. From this set of images, we wish to recover the surface normal and BRDF at each point on the object's surface. Recent work by Alldrin and Kriegman shows how to reliably recover the azimuthal component of each surface normal (relative to the camera coordinate frame), by assuming that the BRDF at each point is isotropic [1]. The primary advantage to their approach is that the BRDF is allowed to vary arbitrarily in both the spatial and angular domain, so long as the BRDF is isotropic. We seek to recover the elevation angle of the normal by imposing two additional constraints : (1) that the surface be composed of a small set of fundamental materials, and (2) that the BRDF at each point is well approximated by a bi-variate function.

More specifically, suppose the BRDF at each surface point is a linear combination of a small set of basis BRDFs. Then the BRDF at each point can be compactly represented as the product of two rank-constrained matrices,

$$\mathbf{H} = \mathbf{W}\mathbf{B}^\top \tag{1}$$

where $\mathbf{H} \in \mathbb{R}^{n \times d}$ is a discretization of the BRDF at each of n surface points, $\mathbf{B} \in \mathbb{R}^{d \times k}$ contains a discretization of k basis BRDFs, and $\mathbf{W} \in \mathbb{R}^{n \times k}$ weights the contribution of each basis BRDF at each surface point. For this decomposition to be physically valid, $\mathbf{W}$ and $\mathbf{B}$ should be non-negative and $\mathbf{B}$ should satisfy BRDF constraints such as energy conservation and reciprocity.

3.1. Bivariate BRDF Assumption

A general isotropic BRDF is a function of three dimensions, and is typically written $\rho(\theta_i, \theta_o, |\phi_i - \phi_o|)$, where (θ_i, ϕ_i) and (θ_o, ϕ_o) are the spherical coordinates of the directions of incident and reflected flux relative to a local coordinate system. (The absolute value, $|\phi_i - \phi_o|$, is sometimes discussed as a separate property called bilateral symmetry, but we do not do so here.) In what follows, it will also be convenient to represent the incident and exitant directions using unit vectors $\mathbf{s}$ and $\mathbf{v}$ in the same coordinate system.

An alternative parameterization is the halfway/difference parameterization of Rusinkeiwicz [18]. Here, an isotropic BRDF is expressed as $\rho(\theta_h, \theta_d, \phi_d)$, where θ_h (the *half-angle*) is the angle between the surface normal and the bisector vector $\mathbf{s} + \mathbf{v}$, and (θ_d, ϕ_d) are the spherical coordinates of the source vector computed relative to the bisector. In particular, θ_d (the *difference angle*) is the angle between the source vector and the bisector vector.

Both of these parameterizations represent all three dimensions of the isotropic BRDF domain. The possibility that general isotropic BRDFs might be well-represented by simpler bi-variate functions was first formally studied by Stark, Arvo and Smits [23]. Their work is motivated by the observation that a number of parametric BRDF models (Lafortune, Phong, Blinn, and Ward) are inherently bivariate functions. Drawing from a combination of empirical observations and theoretical insights, they propose the '$\alpha\sigma$-parameterization' for bivariate BRDFs and show this to represent a small number of measured BRDFs [26] with high fidelity. In this paper, we use an alternative bivariate parameterization based on the half-way and difference angles, $\rho(\theta_h, \theta_d)$. One can show that there is a bijective mapping between (θ_h, θ_d) and (α, σ).

3.2. Image Formation Model

Suppose we know the true surface normal at each surface point. Then this imputes a half-angle for each surface point and light source direction from which we form a data matrix, $\mathbf{E} \in \mathbb{R}^{n \times m}$. The i, jth entry is simply the image intensity at the ith surface point illuminated by the jth light source. If we assume the BRDF at each point is a linear combination of a set of basis BRDFs, then the BRDF of the ith point can be expressed as $\mathbf{H}_i^\top = \mathbf{w}_i^\top \mathbf{B}^\top \in \mathbb{R}^{1 \times d}$, where $\mathbf{w}_i \in \mathbb{R}^{k \times 1}$ is a set of material weights and $\mathbf{B} \in \mathbb{R}^{d \times k}$ contains a discretization of the basis BRDFs. The image intensity at the ith point under the jth illuminant is then modeled as,

$$\begin{aligned}
e_{ij} &= \mathbf{H}_i^\top \mathbf{\Phi}_{ij} \max\{0, \mathbf{n}_i^\top \mathbf{s}_j\} \\
&= \mathbf{H}_i^\top \widetilde{\mathbf{\Phi}}_{ij} \\
&= \mathbf{w}_i^\top \mathbf{B}^\top \widetilde{\mathbf{\Phi}}_{ij}
\end{aligned} \tag{2}$$

where $\max\{0, \mathbf{n}_i^\top \mathbf{s}_j\}$ accounts for shading and $\mathbf{\Phi}_{ij} \in \mathbb{R}^{d \times 1}$ is an interpolation vector mapping the domain of BRDF $\mathbf{H}_i$ to the half-angle / difference angle of the ijth measurement.

Equation 2 is easily extended to multiple color channels by slightly altering the BRDF matrix $\mathbf{B}$ and interpolation matrices, $\widetilde{\mathbf{\Phi}}_{ij}$. Suppose we wish to handle c color channels; then we simply fold each color channel into the BRDF discretization (e.g., $\mathbf{B} \in \mathbb{R}^{dc \times k}$) and modify the interpolation matrices appropriately. Alternatively, color can be encoded in the weight matrix $\mathbf{W}$, which allows arbitrary color scaling per point. This may be useful for surfaces that vary in color, but not in monochromatic reflectance.

4. Alternating Constrained Least Squares

If $\mathbf{W} = (\mathbf{w}_1, ..., \mathbf{w}_n)^\top$ and $\mathbf{B}$ are unknown, then we can estimate them using the method of alternating constrained least squares (ACLS), as described by Lawrence et al. [12].

Figure 2: Surface plot showing discretization of one color channel of a basis BRDF (red channel of the 2nd basis BRDF recovered from the APPLE dataset).

ACLS works by alternately updating $\mathbf{W}$ and $\mathbf{B}$ to minimize the residual between measured intensities and predicted intensities. In each iteration, the material weights $\mathbf{W}$ are updated by fixing $\mathbf{B}$ and solving the resulting constrained convex optimization problem after which $\mathbf{B}$ is updated by fixing $\mathbf{W}$ and solving another constrained convex optimization problem. ACLS is guaranteed to find a local minimum since each update step is guaranteed to not increase the residual. While this means the algorithm may not converge to a global minimum, in practice one can perform multiple trials with random initialization or use domain knowledge to initialize $\mathbf{W}$ and $\mathbf{B}$ near the optimal solution.

Since the elevation angles of the surface normals are also unknown, we also need to incorporate this into our optimization procedure. The simplest thing to do is to simply alternate between all three sets of parameters. However, since the normals are constrained to a single degree of freedom, it's possible to find a global minimum over material weights and surface normals simultaneously. This vastly improves convergence over three-way alternating optimization. We cover each step of our optimization procedure in the following subsections.

4.1. Initialization and Pre-Processing

The first step of our algorithm is to recover the azimuth angle of the surface normals using the technique of Alldrin and Kriegman [1]. Their algorithm is based on the fact that the 2D reflectance field (image intensity as a function of source direction) is symmetric about the plane spanned by the normal and viewing direction. This plane, which corresponds to the azimuth angle of the surface normal, can be estimated from a cone of light source directions parallel to

and centered about the image plane. Thus, our algorithm also requires at least a cone of light source directions centered about the optical axis. More details on this step can be found in their paper [1].

Before starting the optimization process we also randomly initialize $\mathbf{W}$, $\mathbf{B}$, and θ_n (the elevation component of the surface normals).

4.2. Update $\mathbf{B}$ with Fixed $\mathbf{n}$ and $\mathbf{W}$

In this step, we solve for the BRDF matrix $\mathbf{B}$ that minimizes the L_2 error between image measurements e_{ij} and our image formation model $\mathbf{w}_i^\top \mathbf{B}^\top \widetilde{\mathbf{\Phi}}_{ij}$. From equation 2, we set up the following constrained least squares problem,

$$\underset{\mathbf{x}}{\arg\min} \quad \|\mathbf{Ax} - \mathbf{b}\|^2$$

Subject to,
$$\mathbf{x} \geq 0 \qquad (3)$$

where $\mathbf{x} \in \mathbb{R}^{dk \times 1}$ is a vector encoding the entries of $\mathbf{B}$ in column-major order. $\mathbf{A}$ and $\mathbf{b}$ can be constructed as,

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_1 \\ \vdots \\ \mathbf{A}_n \end{pmatrix} \qquad \mathbf{b} = \begin{pmatrix} \mathbf{b}_1 \\ \vdots \\ \mathbf{b}_n \end{pmatrix}$$
$$\mathbf{A}_i = \mathbf{w}_i^\top \otimes \widetilde{\mathbf{\Phi}}_i^\top \qquad \mathbf{b}_i = \mathbf{E}_i^\top \qquad (4)$$

where $\otimes$ denotes the Kronecker product.

4.3. Update $\mathbf{W}$ and $\mathbf{n}$ with Fixed $\mathbf{B}$

For the moment, suppose both $\mathbf{n}$ and $\mathbf{B}$ are fixed. From equation 2, we set up the following constrained least squares problem for each surface point,

$$\underset{\mathbf{w}_i}{\arg\min} \quad \|\mathbf{A}_i \mathbf{w}_i - \mathbf{b}_i\|^2$$

Subject to,
$$\mathbf{w}_i \geq 0 \qquad (5)$$

where,

$$\mathbf{A}_i = \widetilde{\mathbf{\Phi}}_i^\top \mathbf{B} \qquad \mathbf{b}_i = \mathbf{E}_i^\top \qquad (6)$$

with $\mathbf{E}_i = (e_1, ..., e_m)$ the set of measurements at the ith surface point and $\widetilde{\mathbf{\Phi}}_i \in \mathbb{R}^{d \times m}$ the corresponding interpolation matrix. The solution to this optimization problem is the set of weights that minimizes the L_2 error of image measurements to intensities predicted by the image formation model, subject to non-negativity.

Since the weights for each point are estimated independently, the size of each constrained least squares problem is quite small (k variables). The global minimum with respect to both the weights and surface normal can be obtained by exhaustively searching over all possible $\mathbf{n}$ (tractable since there is only one degree of freedom).

5. Additional Constraints

In practice, we found it necessary to impose additional regularization constraints based on domain knowledge of our problem. Specifically, we impose smoothness and monotonicity over the BRDF domain, and we re-weight the constraints in Equations 3 and 5 to prevent specular highlights from dominating the solution. Empirically, these constraints improved convergence as well as the visual quality of the recovered basis BRDFs.

The need for regularization is caused by a number of factors. First, specularities usually occur in a very compact region of the BRDF domain, and within this region the BRDF value can vary by orders of magnitude. As a result, these regions of the BRDF domain are very sensitive to misalignment of light sources; a very small misregistration can lead to large changes in predicted intensity. This is exacerbated by the fact that memory constraints prevent us from using all available pixel measurements when updating the BRDF matrix $\mathbf{B}$.

To introduce a bias toward smooth BRDFs, let $\mathbf{D}_1 \in \mathbb{R}^{d \times d}$ be a discrete operator approximating the gradient over the BRDF domain. We add the following quadratic penalty term to our objective function:

$$\lambda_{D1}(\mathbf{D}_1\mathbf{B}_l)^\top(\mathbf{D}_1\mathbf{B}_l), \text{ for } l = 1...k. \tag{7}$$

This can be incorporated into Equation 3 by augmenting $\mathbf{A}$ and $\mathbf{b}$ with rows,

$$\mathbf{A}_{D1} = \sqrt{\lambda_{D1}}\left(\mathbf{I}_k \otimes \mathbf{D}_1^\top\right) \qquad \mathbf{b}_{D1} = \mathbf{0} \tag{8}$$

where $\mathbf{I}_k$ is a $k \times k$ identity matrix and $\otimes$ denotes the Kronecker product. In our experiments, we non-linearly weight the smoothness penalty so that specular regions (i.e., near $\theta_h = 0$) are penalized less strongly than non-specular regions.

Monotonicity can be enforced by adding the following inequality constraints:

$$B_{h,l} \geq B_{h+1,l}, \text{ for } l = 1...k. \tag{9}$$

It is also quite simple to enforce monotonicity over a portion of the domain (e.g., $\theta_h \in [0, \pi/4]$) by only including inequalities from the desired subset of the domain. Monotonicity is particularly important in specular portions of the BRDF domain, where undersampling and registration errors could otherwise cause unnatural visual artifacts in recovered BRDFs.

5.1. Confidence Weights

While there are relatively few measurements of specularities, such measurements carry a lot of weight since specular pixels typically have intensities more than an order of magnitude stronger than other pixels. To prevent such measurements from overly biasing the final solution, we weight each constraint in Equations 3 and 5 according to the intensity of the corresponding measurement. In our experiments, we found the following ad-hoc weights to work well,

$$c_{ij} = (\log(1 + e_{ij})/e_{ij})^3. \tag{10}$$

6. Discussion on ACLS Procedure

While our optimization procedure is computationally similar to that of Lawrence et al. [12], our methods differ in important ways. At a high level, our primary goal is to recover shape and reflectance in order to extrapolate appearance to novel viewpoints. Lawrence et al., on the other hand, assume they have data from multiple viewpoints as input and seek to obtain compact and separable representations of SVBRDFs for editing purposes. Our data is also very different from [12] in that we consider rather arbitrary geometry instead of focusing on near-planar surfaces.

The two approaches also differ at a more technical level. In their optimization, Lawrence et al. alternate between three sets of variables : BRDF basis, material weights, and surface normals. In this paper, we alternate over only two sets of variables because we find globally optimal material weights and surface normals in each iteration of our optimization algorithm. As a result, our method should be less prone to local minima. In addition, in order to bootstrap their reconstruction, Lawrence et al. use a parametric BRDF model (the Ward model), while in our work we have purposefully avoided the use of parametric BRDF models at any stage of the process. This yields an acquisition system for isotropic surfaces that is as general as possible. Another difference is how scattered data is handled. In [12], measurements are interpolated into the BRDF domain, while in our method, the BRDF domain is interpolated onto the measurements. The effects of this change are twofold : (1) each measurement counts equally in our method, and (2) interpolation of the basis BRDFs is more numerically stable than interpolation of the measured data. A similar interpolation strategy is described in [25], although our method was derived independently.

7. Experimental Validation

To validate our approach, we ran experiments on two datasets consisting of images of a gourd and an apple, respectively. For each dataset, we acquired high-dynamic range images in a dark room (see Figure 5) with the camera and light sources placed between 1.5 and 2 meters from the test object (both test objects have diameter between 5 and 10 centimeters). Light source directions and intensities were measured from specular and diffuse spheres placed in the scene with sources spanning much of the upper hemisphere of lighting directions. 102 images were acquired for
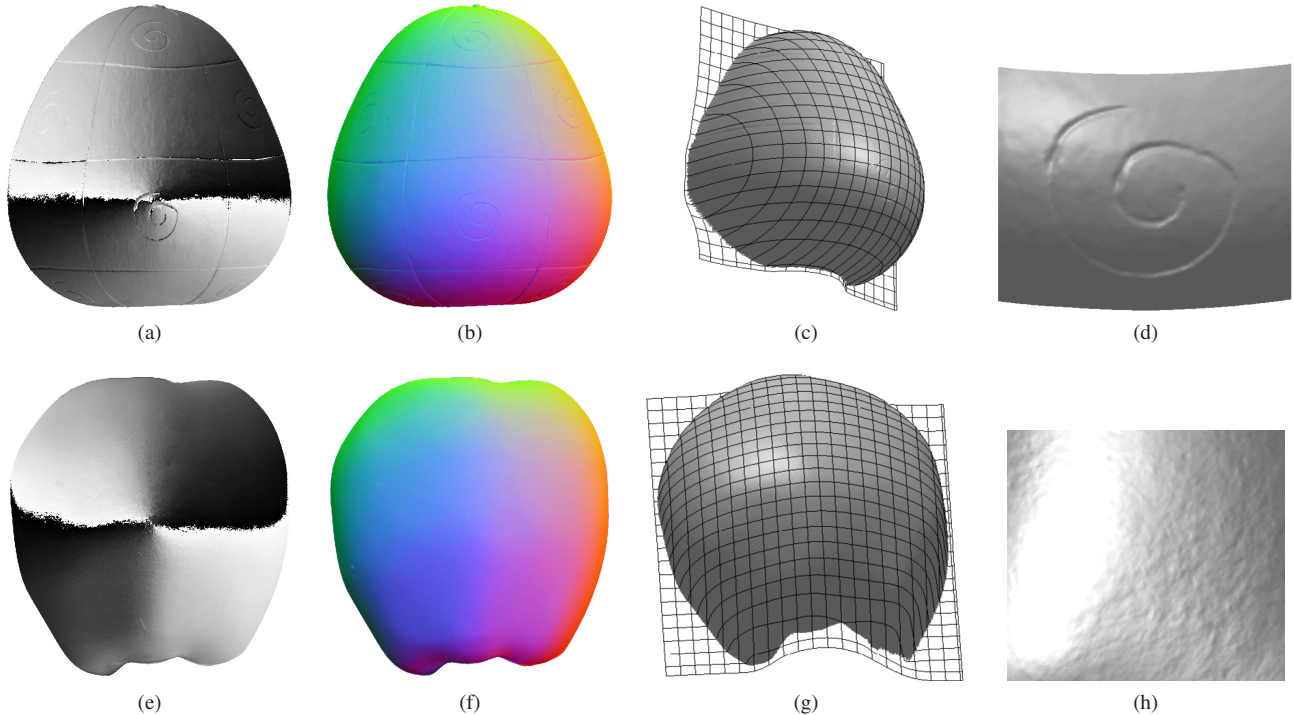
Figure 3: GOURD (top) / APPLE (bottom) shape reconstruction results. (a,e) Phase map showing the azimuthal components of the surface normal field, recovered as in [1]. (b,f) Recovered normal map, encoded to RGB as $r = (n_x + 1)/2$, $g = (n_y + 1)/2$, $b = n_z$. (c,g) Surface obtained by integrating the recovered normal field. (d,h) Detail of the surface; note the recovered mesostructure.

the GOURD dataset and 112 for the APPLE dataset. For both datasets, we assumed three basis BRDFs during reconstruction.

Figure 3 shows the shapes recovered by our algorithm on the GOURD and APPLE datasets. While the overall shape of each surface is simple (we sought to avoid cast shadows and interreflections which are not modeled by our algorithm), note that we accurately recover both the coarse and fine-scale geometric structure (i.e., macrostructure and mesostructure) of the object. In terms of appearance capture, recovery of surface mesostructure plays an important role (observe specular highlights in Figures 5 and 6).

Figure 4 shows the recovered basis BRDFs and material weight maps for the GOURD and APPLE datasets. Note the clear separation of materials visible in the material weight maps as well as the varying shape of specular lobes and body color in the recovered BRDFs.

The most important test of our algorithm is the ability to accurately generate novel views of the test objects. As seen in Figure 5, we are capable of rendering novel views that closely match real photographs. In particular, note the accurate reproduction of specular highlights which depend strongly on both the BRDF at each surface point as well

as the surface mesostructure. As a final test, we rendered each object from a variety of viewpoints under complex illumination conditions (see Figures 1 and 6 as well as the supplementary material). While this is purely qualitative, the resulting images are convincing.

8. Conclusion

Simultaneously estimating shape and reflectance of a surface from a limited set of images is a challenging problem that has traditionally been solved by restricting the reflectance function to a limited, parametric model. While this works as long as the surface does not deviate from the assumed model, it is clearly desirable to relax these restrictions. In this paper we demonstrated a technique which is truly non-parametric, and can yield more "data-driven" solutions. Our approach combines and builds upon recent work in photometric stereo and related fields, and advances the state-of-the-art in appearance capture from a single viewpoint. We also demonstrate the utility of bi-variate approximations of reflectance functions for appearance capture.

|	
(a)	(b)
(c)	(d)

Figure 4: (Top) Material weight maps recovered from the GOURD and APPLE datasets. Red, green, and blue channels correspond to (normalized) weights of the first, second, and third basis BRDFs, respectively. (Bottom) Spheres rendered with the first, second, and third basis BRDFs recovered from the GOURD and APPLE datasets.

Acknowledgments

References

[1] N. Alldrin and D. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance : A stratified photometric stereo approach. *ICCV*, 2007.

[2] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *PAMI*, 25(10):1239–1252, 2003.

[3] V. Blanz and T. Vetter. Face recognition based on fitting a 3d morphable model. *PAMI*, 25(9):1063–1074, 2003.

[4] E. Coleman, Jr. and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *CGIP*, 18(4):309–328, April 1982.

[5] CVX. Matlab software for disciplined convex programming. http://www.stanford.edu/~boyd/cvx/.

[6] FWGrid Project. http://fwgrid.ucsd.edu.

[7] D. Goldman et al. Shape and spatially-varying brdfs from photometric stereo. In *ICCV*, 2005.

[8] A. Hertzmann and S. M. Seitz. Shape and materials by example: A photometric stereo approach. In *CVPR*, 2003.

[9] A. Hertzmann and S. M. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *PAMI*, 27(8):1254–1264, 2005.

[10] High-Resolution Light Probe Image Gallery. http://gl.ict.usc.edu/Data/HighResProbes/.

[11] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *PAMI*, 3(6):661–669, November 1981.

[12] J. Lawrence et al. Inverse shade trees for non-parametric material representation and editing. *SIGGRAPH*, 2006.

[13] H. P. A. Lensch et al. Image-based reconstruction of spatially varying materials. In *Eurographics*, 2001.

[14] S. P. Mallick et al. Beyond lambert: Reconstructing specular surfaces using color. In *CVPR*, 2005.

[15] S. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Trans. on Robotics and Automation*, 6(4):418–431, August 1990.

[16] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of BRDF models. *Eurographics*, 2005.

[17] PBRT Renderer. http://www.pbrt.org.

[18] S. Rusinkiewicz. A new change of variables for efficient BRDF representation. In *Eurographics Rendering Workshop*, 1998.

[19] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. *JOSA A*, 11(11):2990–3002, November 1994.

[20] K. Schlüns and O. Wittig. Photometric stereo for non-lambertian surfaces using color information. In *ICIAP*, 1993.

[21] SDPT3. http://www.math.nus.edu.sg/~mattohkc/sdpt3.html.

[22] W. M. Silver. Determining Shape and Reflectance Using Multiple Images. Master's thesis, MIT, 1980.

[23] M. M. Stark, J. Arvo, and B. Smits. Barycentric parameterizations for isotropic BRDFs. *IEEE Transactions on Visualization and Computer Graphics*, 11(2):126–138, 2005.

[24] P. Tan et al. Isotropy, reciprocity and the generalized bas-relief ambiguity. In *CVPR*, 2007.

[25] R. P. Weistroffer et al. Efficient basis decomposition for scattered reflectance data. In *EGSR*, 2007.

[26] S. Westin. Measurement data, Cornell university program of computer graphics, 2003. http://www.graphics.cornell.edu/online/measurements/.

[27] R. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, January 1980.

[28] T. Zickler, P. N. Belhumeur, and D. J. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In *ECCV*, 2002.

[29] T. E. Zickler, P. N. Belhumeur, and D. J. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *IJCV*, 49(2-3):215–227, 2002.

Figure 5: (Top) Real images of the GOURD and APPLE test objects. (Bottom) Images rendered using recovered shapes and BRDFs. Images in columns 1 and 3 are taken from the training data. Images in columns 2 and 4 are from novel viewpoints.


Figure 6: Images rendered in novel view and illumination conditions using shape and reflectance acquired by our algorithm.