

A Theory of Photometric Reconstruction for Unknown Isotropic Reflectances

Manmohan Chandraker
Jiamin Bai
Ravi Ramamoorthi



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2010-176

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-176.html>

December 21, 2010

Copyright © 2010, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

A Theory of Photometric Reconstruction for Unknown Isotropic Reflectances

Manmohan Chandraker

Jiamin Bai

Ravi Ramamoorthi

EECS Department, University of California, Berkeley

{manukc, bjiamin, ravir}@eecs.berkeley.edu

Abstract

This paper presents a comprehensive theory of surface reconstruction from image derivatives in photometric stereo. For an object with an unknown, general isotropic BRDF, we show that just two measurements of the spatial and temporal image derivatives, under unknown light source positions 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 this 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, for the entire surface. Further, we prove that specification of the surface normal at a single point completely determines the surface depth from these isocontours.

In addition, our theory also suggests practical algorithms that require additional initial or boundary information, but allow reconstruction of depth from lower order derivatives. The theoretical results of the paper are illustrated with examples on synthetic and real data.

1. Introduction

The image formation process is an 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. For Lambertian scenes, there exist wide-ranging theories of photometric stereo and shape-from-shading to recover shape from image intensities and gradients, respectively. However, the theoretical relationships between image gradients and surface geometry, for general BRDFs and unknown light source motions, remains 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 several differential equations when spatial and temporal derivatives of the image are considered. Moreover, these differential equations are linear in the BRDF

derivatives, so under appropriate conditions (such as circular motion of the light source), they may be eliminated to yield a photometric invariant that relates image derivatives, surface geometry and light 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 $z(x, y) : \mathbb{R}^2 \rightarrow \mathbb{R}$, 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, respectively, 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 synthetic and real data examples.

To summarize, this paper introduces a comprehensive theory of photometric stereo using image derivatives, resting on several fundamental contributions:

- A novel photometric invariant that relates image derivatives to the geometry of a surface with arbitrary isotropic BRDF, without requiring light positions.
- Two differential measurements suffice to extract surface information by solving a small linear system.
- Isocontours of constant depth and constant magnitude of the gradient may be recovered from the invariant.
- Theoretical recovery of surface depth by specifying the surface normal at a single point.
- Practical algorithms for surface normal or depth estimation that require only lower order derivatives.

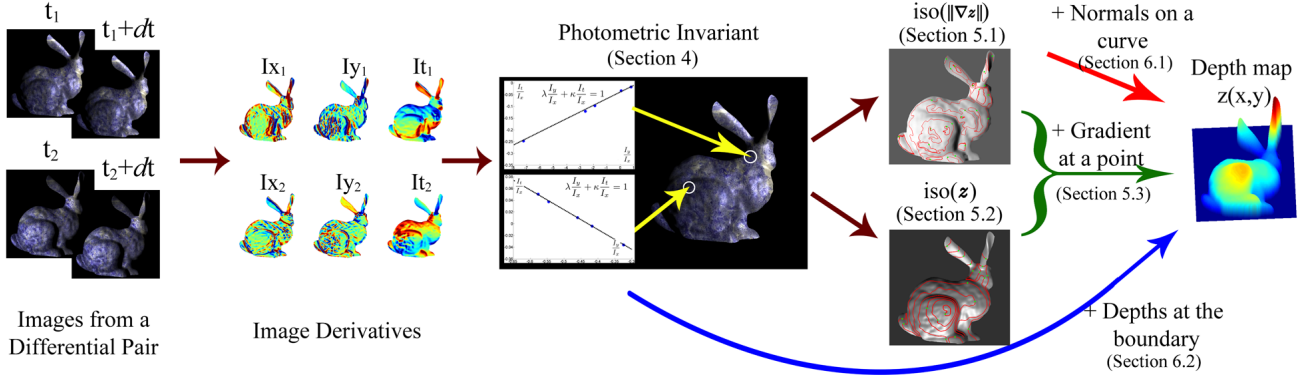


Figure 1. An outline of our differential theory of photometric stereo. Two or more differential pairs lead to an uncalibrated photometric invariant. Isocontours of constant $\|\nabla z\|$ and constant z may be recovered from the invariant. Additional information in the form of gradient at a single point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

2. Related Work

This work differs from most earlier computer vision studies in considering differential information within the context of photometric stereo, with unknown BRDF and light directions. Our theoretical results also distinguish this work by establishing a minimal requirement of two differential pairs for 3D reconstruction.

Most prior work in photometric stereo has dealt with Lambertian surfaces, for which the surface normal can be recovered from three images [13]. In some cases, specularity removal has been used as a pre-processing step for using Lambertian photometric stereo [3].

The example-based photometric stereo of [6] is valid for general BRDFs, but requires a reference object of the same material. With multiple views, Helmholtz stereopsis [15] is an elegant approach to eliminate the BRDF by exploiting reciprocity when the camera and light source are swapped. With a color-based separation of the Lambertian and specular components of a dichromatic reflectance model, shape is recovered from two images in [12]. Further empirical properties of the BRDF can be exploited to recover 3D shape, but require a dense coverage of the lighting hemisphere [2, 7].

The method of [1] uses bilateral symmetry of a spatially varying BRDF to recover isocontours of constant depth. It requires a dense configuration of known lights on a circle and initial information on an entire curve to recover depth. In contrast, our theory is based on derivative information and requires as few as two differential pairs of light sources, at unknown positions on a circle. Further, we recover isocontours of constant depth as well as constant magnitude of the gradient, from which surface depth can be obtained with initial information at a single point.

A related work that uses differential information to derive photometric invariants is the active photometric stereo of [4]. However, it recovers depth using a distance-dependent imaging model and requires calibrated light source positions.

Besides photometric stereo, derivative information in images has been considered in other contexts. Shape from shading seeks to recover depth from a single image of a Lambertian surface [10, 14]. In computer graphics, the first-order behavior of reflection as a convolution and visibility under area lighting have been studied in [11]. Optical flow uses




Figure 2. An illustration of two possible parameterizations for the surface normal. The gnomonic projection (red line) is denoted $\mathbf{n} = \pi_g(\hat{\mathbf{n}})$ and the stereographic projection (purple line) is denoted $\mathbf{n} = \pi_s(\hat{\mathbf{n}})$.

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 does not rely on brightness constancy assumptions, neither 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 2-sphere and its representation \mathbf{n} is determined by the projection mapping $\pi : S^2 \rightarrow \mathbb{R}^2$ [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, $\mathbf{n} = [-z_x, -z_y]^\top$. For a stereographic projection centered on the south pole, $\hat{\mathbf{n}} = \frac{(2\mathbf{n}^\top, l(\mathbf{n})^2 - 2)^\top}{l(\mathbf{n})^2}$. Similarly, we represent a directional point light source $\hat{\mathbf{s}}$ by a 2-vector, \mathbf{s} . Figure 2 illustrates the two parameterizations.

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). For practical implementations, we will choose one of the above two mappings.

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 $\hat{\rho}(\hat{\mathbf{n}}^\top \hat{\mathbf{s}}, \hat{\mathbf{n}}^\top \hat{\mathbf{v}}, \hat{\mathbf{s}}^\top \hat{\mathbf{v}})$. For gnomonic projection, the angles that determine the isotropic BRDF are

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

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 $\hat{\rho}$:

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

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 assumptions, \mathbf{n} depends on space variables \mathbf{x} and \mathbf{s} depends on time variables t .

3.3. Ratio images

The image formation model can be easily generalized to:

$$E(\mathbf{x}, t) = a(\mathbf{x}) \hat{\rho}(\alpha(\mathbf{x}, t), \beta(\mathbf{x}), \gamma(t)), \quad (3)$$

where $a(\mathbf{x})$ is a surface albedo. This is a reasonable model and in particular, subsumes the Lambertian assumption of traditional photometric stereo, while allowing for more general isotropic BRDFs, possibly modulated by a spatially-varying texture or albedo. In fact, 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 collocated 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

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

where $\rho(\cdot)$ is the appropriately defined function. The exact form of $\hat{\rho}(\cdot)$, $\hat{\rho}(\cdot)$ or $\rho(\cdot)$ is not important for us, 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 recovering shape from photometric information is to derive invariants that relate the image intensities to surface geometry and light source directions. Most prior work focuses on more restricted classes of BRDF (such as Lambertian) to avoid the complexities of general BRDFs. 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 space or time variables. These may then be related by eliminating any terms that depend on the functional form of the BRDF. In this section, we use this intuition to derive a novel photometric invariant.

The space and time derivatives of the images in (4) are

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

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

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 $\rho_{\alpha}, \rho_{\beta}, \rho_{\gamma}, \mathbf{n}$ and \mathbf{s} . However, one way to extract constraints on the normal might be to eliminate the BRDF derivatives, which can be done from these three equations if one of $\rho_{\alpha}, \rho_{\beta}$ or ρ_{γ} terms can be made to identically vanish. This can be achieved 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, $\mathbf{s}_t^\top \mathbf{s} = 0$ for lights on a circle, so equation (6) yields $\rho_{\alpha} = \frac{I_t}{\mathbf{s}_t^\top \mathbf{n}}$. Substituting in the two equations in (5):

$$I_x = \left(\frac{I_t}{\mathbf{s}_t^\top \mathbf{n}} \right) \mathbf{s}_t^\top \mathbf{n}_x + \rho_{\beta} \frac{1}{l(\mathbf{n})} \mathbf{n}^\top \mathbf{n}_x \quad (7)$$

$$I_y = \left(\frac{I_t}{\mathbf{s}_t^\top \mathbf{n}} \right) \mathbf{s}_t^\top \mathbf{n}_y + \rho_{\beta} \frac{1}{l(\mathbf{n})} \mathbf{n}^\top \mathbf{n}_y \quad (8)$$

Eliminating ρ_{β} from the two equations leads to:

$$\frac{I_x \mathbf{s}_t^\top \mathbf{n} - I_t \mathbf{s}_t^\top \mathbf{n}_x}{I_y \mathbf{s}_t^\top \mathbf{n} - I_t \mathbf{s}_t^\top \mathbf{n}_y} = \frac{\mathbf{n}^\top \mathbf{n}_x}{\mathbf{n}^\top \mathbf{n}_y}. \quad (9)$$

Clearly, the right hand side depends only on position, while the left hand side contains time-dependent entities. Thus, a photometric invariant that relates image derivatives, surface geometry and light directions can be expressed as

$$\frac{I_x \mathbf{s}_t^\top \mathbf{n} - I_t \mathbf{s}_t^\top \mathbf{n}_x}{I_y \mathbf{s}_t^\top \mathbf{n} - I_t \mathbf{s}_t^\top \mathbf{n}_y} = \text{constant across time.} \quad (10)$$

4.3. The Invariant Is Independent of Light Positions

It might seem at a first glance that, given light source positions, a non-linear minimization framework can be used to solve for the unknowns $\{\mathbf{n}, \mathbf{n}_x, \mathbf{n}_y\}$, up to scale, using the constraint in (9). However, as we show below, the light source directions are not required at all. That is, the invariant (10) directly relates image derivatives to surface geometry.

Proposition 1. *The entire information in the photometric invariant (10) is encapsulated by two entities, which depend only on surface geometry and not on source positions \mathbf{s} .*

While the above may seem surprising given the presence of light source terms in (10), a closer look at the structure of (10) immediately proves the above proposition.

Proof. Defining

$$\lambda = \frac{\mathbf{n}^\top \mathbf{n}_x}{\mathbf{n}^\top \mathbf{n}_y}, \quad \mathbf{u} = \mathbf{n}_x - \lambda \mathbf{n}_y, \quad (11)$$

we can rewrite (9) as

$$\frac{I_x - \lambda I_y}{I_t} = \frac{\mathbf{s}^\top \mathbf{u}}{\mathbf{s}_t^\top \mathbf{n}} = \kappa. \quad (12)$$

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 \mathbf{u} in (11), we have $\mathbf{u}^\top \mathbf{n} = 0$, that is, \mathbf{u} is orthogonal to \mathbf{n} . Since the lights lie on a circle, \mathbf{s}_t is also orthogonal to \mathbf{s} . Moreover, at time t (or equivalently, angular position on the circle of sources), the light source is $\mathbf{s} = (r \cos t, r \sin t)^\top$. Then, $\mathbf{s}_t = (-r \sin t, r \cos t)^\top$, 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:

$$\mathbf{u} = -\kappa \mathbf{n}_\perp. \quad (13)$$

Thus, we have the following constraint, equivalent to the invariant in (9), but which *does not depend on knowledge of light source positions*:

$$I_x - \lambda I_y - \kappa I_t = 0. \quad (14)$$

Clearly, all the information in the invariant is encapsulated by λ and κ , which depend only on surface geometry. \square

It is tempting to compare equation (14) to the optical flow relation [9]. However, we note that (14) is derived without resorting to a 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.

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 just an unmarked piece of string may suffice to create any integral ratio.

4.4. Importance of the Invariant

This result of Proposition 1 is a surprising one and in fact, can be understood as a more 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.*

The above is also useful practically, since it raises the possibility that **surface information may be recovered from image derivatives in photometric stereo for isotropic BRDFs, without knowledge of light source positions**. In fact, as few as two differential image pairs suffice to estimate λ and κ by solving a small linear system. This is an important observation:

Corollary 3. *Two pairs of differential images suffice to recover surface information contained in the invariant of (10).*

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.

At this stage, it is natural to seek a characterization of the exact surface information recoverable from λ and κ . In the following section, we will show that knowledge of λ and κ , together with initial information at a single point, suffices to completely determine the surface depths.

4.5. Experimental Evaluation

Here, we empirically illustrate the validity of the relation in (14). 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 3 illustrates that the ratios $\frac{I_y}{I_x}$ and $\frac{I_t}{I_x}$ computed for the 11 different light positions lie close to a straight line. The entities λ and κ are given by the best-fit straight line.

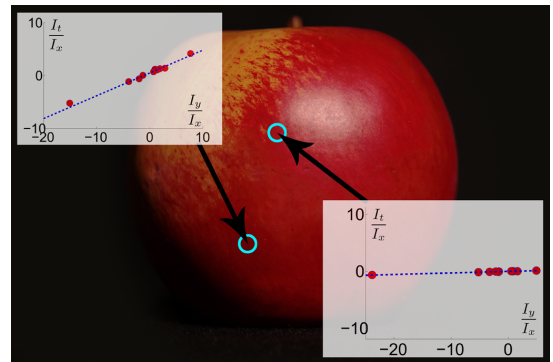


Figure 3. The ratio of image derivatives at a pixel, recorded for various light positions, lie on a straight line given by the equation $\lambda \frac{I_y}{I_x} + \kappa \frac{I_t}{I_x} = 1$.

5. Reconstructibility Using the Invariant

The previous section demonstrates that the entire information in the invariant can be expressed (and recovered) in terms of the two entities λ and κ . However, it is not immediately clear what this means in terms of surface reconstruction. We answer this conclusively in the present 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, 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 gradient associated with this field is $\nabla g = (\|\mathbf{n}\|_x, \|\mathbf{n}\|_y)^\top$. It is well-known that the level curves of a scalar field, $g(x, y) = c$, for constant c , are orthogonal to ∇g . The direction of the tangent to the level curves is $\frac{\|\mathbf{n}\|_y}{\|\mathbf{n}\|_x}$, which is the same as $\frac{1}{\lambda}$, by definition (11). Thus, knowing the value of λ , one may trace the isocontours of constant $\|\mathbf{n}(x, y)\|$, using the knowledge of their tangent directions. The statement of the theorem follows as a special case for the gnomonic projection, where $\mathbf{n} = -\nabla z$. \square

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

5.1.1 Experimental Evaluation

In Figure 4, we illustrate the recovery of isocontours of constant $\|\nabla z\|$ using synthetic data. For the images of a synthetic sphere of varying albedo and the synthetic bunny, we use a simplified Torrance-Sparrow model

$$\bar{\rho} = \frac{1}{4\pi\sigma^2} \exp \left[- \left(\frac{\cos^{-1} \hat{\mathbf{n}}^\top \hat{\mathbf{h}}}{\sigma} \right)^2 \right], \quad \hat{\mathbf{h}} = \frac{\hat{\mathbf{s}} + \hat{\mathbf{v}}}{\|\hat{\mathbf{s}} + \hat{\mathbf{v}}\|}, \quad (15)$$

with $\sigma = 0.3$. For images of the vase of varying albedo, we use a constant coefficient Phong-Blinn model: $\bar{\rho} = \hat{\mathbf{n}}^\top \hat{\mathbf{s}} + (\hat{\mathbf{n}}^\top \hat{\mathbf{h}})^\sigma$, with a typical value of $\sigma = 5$.

In each case, we observe that the recovered isocontours of constant $\|\nabla z\|$ match the ground truth very closely.

5.2. Isocontours of Constant Depth

To recover isocontours of constant $\|\mathbf{n}\|$, we required knowledge only of λ . However, we should be able to further disambiguate the surface using κ . The following proposition gives a constructive proof that λ and κ together determine the isocontours of constant depth (in addition to the isocontours of constant $\|\nabla z\|$).

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

Proof. Let us denote $\mathbf{n} = (p, q)^\top$, thus, $\mathbf{n}_x = (p_x, q_x)^\top$ and $\mathbf{n}_y = (p_y, q_y)^\top$. Given two or more pairs of differential images, one may estimate $\lambda(x, y)$ and $\kappa(x, y)$ using the linear relation in (14). Thereby, one obtains two linear, first order PDEs from (13) at every pixel:

$$p_x - \lambda(x, y)p_y = \kappa(x, y)q \quad (16)$$

$$q_x - \lambda(x, y)q_y = -\kappa(x, y)p. \quad (17)$$

Further, since we are dealing with a surface, it must satisfy the integrability condition:

$$p_y - q_x = 0. \quad (18)$$

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, we get

$$\begin{aligned} h_x &= p_x - \lambda q_x - \lambda_x q \\ &= p_x - \lambda p_y - \lambda_x q = (\kappa - \lambda_x)q \end{aligned} \quad (19)$$

$$\begin{aligned} h_y &= p_y - \lambda q_y - \lambda_y q \\ &= q_x - \lambda q_y - \lambda_y q = -\kappa p - \lambda_y q. \end{aligned} \quad (20)$$

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

$$(\lambda_x - \kappa)q_y - \lambda_y q_x - \kappa p_x = \kappa_x p + \kappa_y q. \quad (21)$$

Thus, we have a linear system in $\{p_x, p_y, q_x, q_y\}$, given by

$$\begin{bmatrix} 1 & -\lambda & 0 & 0 \\ 0 & 0 & 1 & -\lambda \\ 0 & 1 & -1 & 0 \\ -\kappa & 0 & -\lambda_y & (\lambda_x - \kappa) \end{bmatrix} \begin{bmatrix} p_x \\ p_y \\ q_x \\ q_y \end{bmatrix} = \begin{bmatrix} \kappa q \\ -\kappa p \\ 0 \\ \kappa_x p + \kappa_y q \end{bmatrix}, \quad (22)$$

whereby expressions for $\{p_x, p_y, q_x, q_y\}$ are obtained in terms of linear functions of $\{p, q\}$:

$$p_x = \nu_1 p + \eta_1 q, \quad p_y = \nu_2 p + \eta_2 q \quad (23)$$

$$q_x = \nu_2 p + \eta_2 q, \quad q_y = \nu_3 p + \eta_3 q. \quad (24)$$

with

$$\begin{bmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{bmatrix} = \frac{1}{\Delta} \begin{bmatrix} \lambda(\kappa^2 - \lambda_x \kappa + \lambda \kappa_x) \\ \kappa^2 - \lambda_x \kappa + \lambda \kappa_x \\ -(\lambda \kappa^2 + \lambda_y \kappa - \kappa_x) \end{bmatrix}, \quad (25)$$

$$\begin{bmatrix} \eta_1 \\ \eta_2 \\ \eta_3 \end{bmatrix} = \begin{bmatrix} \kappa \\ 0 \\ 0 \end{bmatrix} + \frac{\kappa_y + \kappa^2}{\Delta} \begin{bmatrix} \lambda^2 \\ \lambda \\ 1 \end{bmatrix}, \quad (26)$$

where $\Delta = -(\kappa \lambda^2 + \lambda \lambda_y - \lambda_x + \kappa)$ is the $(3, 2)$ -minor of the 4×4 matrix in (22). At this stage, the reader may verify




Figure 4. 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. (d) Ground truth isocontours for the bunny.

by substitution that the expressions above are consistent with the definitions of λ and κ in (11).

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

$$(\nu_1 p + \eta_1 q)_y = (\nu_2 p + \eta_2 q)_x \quad (27)$$

$$(\nu_2 p + \eta_2 q)_y = (\nu_3 p + \eta_3 q)_x, \quad (28)$$

These are linear, first order PDEs, where we can again replace the first order derivatives $\{p_x, p_y, q_x, q_y\}$ using (23) and (24) to get two (dependent) homogeneous linear equations in $\{p, q\}$. Using, say, the first equation, we can get the ratio of p and q as

$$\frac{q}{p} = \frac{\nu_1 \nu_3 + \eta_1 \nu_3 - \nu_2 \nu_3 - \nu_2 \eta_2}{\nu_2 \eta_1 + \eta_2 \nu_3 + \eta_2^2 - \nu_1 \eta_2 - \eta_1 \nu_3} \quad (29)$$

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: using just Gaussian elimination and repeated use of the integrability condition, **just 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 isotropic BRDF**, without requiring any additional information.

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] that theoretically require a dense configuration of lights at known positions to recover the same information.

5.2.1 Experimental evaluation

The recovery of the direction of the gradient and the level curves of constant depth is illustrated for synthetic data in Figure 5. To display the direction of the gradient, the arrows are plotted with length normalized to one. The isocontours of constant depth are shown in red. All points on a red curve have the same depth as the green point on that curve.

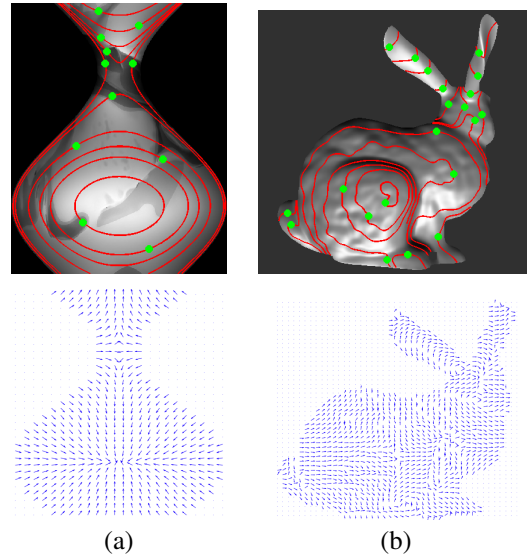


Figure 5. Recovery of isocontours of constant depth. (Top row) Red curves plot the isocontours, while the green dots represent points chosen to start tracing the curves. (Bottom row) Plots of the direction of the gradient, with the length of the vector normalized to one for display. (a) A vase with variable albedo and Blinn-Phong BRDF. (b) A bunny with Torrance-Sparrow BRDF.

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 (10) 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.*

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.

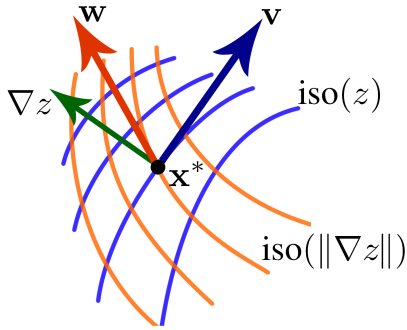


Figure 6. Transversality of isocontours of constant z and constant $\|\nabla z\|$ is sufficient to recover depth, given normal at \mathbf{x}^* . See Proposition 7.

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 6)

$$\mathbf{w} \cdot \frac{\nabla z}{\|\nabla z\|} = \pm \sqrt{1 - (\mathbf{v} \cdot \mathbf{w})^2}. \quad (30)$$

Note the sign ambiguity, which arises since the directions of \mathbf{v} and \mathbf{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 (30) corresponds to a global convex-concave 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 $\|\nabla z\| = \sqrt{\left(\frac{1}{n_3}\right)^2 - 1}$. Thus, specifying the surface normal at a single point is sufficient to determine $\|\nabla z\|$ and seed the above depth reconstruction. \square

Again, we contrast with calibrated methods [1], which use dense sources to recover isocontours of constant z and require additional information on an entire curve to resolve the depth. As the above proof shows, **incorporating gradients has the advantage of reducing this ambiguity, in an uncalibrated framework with only two differential images, to the specification of information at a single point.**

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. 7(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. 7(c)) and recover the depth map (Fig. 7(d)). The recovered depth is nearly the same as ground truth, except in regions where the chosen isocontour (dotted green in Fig. 7(c)) does not cross level curves of z .




Figure 7. Depth recovery for a monkey saddle surface. (a) The ground truth surface. (b) Recovered isocontours of constant $\|\nabla z\|$ (red) and constant z (cyan), from images under a Torrance-Sparrow BRDF model. All points on the red and cyan curves have same depths as the respective green and magenta points. (c) Specifying the surface normal at one point (marked in blue) on the surface determines all depths along the corresponding isocontour of constant $\|\nabla z\|$ (green points), from which depth at every other point can be determined by tracing isocontours of constant z (red curves). (d) The reconstructed surface by tracing all the isocontours.

6. Practical Algorithms for Reconstruction

At this stage, we recall that given image derivatives, one may estimate the entities λ and κ that contain sufficient information for disambiguating the surface, given the surface normal at a single point. However, in practice, the equation (29) involves third-order derivatives of the surface normal, which can lead to noisy estimation.

One alternative to recover the surface normals or depth is to incorporate additional information in the form of initial or boundary conditions. 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.9(a)), 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 teflon ball (Fig.12), apple (Fig.10) and toy dog (Fig.13) 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.




Figure 8. Reconstruction for a synthetic sphere using Algorithm 1 of Section 6.1. (a) Isocontours of constant $r = \|\nabla z\|$ computed using the estimated $\lambda(x, y)$, followed by estimation of $\theta = \tan^{-1}\left(\frac{z_x}{z_y}\right)$ using seed values of the surface normal specified on the brown dots. The values of θ are shown using a color-code that maps the interval $\left[-\frac{\pi}{2}, \frac{\pi}{2}\right]$ between blue and red. (b) Recovered values of θ , with the specified color map. (c) Ground truth θ . (d) Estimated gradient in the x direction, z_x , red indicates large positive value and blue indicates large negative value. (e) Estimated gradient in the y direction, z_y , with the same color map. (f) Side view of the height recovered from integration of estimated z_x and z_y , which matches the ground truth shape. Red indicates higher heights and blue indicates lower heights.

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 (16) and (17).*

Note that including the integrability requirement of (18) 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 (16) and (17) reduces to a pair of ODEs:

$$\dot{p} = \kappa q, \quad \dot{q} = -\kappa p. \quad (31)$$

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:

$$\dot{r} \cos \theta = (r \sin \theta)(\dot{\theta} - \kappa), \quad \dot{r} \sin \theta = (r \cos \theta)(\kappa - \dot{\theta}). \quad (32)$$

Thus, $\dot{r}^2 = r^2(\dot{\theta} - \kappa)^2$. So, a solution to the pair of ODEs can be obtained as

$$\dot{r} = 0, \quad \dot{\theta} = \kappa. \quad (33)$$

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. \square

6.1.1 Experimental Evaluation

We will demonstrate the applicability of this algorithm using the synthetic sphere data. Given the estimated values of $\lambda(x, y)$, Figure 8(a) shows a few contours of constant r (that is, the curves $\dot{r} = 0$), traced with the initial seed point marked by a brown dot. At the brown dots, we specify the value

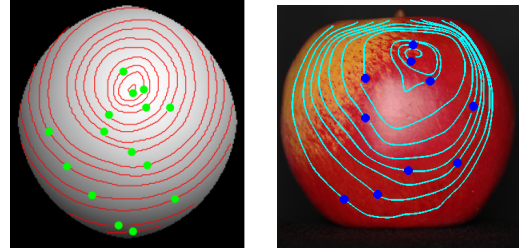


Figure 9. Isocontours of constant $\|\nabla z\|$ for real data corresponding to a hand-moulded clay ball and a plastic apple.

of the surface normal, thus, the value of θ . Then, along the isocontours, we solve the integral equation $\theta = \int \kappa dx$, with the constant of integration being the specified value at the brown dot. The recovered values of θ are shown color-coded on the corresponding isocontours in Figure 8(a).

Figure 8(b) shows the values of θ recovered by tracing isocontours over all the pixels of the image. The artifacts near the center are due to discretization (our isocontours are traced with sub-pixel precision, but the values of θ can be displayed only up to a 1 pixel precision). For comparison, we show the ground truth value of θ in Figure 8(c).

Given the values of r and θ , we can recover the gradients z_x and z_y using the above equations. The recovered values are shown in Figures 8(d) and (e). Finally, using the recovered gradients, one may perform a surface normal integration to estimate the depth map, shown in Figure 8(f).

Note that the algorithm requires knowledge of surface normals along the curve given by the brown dots in Figure 8(a), which is difficult to obtain in practice. But if initial data on a curve is indeed available, Algorithm 6.1 is an elegant method to recover the surface normals directly from first order image derivatives, without resorting to higher order differentiation.

In Figure 9, we show the recovered isocontours of constant $\|\nabla z\|$ on real data for a hand-moulded clay ball and a plastic apple. If the surface normals are specified at a single point on these contours, the shape may be recovered. Note that one may not use the occluding contour as the initial curve, as it can be shown to be non-transversal to isocontours of constant $\|\mathbf{n}\|$.





Figure 10. Input pipeline for the apple dataset. (a) Sample images from the differential pair. (b) Floodlit image. (c) Ratios with respect to the floodlit image. (d) Spatial and temporal derivatives. Blue indicates low values and red indicates high values.



Recovered height map Texture-mapped reconstruction

Figure 11. Reconstruction on real data for the apple dataset using Algorithm 2 of Section 6.2, with Dirichlet boundary conditions. (Top row) Side view. (Bottom row) Top and side view. (Left column) Height map, red indicates higher values and blue stands for lower values. (Right column) Texture-mapped reconstruction.

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 (16), (17) and (18).*

Proof. Using the integrability condition (18), the pair of equations (16) and (17) can be written as a single constraint:

$$p_x = \lambda^2 q_y - \lambda \kappa p + \kappa q. \quad (34)$$

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

$$z_{xx} - \lambda^2 z_{yy} + \lambda \kappa z_x - \kappa z_y = 0. \quad (35)$$

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

While demanding the satisfiability of the constraint in (35) is, in fact, a weaker condition than the coupled constraints in (16) and (17), it leads to a convenient numerical implementation. In practice, depths may be specified at the boundary for scenes where an object rests on a background plane. Indeed, the extensive theory of solutions for hyperbolic PDEs can be used to solve the PDE in (35).

6.2.1 Experimental Evaluation

In Figure 10(a), we show 2 of the 11 differential pair of images for a plastic apple. Note that the object has variable albedo and a non-Lambertian BRDF. These input images are divided by the floodlit image in Figure 10(b) to obtain the ratio images in Figure 10(c). Note that the albedo variations are eliminated in the ratio images. Spatial and temporal derivatives are computed on the ratio images (Figure 10(d)). From these derivatives, we perform a reconstruction using Algorithm 2 of this section, with boundary depths set to 0. As seen in Figure 11, the 3D reconstruction closely resembles the expected shape.

In Fig. 12, as a means of easy comparison to ground truth, we show experimental results for a real teflon sphere with uniform albedo. Note that traditional Lambertian photometric stereo can be performed with just two light sources, or several coplanar lights, for an object with uniform albedo. Fig. 12(a) shows 2 of the 13 differential pairs used for reconstruction, while Fig. 12(b) shows the corresponding spatial and temporal derivatives. It can be easily seen from the images that the material of the sphere is non-Lambertian. Consequently, the height map reconstructed by traditional photometric stereo using 13 lights is clearly sheared (Fig. 12(c)). In contrast, the reconstructed height map using the algorithm presented in this section closely resembles a sphere (Fig. 12(d)).

In Figure 13(a), we show input images for a toy dog, painted with water color. Note the coarse and fine scale variations in the albedo, as well as some non-Lambertian effects. More importantly, the object surface is not differentiable, so it presents a challenging scenario for the theory of this paper. We again acquire a floodlit image (Figure 13(b)) and compute ratio images to eliminate the albedo and obtain images whose intensities depend only on the BRDF (Figure 13(c)).

The reconstruction using Algorithm 2 of this section, with boundary depths set to 0, is shown in Figure 14(a). A depth map for the reconstruction is shown in Figure 14(b). A texture-mapping on the same reconstruction is shown in Figure 14(c). In Figure 14(d), we show a close-up of the reconstruction for the head.

Note the fine scale structure recovered by the algorithm, such as the eyes and the smiling mouth. There is a minor loss of detail near the feet of the toy dog, which might be

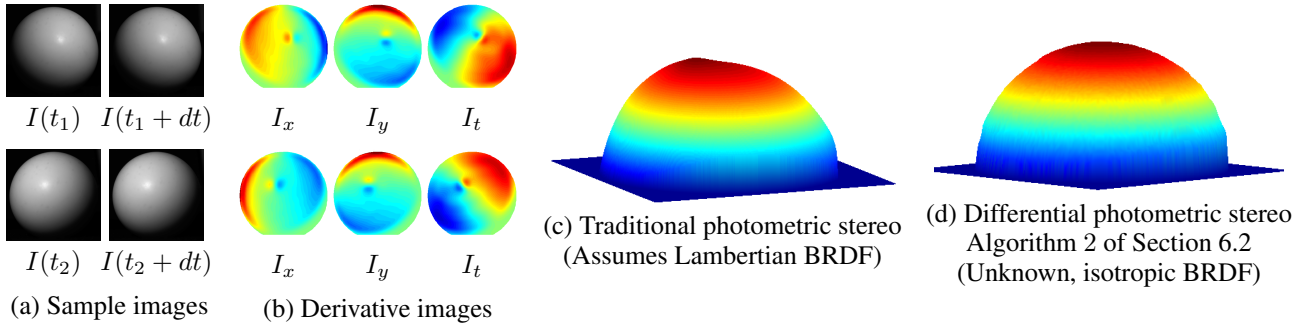


Figure 12. Comparison with traditional Lambertian photometric stereo, using real data. (a) Two sample differential pairs of images of a teflon ball. (b) Spatial and temporal image derivatives. Red stands for positive values and blue stands for negative values. (c) Height recovered using traditional Lambertian photometric stereo. Note the shearing of the surface which is typically due to ignoring non-Lambertian effects. (d) Height recovered using the theory of Section 6.2.

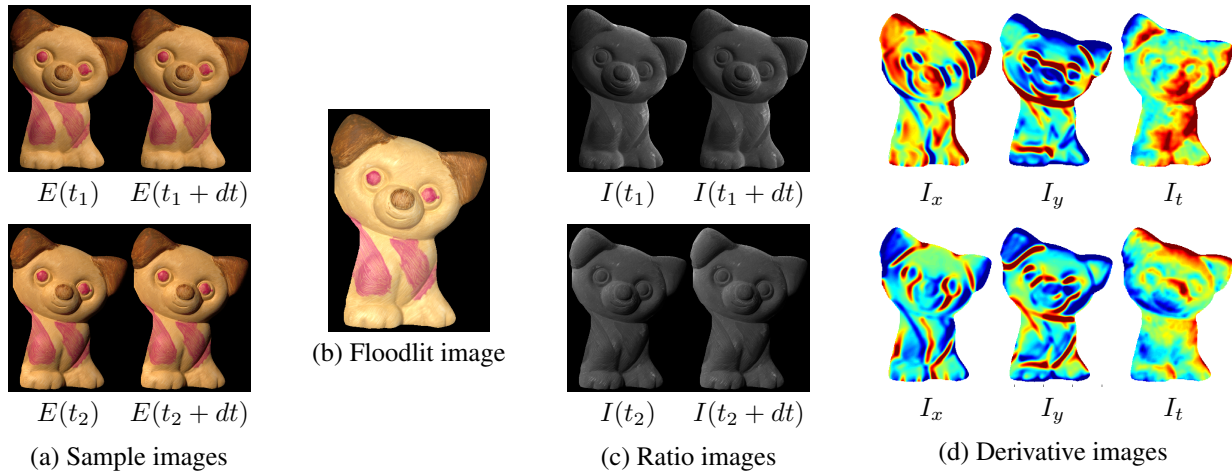


Figure 13. Input real images for the toy dog dataset. (a) Sample images from the differential pair. (b) Floodlit image. (c) Ratios with respect to the floodlit image. (d) Spatial and temporal derivatives. Red stands for positive values and blue stands for negative values.

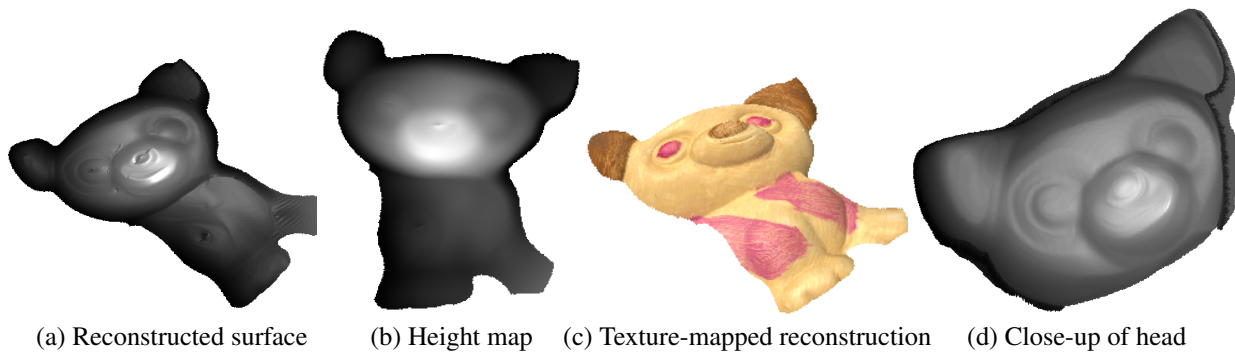


Figure 14. Reconstruction for the toy dog using Algorithm 2 of Section 6.2, with Dirichlet boundary conditions. Note that the surface is not continuously differentiable, so does not meet the strict requirements of our theory. Yet, the PDE solution is robust enough to recover a high quality depth map. (a) A top-and-side view of the recovered surface. (b) Recovered height map. Brighter shading indicates higher values. (c) A texture-mapped display of the 3D reconstruction. (d) A close-up reconstruction of the head.

due to shadows and low levels of illumination (all the light source positions are above the level of the head). While the theory of the paper is derived for differentiable surfaces, this example shows robustness to minor non-differentiability, provided the surface is continuous. We remind the reader

that all the reconstructions in this paper are obtained from unknown light source positions.

7. Implementation Details

In this section, we provide some implementation details relevant to the real data experiments.

7.1. PDE Solution

We take central differences to numerically solve the system of equations arising from equation (35) at every pixel:

$$z_x = \frac{z_r - z_l}{2\Delta x}, \quad z_{xx} = \frac{z_r - 2z + z_l}{(\Delta x)^2},$$

$$z_y = \frac{z_t - z_b}{2\Delta y}, \quad z_{yy} = \frac{z_t - 2z + z_b}{(\Delta y)^2},$$

where $\{\Delta x, \Delta y\}$ are the step lengths in the x and y directions and $\{z_l, z_r, z_b, z_t\}$ are the pixels on the left, right, bottom and top of z , respectively, which are one step length away. The depths of the boundary pixels are specified to be 0. Thus, we obtain a linear system of the form $\mathbf{Az} = \mathbf{b}$, where \mathbf{A} is a large, but extremely sparse matrix. Note that the reconstructed depths at the boundaries may be non-zero, since the PDE constraints and the boundary constraints are solved in a single linear system, in a least squares sense.

Points where $|\lambda|$ and $|\kappa|$ cross a certain threshold (set to 50 in our implementation) are detected. We do not write the PDE constraint at these points, instead, we impose a condition that the depth must be continuous at these points.

Optionally, a regularization term (such as an isotropic prior on the norm-squared gradient, or the Laplacian) may be added to enhance smoothness of the solution. The regularized objective function is

$$\min_{\mathbf{z}} \|\mathbf{Az} - \mathbf{b}\|_2^2 + \mu \|\nabla \mathbf{z}\|_2^2. \quad (36)$$

When regularization is used, we minimize (36) using the CVX convex optimization software [5]. For the apple dataset, we use a value of $\mu = 0.01$. For the teflon ball dataset, no regularization is required ($\mu = 0$).

For the toy dog dataset, instead of the above regularization, we weighted each PDE by the quantity $\frac{1}{\min\{1 + |\lambda|, 1 + |\kappa|, \tau\}}$. This downweights the contribution of regions where λ and κ tend to infinite values. We use a value of $\tau = 5$, although the optimization can tolerate a large range of values.

7.2. Acquisition Setup

The acquisition setup described in Section 6 is pictured in Figure 15. It consists of two light bulbs mounted close together at the end of a rigid shaft. The length of the shaft is set to 8 times the distance between the two lights, which results in a Δt of approximately 7 degrees. The shaft is connected to a wheel, which can be rotated using the crank. A camera is placed facing the object, aligned with the axis of rotation of the wheel.

The crank is moved by hand and rigidly clamped to a table while acquiring images, resulting in different unknown positions of the differential light pair. The differential image pairs are acquired by turning lights 1 and 2 on and off at a few different positions of the crank.




Figure 15. Acquisition setup.

8. Discussions

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. Theoretical analysis of such a framework is already the subject of our ongoing work.

Acknowledgments We gratefully acknowledge the resources provided by Prof. Marc Levoy towards data acquisition using the Stanford Spherical Gantry. We are thankful to Andrew Adams, at the Stanford Computer Graphics Laboratory, for his immense help with setting up the gantry. Many thanks to Miloš Hašan, at UC Berkeley, for help with rendering the bunny and Neil Alldrin, at Tandent Vision Science, for providing the isocontour tracing code of [1].

References

- [1] 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 high-lights and shadows. *PAMI*, 25(10):1239–1252, 2003.
- [4] J. Clark. Active photometric stereo. In *CVPR*, pages 29–34, 1992.
- [5] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming. <http://cvxr.com/cvx>, 2010.

- [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. In *SIGGRAPH ASIA*, 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*, 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. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Engg.*, 19(1):139–144, 1980.
- [14] R. Zhang, P. Tsai, J. Cryer, and M. Shah. Shape from shading: A survey. *PAMI*, 21(8):690–706, 1999.
- [15] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *IJCV*, 49(2/3):1215–1227, 2003.