

A Dictionary-based Approach for Estimating Shape and Spatially-Varying Reflectance

Zhuo Hui and Aswin C. Sankaranarayanan*
ECE Department, Carnegie Mellon University, Pittsburgh, PA

Abstract

We present a technique for estimating the shape and reflectance of an object in terms of its surface normals and spatially-varying BRDF. We assume that multiple images of the object are obtained under fixed view-point and varying illumination, i.e., the setting of photometric stereo. Assuming that the BRDF at each pixel lies in the non-negative span of a known BRDF dictionary, we derive a per-pixel surface normal and BRDF estimation framework that requires neither iterative optimization techniques nor careful initialization, both of which are endemic to most state-of-the-art techniques. We showcase the performance of our technique on a wide range of simulated and real scenes where we outperform competing methods.

Keywords. Photometric stereo, BRDF estimation, Dictionaries, Spatially varying BRDF.

1. Introduction

Photometric stereo [27] seeks to estimate the shape of an object from images which are obtained from a static camera and under varying lighting. While there has been remarkable progress in photometric stereo, the vast majority of techniques are devoted to scenes that exhibit simple reflectance properties. Yet, this creates a significant disconnect between theory and practice since the vast majority of real-life scenes involve materials with complex reflectance properties that interact with light in myriad number of ways.

In this paper, we present a method for recovering the surface normals and the reflectance of opaque objects with complex spatially-varying reflectance. The key challenge here is that the reflectance, characterized in terms of its spatially-varying bidirectional reflectance distribution function (SV-BRDF), and the shape, characterized in terms of surface normals, are coupled and need to be jointly estimated. Further, the SV-BRDF is a very high-dimensional

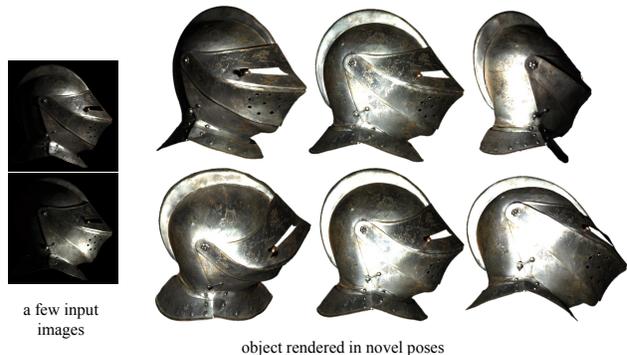


Figure 1. **Recovery of surface normals and spatially-varying BRDF.** We propose a framework for per-pixel estimation of surface normal and BRDF in the setting of photometric stereo. Shown above are the estimated shape and rendered images of a visually-complex object. The results were obtained from 250 images.

signal and, in the absence of additional assumptions, requires a large number of input images for robust estimation.

A common assumption for computational tractability is that the BRDF at each pixel is a weighted combination of a few reference BRDFs [14]. We now need to estimate only the reference BRDFs and their abundances at each pixel which is a significant reduction in the dimensionality of the unknowns. In Goldman et al. [8], the isotropic Ward model, a parametric model is used to characterize the reference BRDFs. Alldrin et al. [2] assume that the reference BRDFs are approximated by the so-called bivariate model, a non-parametric model that approximates the 4D BRDF as a 2D signal. In both cases, the problem of shape and SV-BRDF estimation reduces to alternating minimization over the surface normals, the reference BRDFs, and abundances of the reference BRDFs at each pixel. The drawback of these approaches is that the optimization is not just computationally expensive but also has a critical dependence on the ability to find a good initial solution since the underlying problem is non-convex and riddled with local minima.

An alternate approach called example-based photometric stereo [11, 20] introduces reference objects — objects with known shape — in the scene. These techniques rely on the

*The authors were supported in part by the NSF grant CCF-1117939.
Email: {zhui, saswin}@andrew.cmu.edu

concept of *orientation consistency* [11] which suggests that two surface elements with identical normal and BRDF will take the same appearance when placed in the same illumination. If the reference object had the same BRDF as the target object, we can obtain the shape of the target simply by comparing the intensity profile observed at a pixel on the target to those observed on the reference object. When the target’s BRDF is spatially-varying, it can be shown that two reference objects — one diffuse and the one specular — are sufficient to recover the surface normals of the target by approximating the unknown BRDF at each pixel as a non-negative linear combination of the reference BRDFs [11]. While introducing reference objects is not always desirable, example-based photometric stereo produces precise shape estimates without requiring the knowledge of lighting.

The technique proposed in this paper relies on the core principle of example-based photometric stereo *without actually introducing reference objects into the scene*. Instead, given a dictionary whose atoms are BRDFs associated with a wide range of materials, we can render virtual spheres, one for each atom in the dictionary, under the knowledge of the scene illumination (typically a distant point light source). This provides a set of “virtual examples” that can be used to obtain a per-pixel estimate of the shape and reflectance of the scene with arbitrary spatially-varying BRDF (see Figure 2). The assumption that we make is that the unknown BRDF at each pixel lies in the non-negative span of the dictionary atoms. We show that the surface normals and the BRDFs can be estimated via a sequence of tractable linear inverse problems. This obviates the need for complex iterative optimization techniques as well as careful initialization required to avoid convergence to local minima. The interplay of these ideas leads to a robust surface normal and SV-BRDF estimation technique that provides state-of-the-art results on challenging real-life scenes (see Figure 1).

Contributions. We make the following contributions.

[Model] We propose the use of a dictionary of BRDFs to regularize the surface normal and SV-BRDF estimation. The BRDF at each pixel of an object is assumed to lie in the non-negative span of the dictionary atoms.

[Normal estimation] We show that the surface normal at each pixel can be efficiently estimated using a coarse-to-fine search.

[SV-BRDF estimation] Given the surface normals, we recover the BRDF at each pixel independently by solving a linear inverse problem that enforces sparsity in the occurrence of the reference BRDFs at the pixel.

[Validation] We showcase the accuracy of the shape and SV-BRDF estimation technique on a wide range of simulated and real scenes and demonstrate that the proposed technique outperforms state-of-the-art.

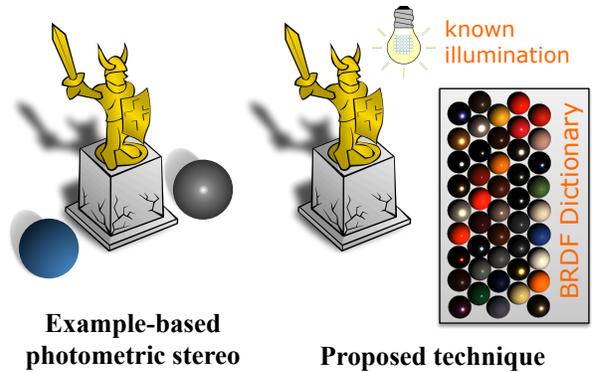


Figure 2. **Virtual examples.** In example-based photometric stereo [11], objects with known shape and reflectance are introduced into a scene. In contrast, we use a dictionary of BRDFs to render “virtual examples” that guide the normal estimation problem. This enables us to handle scenes with very complex reflectance, since we can use a larger collection of virtual examples.

2. Prior work

In this section, we review some of the key techniques for non-Lambertian shape estimation.

The diffuse + specular model. It is well known that the collections of images of a convex Lambertian object typically lies close to a low-dimensional subspace [3, 19]. This naturally leads to techniques [13, 28–30] that robustly fit a low-dimensional subspace, capturing the Lambertian component while isolating non-Lambertian components, such as specularities, as sparse outliers. However, these techniques have restrictive assumptions on the range of BRDFs to which they are applicable, and more importantly, miss out on powerful cues to the shape of the object that is often present in specular highlights.

Parametric BRDFs. Parametric models such as the Blinn-Phong [4], Ward [25], Oren-Nayar [16], and Cook-Torrance model [6] are based on macro-behavior established using specific micro-facet models on the materials, and have been widely used in computer graphics. In the context of shape and SV-BRDF estimation, Goldman et al. [8] utilize the isotropic Ward model [25] to reduce the dimensionality of the inverse problem. Oxholm and Nishino [17, 18] further extend this idea by introducing a probabilistic formulation to estimate the BRDFs from a single image under natural lighting conditions. However, parametric models are inherently limited in their ability to provide precise approximations to the true BRDFs and further, lead to challenging and ill-conditioned optimization problems.

Isotropic BRDFs. Isotropic materials exhibit a form of symmetry, wherein the reflectance of the material is unchanged when the incident and outgoing directions are jointly rotated about the surface normal. This enables the

representation of isotropic BRDFs as the function over three as opposed to four angles. In the context of photometric stereo, Alldrin and Kriegman [1] observe that, for isotropic materials, the surface normal at each point can be restricted to lie on a plane. When the isotropic BRDFs has a single dominant lobe, Shi et al. [24] resolve the planar ambiguity and show that the surface normals can be uniquely determined. Higo et al. [12] utilize properties of isotropy, visibility and monotonicity to restrict the solution space of the surface normal at each pixel. This enables a framework for shape estimation without the need for radiometric calibration. Finally, a bivariate approximation for isotropic materials is used in Romeiro et al. [21,22] to estimate the BRDF of a known shape from a single image and without knowledge of the scene illumination.

Reference basis model. As mentioned in the introduction, a common assumption for scenes with SV-BRDF is that the per-pixel BRDF is generated from a few unknown reference BRDFs [2, 5, 8, 14]. Invariably, this leads to a multi-linear optimization in high-dimensional variables (the reference BRDFs) that is highly dependent on initial conditions. In contrast, our proposed technique avoids the need to estimate high-dimensional optimization by evoking knowledge of a dictionary of BRDFs.

3. Problem setup

Setup. We make the following assumptions. First, the camera is orthographic and hence, the viewing direction $\mathbf{v} \in \mathbb{R}^3$ is constant across all scene points. Second, the scene illumination is assumed to be from a distant point light source. The light sources are assumed to be of constant brightness (equivalently, that calibration is known) and their direction is known. We denote $\mathbf{l}_k \in \mathbb{R}^3$ to refer to the lighting direction in the k -th image I^k . For a light-stage, this information is typically obtained by a one-off calibration. Third, the effects of long-range illumination such as cast shadows and inter-reflections are assumed to be negligible; this is satisfied for objects with a convex shape. Finally, the radiometric response of the camera is assumed to be known.

BRDF representation. We follow the isotropic BRDF representation used in [23] in which a three-angle coordinate system based on half angles is used. Specifically, the BRDF is expressed as a function $\rho(\theta_h, \theta_d, \phi_d)$ with $\theta_h, \theta_d \in [0, \pi/2)$ and $\phi_d \in [0, 2\pi)$. However, by Helmholtz’s reciprocity, the BRDF exhibits the following symmetry: $\rho(\theta_h, \theta_d, \phi_d) = \rho(\theta_h, \theta_d, \phi_d + \pi)$, and hence it is sufficient to express $\phi_d \in [0, \pi)$. Following [26], we use a 1° sampling of each angle. As a consequence, a BRDF is represented as a point in a $T = 90 \times 90 \times 180 = 1,458,000$ -dimensional space. When we deal with color images, we have a BRDF for each color channel and hence, the dimensionality of the BRDF goes up proportionally.

Consider a scene element with BRDF $\rho \in \mathbb{R}^T$, surface normal \mathbf{n} , illuminated from a point light source from a direction \mathbf{l} and viewed from a direction \mathbf{v} . For this configuration of normal, incident light and viewing direction, the BRDF value is simply a linear functional of the vector ρ :

$$\mathbf{s}_{\{\mathbf{l}, \mathbf{v}; \mathbf{n}\}}^\top \rho,$$

where $\mathbf{s}_{\{\mathbf{l}, \mathbf{v}; \mathbf{n}\}}$ is a vector that encodes the geometry of the configuration. In essence, the vector samples the appropriate entry from ρ , allowing for the appropriate interpolation if the required value is off the sampling-grid.

Problem formulation. Our goal is to recover the surface normals and the SV-BRDF in the context of photometric stereo; i.e, multiple images of an object $\{I^1, \dots, I^Q\}$ obtained from a static camera under varying lighting. The intensity value $I_{\mathbf{p}}^i$ observed at pixel $\mathbf{p} = (x, y)$ with lighting \mathbf{l}_i can be written as

$$I_{\mathbf{p}}^i = (\mathbf{s}_{\{\mathbf{l}_i, \mathbf{v}; \mathbf{n}_{\mathbf{p}}\}}^\top \rho_{\mathbf{p}}) \cdot \max\{0, \mathbf{n}_{\mathbf{p}}^\top \mathbf{l}_i\}, \quad (1)$$

where $\rho_{\mathbf{p}}$ is the BRDF and $\mathbf{n}_{\mathbf{p}}$ is the surface normal at pixel \mathbf{p} , respectively, and $\max\{0, \mathbf{n}_{\mathbf{p}}^\top \mathbf{l}_i\}$ accounts for shading.

Given multiple intensity values at pixel \mathbf{p} , one for each lighting direction $\{\mathbf{l}_1, \dots, \mathbf{l}_Q\}$, we can write

$$\begin{aligned} \mathbf{I}_{\mathbf{p}} &= \begin{pmatrix} I_{\mathbf{p}}^1 \\ \vdots \\ I_{\mathbf{p}}^Q \end{pmatrix} = \begin{bmatrix} \max\{0, \mathbf{n}_{\mathbf{p}}^\top \mathbf{l}_1\} \cdot \mathbf{s}_{\{\mathbf{l}_1, \mathbf{v}; \mathbf{n}_{\mathbf{p}}\}}^\top \\ \vdots \\ \max\{0, \mathbf{n}_{\mathbf{p}}^\top \mathbf{l}_Q\} \cdot \mathbf{s}_{\{\mathbf{l}_Q, \mathbf{v}; \mathbf{n}_{\mathbf{p}}\}}^\top \end{bmatrix} \rho_{\mathbf{p}} \\ &= A(\mathbf{n}_{\mathbf{p}}) \rho_{\mathbf{p}} \end{aligned} \quad (2)$$

Given the intensities, $\mathbf{I}_{\mathbf{p}}$, observed at a pixel \mathbf{p} and knowledge of lighting directions $\{\mathbf{l}_1, \dots, \mathbf{l}_Q\}$, we seek to estimate the surface normal $\mathbf{n}_{\mathbf{p}}$ and the BRDF $\rho_{\mathbf{p}}$ at the pixel. This problem is intractable without additional assumptions that constrain the BRDF to a lower-dimensional space.

Model for BRDF. The key assumption that we make is that the BRDF at a pixel \mathbf{p} lies on the non-negative span of the atoms of a BRDF dictionary. Specifically, given dictionary $D = [\rho^1, \rho^2, \dots, \rho^M]$, we assume that the BRDF at pixel \mathbf{p} can be written as

$$\rho_{\mathbf{p}} = D \mathbf{c}_{\mathbf{p}}, \quad \mathbf{c}_{\mathbf{p}} \geq 0,$$

where $\mathbf{c}_{\mathbf{p}} \in \mathbb{R}^M$ are the abundances of the dictionary atoms.

In essence, we have now constrained the BRDF to lie in an M -dimensional cone.¹ This provides immense reduction in the dimensionality of the unknowns at the expense of introducing a model misfit error. Indeed the success of this model relies on having a dictionary that is sufficiently

¹A more appropriate model for the BRDF is that $(D\mathbf{c}) \geq 0$. However, this leads to significantly higher-dimensional constraints. We instead use a sufficient condition to achieve this, $\mathbf{c} \geq 0$.

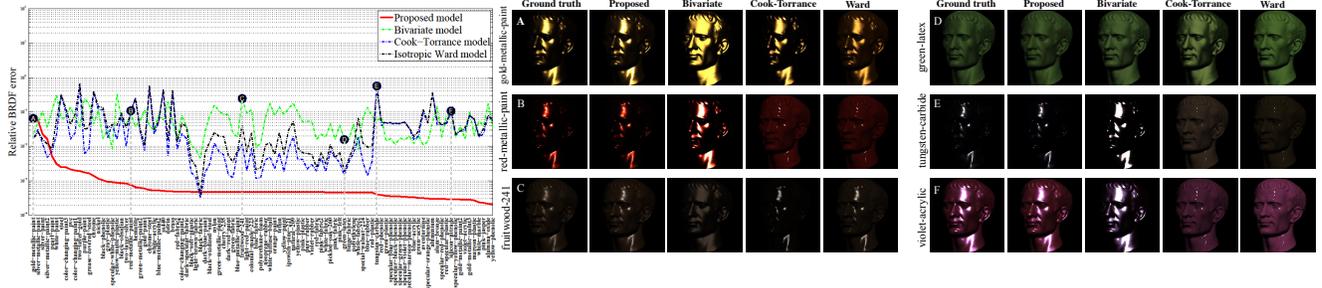


Figure 3. **Accuracy of BRDF models on the MERL database [26].** For the 100 materials in the database, we plot the approximation accuracy in relative RMS error [15] (also see (6)) for the proposed, bivariate [21], Cook-Torrance [6], and the Ward [25] models. For the proposed model, we use a leave-one-out scheme, wherein for each BRDF the remaining 99 BRDFs in the database are used to form the dictionary. The proposed model outperforms competing models both quantitatively (left) as well as in visual perception (right).

rich to cover a wide range of interesting materials. Figure 3 shows the accuracy of various BRDF models on the MERL BRDF database [26].

We also assume that \mathbf{c}_p is sparse, suggesting that BRDF at each pixel is the linear combination of a *few* dictionary atoms. The sparsity constraint avoids over-fitting to the intensity measurements \mathbf{I}_p as well as provides a regularization for under-determined problems.

Solution outline. An estimate of the surface normal and BRDF at a pixel \mathbf{p} can be obtained by solving

$$\{\hat{\mathbf{n}}_p, \hat{\mathbf{c}}_p\} = \arg \min_{\mathbf{n}, \mathbf{c}} \|\mathbf{I}_p - A(\mathbf{n})D\mathbf{c}\|_2^2 + \lambda \|\mathbf{c}\|_1 \quad (3)$$

$$\text{s.t. } \mathbf{c} \geq 0, \|\mathbf{n}\|_2 = 1$$

The ℓ_1 -penalty serves to enforce sparse solutions, with $\lambda \geq 0$ determining the level of sparsity in the solution. The optimization problem in (3) is non-convex due to unit-norm constraint on the surface normal \mathbf{n} as well as the term $A(\mathbf{n})D\mathbf{c}$. Our solution methodology consists of two steps: (i) *Surface normal estimation*. We perform an efficient multi-scale search that provides us with a precise estimate of the surface normal at pixel \mathbf{p} (see Section 4); and, (ii) *BRDF estimation*. We solve (3) only over \mathbf{c} with the normal fixed to obtain the BRDF at \mathbf{p} (see Section 5).

4. Surface normal estimation

In this section, we describe an efficient per-pixel surface normal estimation algorithm.

4.1. Virtual example-based normal estimation

Our surface normal estimation is an extension of the method proposed in [11], where two spheres — one diffuse and one specular — are introduced in a scene along with the target object. To obtain the surface normals at a pixel \mathbf{p} on the target, the intensity observed at pixel \mathbf{p} , \mathbf{I}_p , is matched to those on the reference spheres. The reference

spheres provide a sampling of the space of the normals and hence, we can simply treat them as a collection of candidate normals \mathcal{N} . By orientation consistency, the surface normal estimation now reduces to finding the candidate normal that can best explain the intensity profile \mathbf{I}_p . Given a candidate normal $\tilde{\mathbf{n}}$, we have two intensity profiles, $\mathbf{I}_D(\tilde{\mathbf{n}})$ and $\mathbf{I}_S(\tilde{\mathbf{n}})$, one each for the diffuse and specular sphere, respectively. The estimate of the surface normal at pixel \mathbf{p} is given as

$$\hat{\mathbf{n}}_p = \arg \min_{\tilde{\mathbf{n}} \in \mathcal{N}} \min_{a_1, a_2 \geq 0} \|\mathbf{I}_p - a_1 \mathbf{I}_D(\tilde{\mathbf{n}}) - a_2 \mathbf{I}_S(\tilde{\mathbf{n}})\|.$$

In [11], this is solved by scanning over all the pixels/candidate normals on the reference spheres.

Rendering virtual spheres. We rely on the same approach as [11] with the key difference that we virtually render the reference spheres. The virtual spheres are rendered as follows. Given the lighting directions $\{\mathbf{l}_1, \dots, \mathbf{l}_Q\}$ and the BRDF dictionary $D = [\rho^1, \dots, \rho^M]$, for each candidate normal $\tilde{\mathbf{n}} \in \mathcal{N}$, we render a matrix $B(\tilde{\mathbf{n}}) = [b_{ij}(\tilde{\mathbf{n}})] \in \mathbb{R}^{Q \times M}$ such that $b_{ij}(\tilde{\mathbf{n}})$ is the intensity observed at a surface with normal $\tilde{\mathbf{n}}$ and BRDF ρ^j , under lighting \mathbf{l}_i .

$$b_{ij}(\tilde{\mathbf{n}}) = \max\{0, \tilde{\mathbf{n}}^\top \mathbf{l}_i\} \cdot \mathbf{s}_{\{\mathbf{l}_i, \mathbf{v}; \tilde{\mathbf{n}}\}}^\top \rho^j,$$

We render one such matrix $B(\cdot)$ for each candidate normal in \mathcal{N} . Given these virtually rendered spheres, we can solve (3) by searching over all candidate normals.

Brute-force search. For computational efficiency, we drop the sparsity constraint in (3). We empirically observed that dropping the sparsity constraint made little difference in the estimated surface normals. Now, given the intensity profile \mathbf{I}_p at pixel \mathbf{p} and noting that $B(\tilde{\mathbf{n}}) = A(\tilde{\mathbf{n}})D$, solving (3) reduces to:

$$\hat{\mathbf{n}}_p = \arg \min_{\tilde{\mathbf{n}} \in \mathcal{N}} \min_{\mathbf{c} \geq 0} \|\mathbf{I}_p - B(\tilde{\mathbf{n}})\mathbf{c}\|. \quad (4)$$

The unit-norm constraint on the surface normals is absorbed into the candidate normals being unit-norm. The optimization problem in (4) requires solving a set of a non-negative

least squares (NNLS) sub-problems, one for each element of \mathcal{N} . For the results in the paper, we used the `lsnonneg` function in MATLAB to solve the NNLS sub-problems.

The accuracy and the computational cost in solving (4) depends solely on the cardinality of the candidate set \mathcal{N} , $|\mathcal{N}|$. We obtain \mathcal{N} by uniform or equi-angular sampling on the sphere [10]. As a consequence, the accuracy of the normal estimates, on an average, cannot be better than the half the angular spacing of the candidate set. Yet, the smaller the angular spacing of \mathcal{N} , the larger is its cardinality. For example, a 5° equi-angular sampling over the hemisphere requires approximately 250 candidates while a 0.5° requires 20000 candidates. Given that the time-complexity of the brute-force search is linear in $|\mathcal{N}|$, the computational costs for obtaining very precise normal estimates can be overwhelming. To alleviate this, we outline a coarse-to-fine search strategy that is remarkably faster than the brute-force approach with little loss in accuracy.

4.2. Coarse-to-fine search

Figure 4 shows the value of

$$E(\tilde{\mathbf{n}}) = \min_{\mathbf{c} \geq 0} \|\mathbf{I}_{\mathbf{p}} - B(\tilde{\mathbf{n}})\mathbf{c}\|$$

as a function of the candidate normal $\tilde{\mathbf{n}}$ for a few examples. In our simulations, we observed that there is a gradual increase in error value as we moved away from the global minima of $E(\cdot)$. We exploit this to design a coarse-to-fine search strategy where we first evaluate the candidate normals at a coarse sampling and subsequently search in the vicinity of this solution but at a finer sampling.

Specifically, let \mathcal{N}_θ be the set of equi-angular sampling on the unit-sphere where the angular spacing is θ degrees. Given a candidate normal $\tilde{\mathbf{n}}$, we define

$$C_\theta(\tilde{\mathbf{n}}) = \{\mathbf{n} \mid \langle \mathbf{n}, \tilde{\mathbf{n}} \rangle \geq \cos \theta, \|\mathbf{n}\|_2 = 1\}$$

as the set of unit-norm vectors within θ -degrees from $\tilde{\mathbf{n}}$,

In the first iteration, we initialize the candidate normal set $\mathcal{N}^{(1)} = \mathcal{N}_{\theta_1}$. Now, at the j -th iteration, we solve (4) over a candidate set $\mathcal{N}^{(j)}$. Suppose that $\hat{\mathbf{n}}^{(j)}$ is the candidate normal where the minimum occurs at the j -th iteration. The candidate set for the $(j+1)$ -th iteration is constructed as

$$j \geq 1, \quad \mathcal{N}^{(j+1)} = C_{\theta_j}(\hat{\mathbf{n}}^{(j)}) \cap \mathcal{N}_{\theta_{j+1}},$$

with $\theta_{j+1} < \theta_j$. That is, the candidate set is simply the set of all candidates at a finer angular sampling that are no greater than the current angular sampling from the current estimate. This is repeated till we reach the finest resolution at which we have candidate normals. For the results in this paper, we use the following values: $\theta_1 = 10^\circ$, $\theta_2 = 5^\circ$, $\theta_3 = 3^\circ$, $\theta_4 = 1^\circ$, and $\theta_5 = 0.5^\circ$. For efficient implementation, we pre-render $B(\tilde{\mathbf{n}})$ for $\tilde{\mathbf{n}} \in \mathcal{N}_{\theta_1} \cup \dots \cup \mathcal{N}_{\theta_5}$.

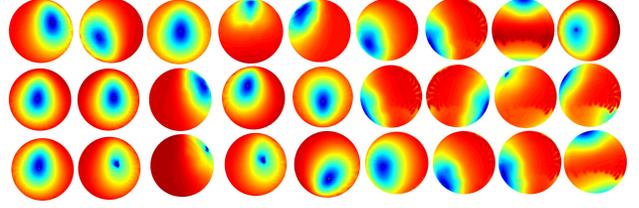


Figure 4. **The error as a function of candidate normals for a few test examples.** We can observe that the global minima is compact and the error increases largely monotonically in its vicinity. This motivates our coarse-to-fine search strategy.

		θ_1	θ_2	θ_3	θ_4	θ_5
Brute force	time	0.18s	0.77s	4.25s	27.3s	74.1s
	ang. error	7.07°	3.99°	1.56°	0.60°	0.42°
Coarse to fine	time	0.18s	0.19s	0.23s	0.34s	0.41s
	ang. error	7.07°	4.99°	2.56°	1.23°	0.82°

Table 1. Comparison of brute-force and coarse-to-fine normal estimation for different angular samplings in the candidate normals. Shown are aggregate statistics over 100 randomly generated trials.

The computational gains obtained via this coarse-to-fine sampling strategy are immense. Table 1 shows the run-time and precision of both brute force and coarse-to-fine normal estimation strategy for different levels of angular sampling in the generation of the candidate normal set. As expected the run-time of the brute force algorithm is linear in the number of candidates. In contrast, the coarse-to-fine strategy requires a tiny fraction of this time while nearly achieving the same precision as the brute force strategy.

While the solution to (4) also produces an estimate of the BRDF at the pixel, this estimate is often poor due to lack of the sparse-regularizer that serves to avoid over-fitting to the observed intensities. In the next section, we use this normal estimate to obtain a per-pixel BRDF estimate.

5. Reflectance estimation

Given the surface normal estimate $\hat{\mathbf{n}}_{\mathbf{p}}$, we obtain an estimate of the BRDF at each pixel, individually, by solving

$$\hat{\mathbf{c}}_{\mathbf{p}} = \arg \min_{\mathbf{c} \geq 0} \|\mathbf{I}_{\mathbf{p}} - B(\hat{\mathbf{n}}_{\mathbf{p}})\mathbf{c}\|_2^2 + \lambda \|\mathbf{c}\|_1. \quad (5)$$

The use of the ℓ_1 -regularizer promotes sparse solutions and primarily helps in avoiding over-fitting to the observed intensities. The optimization problem in (5) is convex and we used CVX [9], a general purpose convex solver, to obtain solutions. The estimate of the BRDF at pixel \mathbf{p} is given as $\hat{\rho}_{\mathbf{p}} = D\hat{\mathbf{c}}_{\mathbf{p}}$. The value of λ was manually tuned for best performance. For color-imagery, we solve for the coefficients associated with each color channel separately.

When we know a priori that multiple pixels share the same BRDF, then we can solve (5) simply by concatenating their corresponding intensity profiles and their respective $B(\cdot)$ matrices. As is to be expected, pooling intensities observed at multiple pixels significantly improves the quality of the estimates. Yet, while spatial averaging or spatial priors improve the quality of the estimate, inherently they require the object to exhibit smooth spatial-variations in its BRDF. The advantage of our per-pixel BRDF estimation framework is the ability to handle arbitrarily complex spatial variations in the BRDF at cost of noisier estimates. In the next section, we carefully characterize the performance of our proposition using synthetic and real examples.

6. Results

We characterize the performance of our technique using both synthetic and real datasets.

6.1. Synthetic experiments

We use the BRDFs in the MERL database [26] in a leave-one-out scheme for testing the accuracy of our proposed algorithms for surface normal and BRDF estimation. Specifically, when we simulate a test object using a particular material, the dictionary is comprised of BRDFs of the remaining $M = 99$ materials from the database. We used the configuration in the light-stage described in [7] for our collection of lighting directions.

Varying number of images. Figure 5 characterizes the errors in surface normal and BRDF estimation for varying number of input images or equivalently, lighting directions. We report the average error computed by randomly generating 20,000 normals per material and varying across all 100 material BRDFs in the database. This experiment is similar in setup to the one reported in [24] which, to our knowledge, is one of the most accurate techniques for photometric stereo on isotropic BRDFs. In [24], for 200 images, the error in estimating the elevation angle *when the azimuth is known* is reported as 0.88° ; in contrast, the proposed technique has an error of 0.82° in estimating the surface normal without any prior knowledge of the azimuth.

Varying BRDF. In Figure 6, we evaluate performance of surface normal estimation for varying material BRDFs. We fixed the number of images at $Q = 253$. Shown are aggregate statistics computed over 50,000 randomly generated surface normals. The worst case error is less than 2° and the error tapers down to 0.5° which is the finest sampling that we used for generating candidate normals. This can presumably be reduced by either choosing a finer sampling grid or using gradient descent techniques.

Comparisons. Figure 7 showcases the performance of many photometric stereo techniques for different objects:

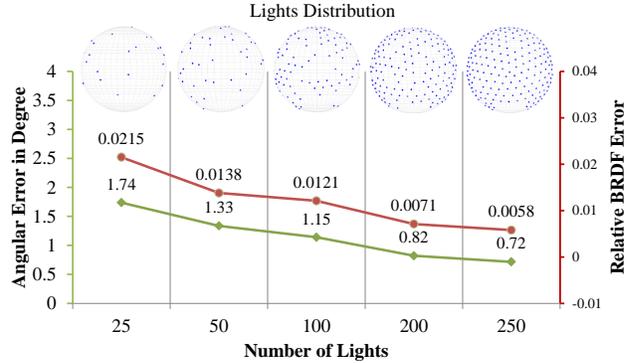


Figure 5. **Normal and BRDF estimation with varying number of images.** Given an input number of images, the angular errors (in green) and relative BRDF errors (in red) were obtained by averaging across all 100 BRDFs and across 20,000 randomly-generated normals per material.

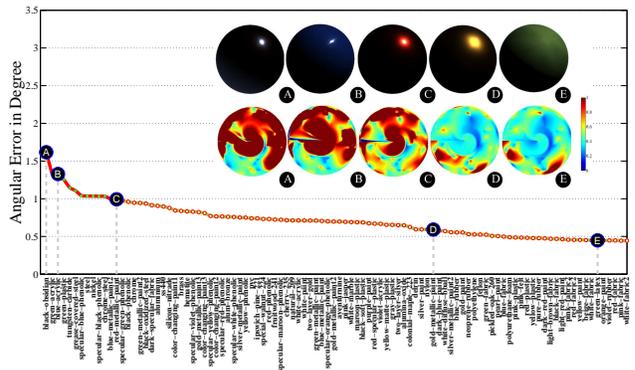


Figure 6. **Normal estimation for different materials.** We fix the number of input images/lighting directions to 253. For each material BRDF, we compute average error over 50,000 randomly-generated surface normals. Inset are the angular error distribution for a few select materials.

a black-obsidian bunny and a gold-painted elephant. We used 253 input images for each object. Photometric stereo under Lambertian model fails to recover precise normal maps indicating the presence of non-Lambertian components. The robust PCA-based solver [28] produces better normal maps as compared to the traditional photometric stereo, however it produces overly smoothed estimates; this can be attributed to removal of non-Lambertian cues which are invaluable for precise normal estimation. The solution of Alldrin et al. [2] while significantly better than Lambertian technique produces errors greater than 1° . In contrast, the proposed method returns reliable normal estimates for both scenes indicating the robustness of the underlying solution. We also simulated the performance of example-based photometric stereo which is identical to the proposed technique when applied to a two-

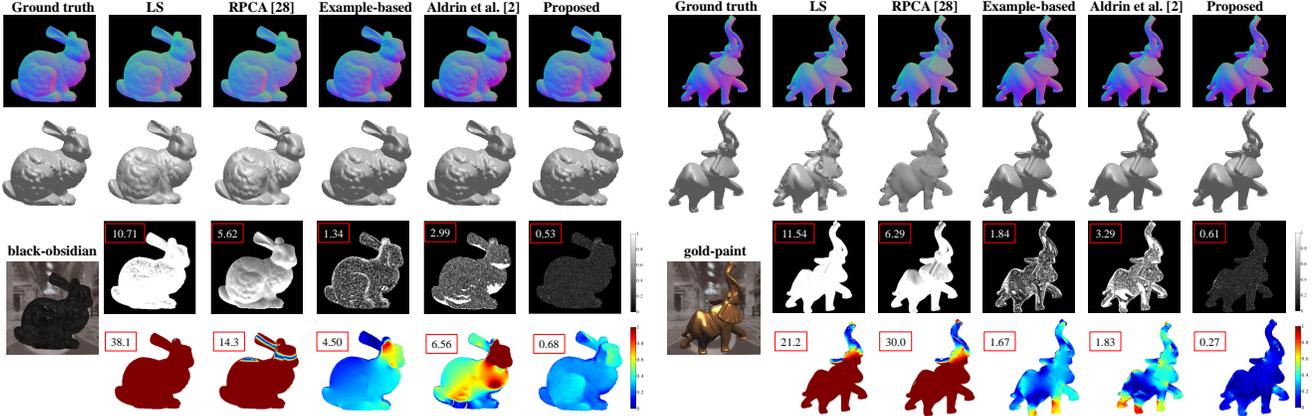


Figure 7. **Normal estimation across algorithms.** We compare the performance of photometric stereo under Lambertian model (LS), robust PCA-based approach [28] (RPCA), simulated example-based [11], Aldrin et al. [2] on two objects using 253 images each. Shown are (top-bottom) the estimated surface normals, recovered 3D surface, angular error in normal estimation in degrees and relative error in depth map based on different approaches. The insets in rows 3 and 4 are the average normal errors in degrees and the relative depth errors.

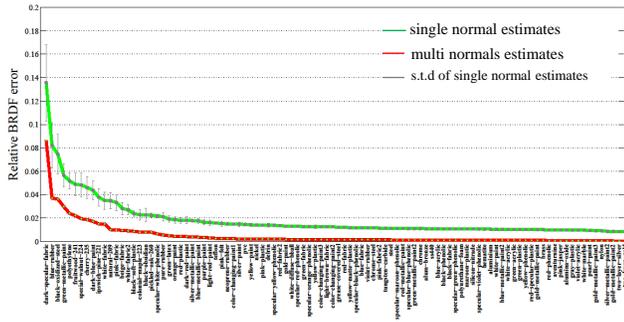


Figure 8. **Quantitative BRDF evaluation.** Given 253 lighting directions, we evaluate accuracy of BRDF estimation across different materials. For each material, we generated 100 normals with random orientations and estimated the BRDF for each instance individually (per-pixel) as well as collectively. For the per-pixel estimates, we plot average and standard deviation of the errors.

material (white-diffuse and chrome) dictionary. As expected, having a larger dictionary of BRDFs as in the proposed technique does provide significant improvements in surface normal estimation.

Performance of BRDF estimation. Given a test BRDF, we generated 100 surface normals with random orientations and rendered their appearance for 253 lighting directions. Assuming the knowledge of the true surface normals, we estimate the BRDF using the optimization in Section 5.

We characterize the performance of the per-pixel BRDF estimate as well as the error in the BRDF estimate when the information at the 100 normals are pooled. We use the relative BRDF error [15] to quantify the accuracy of the estimate. Given true BRDF ρ and estimated value $\hat{\rho}$, the relative

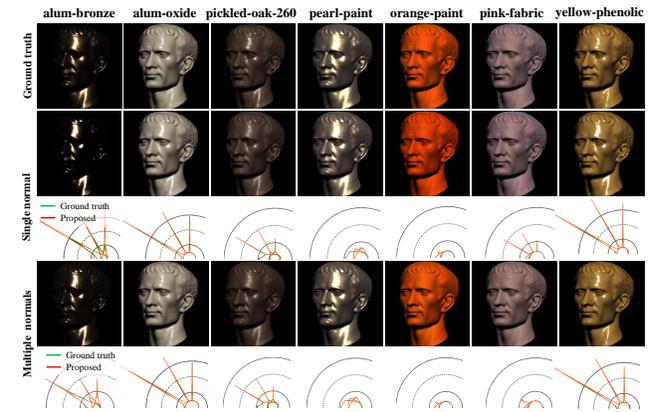


Figure 9. **Qualitative BRDF evaluation.** Shown are rendered BRDF for the *Caesar* statue for a few select material from MERL database [26]. (row 1) The rendered image based on ground truth BRDF; (rows 2 and 4) rendered images based on estimated BRDF from a single normal and 100 randomly generated normals, respectively; (rows 3 and 5) the polar plot for the reflectance function for the incident light angles $[0^\circ, 30^\circ, 60^\circ]$.

BRDF error is given as

$$\sqrt{\frac{\sum_i w_i ((\hat{\rho}(i) - \rho(i)) \cdot \max(0, \cos(\theta_i)))^2}{\sum_i w_i}}, \quad (6)$$

with w_i set equal to 1 for convenience.

Figure 8 shows the errors for different materials in the database — rank-ordered from worst-to-best performance — both for the per-pixel BRDF estimation as well as the joint estimation. In Figure 9, we show relighted images and polar plots for a subset of materials.

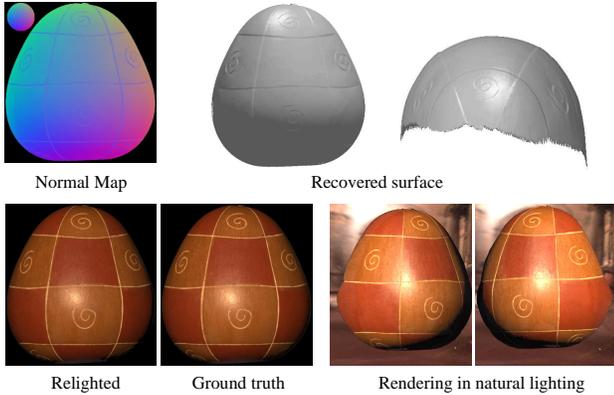


Figure 10. **Results on “gourd1” dataset.** We show the estimated normal map in false color (top-left) and 3D surface (top-right) recovered from it. We also show the relighting results (bottom-left), ground truth under the same lighting direction (bottom-middle), and relighting under natural environment (bottom-right).

6.2. Real data

Real images are present a layer of difficulty well beyond simulations and introduce inter-reflections, sub-surface scattering, cast shadows, and imprecise light source localization. We test the performance of our shape and BRDF recovery algorithm on a wide range of datasets. Specifically, we use images from two sources — the light stage data from [7], and the gourd from [2].

Figures 1, 10, 11 and 12 showcase the performance of our algorithm on the real datasets. The results in Figure 1, 11 and 12 were obtained from 250 input images, and the results of “gourd1” in Figure 10 was obtained from 100 input images. The recovered shape and BRDF (as visualized via rendered images) seem to be in agreement with the results in [2]; however, our algorithm is significantly simpler and employs a per-pixel algorithm that be easily parallelized.

The robustness of the per-pixel BRDF estimate is tested in Figure 11 where there are not just a wide variety of unique materials (the helmet, the breast-plate, the chain, the red scabbard, to name a few) but also significant modeling deviations (inter-reflections, cast-shadows). In spite of this, our approach produces a faithful rendition of the scene. The per-pixel BRDF estimation allows us to handle objects with complex spatial variations. In contrast, methods that assume the presence of just a few reference BRDFs as in [2,8] would not scale easily to such scenes. We refer the reader to the supplemental videos highlighting the relighting results.

7. Discussions

We present a photometric stereo technique for per-pixel normal and BRDF estimation for objects that are visually complex. We demonstrate that the use of a BRDF dictionary

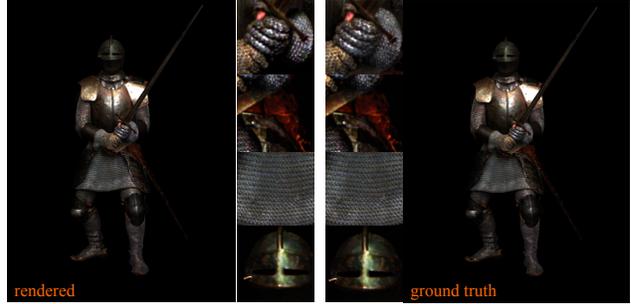


Figure 11. **Relighting results on “knight_fighting” dataset.**

significantly simplifies the inverse problem and provides not just state-of-the-art results in normal and BRDF estimation but also works robustly on a wide range of real scenes. A hallmark of our approach is the ability to obtain per-pixel BRDF estimation without any spatial smoothness assumptions endemic to state-of-the-art SV-BRDF estimation techniques [2, 8]; this makes it applicable to scenes with a large number of unique materials. Finally, our per-pixel framework is ripe for further speed-ups by solve for the shape and reflectance at each pixel in parallel.

Limitations. While the use of virtual examples provides flexibility beyond [11], we require light calibration and hence, our method is most suited to shape and reflectance acquisition from light-stages where the light sources are fixed and the calibration is a one-time effort. The accuracy of our coarse-to-fine normal estimation is lower bounded by finest sampling of our candidate normals. This can potentially be improved by refine the estimates using a gradient descent scheme starting with our solution; however, this approach could be computationally intensive. The SV-BRDF produced by our approach can be noisy especially since we independently recover the BRDF at each pixel. If we have a priori knowledge that the scene has a limited number of unique materials, then enforcing this could lead to robust SV-BRDF estimates. This can be easily incorporated into our framework by enforcing the matrix of sparse coefficients $[c_1 \dots c_p \dots]$ to be low-rank. Finally, it is also important that the scene lies in the linear span of our dictionary. In the failure of this, our results can be unpredictable. Here, the need for a larger dictionary encompassing hundreds, if not thousands, of materials would be invaluable for the broader applicability of our method.

References

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



Figure 12. Recovered surfaces on several real scenes with complex, spatially varying reflectance.

- [2] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *CVPR*, 2008.
- [3] R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *IEEE Trans. Pattern Analysis and Machine Intelligence (PAMI)*, 25:218–233, 2003.
- [4] J. Blinn and M. Newell. Texture and reflection in computer generated images. *Comm. ACM*, 19:542–547, 1976.
- [5] M. Chandraker and R. Ramamoorthi. What an image reveals about material reflectance. In *ICCV*, 2011.
- [6] R. Cook and K. Torrance. A reflectance model for computer graphics. *ACM Trans. Graphics (TOG)*, 1:7–24, 1982.
- [7] P. Einarsson, C. Chabert, A. Jones, W. Ma, B. Lamond, T. Hawkins, M. Bolas, S. Sylwan, and P. Debevec. Relighting human locomotion with flowed reflectance fields. In *Rendering techniques*, 2006.
- [8] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz. Shape and spatially-varying BRDFs from photometric stereo. In *ICCV*, 2005.
- [9] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, 2014.
- [10] R. Harman and V. Lako. On decompositional algorithms for uniform sampling from n-spheres and n-balls. *Journal of Multivariate Analysis*, 101:2297–2304, 2010.
- [11] A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *PAMI*, 27:1254–1264, 2005.
- [12] T. Higo, Y. Matsushita, and K. Ikeuchi. Consensus photometric stereo. In *CVPR*, 2010.
- [13] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *CVPR*, 2012.
- [14] J. Lawrence, A. Ben-Artzi, C. DeCoro, W. Matusik, H. Pfister, R. Ramamoorthi, and S. Rusinkiewicz. Inverse shade trees for non-parametric material representation and editing. *TOG*, 25:735–745, 2006.
- [15] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of brdf models. In *Euro. Conf. Rendering Tech.*, 2005.
- [16] M. Oren and S. Nayar. Generalization of the lambertian model and implications for machine vision. *Intl. J. Computer Vision*, 14:227–251, 1995.
- [17] G. Oxholm and K. Nishino. Shape and reflectance from natural illumination. In *ECCV*, 2012.
- [18] G. Oxholm and K. Nishino. Multiview shape and reflectance from natural illumination. In *CVPR*, 2014.
- [19] R. Ramamoorthi. Analytic PCA construction for theoretical analysis of lighting variability in images of a Lambertian object. *PAMI*, 24:1322–1333, 2002.
- [20] P. Ren, J. Wang, J. Snyder, X. Tong, and B. Guo. Pocket reflectometry. *TOG*, 30:45, 2011.
- [21] F. Romeiro, Y. Vasilyev, and T. Zickler. Passive reflectometry. In *ECCV*, 2008.
- [22] F. Romeiro and T. Zickler. Blind reflectometry. In *ECCV*, 2010.
- [23] S. Rusinkiewicz. A new change of variables for efficient brdf representation. In *Rendering techniques*, pages 11–22. 1998.
- [24] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. In *ECCV*. 2012.
- [25] G. Ward. Measuring and modeling anisotropic reflection. *TOG*, 26:265–272, 1992.
- [26] M. . Wojciech, P. Hanspeter, B. Matt, and M. Leonard. A data-driven reflectance model. *TOG*, 22:759–769, 2003.
- [27] R. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Eng.*, 1980.
- [28] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *ACCV*. 2011.
- [29] C. Yu, Y. Seo, and S. Lee. Photometric stereo from maximum feasible Lambertian reflections. In *ECCV*. 2010.
- [30] L. Yu, S. Yeung, Y. Tai, D. Terzopoulos, and T. Chan. Outdoor photometric stereo. In *ICCP*, 2013.