Shape and Spatially-Varying BRDFs From Photometric Stereo

Dan B Goldman[1] Brian Curless[1] Aaron Hertzmann[2] Steven M. Seitz[1]
[1]University of Washington [2]University of Toronto

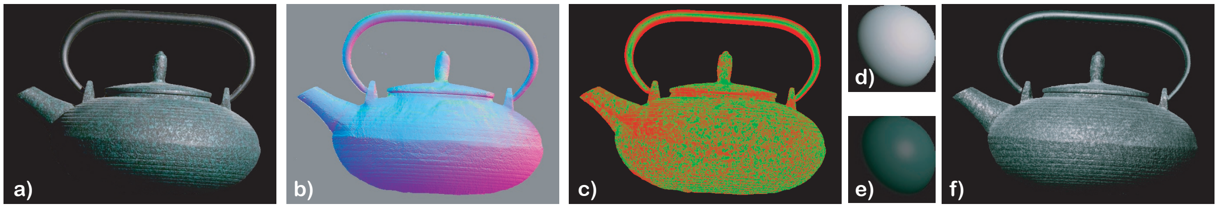

Figure 1 From a) photographs of an object taken under varying illumination (one of ten photographs is shown here), we reconstruct b) its normals and materials, represented as c) a material weight map controlling a mixture of d,e) fundamental materials. Using this representation we can f) re-render the object under novel lighting.

Abstract

This paper describes a photometric stereo method designed for surfaces with spatially-varying BRDFs, including surfaces with both varying diffuse and specular properties. Our method builds on the observation that most objects are composed of a small number of fundamental materials. *This approach recovers not only the shape but also material BRDFs and weight maps, yielding compelling results for a wide variety of objects. We also show examples of interactive lighting and editing operations made possible by our method.*

1 Introduction

Reconstructing scenes with realistic materials from images is a challenging, open research problem. In many real-world scenarios, both the shape *and* the materials are unknown and need to be recovered. Although the problem is very difficult in this general setting, good solutions exist when either the materials or shape are known. In particular, if the shape is known, several BRDF estimation methods can be applied from the graphics literature [12, 14, 15, 18, 27] to estimate material properties. On the other hand, when the BRDF is arbitrary but known or can be measured using reference objects, *example-based* photometric stereo methods [7, 19] enable reconstructing shape models.

In this paper, we address the problem of computing both shape and spatially-varying BRDFs of objects using a novel photometric stereo approach. We seek to achieve much of the generality of the example-based approach to photometric stereo of Hertzmann and Seitz [7], while removing the need for reference objects. In essence, we solve for the reference objects as part of the reconstruction process. Our approach alternates between solving for shape given a model for the BRDFs, and solving for the spatially varying BRDFs given a shape model.

Our problem of shape and material reconstruction is highly underconstrained if the BRDF at each pixel is arbitrary. Moreover, many typical regularization approaches are not applicable. For example, a simple smoothness term would prevent reconstruction of high-frequency texture. We believe the construction of the BRDF model is at the crux of the problem, and our model is both narrow enough to eliminate many ambiguities of interpretation, and flexible enough to produce a visually plausible approximation of a wide variety of real-world objects.

Researchers have observed that many objects, whether man-made or natural, can be decomposed into a small number of materials (e.g., [12]). We contribute the additional observation that, even for surfaces composed of many distinct materials, most pixels can be well-represented by mixtures of just one or two of those materials. Therefore, we impose the novel constraint that each pixel has a mixture of pairwise combinations of the materials. This additional constraint eliminates many sources of ambiguity that would occur with arbitrary linear combinations of multiple materials, and thereby allows us to resolve surface shape using photometric stereo.

Given these constraints, we can now solve for a set of global material parameters, per-pixel normals, and per-pixel material weights for each material. Our approach is to alternate between optimizing global parameters and optimizing per-pixel weights and normals. First, we hold per-pixel normal and material weight estimates fixed, and optimize over BRDF parameters using nonlinear least-squares optimization. Then we hold the BRDF parameters fixed, and jointly optimize normals and material weights for each pixel using a combination of brute force search and linear projection. With this scheme we can minimize a highly nonlinear objective function with tens of thousands of unknowns.

Using our reconstructions, many novel object editing techniques are made possible. For example, in addition to changes of lighting and limited changes of viewpoint, the fundamental materials provide a convenient representation for editing materials globally over a surface. The materials can be painted or adjusted using operations similar to those available in commercial image editing software. Finally, a number of surface property transfer operations are made possible by our representation.

2 Related Work

In the seminal work on photometric stereo [19, 24], the surface materials are assumed to be Lambertian or spatially uniform, but more recent work has begun to extend the range of this approach to surfaces with more complex, spatially-varying BRDFs [4, 7, 13, 16]. In contrast to the approach taken by Hertzmann and Seitz [7], we recover shape and material without requiring sample objects composed of the same material as the target. This allows us to reconstruct natural objects for which samples of known shape are not always available. Furthermore, our approach reconstructs BRDFs as well as shape, which allows rendering from other illuminants and viewing directions. Georghiades [4] describes a method that accomodates variation of diffuse reflectance across a surface, but models the specular properties as being constant, whereas our linear combination model allows for variation of specular properties across the surface. Our approach permits reconstruction of more complex surface properties, such as natural objects comprised of materials with different surface roughness, or mixtures of man-made materials such as metallic and non-metallic paints. However, Georghiades also recovers lighting directions whereas we assume they are known.

BRDF acquisition from photographic data has been widely researched in the computer graphics community. Most methods assume the scene geometry is known [14, 15, 18, 27]. Lensch et al. [12] cluster material estimates over a known surface. In this work they also refine the scanned geometry using extracted normal maps, as in our present work. However, they require approximate scanned geometry as input, whereas our method generates a surface model solely from the input images.

Classical work in physics-based vision includes segmentation based on physical models of surface reflectance [1, 6, 10]. We use a more complete reflectance model, making it possible to recover smoothly varying materials from natural objects. In addition, we advantage of multiple images of the same objects, enabling segmentation of pixels which may be saturated in one or more images.

We also note the success of Helmholtz stereopsis in shape reconstruction on arbitrary surfaces [23, 28]. These methods avoid reconstructing BRDFs explicitly, by exploiting the reciprocity property of BRDFs. However, our approach recovers not only shape but also a BRDF representation that can be edited, relit and viewed from different viewpoints. Our method also uses a simpler capture mechanism, requiring only a single uncalibrated camera position.

Although our method uses only a single viewpoint, some recent work utilizes information in multiple views to reconstruct both shape and materials from images [8, 9, 20, 26]. These works use only a single lighting condition as input, and therefore cannot recover BRDFs. Treuille et al. [21] use a voxel representation to carve away regions of space which are not consistent with the images. In contrast, our approach requires only a single viewpoint, and does not require reference objects.

3 Problem Statement

The input to our system is a set of images of a static target object taken from a distant camera under a different distant illuminant in each image. We assume the lighting is known, and we do not model the effects of cast shadows, interreflections, transparency, or translucency. From these inputs, we seek to reconstruct shape and BRDFs with a constrained material model.

Our material model is motivated by the observation that real world variations in BRDF across a surface are often a result of the surface's composition from a small number of substances. For example, a block of wood with light and dark grain can be viewed as having two different substances that are blended in different amounts across the surface. We call these substances *fundamental materials*, and the mixtures of these materials at each pixel are specified by *material weight maps*.

To reconstruct the surface and materials using this model, the user provides the number of fundamental materials for which to solve. The output is a set of BRDF parameters for each of the fundamental materials, and a surface normal and material weights at each pixel. We can then reconstruct the surface by integrating the normal field.

3.1 Model

We model the lights as distant directional sources, so the lighting direction $\mathbf{L}_i$ is constant over each image (indexed by i). We also model the camera as orthographic, so the view direction $\mathbf{V}$ is constant for all samples (and is therefore elided in the formulation which follows). We model the color at pixel p as if generated from a convex combination of fundamental materials:

$$I_{i,p,c} \leftarrow \sum_m \gamma_{p,m} f_c(\mathbf{n}_p, \mathbf{L}_i, \alpha_m) \tag{1}$$

We use the symbol $I_{i,p,c}$ to represent the intensity of channel c of pixel p in image i. The function f_c represents color channel c of the parameterized lighting model with normal $\mathbf{n}_p$, lighting condition $\mathbf{L}_i$ and BRDF parameter vector α_m; there is one α_m for each fundamental material. In this paper, we have used the isotropic Ward reflectance model [11] because it has low dimensionality, but other parametric models could be substituted.

Since our light sources are directional, f_c is the product of the BRDF and the light intensity, with the viewing direction held constant. (In general, f_c could represent an integral over an arbitrary distant lighting distribution.) We normalize the input images by dividing each by its light intensity. This normalization avoids giving undue weight to some images over others. Note that this also means that we drop the intensity factor when computing f_c.

Although the fundamental materials are constant over the image, the material weight maps γ vary spatially. For example, Figure 1a shows a cast iron teapot with speckled green paint. Figure 1c shows two material weight maps (encoded in the red and green channels), and Figures 1d and 1e are renderings of the two fundamental materials, corresponding to the cast iron and green paint respectively.

Based on this model, we formulate the following objective function to solve for shape and materials:

$$Q(\mathbf{n}, \alpha, \gamma) = \sum_{i,p,c} \left(I_{i,p,c} - \sum_m \gamma_{p,m} f_c(\mathbf{n}_p, \mathbf{L}_i, \alpha_m) \right)^2 + Q_{sp}(\alpha) \tag{2}$$

Q is minimized with respect to the normals, material parameters, and material weight maps (denoted by $\mathbf{n}$, α, and γ, respectively). Here $I_{i,p,c}$ refers to normalized measured pixel values. Not all normal fields correspond to real surfaces, but we constrain the normals $\mathbf{n}$ to be an integrable field.

The term Q_{sp} expresses a prior on the similarity of fundamental materials. Using the data term alone, it is possible for the fundamental materials to extrapolate far beyond any observed materials in the scene. For example, consider the case of a surface with grey matte paint with albedos in the range 0.4 to 0.6. We would like to restrict the individual albedo estimates for the surface to lie in this range, to prevent errors such as setting a fundamental material albedo to 0 (black). However, we don't know the range of materials in advance, and, without an additional penalty term, this surface could easily be described by basis materials with albedos 0 and 1. In order to avoid this problem, we add the term Q_{sp}, which we describe in more detail in Section 4.3.

Although this objective function is easy to express, it can lead to gross overfitting if the material weights γ are left unconstrained. For each additional fundamental material the dimensionality of the mixture increases, but in reality only one or two materials are found at most points on a surface. Therefore, we use pairwise convexity constraints for these material weights, which mitigates overfitting:

$$\begin{gathered} \gamma_{p,m} \geq 0, \\ \textstyle\sum_m \gamma_{p,m} = 1, \\ \exists m_1, m_2 \quad \text{such that} \quad \gamma_{p,m_1} + \gamma_{p,m_2} = 1 \end{gathered} \tag{3}$$

Optimization of Equation (2) is non-trivial; many optimization algorithms become easily trapped in local minima due to the nonlinear terms of f_c and integrability constraints for the normal field $\mathbf{n}$. Our optimization approach, described in Section 4, is designed to avoid these problems.

4 Algorithm

Our approach has five components, each of which is described in detail in the sections which follow.

Light calibration. The light source direction and intensity is estimated using diffuse grey and chrome spheres captured under the same illumination as the target object. The calibration method is described in Section 4.1.

Initialization. Our normal maps are initialized using Lambertian photometric stereo, with threshholds to reject specular highlights. This also gives us an initial estimate of diffuse albedo, which is used to initialize the material weight maps (Section 4.2).

After initialization, the system optimizes the objective function iteratively by repeating the following three steps:

1. Optimize BRDF parameters. The BRDF parameters are optimized while holding the normals and material weights constant (Section 4.3).

2. Compute surface normals and material weight maps. To compute the normals and material weight maps, while holding the BRDF parameters constant, Equation (2) is optimized jointly over normals and material weights. The normal optimization is performed as a discrete search, and the material weights are optimized by linear projection (Section 4.4). For the first several iterations of the algorithm, the material weights are held at their initial values. Once the other parameters have converged, they are allowed to vary freely in this step.

3. Enforce integrability. The normals generated in the previous step are not guaranteed to be consistent with a 3D surface, so integrability is enforced by solving a Poisson equation to obtain a least-squares surface reconstruction, and subsequently the normals are recomputed (Section 4.5).

Termination. Steps 1 - 3 are iterated until the objective function no longer decreases in successive iterations of the outer loop. Each step in our algorithm is guaranteed to monotonically decrease the objective function, except the enforcement of integrability (Section 4.5). However, in our experience, this projection step makes only very minor changes to the surface. Therefore, our optimization approach is likely – although not guaranteed – to find a solution near a local optimum.

4.1 Light Calibration

We begin by calibrating the light source directions and relative intensities by photographing calibration objects under the various lighting conditions. In principle, images of a chrome sphere alone suffice to recover both light direction and intensity: the direction is recovered by reflecting the viewing vector about the normal at the point of greatest brightness, and the relative intensity by integrating the measured radiance. However, we have found that a large number of exposures may be necessary to obtain an intensity estimate with low variance from an image of a chrome sphere alone. This is because all of the intensity is concentrated in a small number of pixels, so even a small amount of Gaussian pixel noise gives a high variance to the resulting integral estimate.

Instead, we use the following method employing two calibration objects: we determine the illumination direction from the image of the chrome sphere, and then determine the light intensity from an image of a diffuse sphere. Thus, only one low-dynamic-range exposure for each calibration object is required to reconstruct both intensity and direction.

Given the correct lighting direction $\mathbf{L}_i$, the intensity of a diffuse sphere in image i is $I_{i,p} = \ell_i \rho \mathbf{n}_p^T \mathbf{L}_i$, where ρ is the diffuse albedo of the sphere and ℓ_i is the intensity of the light in image i. So, given the known normals, we solve for the relative light intensity $\ell_i \rho = \sum_p I_{i,p} / \sum_p \mathbf{n}_p^T \mathbf{L}_i$. This is performed independently for each color channel.

If ρ is known, the absolute intensity can be recovered. However, in this paper we have used the relative intensities $\ell_i \rho$ to solve for materials, so all the BRDFs we have recovered are actually scaled by the unknown constant scale factor ρ.

These calibration objects are similar to those used by Hertzmann and Seitz [7], but they are used in a different way. In particular, we do not require or expect that the BRDFs of our target objects are composed of a linear combination of the calibration materials.

4.2 Initialization

We initialize normals using Lambertian photometric stereo [24]. Since this method fails in the presence of specular highlights, we employ manually-selected intensity threshholds to reject shadow and specular highlight pixels from consideration. For particular configurations of lights and shiny objects, some pixels may have fewer than 4 inlier samples with which to estimate a normal, and for these pixels we simply choose an arbitrary plausible normal. Although the resulting normal map is quite poor (Figure 2) it suffices as an initial guess.

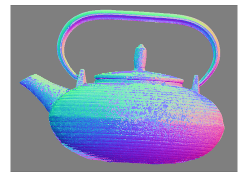

Figure 2 RGB-encoded normals acquired using the Lambertian photometric stereo method, with threshholding to exclude shadows and highlights.

Lambertian photometric stereo also provides an estimate of the diffuse albedo. To compute our initial material weight maps, we first transform the diffuse albedo into HSV colorspace [1] and discard the V channel. This transformation reduces distortions of the diffuse albedo estimate due to specular highlights and shadows. We then cluster the pixels of the image using the parameterization $(\cos(2\pi \mathrm{H}), \sin(2\pi \mathrm{H}), \mathrm{S})$. We use an EM optimization for mixtures-of-Gaussians to segment this image into separate regions [25].

4.3 BRDF Parameter Optimization

The BRDFs of the fundamental materials, denoted as f_c in Equation (2), are generally nonlinear functions of their parameters α. We optimize the objective function over all α simultaneously, using the Levenberg-Marquardt nonlinear optimization algorithm [17].

To keep the solution space highly constrained, we use the isotropic Ward model [11] as our parametric reflectance model:

$$\rho_{bd,iso}(\theta_i, \phi_i; \theta_r, \phi_r) = \frac{\rho_d}{\pi} + \frac{\rho_s}{\sqrt{\cos\theta_i \cos\phi_r}} \frac{\exp[-\tan^2\delta/\beta^2]}{4\pi\beta^2} \tag{4}$$

where ρ_d and ρ_s are the diffuse and specular reflectance coefficients, β is a measure of roughness, and δ is the angle between $\mathbf{n}$ and the halfway vector $\mathbf{h} = (\mathbf{V}+\mathbf{L})/||\mathbf{V}+\mathbf{L}||$. This model has only seven parameters (in our system, ρ_d and ρ_s are RGB vectors), and is thus well-suited to our problem. Each fundamental material thus has a parameter vector α_m comprised of $\rho_{d,m}$, $\rho_{s,m}$, and β_m.

As previously stated, we include a small penalty term between each pair of fundamental materials in our objective function: Q_{sp} in Equation (2). This takes the form of a spring term that works together with the pairwise convexity constraint to constrain the range of materials to those observed in the data. This term need not be very strong, as it is intended only to disambiguate between solutions which would have the same or nearly the same energy under the original objective function.

For the Ward model, we use the following spring term:

$$Q_{sp}(\alpha) = \sum_{i \neq j} \varepsilon_d ||\rho_{d,i} - \rho_{d,j}||^2 + \varepsilon_s ||\rho_{s,i} - \rho_{s,j}||^2 + \varepsilon_\beta ||\beta_i - \beta_j||^2 \tag{5}$$

where (i, j) are all pairs of distinct materials. In practice, ρ_d and ρ_s tend to vary over a similar range, whereas the roughness β typically varies an order of magnitude more. Accordingly, we set $\varepsilon_d = \varepsilon_s = 1.0$ and $\varepsilon_\beta = 0.1$ for all of our examples. Note that these values are very small relative to the data term, which sums over many pixels, so that the spring term has an effect only when the data term alone does not constrain the solution.

4.4 Computing Normals and Material Weight Maps

Next, we jointly optimize the normals ($\mathbf{n}_p$) and material weights ($\gamma_{p,m}$) of Equation (2). We first precompute the function f_c over a discrete sampling of normals $\mathbf{n}$ for each of the lighting samples $\mathbf{L}_i$ and fundamental material parameters α_m. In practice, this simply means rendering a small "virtual sphere" (Figures 1d, 1e, 5d) of each fundamental material under each lighting condition of the input set.

Given these samples of the appearance functions f, and the pairwise convex combination constraint, weights $\gamma_{p,m}$ are computed by linear projection and brute force search over all normals and all pairwise combinations of fundamental materials. Specifically, let $\phi^c_{m,i}(\mathbf{n}_p) = f_c(\mathbf{n}_p, \mathbf{L}_i, \alpha_m)$ denote the virtual sphere images for material m. For a given choice of normal $\mathbf{n}_p$ and pair of fundamental materials m_1, m_2, the objective function reduces to

$$Q_p = \sum_{i,c} \left(I_{i,p,c} - \gamma_{p,m_1} \phi^c_{m_1,i}(\mathbf{n}_p) - \gamma_{p,m_2} \phi^c_{m_2,i}(\mathbf{n}_p) \right)^2 \tag{6}$$

This is minimized by substituting in the constraint $\gamma_{p,m_2} = 1 - \gamma_{p,m_1}$, and solving $dQ_p/d\gamma = 0$. After some rearrangement we have:

$$\gamma_{p,m_1} \leftarrow \frac{\sum_{i,c} \left(I_{i,p,c} - \phi^c_{m_2,i}(\mathbf{n}_p) \right) \left(\phi^c_{m_2,i}(\mathbf{n}_p) - \phi^c_{m_1,i}(\mathbf{n}_p) \right)}{\sum_{i,c} \left(\phi^c_{m_2,i}(\mathbf{n}_p) - \phi^c_{m_1,i}(\mathbf{n}_p) \right)^2} \tag{7}$$

Since we wish to constrain the solution to convex combinations, we also clamp γ_{p,m_1} to lie between 0 and 1. We solve for these optimal convex weights for each normal and pair of materials, and select the pair and normal with the lowest objective Q_p.

Depending on the resolution of the virtual sphere, the full brute force normal search described in Section 4.4 can be quite slow; a single pass over the image may take several hours to complete. In order to accelerate the computation, we limit the brute-force search to normals which lie close to the previous normal for each pixel. This can result in small areas which become "trapped" at the wrong normal values, so after the algorithm converges, we perform a final pass of normal/weight optimization using the full global normal search. Although in principle this strategy could cause the algorithm to converge to a sub-optimal solution, in practice we have found it gives good results with dramatically less computation than the full normal search.

4.5 Enforcing Integrability

To compute a 3D surface from the estimated surface orientations, given the normal $\{n_x, n_y, n_z\}$ for each point, we solve for the height field $z(x, y)$ that minimizes

$$\Psi(z) = \sum_{x,y} \left(n_z \frac{\partial z(x,y)}{\partial x} + n_x \right)^2 + \left(n_z \frac{\partial z(x,y)}{\partial y} + n_y \right)^2 \tag{8}$$

using the approximations $\frac{\partial z(x,y)}{\partial x} = (z(x+1, y) - z(x, y))$, $\frac{\partial z(x,y)}{\partial y} = (z(x, y+1) - z(x, y))$ [3, 22]. This amounts to integrating the normal field. The minimization gives rise to a large but sparse system of linear equations which we solve using the conjugate gradient method [17].

The normals are then recomputed from this surface approximation. This step can be viewed as projecting the normal field into the subspace of feasible normal fields.

5 Results and Applications

To capture our source images, we programmed a Lutron lighting control system and Canon 10D camera with a 400mm telephoto lens to automatically capture multiple exposures of each lighting direction. (Although our algorithm assumes an orthographic camera and parallel light rays, we have obtained good results with distances from target to camera and lights of only 5 feet. The targets were all 6 inches in diameter or smaller.) Our images are captured at the full resolution of the camera (3072 × 2048), but most of the examples in this paper were computed at a downsampled resolution of 768 × 512. A typical capture session with 12 light sources takes about 25 minutes, most of which is simply the time to download the high-resolution images to disk via the camera's USB 1.0 interface. The multiple exposures of each lighting direction are then combined into high-dynamic range images using the technique of Debevec and Malik [2]. Since we use multiple fixed light sources, we only calibrate the light sources once each capture session.

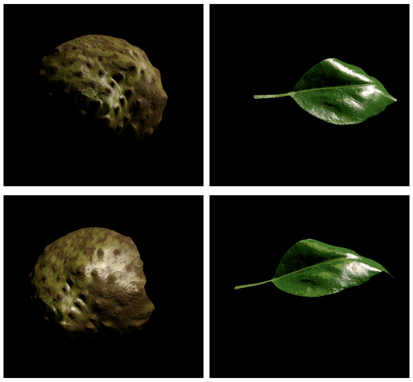

Figure 3 Synthesized views of our reconstructions of a cherimoya (a tropical fruit) and a leaf.

Most of the examples shown in this paper converged after 10-20 iterations of the outer loop of the algorithm (steps 1, 2 and 3 from Section 3), in about 5-10 hours on a 2.8GHz Xeon processor.

A few of our reconstructions, with different viewpoints and illuminations, are shown in Figures 1, 3, 4, and 5. Note that, for each object, the algorithm estimates a detailed normal map, plausibly segments the surface materials, estimates the reflectance properties of the materials, and reproduces the input imagery fairly accurately. A few artifacts still occur in regions of the surface that had highlights in most of the views, such as the frontal portion of the candlestick, and the lower part of the leaf. The presence of highlights in all of the images causes the algorithm to overesimate the diffuse component. These artifacts do not have a significant impact on rerendering and relighting of these objects.

Since our goal is to produce a visually plausible model, we compare to ground truth by considering the reconstruction error for novel lighting conditions. In Figure 4, we show our relighting results compared with real photos taken under lighting conditions that were held out from the input dataset. The RMS error is 0.0290 for the leaf image pair and 0.0303 for the teapot image pair (with source images scaled to a maximum pixel value of 1.0).

5.1 Editing Operations

Direct BRDF Manipulation. The Ward BRDF model has a small number of parameters that can be directly manipulated to change one or more of the fundamental materials, without modifying the others. For example, in Figure 6 (top), we have manually edited the BRDF parameters of the green paint of our teapot to appear as gold leaf.

BRDF Transfer. BRDF parameters captured from one object can easily be transferred to another. Figure 6 (bottom) illustrates materials captured from our cherimoya, applied to the re-rendering of our teapot.

Figure 6 Teapot edited to appear as gold leaf by direct manipulation of BRDF parameters, and as a cherimoya by transfer from another model.

Material/shape extraction and transfer. One application of our method is to extract surface texture and detail for transfer to new objects; we show an example in which texture and normals were extracted from a portion of the cherimoya in Figure 5, and reapplied to a computer-graphic model of a head. Details of this procedure are given in [5], and the results are shown in Figure 7.

6 Discussion and Future Work

We have demonstrated a method that acquires both shape and spatially-varying BRDFs from a set of photographs captured under varying illumination. Although our shape and material reconstructions are lower fidelity than those attainable using methods that assume one or the other is given and use large numbers of observations, we can nonetheless acquire a wide range of models, which can be reused under various lighting and viewing conditions. The spatially varying BRDFs that we acquire enable a set of useful and interesting editing operations. We believe this approach represents an important step towards acquisition and reconstruction of both shape and material from a single set of photographic data.

Our approach is able to capture shape and BRDFs of reflective objects using a small number of photos and without

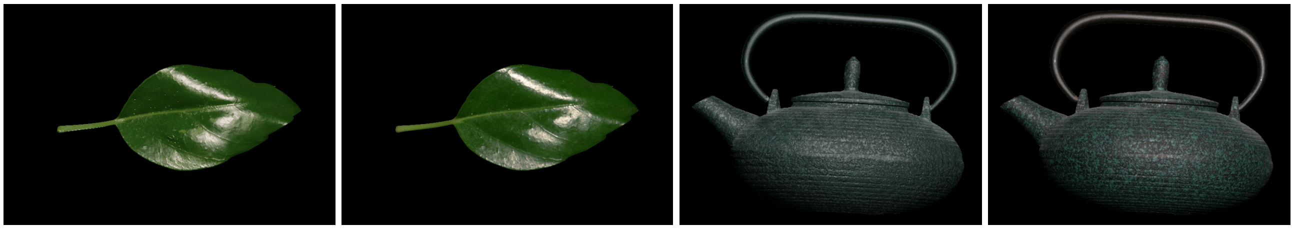

Figure 4 Comparison to ground truth. Left in each pair: models rendered under novel lighting conditions not in the training set. Right in each pair: images of real object under same lighting condition.

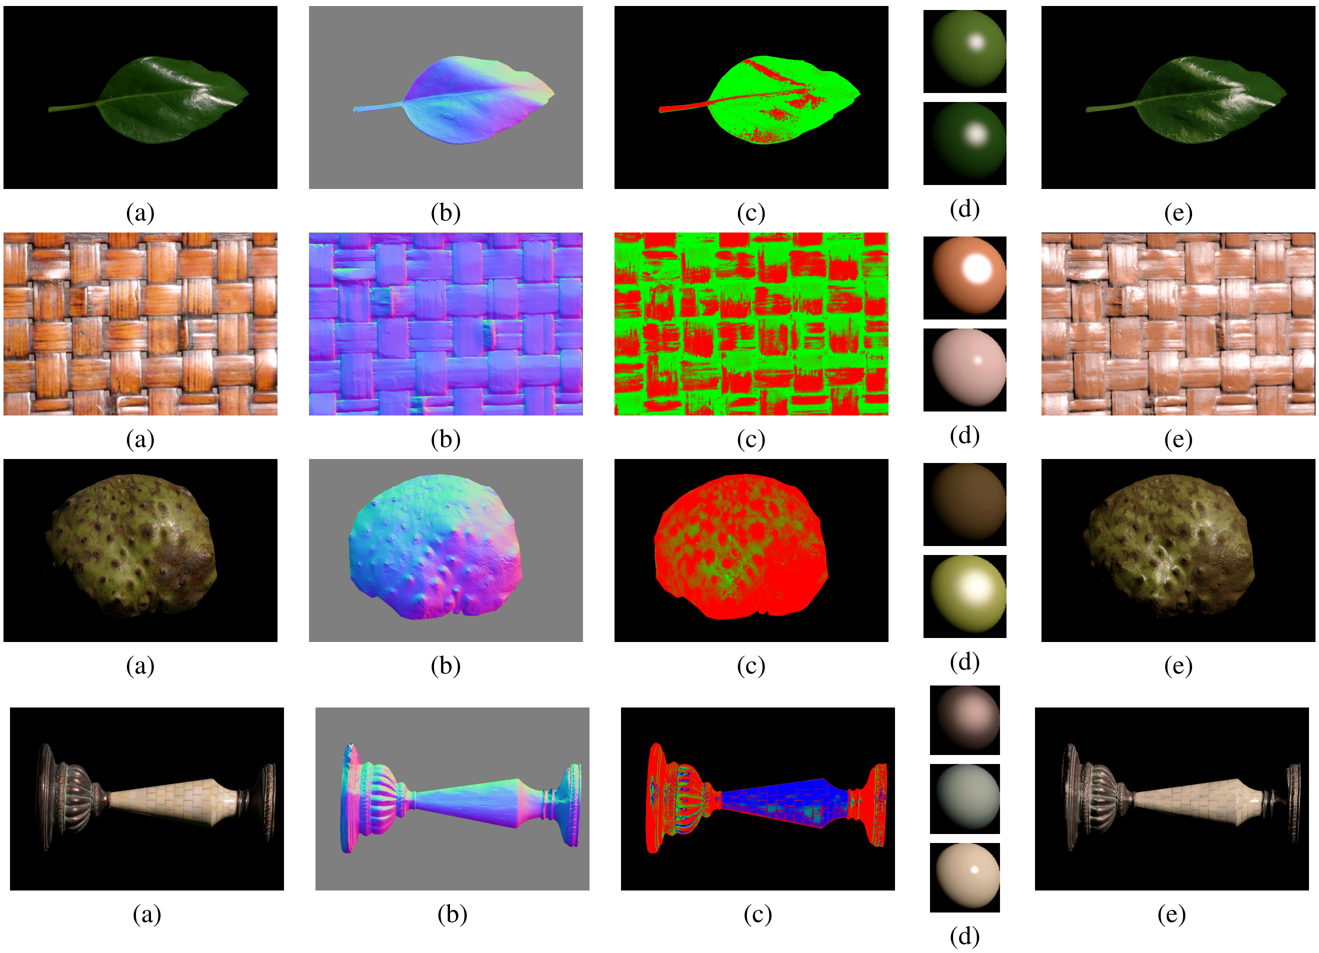

Figure 5 Input, model, and reconstructions for leaf, woodweave, cherimoya, and candlestick. a) Source image (1 of 10). b) Recovered normal map, RGB-encoded. c) Recovered material weights, in false color. d) Virtual spheres for fundamental materials. e) Model rendered under original lighting condition (compare with a).

repainting the objects gray, as is typically required for high-quality laser range scanning. However, the objective function in Equation 2 is still subject to some overfitting, usually in the case where pixels appear in specular highlight in most or all of the image samples. Such pixels may be assigned a bright diffuse material instead of a dark specular material, because there is too little data to distinguish diffuse from specular color. However, these pixels often have neighbors which are assigned properly. It is therefore possible that adding a smoothness term for material weight maps to our objective function will ameliorate these artifacts. Such an objective function will require a modified optimization approach.

Also, because we employ a local reflectance model, our algorithm does not properly account for shadows, interreflections and subsurface scattering. Few real-world objects are free of these effects, therefore future work must address techniques to compensate for their appearance.

Figure 7 A head model textured using material maps and bump maps synthesized from a cherimoya skin.

Despite these limitations, we believe our method will enable more rapid acquisition of computer models suitable for applications in computer graphics and other domains, from source material at a range of scales which are inaccessible to laser range scanners.

Acknowledgements

This work was supported in part by National Science Foundation grants CCR-0098005, IIS-0413198, NSERC, the Washington Research Foundation, an Office of Naval Research YIP award, the Animation Research Labs, and Microsoft Corporation.

References

[1] R. Bajcsy, S. Lee, and A. Leonardis. Color image segmentation with detection of highlights and local illumination induced by inter-reflections. In *Proc. ICPR*, pages 785–790, 1990.

[2] P. E. Debevec and J. Malik. Recovering High Dynamic Range Radiance Maps from Photographs. In *Proc. SIGGRAPH 97*, pages 369–378, 1997.

[3] D. A. Forsyth and J. Ponce. *Computer Vision: A Modern Approach*. Prentice Hall, 2003.

[4] A. S. Georghiades. Recovering 3-D shape and reflectance from a small number of photographs. In *Proc. EGWR*, pages 230–240, 2003.

[5] D. B. Goldman, B. Curless, A. Hertzmann, and S. Seitz. Shape and spatially-varying BRDFs from photometric stereo. Technical Report UW-CSE-04-05-03, University of Washington CSE, 2004.

[6] G. Healey. Segmenting images using normalized color. *IEEE Trans. on Systems, Man and Cybernetics*, 22(1):64–73, 1992.

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

[8] H. Jin, S. Soatto, and A. J. Yezzi. Multi-view stereo beyond Lambert. In *Proc. CVPR 2003*, pages 171–178, 2003.

[9] H. Jin, A. J. Yezzi, and S. Soatto. Region-based segmentation on evolving surfaces with application to 3D shape and radiance estimation. In *Proc. ECCV*, pages 114–125, 2004.

[10] G. Klinker, S. Shafer, and T. Kanade. A physical approach to color image understanding. *IJCV*, 4(1):7–38, 1990.

[11] G. J. W. Larson. Measuring and Modeling Anisotropic Reflection. In *Proc. SIGGRAPH*, pages 265–272, 1992.

[12] H. P. A. Lensch, J. Kautz, M. Goesele, W. Heidrich, and H.-P. Seidel. Image-based reconstruction of spatial appearance and geometric detail. *ACM Trans. Graph.*, 22(2):234–257, 2003.

[13] S. P. Malick, T. E. Zickler, D. J. Kriegman, and P. N. Belhumeur. Beyond Lambert: Reconstructing specular surfaces using color. In *Proc. CVPR*, 2005.

[14] 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 *Proc. EGWR*, pages 139–152, 1999.

[15] W. Matusik, H. Pfister, M. Brand, and L. McMillan. Efficient isotropic BRDF measurement. In *Proc. EGWR*, pages 241–248, 2003.

[16] S. K. 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, 1990.

[17] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery. *Numerical Recipes in C: The Art of Scientific Computing (2nd ed.)*. Cambridge University Press, 1992.

[18] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *Proc. SIGGRAPH 97*, pages 379–387, 1997.

[19] W. M. Silver. Determining shape and reflectance using multiple images. Master's thesis, MIT, Cambridge, MA, 1980.

[20] S. Soatto, A. J. Yezzi, and H. Jin. Tales of shape and radiance in multi-view stereo. *IJCV*, pages 947–981, October 2003.

[21] A. Treuille, A. Hertzmann, and S. Seitz. Example-based stereo with general BRDFs. In *Proc. ECCV*, 2004.

[22] E. Trucco and A. Verri. *Introductory Techniques for 3-D Computer Vision*. Prentice Hall, 1998.

[23] P. Tu and P. R. S. Mendonça. Surface reconstruction via helmholtz reciprocity with a single image pair. In *Proc. CVPR*, pages 541–547, 2003.

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

[25] T. Yamazaki. Introduction of EM algorithm into color image segmentation. In *Proc. ICIPS'98*, pages 368–371, 1998.

[26] T. Yu, N. Xu, and N. Ahuja. Recovering shape and reflectance model of non-Lambertian objects from multiple views. In *Proc. CVPR 2004*, volume 2, pages 226–233, 2004.

[27] Y. Yu, P. Debevec, J. Malik, and T. Hawkins. Inverse global illumination: Recovering reflectance models of real scenes from photographs. In *Proc SIGGRAPH*, pages 215–224, 1999.

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