

What Camera Motion Reveals About Shape With Unknown BRDF

Manmohan Chandraker
NEC Labs America, Cupertino, CA

Abstract

Psychophysical studies show motion cues inform about shape even with unknown reflectance. Recent works in computer vision have considered shape recovery for an object of unknown BRDF using light source or object motions. This paper addresses the remaining problem of determining shape from the (small or differential) motion of the camera, for unknown isotropic BRDFs. Our theory derives a differential stereo relation that relates camera motion to depth of a surface with unknown isotropic BRDF, which generalizes traditional Lambertian assumptions. Under orthographic projection, we show shape may not be constrained in general, but two motions suffice to yield an invariant for several restricted (still unknown) BRDFs exhibited by common materials. For the perspective case, we show that three differential motions suffice to yield surface depth for unknown isotropic BRDF and unknown directional lighting, while additional constraints are obtained with restrictions on BRDF or lighting. The limits imposed by our theory are intrinsic to the shape recovery problem and independent of choice of reconstruction method. We outline with experiments how potential reconstruction methods may exploit our theory. We illustrate trends shared by theories on shape from motion of light, object or camera, relating reconstruction hardness to imaging complexity.

1. Introduction

Image formation is an outcome of the interaction between shape, lighting and camera, governed by material reflectance. Motion of the object, light source or camera are important cues for recovering object shape from images. Each of those cues have been extensively studied in computer vision, under the umbrellas of optical flow for object motion [7, 9], photometric stereo for light source motion [19] and multiview stereo for motion of the camera [15]. Due to the complex and often unknown nature of the bidirectional reflectance distribution function (BRDF) that determines material behavior, simplifying assumptions like brightness constancy and Lambertian reflectance are often employed.

However, psychophysical studies have established that complex reflectance does not impede shape perception [16]. Correspondingly, recent works in computer vision show that differential motion of the light source [2] or the object [4] inform about shape even with unknown BRDFs. This paper solves the remaining problem of characterizing shape recovery for unknown BRDFs, using differential motion of the camera.

We begin in Sec. 3 by proposing a differential stereo relation that relates depth to camera motion, while accounting for general material behavior in the form of an isotropic BRDF. Diffuse photoconsistency of traditional Lambertian stereo follows as a special case. Surprisingly, it can be shown that considering a sequence of motions allows eliminating the BRDF dependence of differential stereo.

The mathematical basis for differential stereo is outwardly similar to differential flow for object motion [4]. However, while BRDFs are considered black-box functions in [4], our analysis explicitly considers the angular dependencies of isotropic BRDFs to derive additional insights. A particular benefit is to show that ambiguities exist for the case of camera motion, which render shape recovery more difficult.

Consequently, for orthographic projection, Sec. 4 shows a negative result whereby constraints on the shape of a surface with general isotropic BRDF may not be derived using camera motion as a cue. But we show the existence of an orthographic invariant for several restricted isotropic BRDFs, exhibited by common materials like plastics, metals, some paints and fabrics. The invariant is characterized as a quasilinear partial differential equation (PDE), which specifies the topological class up to which reconstruction may be performed.

Under perspective projection, Sec. 5 shows that depth for a surface with unknown isotropic BRDF, under unknown directional lighting, may be obtained using differential stereo relations from three or more camera motions. Further, we show that an additional linear constraint on the surface gradient is available for the above restricted families of BRDFs. These results substantially generalize Lambertian stereo, since depth information is obtained without assuming diffuse photoconsistency, while a weak assumption on material type yields even richer surface information. Table 1 summarizes the main theoretical results of this paper.

Finally, Sec. 6 explores relationships between differential theories pertaining to motion of the object, light source or camera. We discuss their shared traits that allow shape recovery and common trends relating hardness of surface reconstruction to complexity of imaging setup. Those perspectives on shape from motion are summarized in Table 2.

2. Related Work

Relating shape to intensity variations due to differential motion has a significant history in computer vision, dating to studies in optical flow [7, 9]. The limitations of the Lambertian assumption have been recognized by early works [11, 18].

Camera	Light	BRDF	#Motions	Surface Constraint	Shape Recovery	Theory
Persp.	Unknown	Lambertian	1	Linear eqn.	Depth	Stereo
Orth.	Unknown	General, unknown	-	Ambiguous for reconstruction		Prop. 3
Orth.	Unknown	View-angle, unknown	2	Homog. quasilinear PDE	Level curves	Rem. 1, Prop. 6
Orth.	Known	Half-angle, unknown	2	Inhomog. quasilinear PDE	Char. curves	Props. 4, 7
Persp.	Unknown	General, unknown	3	Linear eqn.	Depth	Prop. 8
Persp.	Unknown	View-angle, unknown	3	Lin. eqn. + Homog. lin. PDE	Depth + Gradient	Sec. 5.2
Persp.	Known	Half-angle, unknown	3	Lin. eqn. + Inhomog. lin. PDE	Depth + Gradient	Prop. 9

Table 1. Summary of main results of this paper. Note that k camera motions result in $k + 1$ images. BRDF is described by angular dependences and functional form. In general, more constrained BRDF or lighting yields richer shape information. Also, see discussion in Sec. 6.

Several stereo methods have been developed for non-Lambertian materials. Zickler et al. exploit Helmholtz reciprocity for reconstruction with arbitrary BRDFs [20]. Reference shapes of known geometry and various material types are used in example-based methods [17]. Savarese et al. study stereo reconstructions for specular surfaces [14]. In contrast, this paper explores how an image sequence derived from camera motion informs about shape with unknown isotropic BRDFs, regardless of reconstruction method.

Light source and object motions have also been used to understand shape with unknown BRDFs. Photometric stereo based on small light source motions is presented by [5], while generalized notions of optical flow are considered in [6, 12]. An isometric relationship between changes in normals and radiance profiles under varying light source is used by Sato et al. to recover shape with unknown reflectance [13].

Closely related to this paper are the works of Chandraker et al. that derive topological classes up to which reconstruction can be performed for unknown BRDFs, using differential motion of the source [2] or object [4]. This paper derives limits on shape recovery using the third cue, namely camera motion. The similarities and differences between the frameworks are explored throughout this paper and summarized in Sec. 6.

3. Differential Stereo for General BRDFs

In this section, we state our assumptions and derive the relationship between camera motion and surface depth, for unknown isotropic BRDFs. We also provide some intuition into that relationship, which will be used for subsequent shape recovery results. To facilitate presentation, we will occasionally point the reader to an extended version [1] for details.

3.1. Derivation of the Differential Stereo Relation

Assumptions and setup We assume static object and lighting, while the camera moves. For rigid body motion, our analysis equivalently considers a fixed camera, with the object and light source undergoing the inverse motion. The illumination is assumed directional and distant. The object BRDF is assumed to be homogeneous and isotropic, with unknown functional form. We make a technical assumption that the camera direction is constant over the entire object.¹ Global illumination effects are assumed negligible.

¹This assumption is exact for orthographic cameras, but only an approximation for perspective projection where viewing direction may vary over object dimensions. The approximation is reasonable in practical situations where the camera is not too close to the object, relative to object size.

Let the focal length of the camera be f . The camera model is perspective for finite values of f and approaches orthographic as $f \rightarrow \infty$. The principal point on the image plane is defined as the origin of the 3D coordinate system, with the camera center at $(0, 0, -f)^\top$. Denoting $\beta = f^{-1}$, a 3D point $\mathbf{x} = (x, y, z)^\top$ is imaged at $\mathbf{u} = (u, v)^\top$, where

$$(1 + \beta z)u = x, \quad (1 + \beta z)v = y. \quad (1)$$

Motion field Let the camera undergo a rigid body rotation $\tilde{\mathbf{R}}$ and translation $\tilde{\boldsymbol{\tau}}$. We equivalently assume that the object and light source undergo a rotation $\mathbf{R} = \tilde{\mathbf{R}}^{-1}$ and translation $\boldsymbol{\tau} = \tilde{\mathbf{R}}^{-1}\tilde{\boldsymbol{\tau}}$, with a fixed camera. For differential motion, we approximate $\mathbf{R} \approx \mathbf{I} + [\boldsymbol{\omega}]_\times$, where $\boldsymbol{\omega} = (\omega_1, \omega_2, \omega_3)^\top$.

The motion field, $\boldsymbol{\mu} = (\dot{u}, \dot{v})^\top$, is the differential motion of the image obtained by differentiating (1). We refer the reader to prior works like [11, 4] for a derivation and simply express the motion field here in a form similar to [4]:

$$\text{Perspective: } \boldsymbol{\mu} = \left(\alpha_1 + \frac{\alpha_2 + \omega_2 z}{1 + \beta z}, \alpha_3 + \frac{\alpha_4 - \omega_1 z}{1 + \beta z} \right)^\top, \quad (2)$$

$$\text{Orthographic: } \boldsymbol{\mu} = (\alpha_5 + \omega_2 z, \alpha_6 - \omega_1 z)^\top, \quad (3)$$

where $\alpha_i, i = 1, \dots, 6$, are known functions of $\boldsymbol{\omega}, \boldsymbol{\tau}, \mathbf{u}$ and β , whose algebraic forms are shown in [1].

Differential stereo relation Let \mathbf{s} be the light source direction and $\mathbf{v} = (0, 0, -1)^\top$ be the camera direction. For a 3D point $\mathbf{x} = (x, y, z(x, y))^\top$ on the surface, the unit normal is

$$\mathbf{n} = (n_1, n_2, n_3)^\top = (z_x^2 + z_y^2 + 1)^{-\frac{1}{2}}(z_x, z_y, -1)^\top, \quad (4)$$

where $\nabla z = (z_x, z_y)^\top$ is the surface gradient. For a homogeneous isotropic BRDF ρ , with distant light source, the image intensity at pixel \mathbf{u} of a 3D point \mathbf{x} is assumed to be

$$I(\mathbf{u}, t) = \sigma(\mathbf{x})\rho(\mathbf{x}, \mathbf{n}, \mathbf{s}, \mathbf{v}), \quad (5)$$

where σ is the albedo and the cosine fall-off is absorbed in ρ . This is a reasonable imaging model that subsumes traditional ones like Lambertian and allows general isotropic BRDFs modulated by spatially varying albedo. We do not make any assumptions on the functional form of ρ , except smoothness.

Taking the *total* derivative on both sides of (5), we get

$$I_u \dot{u} + I_v \dot{v} + I_t = \sigma \frac{d}{dt} \rho(\mathbf{x}, \mathbf{n}, \mathbf{s}, \mathbf{v}) + \rho \frac{d\sigma}{dt}. \quad (6)$$

Since σ is intrinsically defined on the surface coordinates, its total derivative vanishes. For the rigid body motion we

consider, $\dot{\mathbf{n}} = \boldsymbol{\omega} \times \mathbf{n}$ and $\dot{\mathbf{s}} = \boldsymbol{\omega} \times \mathbf{s}$, while the camera direction remains unchanged. Using chain rule differentiation and noting that $\boldsymbol{\mu} = (\dot{u}, \dot{v})^\top$ is the motion field, we have

$$(\nabla_{\mathbf{u}} I)^\top \boldsymbol{\mu} + I_t = \sigma [(\nabla_{\mathbf{x}} \rho)^\top \boldsymbol{\nu} + (\nabla_{\mathbf{n}} \rho)^\top (\boldsymbol{\omega} \times \mathbf{n}) + (\nabla_{\mathbf{s}} \rho)^\top (\boldsymbol{\omega} \times \mathbf{s})], \quad (7)$$

where $\boldsymbol{\nu} = \dot{\mathbf{x}}$ is the linear velocity. While the above discussion gives intuition for the differential relation in (7), we refer the reader to [1] for a rigorous derivation.

For distant lighting and homogeneous reflectance in our setup, we may assume that $\nabla_{\mathbf{x}} \rho$ is negligible. Further, dividing the two sides of (7) with those of (5), we get

$$\boxed{(\nabla_{\mathbf{u}} E)^\top \boldsymbol{\mu} + E_t = (\mathbf{n} \times \nabla_{\mathbf{n}} \log \rho + \mathbf{s} \times \nabla_{\mathbf{s}} \log \rho)^\top \boldsymbol{\omega}} \quad (8)$$

where we use the notation $E = \log I$ and the identities $(\nabla_{\mathbf{a}} \rho)^\top (\boldsymbol{\omega} \times \mathbf{a}) = (\mathbf{a} \times \nabla_{\mathbf{a}} \rho)^\top \boldsymbol{\omega}$ and $\nabla_{\mathbf{a}} \log \rho = \rho^{-1} \nabla_{\mathbf{a}} \rho$, for some $\mathbf{a} \in \mathbb{R}^3$. We call (8) the *differential stereo relation*.

3.2. Understanding the Differential Stereo Relation

Generalization of Lambertian stereo Initial intuition into the differential stereo relation of (8) may be derived by noting how it generalizes traditional Lambertian stereo. For two images I^1 and I^2 related by a known motion, Lambertian stereo seeks the depth z that best satisfies brightness constancy: $I^1(\mathbf{u}) = I^2(\mathbf{u} + \boldsymbol{\mu}(z))$. Substituting a Lambertian reflectance $\rho(\mathbf{n}, \mathbf{s}, \mathbf{v}) = \mathbf{n}^\top \mathbf{s}$ in (8), we get

$$(\nabla_{\mathbf{u}} E)^\top \boldsymbol{\mu} + E_t = (\mathbf{n} \times (\mathbf{n}^\top \mathbf{s})^{-1} \mathbf{s} + \mathbf{s} \times (\mathbf{n}^\top \mathbf{s})^{-1} \mathbf{n})^\top \boldsymbol{\omega} = \mathbf{0}^\top \boldsymbol{\omega} = 0, \quad (9)$$

which is precisely brightness constancy (total derivative of image intensity is zero). Thus, diffuse photoconsistency imposed by traditional stereo is a special case of our theory.

Relation to object motion For object motion, a differential flow relation is derived in [4]. The differential stereo relation of (8) has an additional dependence on BRDF derivatives with respect to the light source. This stands to reason, since both the object and lighting move relative to camera in the case of camera motion, while only the object moves in the case of object motion. Thus, intuitively, camera motion leads to a harder reconstruction problem. Indeed, as we will see next, the dependence of (8) on lighting leads to a somewhat surprising additional ambiguity.

Ambiguity for camera motion Now we derive some additional insights by making a crucial departure from the analysis of [4]. Namely, instead of treating isotropic BRDFs as entirely black box functions, we explicitly consider their physical property in the form of angular dependencies between the normal, light source and camera directions. We define

$$\boldsymbol{\pi} = \mathbf{n} \times \nabla_{\mathbf{n}} \log \rho + \mathbf{s} \times \nabla_{\mathbf{s}} \log \rho, \quad (10)$$

which allows rewriting the differential stereo relation (8) as:

$$(\nabla_{\mathbf{u}} E)^\top \boldsymbol{\mu} + E_t = \boldsymbol{\omega}^\top \boldsymbol{\pi}. \quad (11)$$

The entity $\boldsymbol{\pi}$ is central to our theory, since it captures the dependence of differential stereo on BRDF. Its practical significance is that any shape recovery method that seeks invariance to material behavior must either accurately model $\boldsymbol{\pi}$, or eliminate it. Our work adopts the latter approach.

This definition of $\boldsymbol{\pi}$ leads to an observation intrinsic to shape recovery with isotropic BRDFs:

Proposition 1. *The BRDF dependence of differential stereo is captured by a 2D vector in the principal plane of the camera.*

Proof. Since an isotropic BRDF depends on the three angles between normal, camera and light directions, we may write

$$\tilde{\rho}(\mathbf{n}^\top \mathbf{s}, \mathbf{s}^\top \mathbf{v}, \mathbf{n}^\top \mathbf{v}) = \log \rho(\mathbf{n}, \mathbf{s}, \mathbf{v}). \quad (12)$$

Denote $\theta = \mathbf{n}^\top \mathbf{s}$, $\phi = \mathbf{s}^\top \mathbf{v}$ and $\psi = \mathbf{n}^\top \mathbf{v}$. Then, applying chain-rule differentiation, we may write (10) as

$$\begin{aligned} \boldsymbol{\pi} &= \mathbf{n} \times \nabla_{\mathbf{n}} \tilde{\rho} + \mathbf{s} \times \nabla_{\mathbf{s}} \tilde{\rho} \\ &= \tilde{\rho}_\theta (\mathbf{n} \times \mathbf{s}) + \tilde{\rho}_\psi (\mathbf{n} \times \mathbf{v}) + \tilde{\rho}_\theta (\mathbf{s} \times \mathbf{n}) + \tilde{\rho}_\phi (\mathbf{s} \times \mathbf{v}) \\ &= \tilde{\rho}_\psi (\mathbf{n} \times \mathbf{v}) + \tilde{\rho}_\phi (\mathbf{s} \times \mathbf{v}). \end{aligned} \quad (13)$$

From the form of $\boldsymbol{\pi}$ in (13), it is evident that $\boldsymbol{\pi}^\top \mathbf{v} = 0$. For our choice of coordinate system, $\mathbf{v} = (0, 0, -1)^\top$. Thus,

$$\pi_3 = 0. \quad (14)$$

It follows that the BRDF-dependent entity $\boldsymbol{\pi} = (\pi_1, \pi_2, 0)^\top$ lies on the principal plane of the camera. \square

This is an important result that limits the extent to which shape may be recovered from differential stereo. The following sections explore the precise nature of those limits.

4. Orthographic Projection

Estimating the motion field, $\boldsymbol{\mu}$, is equivalent to determining dense correspondence and thereby, object shape. In this section, we consider shape recovery with unknown BRDF under orthographic projection, using a sequence of differential motions. Our focus will continue to be on providing intuitions, while referring the reader to [1] for details.

4.1. Rank Deficiency and Depth Ambiguities

It is clear that just one differential motion of the camera is insufficient to extract depth, since (11) is a linear relation in the multiple unknowns $\{z, \boldsymbol{\pi}\}$. Consequently, we consider a sequence of differential motions. We start by observing that, in the case of orthography, a rank deficiency similar to the case of object motion exists for camera motion too.

Under orthography, the motion field $\boldsymbol{\mu}$ is given by (3). Noting the μ_1 and μ_2 are linear in z , we observe that the differential stereo relation of (11) reduces to:

$$pz + q = \boldsymbol{\omega}^\top \boldsymbol{\pi} \quad (15)$$

where, using (3), p and q are known entities given by

$$p = \omega_2 E_u - \omega_1 E_v \quad (16)$$

$$q = \alpha_5 E_u + \alpha_6 E_v + E_t. \quad (17)$$

Consider $m > 0$ differential motions of the camera about a base position, given by $\{\omega^i, \tau^i\}$, for $i = 1, \dots, m$. Let $E^0 = \log I^0$ be the logarithm of the base image, with E^i being the log-image for each motion $\{\omega^i, \tau^i\}$. Note that the spatial gradient of the image is independent of motion and corresponds to the derivative of E^0 with respect to \mathbf{u} . We will simply denote it as $\nabla_{\mathbf{u}} E = (E_u, E_v)^\top$. The temporal derivative, E_t , as well as p and q , depend on the motion.

To recover the unknown depth z , an initial approach may consider $m \geq 3$ relations of the form (11) as a linear system:

$$\tilde{\mathbf{A}} \begin{bmatrix} z \\ \pi_1 \\ \pi_2 \end{bmatrix} = \mathbf{q}, \text{ with } \tilde{\mathbf{A}} = \begin{bmatrix} -p^1 & \omega_1^1 & \omega_2^1 \\ \vdots & \vdots & \vdots \\ -p^m & \omega_1^m & \omega_2^m \end{bmatrix}, \quad (18)$$

where $\mathbf{q} = (q^1, \dots, q^m)^\top$. Note that Proposition 1 allows us to drop any dependence on ω_3 , since $\pi_3 = 0$ is not an unknown. But observe the form of $p^i = \omega_2^i E_u - \omega_1^i E_v$ from (16), which makes $\tilde{\mathbf{A}}$ rank-deficient. Thus, we have shown:

Proposition 2. *Under orthographic projection, surface depth under unknown BRDF may not be unambiguously recovered using solely camera motion as the cue.*

While depth cannot be directly recovered from differential stereo under orthography, a natural next step is to consider the possibility of any constraints on the depth. However, the result of Proposition 1 makes this challenging for the case of camera motion. To see this, we note that the null vector of the rank-deficient matrix $\tilde{\mathbf{A}}$ is $(1, -E_v, E_u)^\top$. With $\tilde{\mathbf{A}}^+$ denoting the Moore-Penrose pseudoinverse of $\tilde{\mathbf{A}}$, for any $k \neq 0$, we have a parameterized solution to (18):

$$\begin{bmatrix} z \\ \pi_1 \\ \pi_2 \end{bmatrix} = \gamma + k \begin{bmatrix} 1 \\ -E_v \\ E_u \end{bmatrix}, \quad (19)$$

where $\gamma = (\gamma_1, \gamma_2, \gamma_3)^\top = \tilde{\mathbf{A}}^+ \mathbf{q}$. From the first equation in the above system, we have $k = z - \gamma_1$. Thereby, we get the following two relations between z and π :

$$\pi_1 = (\gamma_2 + E_v \gamma_1) - E_v z \quad (20)$$

$$\pi_2 = (\gamma_3 - E_u \gamma_1) + E_u z \quad (21)$$

Now, any relationship between π_1 and π_2 gives a constraint on the depth, z . But from Prop. 1, π is an arbitrary vector in the principal plane. That is, from (13), π_1 and π_2 depend on two unknown BRDF derivatives $\tilde{\rho}_\psi$ and $\tilde{\rho}_\phi$. It follows that without imposing any external constraint on $\tilde{\rho}_\psi$ and $\tilde{\rho}_\phi$, one may not derive a constraint on surface depth. Thus, we state:

Proposition 3. *Under orthographic projection, for unknown isotropic BRDF, an unambiguous constraint on surface depth may not be derived using solely camera motion as the cue.*

The above is an example of the comparative limits on shape recovery with object or camera motion. For object motion, only the relative motion between camera and object must be accounted. Thus, the definition of $\pi = \mathbf{n} \times \nabla_{\mathbf{n}} \log \rho$ in [4] depends only on the BRDF-derivative with respect to surface normal. For camera motion, both object and lighting move relative to the camera. Additional dependence of π in (10) on BRDF-derivative with respect to lighting makes it indeterminate without further restrictions on BRDF or lighting.

While Prop. 3 is a negative result, its development provides valuable insight. An $m \times 3$ rank-deficient matrix $\tilde{\mathbf{A}}$ constructed using m differential motions, for any $m \geq 2$, determines $\gamma = \tilde{\mathbf{A}}^+ \mathbf{q}$. From (14), (20) and (21), it follows:

Corollary 1. *Under orthography, two differential motions of the camera suffice to constrain π to a linear relation in z .*

This fact will now be used, along with some restrictions on the BRDF, to derive BRDF-invariant constraints on depth.

4.2. BRDF-Invariance for Certain Material Types

The result of Prop. 3 immediately suggests a possibility to constrain z : consider a BRDF whose dependence on ψ and ϕ is restricted in a way that introduces a constraint on π . Note that the functional form of the BRDF, $\rho(\cdot)$, remains unknown.

4.2.1 BRDFs Dependent on View Angle

Some reflectances depend on the angles subtended by the normal on the source and view directions. Such BRDFs can explain the darkening near image edges for materials like fabrics. Another well-known example is the Minnaert BRDF for lunar reflectance [10]. In such cases, we may define

$$\bar{\rho}(\mathbf{n}^\top \mathbf{s}, \mathbf{n}^\top \mathbf{v}) = \log \rho(\mathbf{n}, \mathbf{s}, \mathbf{v}). \quad (22)$$

Again denoting $\theta = \mathbf{n}^\top \mathbf{s}$ and $\psi = \mathbf{n}^\top \mathbf{v}$, we get from (10):

$$\begin{aligned} \pi &= \mathbf{n} \times \nabla_{\mathbf{n}} \bar{\rho} + \mathbf{s} \times \nabla_{\mathbf{s}} \bar{\rho} \\ &= \mathbf{n} \times (\bar{\rho}_\theta \mathbf{s} + \bar{\rho}_\psi \mathbf{v}) + \mathbf{s} \times \bar{\rho}_\theta \mathbf{n} = \bar{\rho}_\psi \mathbf{n} \times \mathbf{v}. \end{aligned} \quad (23)$$

Noting that $\mathbf{n} \times \mathbf{v} = (-n_2, n_1, 0)^\top$, one may eliminate the BRDF-dependent term $\bar{\rho}_\psi$ using (20) and (21) to obtain a relationship between depths and normals:

$$\frac{\pi_1}{\pi_2} = \frac{-n_2}{n_1} = \frac{(\gamma_2 + E_v \gamma_1) - E_v z}{(\gamma_3 - E_u \gamma_1) + E_u z}. \quad (24)$$

Using (4) to relate the normal to the gradient, this reduces to

$$[(\gamma_2 + E_v \gamma_1) - E_v z] z_x + [(\gamma_3 - E_u \gamma_1) + E_u z] z_y = 0, \quad (25)$$

which is a constraint on surface depth and gradient that is independent of both the BRDF and lighting. We note that $m \geq 2$ differential motions of the camera suffice to determine γ from (19) and yield the constraint in (25). Thus, we state:

Remark 1. *Under orthography, for a BRDF of unknown functional form that depends on light and view directions, two differential motions of the camera suffice to yield a constraint on surface depth independent of BRDF and lighting.*

4.2.2 BRDFs Dependent on Half-angle

For many common materials like metals or plastics, it is reasonable to assume that reflectance depends on the angle between the surface normal and the half-angle between the source and view directions. For a surface of such material type, we can show that a sequence of differential stereo relations yields a BRDF-invariant constraint on surface depth. For this case, we assume a known light source direction.

Proposition 4. *Under orthographic projection, for a BRDF of unknown functional form that depends on known light and half-angle directions, two differential motions of the camera suffice to yield a BRDF-invariant constraint on surface depth.*

Proof. For a BRDF that depends on half-angle \mathbf{h} , we define

$$\log \rho(\mathbf{n}, \mathbf{s}, \mathbf{v}) = \bar{\rho}(\mathbf{n}^\top \mathbf{s}, \mathbf{n}^\top \mathbf{h}), \text{ where } \mathbf{h} = \frac{\mathbf{s} + \mathbf{v}}{\|\mathbf{s} + \mathbf{v}\|}. \quad (26)$$

The definition of $\boldsymbol{\pi}$ in (10) may now be rewritten as

$$\boldsymbol{\pi} = \mathbf{n} \times \nabla_{\mathbf{n}} \bar{\rho} + \mathbf{s} \times \nabla_{\mathbf{s}} \bar{\rho}. \quad (27)$$

We again denote $\theta = \mathbf{n}^\top \mathbf{s}$, $\phi = \mathbf{s}^\top \mathbf{v}$ and define $\eta = \mathbf{n}^\top \mathbf{h}$. Then, using the definition of \mathbf{h} in (26), we have after some algebraic simplification:

$$\boldsymbol{\pi} = \bar{\rho}_\eta \left[\frac{\mathbf{n} \times \mathbf{v}}{\|\mathbf{s} + \mathbf{v}\|} - \frac{(\mathbf{n}^\top \mathbf{h}) \mathbf{s} \times \mathbf{v}}{\|\mathbf{s} + \mathbf{v}\|^2} \right]. \quad (28)$$

The symmetry of $\boldsymbol{\pi}$ with respect to \mathbf{n} and \mathbf{s} means it is independent of ρ_θ . Recall that $\pi_3 = 0$ from (14). After some algebra and using (4) to relate \mathbf{n} to gradient, we get from (28):

$$\frac{\pi_1}{\pi_2} = \frac{\mathbf{a}^\top \mathbf{n}}{\mathbf{b}^\top \mathbf{n}} = \frac{a_1 z_x + a_2 z_y - a_3}{b_1 z_x + b_2 z_y - b_3}, \quad (29)$$

where $a_1 = s_2 h_1$, $a_2 = s_2 h_2 - \sqrt{2(1 + \phi)}$, $a_3 = s_2 h_3$ and $b_1 = \sqrt{2(1 + \phi)} - s_1 h_1$, $b_2 = -s_1 h_2$, $b_3 = -s_1 h_3$. Thus, we have obtained a relationship on π independent of $\bar{\rho}_\eta$.

Finally, for $m \geq 2$ differential camera motions, we have from Corollary 1 that π_1 and π_2 are linear in z . Substituting their expressions from (20) and (21) into (29), we get

$$(\lambda_1 + \lambda_2 z) z_x + (\lambda_3 + \lambda_4) z_y + \lambda_5 = 0, \quad (30)$$

where λ_i , for $i = 1, \dots, 5$, are known entities. Complete expressions for λ_i are provided in [1]. Here, it suffices to note that the λ_i are independent of $\bar{\rho}$, thus, (30) is a BRDF-invariant constraint on surface depth. \square

Intuitively, the ambiguity of Prop. 3 exists since one cannot constrain π_1 and π_2 if they depend on two independent unknowns ρ_ϕ and ρ_ψ . Considering an appropriate BRDF, such as a half-angle one, allows expressing $\boldsymbol{\pi}$ in terms of a single unknown ρ_η . The linearity of differentiation eliminates ρ_η to yield a BRDF-invariant constraint on $\boldsymbol{\pi}$. Thus, Prop. 4 can derive an invariant purely in terms of depth and gradient.

Note that Proposition 1 is a basic property of isotropic BRDFs under camera motion. So, it may be verified that it holds true even in the restricted case of half-angle BRDFs, that is, $\boldsymbol{\pi}^\top \mathbf{v} = -\pi_3 = 0$ even for the $\boldsymbol{\pi}$ defined by (28).

4.2.3 Dependence on Arbitrary Angle in $\{\mathbf{s}, \mathbf{v}\}$ -Plane

Recent works on empirical analysis of measured BRDFs show that reflectance functions often depend on the angles the surface normal makes with the light source and another direction in the plane defined by the source and camera directions [3]. In such cases, we may write

$$\log \rho(\mathbf{n}, \mathbf{s}, \mathbf{v}) = \bar{\rho}(\mathbf{n}^\top \mathbf{s}, \mathbf{n}^\top \mathbf{y}), \text{ with } \mathbf{y} = \frac{\mathbf{s} + \kappa \mathbf{v}}{\|\mathbf{s} + \kappa \mathbf{v}\|}, \quad (31)$$

for some $\kappa \in \mathbb{R}$. Note that half-angle BRDFs for materials like plastics and metals considered in Section 4.2.2 are a special case with $\kappa = 1$. The BRDFs considered in Section 4.2.1 may also be considered a limit case with $\kappa \gg 1$. Empirical examples for BRDFs with finite $\kappa \neq 1$ are shown for materials like paints and fabrics in [3]. We may now state:

Proposition 5. *Under orthographic projection, for a BRDF of unknown functional form that depends on light source and an arbitrary direction in the source-view plane, two differential motions of the camera suffice to yield a BRDF-invariant constraint on surface depth.*

Proof. The proof directly generalizes the development in Proposition 4. We refer the reader to [1] for complete details. We only note here that we obtain:

$$\boldsymbol{\pi} = \frac{\kappa \rho_{\mathbf{n}^\top \mathbf{y}}}{1 + \kappa^2 + 2\kappa\phi} \left[(1 + \kappa^2 + 2\kappa\phi)^{\frac{1}{2}} \mathbf{n} - (\mathbf{n}^\top \mathbf{y}) \mathbf{s} \right] \times \mathbf{v}, \quad (32)$$

thus, dependence on ρ may be eliminated similar to (29) using $\pi_3 = 0$ and considering the ratio of π_1 and π_2 . We then invoke Corollary 1 to constrain π_1 and π_2 to linear functions in z using $m \geq 2$ differential motions, yielding an invariant:

$$(\lambda'_1 + \lambda'_2 z) z_x + (\lambda'_3 + \lambda'_4) z_y + \lambda'_5 = 0, \quad (33)$$


where λ'_i , for $i = 1, \dots, 5$, are again known entities. \square

We urge the reader to observe that the constraint (33) is invariant to the functional form of the BRDF, but is not independent of κ . However, note that κ can be estimated from image data without requiring a full BRDF estimation, for instance, using the methods proposed in [3]. Also, we again note that Proposition 1 is an intrinsic property of isotropic BRDFs and $\pi_3 = 0$ continues to hold even for the $\boldsymbol{\pi}$ in (32).

4.3. Results on Surface Estimation

Recall that Sec. 4.1 establishes that, for general unknown isotropic BRDFs, one may neither estimate the surface depth, nor derive any BRDF-invariant constraints on it. However, Sec. 4.2 derives constraints on depth and gradient for restricted (but unknown) BRDFs. Here, we characterize the PDE defined by those constraints, which directly informs the extent to which shape may be recovered.

For a BRDF of the form (22), an invariant of the form (25) is obtained. We note that (25) is characterized as a *homogeneous first-order quasilinear PDE* [8]. For a surface level curve $z(x, y) = z_0$, the solution to (25) from PDE theory is:



(a) Input [1 of 3] (b) Level curves (c) Reconstruction

Figure 1. (a) One of three simulated orthographic images (two motions), with unknown lighting and arbitrary non-Lambertian BRDF dependent on source and view angles. (b) Level curves estimated using Prop. 6. (c) Surface reconstructed after interpolation.

$$z = z_0, \quad \frac{dy}{dx} = \frac{(\gamma_3 - E_u \gamma_1) + E_u z_0}{(\gamma_2 + E_v \gamma_1) - E_v z_0}. \quad (34)$$

That is, given the depth at a point, the ODE (34) defines a step in the tangent direction, thereby tracing out the level curve through that point. We refer the reader to [1] for a formal proof, while only stating the result here:

Proposition 6. *Under orthography, two or more differential motions of the camera yield level curves of depth for a surface with BRDF dependent on light source and view angles.*

For half-angle BRDFs given by (26), a BRDF-invariant constraint on surface depth is obtained as (30). We note that it is characterized as an *inhomogeneous first-order quasilinear PDE*, whose solution is also available from PDE theory [8]. Here, with λ_i as defined in (30), we again state the result while referring the reader to [1] for a proof:


Proposition 7. *Under orthography, for a surface with half-angle BRDF, two or more differential camera motions yield characteristic surface curves $\mathcal{C}(x(s), y(s), z(s))$ defined by*

$$\frac{1}{\lambda_1 + \lambda_2 z} \frac{dx}{ds} = \frac{1}{\lambda_3 + \lambda_4 z} \frac{dy}{ds} = \frac{-1}{\lambda_5} \frac{dz}{ds} \quad (35)$$

corresponding to depths at some (possibly isolated) points.

Finally, for a BRDF dependent on an arbitrary angle in the source-view plane given by (31), the invariant (33) is also an inhomogeneous quasilinear PDE. So, as a direct generalization of Prop. 7, differential stereo also yields characteristic curves for this case, with λ'_i of (33) instead of λ_i .

Shape recovery Given depths at a few points on a surface with unknown BRDF, the above propositions yield depths along certain characteristic curves. For a smooth surface, one may interpolate the depths between the curves, in order to recover depth for the whole surface. The procedure is shown for synthetic data in Fig. 1 for a BRDF that depends on source and view angles (unknown lighting) and in Fig. 2 for unknown half-angle BRDFs (known lighting). Depth is assumed known at the green points to yield characteristic curves shown in red. We note that reconstruction methods in practice may use tracked feature points as seeds.



(a) Input [1 of 3] (b) Characteristics (c) Reconstruction

Figure 2. (a) One of three synthetic images (two motions), with unknown half-angle BRDF and known lighting, under orthographic projection. (b) Characteristic curves estimated using Prop. 7. (c) Surface reconstructed after interpolation.

5. Perspective Projection

In this section, we consider shape recovery from differential stereo under perspective projection. In particular, we show that unlike the orthographic case, depth may be unambiguously recovered in the perspective case, even when both the BRDF and lighting are unknown.

5.1. Depth from Differential Stereo

In the perspective case, the motion field μ is given by (2). Substituting for μ in the differential stereo relation (8), we obtain its expression for the perspective case:

$$p' \left(\frac{z}{1 + \beta z} \right) + r' \left(\frac{1}{1 + \beta z} \right) + q' = \omega^\top \pi, \quad (36)$$

where π is defined as before by (10) and $p' = E_u \omega_2 - E_v \omega_1$, $q' = \alpha_1 E_u + \alpha_3 E_v + E_t$ and $r' = \alpha_2 E_u + \alpha_4 E_v$ are known entities. Unlike the orthographic case, the differential stereo relation is not linear in $\{z, \pi\}$ for perspective projection.

We are now in a position to show depth recovery:

Proposition 8. *Under perspective projection, three differential motions of the camera suffice to yield depth of a surface with unknown isotropic BRDF and unknown light source.*

Proof. For $m \geq 3$, let images E^1, \dots, E^m be related to a base image E^0 by known differential motions $\{\omega^i, \tau^i\}$. Then, we obtain from (36) a sequence of differential stereo relations:

$$(p'^i + \beta q'^i)z - ((1 + \beta z)\pi)^\top \omega^i + (q'^i + r'^i) = 0, \quad (37)$$

for $i = 1, \dots, m$. Let $\tilde{\mathbf{C}} = [\mathbf{c}^1, \dots, \mathbf{c}^m]^\top$ be the $m \times 3$ matrix with rows $\mathbf{c}^i = [-(p'^i + \beta q'^i), \omega_1^i, \omega_2^i]^\top$. Further, let $\mathbf{q}' = [q'^1, \dots, q'^m]^\top$ and $\mathbf{r}' = [r'^1, \dots, r'^m]^\top$. Then, the system of differential stereo relations (37) may be written as

$$\tilde{\mathbf{C}} \begin{bmatrix} z \\ (1 + \beta z)\pi_1 \\ (1 + \beta z)\pi_2 \end{bmatrix} = \mathbf{q}' + \mathbf{r}', \quad (38)$$

since $\pi_3 = 0$, from Prop. 1. With $\tilde{\mathbf{C}}^+$ as the Moore-Penrose pseudoinverse of $\tilde{\mathbf{C}}$, let $\epsilon = (\epsilon_1, \epsilon_2, \epsilon_3)^\top = \tilde{\mathbf{C}}^+(\mathbf{q}' + \mathbf{r}')$. Then, (38) has the solution:

$$[z, (1 + \beta z)\pi_1, (1 + \beta z)\pi_2]^\top = \epsilon. \quad (39)$$

It follows that $z = \epsilon_1$ yields the surface depth. \square

Thus, Prop. 8 yields surface depth from differential stereo when both the BRDF and lighting are unknown. Again, a comparison is merited to the case of object motion in [4]. With object motion, depth and an additional constraint on the gradient are recovered with unknown BRDF and lighting. While camera motion also recovers depth, following Prop. 1, it is likely that further information on the gradient is recoverable only with additional constraints on the BRDF or lighting.

5.2. Additional Constraints for Certain Materials

As indicated above, now we show that additional constraints on the gradient are available for several material types.

Proposition 9. *Under perspective projection, three or more differential motions of the camera suffice to yield both depth and a linear constraint on the gradient for a surface with BRDF dependent on known light and half-angle directions.*

Proof. Prop. 8 already shows depth recovery from $m \geq 3$ differential motions of the camera. Further, from the form of π for half-angle BRDFs in (28) and using (39), we have

$$\frac{\pi_1}{\pi_2} = \frac{\mathbf{a}^\top \mathbf{n}}{\mathbf{b}^\top \mathbf{n}} = \frac{\epsilon_2}{\epsilon_3}, \quad (40)$$

where \mathbf{a} and \mathbf{b} are known entities dependent on light, defined by (29). Finally, using (4) yields a linear PDE in depth:

$$l_1 z_x + l_2 z_y + l_3 = 0, \quad (41)$$

with $l_1 = \epsilon_2 b_1 - \epsilon_3 a_1$, $l_2 = \epsilon_2 b_2 - \epsilon_3 a_2$, $l_3 = \epsilon_3 a_3 - \epsilon_2 b_3$. Thus, we obtain a linear constraint on the gradient. \square

It is evident from the above result that materials for which a relation between π_1 and π_2 may be derived independent of derivatives of the BRDF ρ , yield an additional constraint on the surface gradient. The form of this constraint will mirror the relationship between π_1 and π_2 , since their ratio is a measurable quantity from (40). In particular, it follows that a linear constraint may also be derived for the more general case considered in Section 4.2.3:

Remark 2. *Under perspective projection, three differential motions of the camera suffice to yield both depth and a linear constraint on the gradient for a BRDF dependent on known light and an arbitrary direction in the source-view plane.*

5.3. Shape Recovery

From Prop. 8, depth is directly available in the perspective case for unknown lighting and unknown BRDF. Fig. 3 illustrates this for synthetic data, where the object is imaged under an unknown light source. The object has diameter 20cm and is placed 1.5m away from the camera of focal length 10cm. Three random motions, of approximately 2° rotation and 5mm translation are imparted to the camera. The recovered shape using the theory of Sec. 5.1 is shown in Fig. 3. No prior knowledge of the BRDF, lighting or depth are required.




Figure 3. (a) One of four synthetic images (three motions), with arbitrary non-Lambertian BRDF and unknown lighting, under perspective projection. (b,c) π_1 and π_2 recovered using (43), as discussed in Sec. 6. (d) Depth estimated using Prop. 8.




Figure 4. Reconstruction for a non-Lambertian ball with an unknown BRDF, under unknown light source, using Proposition 8.

For evaluation with real data, images of the plastic sphere in Fig. 4(a) are obtained against a textured background under unknown directional lighting. Note the clear non-Lambertian effects. The camera focal length is 55mm and the object is approximately 80cm away. Five small motions (and several large ones for stable pose optimization) are imparted to the camera and estimated using bundle adjustment. The surface depth obtained from Prop. 8 is shown in Fig. 4(b) and (c).

While we do not explore the direction in this paper, note that with known lighting, one may use the constraint (41) to solve a joint depth and gradient optimization weighted by λ :

$$\min_z (z - \epsilon_1)^2 + \lambda(l_1 z_x + l_2 z_y + l_3)^2. \quad (42)$$

With standard differencing, the above is a highly sparse linear system in z which may be solved efficiently.

6. Perspectives on Shape from Motion Theories

We now provide a unified perspective on shape from motion theories corresponding to light, object or camera motions. Despite the apparent complexity of material behavior, differential motion of light, object or camera allows shape recovery with unknown BRDF and often unknown lighting. More importantly, theoretical limits on shape from motion are also derivable in these frameworks. Prior works have presented theories for light [2] and object [4] motions. This paper has studied camera motion to complete the differential analysis framework for shape recovery with unknown BRDFs.

In Table 2, we summarize a few results from each theory. In each case, the theories generalize well-known special cases that assume Lambertian BRDF or brightness constancy. Only orthographic projection for light motion (generalizing photometric stereo) and perspective projection for object or camera motion (generalizing optical flow and multiview stereo, respectively) are shown. A complete table is provided in [1].

A few important traits are shared by these theories:

Motion	Light	BRDF	#Motions	Surface Constraint	Shape Recovery	Theory
Light	Known	Lambertian	2	Lin. eqns. on gradient	Gradient	[19]
Light	Unknown	General, unknown	2	Linear PDE	Level curves + $\ \nabla z\ $ isocontours	[2]
Object	Brightness constancy		1	Lin. eqn. (optical flow)	Depth	[7, 9]
Object	Colocated	Uni-angle, Unknown	3	Lin. eqn. + Homog. lin. PDE	Depth + Gradient	[4]
Object	Unknown	Full, unknown	4	Lin. eqn. + Homog. lin. PDE	Depth + Gradient	[4]
Camera	Brightness constancy		1	Lin. eqn. (stereo)	Depth	[15]
Camera	Unknown	General, unknown	3	Linear eqn.	Depth	Prop. 8
Camera	Known	Half angle, unknown	3	Lin. eqn. + Inhomog. lin. PDE	Depth + Gradient	Prop. 9

Table 2. A unified look at BRDF-invariant frameworks on shape recovery from motions of light source, object or camera. In each case, PDE invariants specify precise topological limits on shape recovery. More general BRDFs or imaging setups require greater number of motions to derive the reconstruction invariant. More restricted BRDFs require fewer motions or yield richer shape information. For comparison, the traditional diffuse equivalents are also shown. Only the perspective cases are shown for object and camera motion (the full table is in [1]).

- They all rely on the linearity of chain rule differentiation to eliminate the BRDF from a system of equations.
- The invariant in each case can be characterized as a PDE amenable to solution through standard analysis tools.
- The involved PDEs provide intrinsic limits on the topological class upto which shape may be recovered from each motion cue, regardless of reconstruction method.
- More general imaging setups require greater number of motions for shape recovery. For instance, colocated lighting requires less motions than general lighting, or perspective projections require more motions than orthographic.
- Constraining the BRDF either reduces the minimum requirement on number of motions (colocated or general BRDFs for object motion), or provides richer shape information (half-angle or general BRDFs for camera motion).

The cases of object and camera motion are more closely related, but with important differences due to additional ambiguities entailed by camera motion. Qualitatively, this leads to a harder problem for the case of camera motion. The practical manifestation of this hardness is requiring a more restricted BRDF (although still unknown) to obtain the same shape information. For instance, a half-angle BRDF yields depth and a gradient constraint with camera motion, while the same can be obtained with general BRDFs for object motion. Allowing a general BRDF means only depth may be obtained for camera motion, while object motion yields an additional gradient constraint. Throughout this paper, we have highlighted such distinctions and their impact on shape recovery.

Future work Our theory focuses on shape recovery, but some BRDF information may also be recovered as a by-product. For instance, with perspective projection, Props. 8 and 1 completely define π . That is, from (39),

$$\pi = (1 + \beta\epsilon_1)^{-1} (\epsilon_2, \epsilon_3, 0)^\top. \quad (43)$$

Fig. 3 shows an example recovery of π . In turn, this places constraints on the derivative of BRDF. An avenue for future work is to characterize the extents to which BRDF may be recovered using motions of the light source, object or camera.

Further, while [2, 4] and this paper together provide limits on shape recovery from motions corresponding to each imaging element, an interesting problem is to achieve similar limits when lighting, object and camera all undergo simultaneous

motion. We also anticipate robust estimation methods that exploit our theory to extend traditional implementations of optical flow and multiview stereo to handle general BRDFs.

References

- [1] M. Chandraker. What camera motion reveals about shape with unknown BRDF. Technical Report 2014-TR044r, NEC Labs America, 2014.
- [2] M. Chandraker, J. Bai, and R. Ramamoorthi. A theory of differential photometric stereo for unknown isotropic BRDFs. In *CVPR*, pages 2505–2512, 2011.
- [3] M. Chandraker and R. Ramamoorthi. What an image reveals about material reflectance. In *ICCV*, pages 1076–1083, 2011.
- [4] M. Chandraker, D. Reddy, Y. Wang, and R. Ramamoorthi. What object motion reveals about shape with unknown BRDF and lighting. In *CVPR*, pages 2523–2530, 2013.
- [5] J. Clark. Active photometric stereo. In *CVPR*, pages 29–34, 1992.
- [6] H. W. Haussecker and D. J. Fleet. Computing optical flow with physical models of brightness variation. *PAMI*, 23(6):661–673, 2001.
- [7] B. Horn and B. Schunck. Determining optical flow. *Artificial Intelligence*, 17:185–203, 1981.
- [8] F. John. *Partial Differential Equations*. App. Math. Sci. Springer, 1981.
- [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. Minnaert. The reciprocity principle in lunar photometry. *Astrophysical Journal*, 3:403–410, 1941.
- [11] H.-H. Nagel. On a constraint equation for the estimation of displacement rates in image sequences. *PAMI*, 11(1):13–30, 1989.
- [12] S. Negahdaripour. Revised definition of optical flow: Integration of radiometric and geometric cues for dynamic scene analysis. *PAMI*, 20(9):961–979, 1998.
- [13] I. Sato, T. Okabe, Q. Yu, and Y. Sato. Shape reconstruction based on similarity in radiance changes under varying illumination. In *ICCV*, pages 1–8, 2007.
- [14] S. Savarese, M. Chen, and P. Perona. Local shape from mirror reflections. *IJCV*, 64(1):31–67, 2005.
- [15] S. Seitz, B. Curless, J. Diebel, D. Scharstein, and R. Szeliski. A comparison and evaluation of multiview stereo reconstruction algorithms. In *CVPR*, pages 519–526, 2006.
- [16] J. T. Todd, J. F. Norman, J. J. Koenderink, and A. M. L. Kappers. Effects of texture, illumination and surface reflectance on stereoscopic shape perception. *Perception*, 26(7):807–822, 1997.
- [17] A. Treuille, A. Hertzmann, and S. Seitz. Example-based stereo with general BRDFs. In *ECCV*, pages 457–469, 2004.
- [18] A. Verri and T. Poggio. Motion field and optical flow: Qualitative properties. *PAMI*, 11(5):490–498, 1989.
- [19] P. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Engg.*, 19(1):139–144, 1980.
- [20] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *IJCV*, 49(2/3):1215–1227, 2002.