A Theory of Differential Photometric Stereo for Unknown Isotropic BRDFs

Manmohan Chandraker

Jiamin Bai

Ravi Ramamoorthi

University of California, Berkeley

Abstract

This paper presents a comprehensive theory of photometric surface reconstruction from image derivatives. For unknown isotropic BRDFs, we show that two measurements of spatial and temporal image derivatives, under unknown light sources on a circle, suffice to determine the surface.

This result is the culmination of a series of fundamental observations. First, we discover a photometric invariant that relates image derivatives to the surface geometry, regardless of the form of isotropic BRDF. Next, we show that just two pairs of differential images from unknown light directions suffice to recover surface information from the photometric invariant. This is shown to be equivalent to determining isocontours of constant magnitude of the surface gradient, as well as isocontours of constant depth. Further, we prove that specification of the surface normal at a single point completely determines the surface depth from these isocontours.

In addition, we propose practical algorithms that require additional initial or boundary information, but recover depth from lower order derivatives. Our theoretical results are illustrated with several examples on synthetic and real data.

1. Introduction

Image formation is a complex interplay between the geometry of a scene, its reflectance properties and the illumination conditions under which it is observed. Photometric stereo aims to recover the surface structure, using shading cues from varying illumination. Under assumptions of diffuse reflectance or illumination constancy, there exist wideranging theories of photometric stereo, shape-from-shading and optical flow to recover shape from image intensities and gradients. However, algorithms for general BRDFs and unknown light source motions remain largely unexplored.

This paper presents an intensive theoretical study of the utility of image gradients for recovering scene structure for complex BRDFs and unknown light directions. We assume a homogeneous, isotropic BRDF, while a differential pair of lights undergo unknown motion on a circle. The various aspects of our theory are illustrated in Figure 1.

We begin with the observation that a single image formation equation yields three differential equations upon spatial and temporal differentiation. Moreover, these differential equations are linear in BRDF derivatives, so under appropriate conditions (such as circular motion of the light), they may be eliminated to yield a photometric invariant that relates image derivatives, surface geometry and source directions.

A surprising discovery is that for isotropic BRDFs, it is possible to uncover such an invariant that is independent of lighting directions. Moreover, we show that the $N \times 3$ matrix whose columns are the image derivatives at a pixel under N light source positions, must be rank 2 and its null vector, $(\lambda, \kappa, 1)^{\top}$, is determined by surface geometry alone.

For a surface $z(x, y) : \mathbb{R}^2 \to \mathbb{R}$, the entities λ and κ are space varying functions whose relationship to surface depths, z, or gradients ∇z , is not straightforward. Yet, we show that the information contained in those functions can be succinctly expressed as the direction of the surface gradient and the direction of the gradient of the scalar field $\|\nabla z\|$. Thus, from image information alone, one may determine surface structure up to level curves of constant depth and isocontours of constant magnitude of the gradient.

Further, we show that for general surfaces where these two sets of isocontours intersect transversally, surface normal information at a single point suffices to determine the depth. Thus, in theory, differential photometric stereo allows recovery of surface depths from unknown light positions on a circle, for unknown isotropic BRDFs.

While the theory suggests an elegant algorithm for depth reconstruction, recovering isocontours of constant depth requires higher-order derivatives, which may lead to noisy estimation. For practical applications, we suggest two additional algorithms that can recover surface normals and depths, resepectively, given additional information in the form of normals on a curve, or depths at the boundary.

Throughout the paper, we validate the theoretical results with several examples on synthetic and real data.

2. Related Work

This work differs from most prior studies in considering differential information within the context of photometric stereo, with unknown BRDF and light directions. Our theoretical results distinguish this work by establishing a minimal requirement of two differential pairs for 3D reconstruction. We refer the reader to the extended version [4] for details of proofs, implementations and experiments.



Figure 1. Outline of our differential theory of photometric stereo. Two or more differential pairs lead to an uncalibrated photometric invariant. Isocontours of constant $||\nabla z||$ and constant z may be recovered from the invariant. Additional information in the form of the gradient at a single point suffices to recover the surface. Normals on a curve or depths on the boundary can be used for a more noise robust estimation.

Most prior work in photometric stereo has dealt with Lambertian surfaces, for which the surface normal can be recovered from three images [14]. In some cases, specularity removal has been used as a pre-processing step [3].

Hertzmann and Seitz use example-based methods for general BRDFs, but require a reference object of the same material [6]. Helmholtz stereopsis [16] eliminates the BRDF by exploiting reciprocity, but needs active control of both light and camera positions. Tan et al. use color-based separation of the diffuse and specular components of dichromatic reflectance for reconstruction [12]. Further empirical properties of the BRDF can be exploited to recover shape, but require dense coverage of the lighting hemisphere [2, 7].

Isocontours of constant depth are recovered by Alldrin and Kriegman [1] using bilateral symmetry of a spatially varying BRDF. They require a dense configuration of known lights on a circle and initial information on a curve to recover depth. Tan and Zickler [13] further recover isocontours of constant magnitude of the gradient, whereby initial information at a single point suffices. In contrast, our theory does not depend on pointwise BRDF symmetries, rather uses image derivatives. Instead of the dense, calibrated ring of lights of [1, 13], we need as few as two differential pairs, at unknown positions on a circle. While the results of [13] are derived only for the gnomonic projection, our theory holds for many projections (see Sec. 3.1). Thus, we can also recover isocontours of constant depth as well as constant magnitude of the gradient, but our derivations are more general and may lead to other interesting topological classes in 3D reconstruction.

A related work that uses differential information to derive photometric invariants is Clark's active photometric stereo [5]. However, it recovers depth using a distance-dependent imaging model and calibrated source positions.

Besides photometric stereo, image derivatives have been considered in other contexts. Shape from shading seeks to recover depth from a single image of a Lambertian surface [10, 15]. In computer graphics, Ramamoorthi et al. have studied the first-order behavior of reflection as a convolution and visibility under area lighting [11]. Optical flow uses spatial and temporal derivatives for recovering the motion field [9]. Indeed, the form of our photometric invariant bears a striking resemblance to the optical flow constraint. However, our theory neither relies on assumptions like brightness constancy, nor does it suffer from the aperture problem.

3. Image Formation for Isotropic BRDFs

Throughout this paper, the object and the camera are fixed, while the directional point light source moves around the object. The camera principal axis is oriented along $\hat{\mathbf{v}} = (0, 0, 1)^{\top}$, pointing towards the origin. We assume orthographic projection and the object is represented by a surface z(x, y), where $\mathbf{x} = (x, y)^{\top}$ represents a point on the image plane. A unit 3-vector on $S^2 \subset \mathbb{R}^3$ is represented as $\hat{\mathbf{w}}$, whereas a 2-vector on \mathbb{R}^2 , is represented as \mathbf{w} . For a vector $\mathbf{w} \in \mathbb{R}^2$, we denote $l(\mathbf{w}) = \sqrt{||\mathbf{w}||^2 + 1}$.

3.1. A Note on Surface Normal Parameterizations

The relationship between the unit normal $\hat{\mathbf{n}}$ on the 2sphere and its representation \mathbf{n} is determined by the projection mapping $\pi : S^2 \to \mathbb{R}^2$ [8]. For a gnomonic projection, from the center of the sphere to the tangent plane resting on the north pole, $\hat{\mathbf{n}} = \frac{(\mathbf{n}^\top, 1)^\top}{l(\mathbf{n})}$, thus, $\mathbf{n} = [-z_x, -z_y]^\top$. A stereographic projection is centered on the south pole. Similarly, we represent a directional point light source $\hat{\mathbf{s}}$ by a 2-vector, s. The theory of the paper is independent of this choice and valid for both projections (indeed, for any centered on the line joining the north and south poles).

3.2. Parameterizing Isotropic BRDFs

We develop our theory for homogeneous isotropic BRDFs. Such reflectance functions depend only on the three angles between the unit normal $\hat{\mathbf{n}}$, the light source direction $\hat{\mathbf{s}}$ and the viewing direction $\hat{\mathbf{v}}$, thus, they can be represented as a function of the form $\bar{\rho}(\hat{\mathbf{n}}^{\top}\hat{\mathbf{s}}, \hat{\mathbf{n}}^{\top}\hat{\mathbf{v}}, \hat{\mathbf{s}}^{\top}\hat{\mathbf{v}})$. For gnomonic projection, the angles that determine the isotropic BRDF are

$$\widehat{\mathbf{n}}^{\top}\widehat{\mathbf{s}} = \frac{\mathbf{n}^{\top}\mathbf{s} + 1}{l(\mathbf{n})l(\mathbf{s})}, \quad \widehat{\mathbf{n}}^{\top}\widehat{\mathbf{v}} = \frac{1}{l(\mathbf{n})}, \quad \widehat{\mathbf{s}}^{\top}\widehat{\mathbf{v}} = \frac{1}{l(\mathbf{s})}.$$
 (1)

Thus, the BRDF can be written as a function of the form $\hat{\rho}(\mathbf{n}^{\top}\mathbf{s}, l(\mathbf{n}), l(\mathbf{s}))$, appropriately defined from $\bar{\rho}$:

$$\bar{\rho}(\widehat{\mathbf{n}}^{\top}\widehat{\mathbf{s}}, \widehat{\mathbf{n}}^{\top}\widehat{\mathbf{v}}, \widehat{\mathbf{s}}^{\top}\widehat{\mathbf{v}}) = \widehat{\rho}(\mathbf{n}^{\top}\mathbf{s}, l(\mathbf{n}), l(\mathbf{s})) = \widehat{\rho}(\alpha, \beta, \gamma),$$
(2)

where we denote $\alpha = \mathbf{n}^{\top} \mathbf{s}$, $\beta = l(\mathbf{n})$, $\gamma = l(\mathbf{s})$ for the parameters of an isotropic BRDF. It may be verified that the same is true for a stereographic projection. Note that for our fixed pose and variable illumination setup, \mathbf{n} depends on space variables \mathbf{x} and \mathbf{s} on time variables t.

3.3. Ratio images

The image formation model can be easily generalized to:

$$E(\mathbf{x},t) = a(\mathbf{x}) \,\widehat{\rho} \left(\alpha(\mathbf{x},t), \beta(\mathbf{x}), \gamma(t) \right), \tag{3}$$

where $a(\mathbf{x})$ is a surface albedo. This is a reasonable model and in particular, subsumes traditional Lambertian assumptions, while allowing for more general isotropic BRDFs, possibly modulated by a spatially-varying texture. We make no assumptions on the form of the isotropic BRDF $\hat{\rho}$.

In practice, one may eliminate the pointwise albedo $a(\mathbf{x})$ by considering ratio images. Note that ratios of arbitrarily illuminated images need not preserve the functional dependence of the isotropic BRDF on α , β and γ . But ratios with respect to the image obtained from a light source colocated with the sensor, that is, $\mathbf{s} = (0, 0)^{\top}$, do preserve the desired form of the BRDF. It can be easily seen from (2) that such ratio images, denoted by I, have the form

$$I(\mathbf{x},t) = \frac{a(\mathbf{x})\,\widehat{\rho}(\alpha,\beta,\gamma)}{a(\mathbf{x})\,\widehat{\rho}(\beta)} = \rho(\alpha,\beta,\gamma), \qquad (4)$$

where $\rho(\cdot)$ is the appropriately defined function. The exact form of $\bar{\rho}(\cdot)$, $\hat{\rho}(\cdot)$ or $\rho(\cdot)$ is immaterial, since we will derive our photometric invariant by eliminating it.

Alternatively, one may also take ratios with respect to an image under uniform (floodlit or cloudy sky) illumination.

4. A Novel Differential Photometric Invariant

A common approach to photometric reconstruction derives invariants that relate image intensities to surface geometry and source directions. Most prior work focuses on restricted reflectances (such as Lambertian). However, one may deal with complex BRDFs if they can be eliminated from a system of equations. In this section, we use differential information to derive such an invariant. A surprising and useful result is that **our invariant does not depend on light source positions**, so it is in fact an uncalibrated invariant.

4.1. Differential Images

Intuitively, a single image formation equation leads to independent relations upon differentiation with respect to various variables. These may then be related by eliminating any terms that depend on the functional form of the BRDF. We use this intuition to derive a novel photometric invariant. The space and time derivatives of the images in (4) are

$$\nabla_{\mathbf{x}} I(\mathbf{x}, t) = \rho_{\alpha} \mathbf{J}^{\top}(\mathbf{n}) \mathbf{s} + \rho_{\beta} \frac{1}{l(\mathbf{n})} \mathbf{J}^{\top}(\mathbf{n}) \mathbf{n} \qquad (5)$$

$$I_t(\mathbf{x},t) = \rho_{\alpha} \mathbf{s}_t^{\top} \mathbf{n} + \rho_{\gamma} \frac{1}{l(\mathbf{s})} \mathbf{s}_t^{\top} \mathbf{s}$$
(6)

where α, β, γ are defined previously and $\mathbf{J}(\mathbf{n}) = [\mathbf{n}_x, \mathbf{n}_y]$ is the 2 × 2 Jacobian related to the second fundamental form, *II*. (For the gnomonic projection, $\mathbf{J}(\mathbf{n}) = l(\mathbf{n}) \cdot II$.)

This system of three equations (note that (5) represents two equations) is clearly underconstrained, with the unknowns ρ_{α} , ρ_{β} , ρ_{γ} , **n** and **s**. However, the BRDF derivatives can be eliminated to extract constraints on the surface normal in a setup of circular motion, as discussed next.

4.2. Circular Motion Yields an Invariant

Let us constrain the source to move in a circle around the camera axis. Intuitively, since $||\mathbf{s}||$ is now constant, γ stays constant. Thus, the BRDF reduces to a 2D one, which allows elimination of derivatives with respect to α and β . Mathematically, $\mathbf{s}_t^{\mathsf{T}} \mathbf{s} = 0$ for lights on a circle, so the three equations (5) and (6) can be reduced to

$$\frac{I_{x}\mathbf{s}_{t}^{\top}\mathbf{n} - I_{t}\mathbf{s}^{\top}\mathbf{n}_{x}}{I_{y}\mathbf{s}_{t}^{\top}\mathbf{n} - I_{t}\mathbf{s}^{\top}\mathbf{n}_{y}} = \frac{\mathbf{n}^{\top}\mathbf{n}_{x}}{\mathbf{n}^{\top}\mathbf{n}_{y}} = \text{constant across time},$$
(7)

since the right hand side depends only on position. Thus, we have a photometric invariant that relates image derivatives, surface geometry and light directions, regardless of BRDF.

4.3. Illumination Invariant Photometric Flow

It might seem that, with known sources, non-linear minimization can be used to estimate the unknowns $\{\mathbf{n}, \mathbf{n}_x, \mathbf{n}_y\}$, up to scale, using (7). However, as shown below, the source directions are not required at all. That is, the invariant (7) directly relates image derivatives to surface geometry.

Proposition 1. *The entire information in the photometric invariant* (7) *is encapsulated by two entities, which depend only on surface geometry and not on source positions* **s**.

Proof. With the following series of definitions

$$\lambda = \frac{\mathbf{n}^{\top} \mathbf{n}_x}{\mathbf{n}^{\top} \mathbf{n}_y}, \qquad \mathbf{u} = \mathbf{n}_x - \lambda \mathbf{n}_y, \qquad \kappa = \frac{\mathbf{s}^{\top} \mathbf{u}}{\mathbf{s}_t^{\top} \mathbf{n}}, \qquad (8)$$

we may rewrite (7) as

$$I_x - \lambda I_y - \kappa I_t = 0.$$
(9)

By definition, λ depends only on the surface normal and is constant across time (that is, independent of light source positions). Also, by definition of λ and **u** in (8), we have **u** \perp **n**. Since the lights lie on a circle, we have $\mathbf{s}_t \perp \mathbf{s}$. Moreover, at time t (or equivalently, angular position t on the circle of sources), the light source is $\mathbf{s} = (r \cos t, r \sin t)^{\top}$. Then,



Figure 2. The ratio of image derivatives at a pixel, recorded for various light positions, lie on a straight line given by the equation $\lambda \frac{I_y}{I_x} + \kappa \frac{I_t}{I_x} = 1.$

 $\mathbf{s}_t = (-r \sin t, r \cos t)^{\top}$, thus, $\|\mathbf{s}\| = \|\mathbf{s}_t\|$. It follows that κ is actually the (signed) ratio of the magnitudes of \mathbf{u} and \mathbf{n} .¹ Indeed, with $\mathbf{n}_{\perp} = (-n_2, n_1)^{\top}$, it immediately follows that

$$\mathbf{u} = -\kappa \mathbf{n}_{\perp}.\tag{10}$$

Thus, κ is also independent of illumination, so the constraint (9) is *independent of both the BRDF and the light sources*. Clearly, all the information in the invariant is encapsulated by λ and κ , which depend only on surface geometry.

Noting the similarity of (9) to optical flow [9], we call it the **photometric flow** relation. However, (9) is derived without resorting to any brightness constancy assumption and as we discuss next, recovery of λ and κ in a photometric stereo setup does not suffer from the aperture problem.

4.4. Importance of the Invariant

The result of Proposition 1 is a surprising one and can be understood as a fundamental relationship between spatial and temporal derivatives of images due to isotropic BRDFs:

Corollary 2. For a surface with isotropic BRDF, the $N \times 3$ matrix of spatial and temporal image derivatives at a pixel, recorded for N > 1 unknown light positions on a circle, must be rank 2. In addition, the null-vector, denoted $(\lambda, \kappa, 1)^{\top}$, depends only on the surface geometry.

This is also useful practically, since it establishes that surface information may be recovered from image derivatives with unknown isotropic BRDFs and unknown light sources. In fact, as few as two differential image pairs suffice to estimate λ and κ . This is an important observation:

Corollary 3. *Two pairs of differential images suffice to recover surface information from the photometric flow.*

Note that three light sources can create two differential pairs. It is instructive to recall that traditional photometric stereo for Lambertian surfaces requires three images to completely determine the surface normal.

4.5. Experimental Evaluation

Here, we empirically illustrate the validity of (9). A differential pair of lights is moved on a circle and real images of a plastic apple of varying albedo are acquired at 11 unknown light positions. Ratio images are computed with respect to a floodlit image. Figure 2 shows that the ratios $\frac{I_y}{I_x}$ and $\frac{I_t}{I_x}$ computed for various light positions lie on a straight line. The entities λ and κ are given by the best-fit straight line.

5. Reconstructibility From Photometric Flow

While the previous section proves the existence of a BRDF and illumination-independent photometric flow, its relation to surface reconstruction is not immediately apparent. We conclusively establish that relationship in this section.

5.1. Isocontours of Constant $\|\nabla z\|$

To begin, we show that one may recover isocontours where magnitude of the gradient, $\|\nabla z\|$, stays constant. This follows directly from the definition of λ in Proposition 1:

Corollary 4. From two or more differential images of a surface with unknown isotropic BRDF, obtained from unknown light source positions, one may recover isocontours of constant magnitude of the surface gradient.

Proof. Consider the scalar field $g(x, y) = ||\mathbf{n}(x, y)||$. The associated gradient is $\nabla g = (||\mathbf{n}||_x, ||\mathbf{n}||_y)^{\top}$. The level curves of a scalar field are orthogonal to the gradient. So, the direction of the tangent to the level curves of g is $\frac{||\mathbf{n}||_y}{||\mathbf{n}||_x}$, which is the same as λ^{-1} , by definition (8). Thus, knowing the value of λ , one may trace the isocontours of constant $||\mathbf{n}(x, y)||$. The statement of the theorem follows as a special case for the gnomonic projection, where $\mathbf{n} = -\nabla z$.

Note that the result we have proved is more general than the statement of Corollary 4. It may be easily verified that the isocontours of constant $||\mathbf{n}||$ are the same for any **n** derived from a projection centered on the line joining the poles.

5.1.1 Experimental Evaluation

Figure 3 illustrates the recovery of isocontours of constant $\|\nabla z\|$ using synthetic data. For the sphere of varying albedo and the bunny, we use a simplified Torrance-Sparrow model

$$\bar{\rho} = (2\sqrt{\pi}\sigma)^{-2} e^{-\left(\sigma^{-1}\cos^{-1}\widehat{\mathbf{n}}^{\top}\widehat{\mathbf{h}}\right)^{2}}, \ \widehat{\mathbf{h}} = \frac{\widehat{\mathbf{s}} + \widehat{\mathbf{v}}}{\|\widehat{\mathbf{s}} + \widehat{\mathbf{v}}\|}.$$
 (11)

with $\sigma = 0.3$. For the vase of varying albedo, we use a constant coefficient Phong-Blinn model: $\bar{\rho} = \hat{\mathbf{n}}^{\mathsf{T}} \hat{\mathbf{s}} + (\hat{\mathbf{n}}^{\mathsf{T}} \hat{\mathbf{h}})^{\sigma}$, with a typical value of $\sigma = 5$. In each case, we observe that the recovered isocontours closely match the ground truth.

5.2. Isocontours of Constant Depth

To recover isocontours of constant $||\mathbf{n}||$, we required knowledge only of λ . However, we can further disambiguate the surface using κ . The following constructively proves that λ and κ together determine the isocontours of constant depth (in addition to the isocontours of constant $||\mathbf{n}||$).

¹In practice, the temporal derivative is obtained as a difference between images at time t and $t + \delta t$, so asserting $\|\mathbf{s}\| = \|\mathbf{s}_t\|$ assumes that the angular difference δt between the lights of the rotating differential pair is known. This is the same as the ratio of their distance and the radius of the circle on which they are situated. Note that calibrated measurement is not required for obtaining this ratio (for example, a piece of string suffices).



Figure 3. Recovery of isocontours of constant $\|\nabla z\|$. Red curves plot the isocontours, while the green dots represent points chosen to start tracing the curves. (a) A sphere with variable albedo and simplified Torrance-Sparrow BRDF. (b) A vase with Blinn-Phong BRDF. (c) A bunny with Torrance-Sparrow BRDF.

Proposition 5. From two or more differential images of a surface with unknown isotropic BRDF, obtained from unknown source positions, one may recover the direction of the surface gradient at every point.

Proof. We denote $\mathbf{n} = (p,q)^{\top}$, thus, $\mathbf{n}_x = (p_x,q_x)^{\top}$ and $\mathbf{n}_y = (p_y, q_y)^{\top}$. Using (9), two or more pairs of differential images determine $\lambda(x, y)$ and $\kappa(x, y)$. Thereby, two linear, first order PDEs are obtained from (10) at every pixel:

$$p_x - \lambda(x, y)p_y = \kappa(x, y)q \tag{12}$$

$$q_x - \lambda(x, y)q_y = -\kappa(x, y)p.$$
(13)

Further, a surface must satisfy the integrability condition:

$$p_y - q_x = 0. \tag{14}$$

Thus, we have a coupled first order system of three linear PDEs in the two variables p and q. Note that this is an overdetermined system, which may not be solvable in general. However, in our particular case, we can exploit the special form of the PDEs to derive an unusual solution.

Consider the function $h = p - \lambda q$. Taking partial derivatives and using (12), (13) and (14), we get

$$h_x = (\kappa - \lambda_x)q, \quad h_y = -\kappa p - \lambda_y q.$$
 (15)

By integrability of h, we have $h_{xy} = h_{yx}$, which gives us another first order linear PDE:

$$(\lambda_x - \kappa)q_y - \lambda_y q_x - \kappa p_x = \kappa_x p + \kappa_y q.$$
(16)

From (12), (13), (14) and (16), we have a linear system in $\{p_x, p_y, q_x, q_y\}$, whereby expressions for them may be obtained as linear functions of $\{p, q\}$:

$$p_x = \nu_1 p + \eta_1 q, \quad q_x = p_y = \nu_2 p + \eta_2 q, \quad q_y = \nu_3 p + \eta_3 q, \quad (17)$$

where, with $\Delta = -(\kappa \lambda^2 + \lambda \lambda_y - \lambda_x + \kappa)$,

$$\begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix} = \frac{1}{\Delta} \begin{bmatrix} \lambda(\kappa^2 - \lambda_x \kappa + \lambda \kappa_x) \\ \kappa^2 - \lambda_x \kappa + \lambda \kappa_x \\ -(\lambda \kappa^2 + \lambda_y \kappa - \kappa_x) \end{bmatrix}, \quad (18)$$

$$\begin{bmatrix} \eta_1 \\ \eta_2 \\ \eta_3 \end{bmatrix} = \begin{bmatrix} \kappa \\ 0 \\ 0 \end{bmatrix} + \frac{\kappa_y + \kappa^2}{\Delta} \begin{bmatrix} \lambda^2 \\ \lambda \\ 1 \end{bmatrix}, \quad (19)$$



Figure 4. Recovery of isocontours of constant depth. Red curves plot the isocontours, while the green dots represent points chosen to start tracing the curves. (a) A vase with variable albedo and Blinn-Phong BRDF. (b) A bunny with Torrance-Sparrow BRDF. (c) Plot of the direction of the gradient corresponding to (b), with the length of the vector normalized to one for display.

The reader may verify by substitution that the expressions above are consistent with the definitions of λ and κ in (8).

Now, using integrability of p and q, that is $(p_x)_y = (p_y)_x$ and $(q_x)_y = (q_y)_x$, we get two new PDEs:

$$(\nu_1 p + \eta_1 q)_y = (\nu_2 p + \eta_2 q)_x \tag{20}$$

$$(\nu_2 p + \eta_2 q)_y = (\nu_3 p + \eta_3 q)_x, \tag{21}$$

These are linear, first order PDEs, where we can again replace the first order derivatives $\{p_x, p_y, q_x, q_y\}$ using (17) to get two (dependent) homogeneous linear equations in $\{p, q\}$. Using, say, the first equation, we get

$$\frac{q}{p} = \frac{\nu_{1y} + \eta_1 \nu_3 - \nu_{2x} - \nu_2 \eta_2}{\nu_2 \eta_1 + \eta_{2x} + \eta_2^2 - \nu_1 \eta_2 - \eta_{1y} - \eta_1 \eta_3}$$
(22)

The statement of the theorem follows by considering a gnomonic projection, where $p = -z_x$ and $q = -z_y$. \square

Again, we note that the actual result proved is stronger and holds for several projections besides gnomonic. To emphasize the import of the result: with just Gaussian elimination and repeated use of the integrability condition, two pairs of differential images at unknown light source positions allow us to recover the direction of the gradient at every point of a surface with unknown isotropic BRDF.

Similar to Corollary 4, it immediately follows that:

Corollary 6. From two or more differential images of a surface, obtained from unknown source positions, it is possible to recover the isocontours of constant depth (or the level curves) for the entire surface.

This result may be contrasted with symmetry-based methods [1, 13] that theoretically require a dense configuration of lights at known positions to recover the same information.

5.2.1 Experimental evaluation

The recovery of gradient direction and level curves of constant depth is illustrated for synthetic data in Figure 4. The BRDF parameters used are the same as Figure 3.

2509

5.3. Surface Reconstruction from Isocontours

We note that one may not recover magnitude of the gradient without additional information, since the uncalibrated invariant of (7) is homogeneous in p, q and their derivatives. The following proposition establishes that, with the results of Corollaries 4 and 6, additional information is required for general surfaces only on a set of measure zero.

We assume that the surface can be split into a finite number of regions, each of which satisfies a generality condition, namely, that the isocontours of constant z and constant $\|\nabla z\|$ are not everywhere parallel. Surfaces such as a hemisphere are not general, but most surfaces do satisfy this condition.

Proposition 7. Given the isocontours of constant depth and constant $||\nabla z||$, under certain assumptions of generality for the surface z(x, y), specification of the surface normal at a single point suffices to reconstruct the depth map up to a global convex-concave ambiguity and additive offset.

A similar result is proved in [13]. However, the proof here is more easily generalizable to other classes of isocontours.

Proof. Let \mathbf{x}^* be a point in an open set \mathcal{U} where the isocontours of constant z and constant $\|\nabla z\|$ intersect transversally. Then, within \mathcal{U} , one may define unit vector fields \mathbf{v} and \mathbf{w} that are tangent, respectively, to the isocontours of constant z and constant $\|\nabla z\|$ and thus, are transversal. Then, since ∇z is orthogonal to \mathbf{v} , it is apparent that the following relation must hold at the point $\mathbf{x}^* \in \mathcal{U}$ (see Figure 5)

$$\mathbf{w} \cdot \frac{\nabla z}{\|\nabla z\|} = \pm \sqrt{1 - (\mathbf{v} \cdot \mathbf{w})^2}.$$
 (23)

Note the sign ambiguity, which arises since the directions of **v** and **w** can be specified at most up to a global sign. If the value of $||\nabla z||$ is specified at \mathbf{x}^* , it is also specified at every point on the isocontour of constant $||\nabla z||$. Thus, we have a linear ODE in z along that isocontour, which may be solved up to an additive constant, c^* .

Since the isocontours of constant z and constant $\|\nabla z\|$ are transversal in \mathcal{U} , the values of depths can now be assigned, up to an unknown c^* , along all the isocontours of constant z that intersect the isocontour of constant $\|\nabla z\|$ passing through \mathbf{x}^* and thus, almost everywhere on \mathcal{U} . Note that the sign ambiguity in (23) corresponds to a global convexconcave ambiguity and the unknown constant c^* corresponds to a global additive offset.

Finally, we note that for a unit normal $\hat{\mathbf{n}} = (n_1, n_2, n_3)^{\top}$, we have $\|\nabla z\| = \sqrt{n_3^{-2} - 1}$. Thus, specifying the surface



Figure 5. Transversality of isocontours of constant z and constant $||\nabla z||$ is sufficient to recover depth, given normal at \mathbf{x}^* . See Proposition 7.



Figure 6. Depth recovery for a monkey saddle surface. (a) The ground truth surface. (b) Recovered isocontours of constant $\|\nabla z\|$ (red) and constant z (blue), from images under a Torrance-Sparrow BRDF model. All points on the red and cyan curves have same values as respective green and magenta points. (c) Specifying the surface normal at one point (marked in blue) determines all depths along the corresponding isocontour of constant $\|\nabla z\|$ (green points), from which depth at every other point is determined by isocontours of constant z (red curves). (d) The reconstructed surface.

normal at a single point is sufficient to determine $\|\nabla z\|$ and seed the above depth reconstruction.

Again, we contrast with [1, 13], which use dense, calibrated sources and in the case of [1], additional information on an entire curve to resolve the depth. As the above proof shows, our gradient-based framework has the advantage of requiring **initial information at a single point**, in an uncal**ibrated setting**, with only two differential image pairs.

5.3.1 Experimental Evaluation

To empirically demonstrate reconstruction from isocontours, we simulate a monkey saddle, $z = x^3 - 3xy^2$, for which the isocontours of constant z and $\|\nabla z\|$ are transversal. In Fig. 6(b), we show recovery of isocontours of constant z and $\|\nabla z\|$. Specifying the normal at a single point on one of the isocontours of constant $\|\nabla z\|$ allows us to assign depths to isocontours of constant z (Fig. 6(c)) and recover the depth map (Fig. 6(d)). The recovered depth is nearly the same as ground truth except in regions where the chosen isocontour (dotted green in Fig. 6(c)) does not cross level curves of z.

6. Practical Algorithms for Reconstruction

Recall that, given image derivatives, one may estimate the entities λ and κ that contain sufficient information for disambiguating the surface, given the normal at a single point. However, in practice, the equation (22) involves thirdorder derivatives of the surface normal, which can lead to noisy estimation. In this section, we present algorithms for recovering surface normals given initial information across a curve, or recovering depths given boundary information. These algorithms require only the estimation of λ and κ , which can be performed directly from image derivatives, without resorting to higher order differentiation.

The images used for real data experiments were acquired using a Canon 5D camera. For the clay ball dataset (Fig.7(a)),



Figure 7. Isocontours of constant $||\nabla z||$ for a hand-moulded clay ellipsoid and for a plastic apple.

the differential pair is created by attaching two light bulbs to a shaft, which can be rotated around a wheel using a crank. The axis of the wheel is aligned with the camera principal axis. The time step Δt , computed as the ratio of the distance between the bulbs and the radius of the shaft, was 0.12 (about 7 degrees). For the apple (Fig.7(b)) and the teflon ball (Fig.8) datasets, we used a gantry to acquire images at Δt corresponding to 2 degrees.

Image derivatives are computed using a Savitzky-Golay filter. The PDE solution for Algorithm 2 in Section 6.2 uses central differences, with a smoothness regularizer and can be implemented in standard PDE solution frameworks. Please see [4] for further acquisition and implementation details.

6.1. Algorithm 1: Recovering Surface Normals

Proposition 8. Initial data in the form of known surface normals on a curve suffices to recover the surface normals from the coupled PDEs in (12) and (13).

Note that including the integrability requirement of (14) is ignored here (as is done in traditional Lambertian photometric stereo). Once the surface normals are estimated, one may impose integrability while recovering the depth map.

Proof. Consider the isocontours of constant $\|\nabla z\|$, parameterized by x, which are solutions to the ODE $\dot{y} = -\lambda(x, y)$. For a closed, smooth surface, these characteristic curves will, in general, be non-intersecting.

Along the above curves, the pair of PDEs in (12) and (13) reduces to a pair of ODEs: $\{\dot{p} = \kappa q, \ \dot{q} = -\kappa p\}$. Let $p = r \sin \theta$ and $q = r \cos \theta$. Then, the following pair of relations, obtained by differentiating p and q with respect to the curve parameter (in this case, x) and substituting in the above pair of ODEs, must be true:

$$\dot{r}\cos\theta = (\dot{\theta} - \kappa)r\sin\theta, \quad \dot{r}\sin\theta = (\kappa - \dot{\theta})r\cos\theta.$$
 (24)

Thus, $\dot{r}^2 = r^2(\dot{\theta} - \kappa)^2$. So, a solution to the pair of ODEs can be obtained as $\{\dot{r} = 0, \dot{\theta} = \kappa\}$. Given initial data along a curve, we can solve the above pair of ODEs. These initial conditions amount to specifying r and θ along a curve. \Box

6.1.1 Experimental Evaluation

Figure 7 shows the estimated isocontours of constant $\|\nabla z\|$ on real data for a hand-moulded clay ball and a plastic apple. The shape may be recovered by specifying the normals on a curve. An example reconstruction is shown in [4]. A practical limitation is that specifying such initial information



(a) Sample images (c) Our BRDF-Invariant Method Figure 8. Reconstruction with ground truth. (a) Two differential pairs for a teflon ball. (b) Height map using traditional photometric stereo. Note the typical shearing due to non-Lambertian effects. (c) Height map using the BRDF-invariant theory of Section 6.2.

may be difficult (the occluding contour is not usable as it is non-transversal to isocontours of constant $\|\mathbf{n}\|$).

6.2. Algorithm 2: Recovering Depth

Proposition 9. Boundary data in the form of known depths suffices to recover the surface depth from the PDEs in (12), (13) and (14).

Proof. Using the integrability condition (14), the pair of equations (12) and (13) can be written as a single constraint:

$$p_x = \lambda^2 q_y - \lambda \kappa p + \kappa q. \tag{25}$$

For a gnomonic projection, we have $p = -z_x$ and $q = -z_y$. Thus, we can rewrite the above as

$$z_{xx} - \lambda^2 z_{yy} + \lambda \kappa z_x - \kappa z_y = 0.$$
 (26)

This is a linear, second-oder hyperbolic PDE, which is well-posed given Dirichlet boundary conditions. \Box

While (26) is, in fact, a weaker condition than the coupled constraints in (12) and (13), it leads to a convenient numerical implementation. In practice, boundary depths may be specifiable when an object rests on a background plane.

6.2.1 Experimental Evaluation

Fig. 8(a) shows 2 of the 13 differential images of a teflon sphere. It can be easily seen that the material of the sphere is non-Lambertian. Consequently, the height map reconstructed by traditional photometric stereo (Fig. 8(b)) using 13 lights is clearly sheared. In contrast, the reconstructed height map using the algorithm presented in this section closely resembles a sphere (Fig. 8(c)).

Fig. 9 shows the complete pipeline for a clearly glossy plastic apple. Note the non-Lambertian effects in the derivative images. The recovered height map closely matches the expected shape. Fig. 10 shows reconstruction of an object that does not exactly satisfy the theoretical smoothness requirements. Note the high level of detail recovered, which demonstrates robustness to minor non-differentiability. Please see [4] for a detailed experimental analysis.



7. Conclusion and Future Work

In this paper, we have presented a comprehensive theory that relates image gradients to surface geometry in uncalibrated photometric stereo for isotropic BRDFs. In the process, we have uncovered fundamental insights into the nature of differential information contained in photometric images for isotropic BRDFs. We have presented a novel invariant for surface reconstruction and precisely characterized the extent to which this invariant informs surface reconstruction.

A key observation in our work is the linearity of the differentiation operation, that may be used to derive novel constraints on surface geometry, regardless of the exact form of the BRDF. This insight is of potential relevance in many other domains like shape from shading and optical flow and may provide a unified framework to analyze all of these problems with general BRDFs.

Acknowledgments This work is funded by ONR PECASE grant N00014-09-1-0741, a National Science Scholarship from the



(a) Input images (b) Ratio images (c) Derivative images



(d) Reconstruction (e) Textured (f) Close-up Figure 10. Reconstruction for toy dog. (a) Sample images from the differential pair. (b) Ratios with respect to the floodlit image. (c) Spatial and temporal derivatives. Red stands for positive values and blue stands for negative values. (d,e,f) Reconstruction using Algorithm 2 of Section 6.2, with Dirichlet boundary conditions. Note the high detail in the close-up in spite of minor non-differentiability.



(d) Recovered (e)Textured height map reconstruction

Figure 9. Reconstruction for the apple dataset. (a) Sample images from the differential pair. (b) Ratios with respect to the floodlit image. (c) Spatial and temporal derivatives. Blue indicates low values and red indicates high values. (d,e) Depth recovery.

A*STAR Graduate Academy of Singapore, as well as generous support from Intel, Adobe, NVIDIA and Pixar. We thank Andrew Adams and Marc Levoy for their kind help with the Stanford gantry, Miloš Hašan for the bunny data, Neil Alldrin for code from [1], as well as Todd Zickler, Steve Marschner, Szymon Rusinkiewicz and anonymous reviewers for helpful discussions and comments.

References

- N. Alldrin and D. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance : A stratified photometric stereo approach. In *ICCV*, 2007.
- [2] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *CVPR*, 2008.
- [3] S. Barsky and M. Petrou. The 4-source photometric stereo method for three-dimensional surfaces in the presence of highlights and shadows. *PAMI*, 25(10):1239–1252, 2003.
- [4] M. Chandraker, J. Bai, and R. Ramamoorthi. A theory of photometric reconstruction for unknown isotropic reflectances. Technical Report EECS-2010-176, UC Berkeley, 2010.
- [5] J. Clark. Active photometric stereo. In CVPR, 1992.
- [6] A. Hertzmann and S. Seitz. Example-based photometric stereo: Shape reconstruction with general, varying BRDFs. *PAMI*, 27(8):1254–1264, 2005.
- [7] M. Holroyd, J. Lawrence, G. Humphreys, and T. Zickler. A photometric approach for estimating normals and tangents. *ACM ToG*, 27(5):133:1–133:9, 2008.
- [8] B. Horn. Robot Vision. MIT Press, 1986.
- [9] B. Lucas and T. Kanade. An iterative image registration technique with an application to stereo vision. In *Image Understanding Workshop*, pages 121–130, 1981.
- [10] M.Brooks and B. Horn. Shape and source from shading. In *IJCAI*, pages 932–936, 1985.
- [11] R. Ramamoorthi, D.Mahajan, and P. Belhumeur. A first order analysis of lighting, shading and shadows. TOG, 26(1), 2007.
- [12] P. Tan, S. Mallick, D. Kriegman, L. Quan, and T. Zickler. Isotropy, reciprocity and the GBR ambiguity. In *CVPR*, 2007.
- [13] P. Tan and T. Zickler. A projective framework for radiometric image analysis. In CVPR, 2009.
- [14] P. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Engg.*, 19(1):139–144, 1980.
- [15] R. Zhang, P. Tsai, J. Cryer, and M. Shah. Shape from shading: A survey. *PAMI*, 21(8):690–706, 1999.
- [16] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *IJCV*, 49(2/3):1215–1227, 2003.