

Reflections on the Generalized Bas-Relief Ambiguity

Manmohan Krishna Chandraker

Fredrik Kahl

David J. Kriegman

Department of Computer Science and Engineering
University of California, San Diego

Abstract

Prior work has argued that when a Lambertian surface in fixed pose is observed in multiple images under varying distant illumination, there is an equivalence class of surfaces given by the generalized bas-relief (GBR) ambiguity that could have produced these images. In contrast, this paper shows that for general nonconvex surfaces, interreflections completely resolve the GBR ambiguity. In turn, the full Euclidean geometry can be recovered from uncalibrated photometric stereo for which the light source directions and strengths are unknown. Further, we show that surfaces with a translational symmetry do not lend enough constraints to be disambiguated by interreflections.

1. Introduction

It has long been established that interreflections between surface patches form a vital component of scene inference [7]. The concept of an interreflection kernel, which explains the physics of mutual illumination, has existed in literature pertaining to the field of radiometry for over half a century [13]. Efficient algorithms to find approximate solutions to the rendering equation [10] (which is a manifestation of interreflections in real surfaces), such as radiosity methods [2], form a cornerstone of computer graphics research.

But curiously, an overwhelming majority of computer vision algorithms simply ignore the effect of interreflections. Even strongly intensity dependent techniques such as photometric stereo [17] and shape from shading [9] fail to substantively account for the augmentation to observed image intensity due to mutual illumination. The practical difficulty encountered arises due the complex structure of the kernel - each point in the range of the kernel encodes the geometric and radiometric structure of the entire domain.

Thus, while interreflection effects are deemed significant in real imaging situations and treated as a complex nuisance, it almost seems *non-sequitur* that they can actually




Figure 1. The first column shows a few synthetic input images. The second column shows the result of traditional uncalibrated photometric stereo reconstruction, which can be disambiguated only up to a GBR-transformation. Two such ambiguous reconstructions are shown. The final column shows the true reconstruction obtained by our algorithm where the GBR ambiguity is resolved by considering interreflections. See Section 4.3.

aid the surface reconstruction process in any manner. Yet, we show that incorporating mutual illumination into our image formation model makes Euclidean reconstruction from uncalibrated photometric stereo possible.

It has been shown [1, 12] that without interreflections, uncalibrated photometric stereo can recover a Lambertian surface illuminated by distant point light sources only up to a generalized bas-relief (GBR) transformation. (If an object's true surface, viewed under orthographic projection, is $z = f(x, y)$, then the GBR-transformed surface is $\tilde{f} = \lambda f(x, y) + \mu x + \nu y$, which can be considered as a scaling along the viewing direction composed with an additive plane. See Figure 1.)

The presence of interreflections results in counterintuitive shading effects in images, such as Figure 2. Thus, when interreflections are present but ignored, calibrated photometric stereo with known light sources can only recover a pseudo-normal field. Hence, uncalibrated photo-




Figure 2. Intensity profile across a horizontal cross-section of a concave dihedral angle. One might expect intensity to fall at a greater distance from light source, but note the sharp rise in intensity towards the center due to interreflections.

metric stereo in the presence of interreflections yields a GBR-transformation of the "pseudo-surface" (a term we will use to denote the surface constituted by the integrable normal field closest to the pseudo-normal field).

While a rigorous proof is presented in subsequent sections, the intuition behind our claim that interreflections resolve the GBR ambiguity is a simple one : distance is not an invariant for the group of GBR-transformations, whereas interreflections are strongly distance-dependent phenomena. Though the bas-relief ambiguity is resolved for general 3D surfaces using interreflections, we prove that a non-linear ambiguity remains in the case of surfaces with translational symmetry. Our experiments, designed to accommodate the sheer size of the interreflection kernel, validate these claims.

2. Background

2.1. Previous Work

Formalization of the contribution of interreflections towards the image of a parametrizable surface can be achieved through a Fredholm integral equation of the second kind, which can be solved through the Neumann series [15]. A finite element method to solve the forward problem was presented for surfaces with translational symmetry in [4]. We show that the inverse problem is ill-constrained for such surfaces. The extent to which some vision algorithms like shape from shading and edge detection err in ignoring interreflections is discussed in [5].

An alternate solution approach is to expand the radiance function in terms of eigenfunctions of the interreflection kernel, which are termed "pseudo-facets" in [11]. The discrete representation of the interreflection kernel was presented in [14] where shape was recovered from calibrated photometric stereo in the presence of interreflections. Shape from shading with a simplified (single-bounce) interreflection model is studied in [16].

Further generalization of the GBR ambiguity is achieved by the KGBR viewpoint-lighting ambiguity [18], but interreflections are ignored. Consideration of non-Lambertian reflectance models that include specularities can also be used to resolve the GBR ambiguity [3, 6]. A specularity imposes a constant viewpoint constraint that requires all light sources reflected along the viewing direction which reduces

the ambiguity. The assumption of a constant albedo was shown to resolve the ambiguity in [1, 8]. We resolve the GBR ambiguity using interreflections for a purely Lambertian surface, without any further modifications to the reflectance model, so in some sense, our approach is thoroughly uncalibrated photometric stereo.

2.2. Notation

We denote by (\mathcal{S}, a) the geometrical shape of a surface in \mathbb{R}^3 and its albedo field being imaged by an orthographic camera with its image plane and the global coordinate system mutually aligned. Hence \mathcal{S} consists of points of the form $[x, y, f(x, y)]^\top = [\mathbf{x}, f(\mathbf{x})]^\top$ and the unit normal is given by $\mathbf{n} = [-f_x, -f_y, 1]^\top / \sqrt{f_x^2 + f_y^2 + 1}$.

2.3. Bas-Relief Ambiguity in Uncalibrated Photometric Stereo

The radiance of a convex (interreflection-free) Lambertian surface illuminated by a distant point light source \mathbf{s} is

$$L_{f,a,s}(\mathbf{x}) = \Psi_{f,s}(\mathbf{x})a(\mathbf{x})\mathbf{n}(\mathbf{x})^\top \mathbf{s} \quad (1)$$

where the binary function Ψ , which depends purely on the surface geometry and light source direction, determines shadowing ($\Psi = 0$ for points in shadow and 1 otherwise).

Traditional photometric stereo is based on the assumption that an orthographic camera is used to image a Lambertian surface under direct illumination by distant point light sources, with no interreflections. In the uncalibrated case with no shadowing, three images with linearly independent (unknown) light source directions suffice to simultaneously determine the surface and light source directions up to an invertible 3×3 transformation. If $\mathbf{b}(\mathbf{x}) = a(\mathbf{x})\mathbf{n}(\mathbf{x})^\top$, one recovers some $\bar{\mathbf{b}}(\mathbf{x}) = \mathbf{b}(\mathbf{x})\mathbf{A}^{-\top}$ and $\mathbf{A}\bar{\mathbf{s}}$ since

$$L_{\bar{f},\bar{a},\bar{\mathbf{s}}} = \Psi_{\bar{f},\bar{\mathbf{s}}}\bar{\mathbf{b}}\bar{\mathbf{s}} = \Psi_{f,s}(\mathbf{b}\mathbf{A}^{-\top})(\mathbf{A}\mathbf{s}) = L_{f,a,s}.$$

It can be shown [1] that imposing the integrability constraint on the recovered normal field restricts \mathbf{A} to lie within the group of GBR-transformations. (A 3×3 matrix $\mathbf{G} \in GBR \subset GL(3)$ must have the form

$$\mathbf{G} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{bmatrix} \quad (2)$$

where $\mu, \nu, \lambda \in \mathbb{R}$ and $\lambda > 0$.) A GBR-transformation on a surface-albedo pair, $\mathbf{G} : (\mathcal{S}, a) \rightarrow (\bar{\mathcal{S}}, \bar{a})$ is defined by a transformation on the normal field and albedo

$$\bar{\mathbf{n}} = \frac{\mathbf{G}^{-\top} \mathbf{n}}{\|\mathbf{G}^{-\top} \mathbf{n}\|} \text{ and } \bar{a} = a \|\mathbf{G}^{-\top} \mathbf{n}\| \quad (3)$$

while the corresponding transformation on the light source $\mathbf{s} \in \mathbb{R}^3$ is $\mathbf{G} : \mathbf{s} \rightarrow \bar{\mathbf{s}} = \mathbf{G}\mathbf{s}$.

The following two theorems establish that the generalized bas-relief transformation is the only transformation that preserves both the shading and shadowing of a Lambertian surface.

Theorem 1 (Kriegman and Belhumeur). *The image of the shadow boundaries (cast and attached) of a Lambertian surface \mathcal{S} and light source \mathbf{s} are identical to that of a GBR-transformed surface $\bar{\mathcal{S}}$ and light source $\bar{\mathbf{s}}$. Further, $GBR \subset GL(3)$ is the only group of transformations for which shadow boundaries are invariant.*

Theorem 2 (Belhumeur, Kriegman and Yuille). *The image of a Lambertian surface-albedo pair (\mathcal{S}, a) under orthography with directional point light source \mathbf{s} is identical to that of a GBR-transformed surface-albedo pair $(\bar{\mathcal{S}}, \bar{a})$ and light source $\bar{\mathbf{s}}$.*

These theorems were proved, respectively, in [12] and [1] under assumptions of general surface conditions - we assume the same throughout this paper.

2.4. The Interreflection Kernel

The interreflection equation computes the radiance, in the presence of interreflections, emitted by a point \mathbf{x} on a surface \mathcal{S} with surface normal \mathbf{n} and albedo a , when illuminated by a point light source \mathbf{s} , as

$$L(\mathbf{x}) = \Psi_{f,\mathbf{s}}(\mathbf{x})a(\mathbf{x})\mathbf{n}(\mathbf{x})^\top \mathbf{s} + a(\mathbf{x}) \int_{\mathcal{S}} K(\mathbf{x}, \mathbf{x}')L(\mathbf{x}')dS' \quad (4)$$

Here dS' represents a differential surface element at \mathbf{x}' and K is the symmetric, positive semi-definite interreflection kernel, which is a function of the local surface orientation at points \mathbf{x} and \mathbf{x}' as well as the global surface geometry. More specifically,

$$K(\mathbf{x}, \mathbf{x}') = \frac{(\mathbf{n}^\top(-\mathbf{r}))(\mathbf{n}'^\top \mathbf{r})V(\mathbf{x}, \mathbf{x}')}{(\mathbf{r}^\top \mathbf{r})^2} \quad (5)$$

represents the fraction of light emitted by point \mathbf{x}' that is incident on \mathbf{x} , subject to a visibility criterion V that evaluates to 1 when the surface geometry allows for two points to be visible to each other, else evaluates to 0. The global shape determines the value of V to be 0 when \mathbf{x} and \mathbf{x}' are occluded, else the local orientation at the two points governs V according to

$$V(\mathbf{x}, \mathbf{x}') = \left(\frac{\mathbf{n}^\top(-\mathbf{r}) + |\mathbf{n}^\top(-\mathbf{r})|}{2|\mathbf{n}^\top(-\mathbf{r})|} \right) \left(\frac{\mathbf{n}'^\top \mathbf{r} + |\mathbf{n}'^\top \mathbf{r}|}{2|\mathbf{n}'^\top \mathbf{r}|} \right) \quad (6)$$

Here $\mathbf{r} = \mathbf{x} - \mathbf{x}'$ stands for the displacement from \mathbf{x}' to \mathbf{x} .

2.5. Calibrated Photometric Stereo

For a practical implementation, the surface is considered discretized into, say, m facets, each of constant albedo and radiance. Then, the interreflection kernel is an $m \times m$ matrix \mathbf{K} , whose entries are given by

$$\mathbf{K}_{ij} = \frac{(\mathbf{n}_i^\top \mathbf{r}_{ij})(\mathbf{n}_j^\top \mathbf{r}_{ji})}{(\mathbf{r}_{ij}^\top \mathbf{r}_{ij})^2} \frac{dA_j}{\mathbf{v}^\top \mathbf{n}_j} \mathbf{V}_{ij} \quad (7)$$

where \mathbf{V} is the visibility matrix, A_j is the area of facet j and \mathbf{v} is the viewing direction (which is the same for all facets under orthographic projection). \mathbf{K}_{ii} is set to 0 for all i . The radiosity equation can now be modified to its discrete formulation

$$\mathbf{L} = (\mathbf{I} - \mathbf{P}\mathbf{K})^{-1}\mathbf{L}_s \quad (8)$$

where $\mathbf{L} = [L_1, \dots, L_m]^\top$ is the radiance vector, $\mathbf{L}_s = [L_{s1}, \dots, L_{sm}]^\top$ is the source contribution vector, \mathbf{P} is a diagonal matrix composed of albedo values and \mathbf{K} is the interreflection kernel matrix in (7).

This discrete formulation is used in [14] to recover a "pseudo-normal field" by treating the radiance field as that of a directly illuminated surface. This pseudo-normal field is then iteratively projected on to an integrable normal field to obtain a surface estimate and the interreflection kernel is computed which yields a next estimate for the surface. It can also be shown that the "pseudo-shape" is less concave than the true shape.

3. Euclidean Reconstruction with Uncalibrated Photometric Stereo

3.1. Resolving Bas-Relief Ambiguities

As discussed above, assuming interreflection effects to be absent, two surfaces differing by a GBR-transformation, when viewed under orthographic projection in suitably transformed lighting conditions, would produce exactly the same images. We prove in the following theorem that the presence of interreflections ensures that images of surfaces differing by a GBR are not the same.

Proposition 1. *There does not exist a non-trivial GBR-transformation \mathbf{G} of a Lambertian surface-albedo pair (\mathcal{S}, a) that, with light sources correspondingly transformed, results in the same shading configuration under an orthographic camera when interreflections are present.*

Proof. Let $(\bar{\mathcal{S}}, \bar{a})$ be a GBR-transformed surface-albedo pair associated with the original surface-albedo pair (\mathcal{S}, a) . We note that $\bar{\mathbf{n}}^\top \bar{\mathbf{r}} = (\mathbf{G}^{-\top} \mathbf{n} / \|\mathbf{G}^{-\top} \mathbf{n}\|)(\mathbf{G}\mathbf{r}) = \mathbf{n}^\top \mathbf{r} / \|\mathbf{G}^{-\top} \mathbf{n}\|$, thus, from (5) and (6),

$$\bar{K}(\bar{\mathbf{x}}, \bar{\mathbf{x}}') = \frac{(\mathbf{r}^\top \mathbf{r})^2}{\|\mathbf{G}^{-\top} \mathbf{n}\| \|\mathbf{G}^{-\top} \mathbf{n}'\|} \frac{K(\mathbf{x}, \mathbf{x}')}{(\mathbf{r}^\top \mathbf{G}^\top \mathbf{G}\mathbf{r})^2}. \quad (9)$$

Substituting (9) and (3) into (4), the radiance at a point on the GBR-transformed surface is

$$\begin{aligned}\bar{L}(\bar{\mathbf{x}}) &= \Psi_{\bar{f},\bar{s}} \bar{a} \bar{\mathbf{n}}^\top \bar{\mathbf{s}} + \bar{a} \int_{\bar{S}} \bar{K}(\bar{\mathbf{x}}, \bar{\mathbf{x}}') \bar{L}(\bar{\mathbf{x}}') d\bar{S}' \\ &= \Psi_{f,s} a \mathbf{n}^\top \mathbf{s} + \|\mathbf{G}^{-\top} \mathbf{n}\| a \\ &\quad \times \int_{\bar{S}'} \frac{(\mathbf{r}^\top \mathbf{r})^2}{\|\mathbf{G}^{-\top} \mathbf{n}\| \|\mathbf{G}^{-\top} \mathbf{n}'\|} \frac{K(\mathbf{x}, \mathbf{x}')}{(\mathbf{r}^\top \mathbf{G}^\top \mathbf{G} \mathbf{r})^2} \bar{L}(\bar{\mathbf{x}}') d\bar{S}'\end{aligned}\quad (10)$$

as $\Psi_{\bar{f},\bar{s}} \bar{a} \bar{\mathbf{n}}^\top \bar{\mathbf{s}} = \Psi_{f,s} a \mathbf{n}^\top \mathbf{s}$ from Theorems 1 and 2.

Now, assume to the contrary that the original surface and the GBR-transformed surface produce the same images,

$$\bar{L}(\bar{\mathbf{x}}) = L(\mathbf{x}). \quad (11)$$

We recall the assumption of orthography and for the Cartesian (image coordinate) surface parametrization it provides, the Jacobian for the change of coordinates in a GBR-transformation is identity. Let $\mathbf{N} = [-f_x, -f_y, 1]^\top = \|\mathbf{N}\| \mathbf{n}$ be the *un-normalized* surface normal, then a differential area element on the surface is given by

$$\begin{aligned}dS &= \|\mathcal{S}_x \times \mathcal{S}_y\| dx dy \\ &= \|[1, 0, f_x]^\top \times [0, 1, f_y]^\top\| dx dy = \|\mathbf{N}\| dx dy\end{aligned}\quad (12)$$

Using equations (11) and (12), relation (10) becomes

$$\begin{aligned}L(\mathbf{x}) &= \bar{L}(\bar{\mathbf{x}}) \\ &= \Psi_{f,s} a \mathbf{n}^\top \mathbf{s} \\ &\quad + a \int_{\bar{S}'} \frac{(\mathbf{r}^\top \mathbf{r})^2}{\|\mathbf{G}^{-\top} \mathbf{n}'\|} \frac{K(\mathbf{x}, \mathbf{x}')}{(\mathbf{r}^\top \mathbf{G}^\top \mathbf{G} \mathbf{r})^2} L(\mathbf{x}') \|\bar{\mathbf{N}}'\| d\bar{x}' d\bar{y}' \\ &= \Psi_{f,s} a \mathbf{n}^\top \mathbf{s} \\ &\quad + a \int_S \frac{K(\mathbf{x}, \mathbf{x}') (\mathbf{r}^\top \mathbf{r})^2}{(\mathbf{r}^\top \mathbf{G}^\top \mathbf{G} \mathbf{r})^2} L(\mathbf{x}') \frac{\|\mathbf{G}^{-\top} \mathbf{N}'\|}{\|\mathbf{G}^{-\top} \mathbf{n}'\|} dx' dy'\end{aligned}\quad (13)$$

where we have used the definition of GBR-transformation $\bar{\mathbf{N}} = \mathbf{G}^{-\top} \mathbf{N}$ and the identity Jacobian to change the coordinate system. From (12), the radiance in (13) becomes

$$\begin{aligned}L(\mathbf{x}) &= \Psi_{f,s} a \mathbf{n}^\top \mathbf{s} + a \int_S \frac{K(\mathbf{x}, \mathbf{x}') (\mathbf{r}^\top \mathbf{r})^2}{(\mathbf{r}^\top \mathbf{G}^\top \mathbf{G} \mathbf{r})^2} L(\mathbf{x}') \|\mathbf{N}'\| dx' dy' \\ &= \Psi_{f,s} a \mathbf{n}^\top \mathbf{s} + a \int_S \frac{K(\mathbf{x}, \mathbf{x}') (\mathbf{r}^\top \mathbf{r})^2}{(\mathbf{r}^\top \mathbf{G}^\top \mathbf{G} \mathbf{r})^2} L(\mathbf{x}') dS'\end{aligned}\quad (14)$$

Comparing equations (4) and (14), since the surfaces we assumed are general, it can be easily seen that $\mathbf{G}^\top \mathbf{G}$ is forced to be identity. That is, \mathbf{G} must be a rotation matrix. But \mathbf{G} is a generalized bas-relief transformation and must have the structure outlined in (2). Unit norm on the columns

forces $\mu = \nu = 0$ and $\lambda = \pm 1$. Moreover, as $\lambda > 0$ by definition of a GBR transformation, we must have $\lambda = 1$. Thus, \mathbf{G} can only be the identity matrix, which contradicts our assumption of a non-trivial GBR-transformation. We have shown that the original and GBR-transformed surfaces cannot produce the same images and our proof stands completed. \square

The above proof was contingent on the fact that the visibility between two points does not change due to a GBR transformation. We note that as a consequence of Theorem 1 the occlusion relationship between two points is also invariant under a GBR transformation. So, we have been able to drop the assumption in [14] that required all the surface points be visible to each other.

From Theorem 1, the GBR-transformation is the only surface transformation that preserves the shadowing configuration. For a reconstruction method based solely on shadows, we have just demonstrated that the unknown GBR-transformation of the surface can be resolved by interreflections. So, we can assert that no two surface-albedo pairs can produce the same images in presence of interreflections *and* shadows. In fact, for any technique that can reconstruct a surface up to an unknown GBR, the ambiguity can be resolved in the presence of interreflections.

3.2. Surfaces with No Shadows

In practice, it is hard to deal with shadows. Often, they are not visible everywhere on the surface. Moreover, they inevitably lead to computationally more complex optimization problems. Therefore, we investigate the problem for surfaces with no shadows as well.

Proposition 2. *There does not exist a non-trivial transformation of a Lambertian surface-albedo pair (S, a) , with light sources correspondingly transformed, that results in the same shading configuration under an orthographic camera when interreflections are present.*

Proof. We may consider the surface-albedo pair (S, a) and its light sources \mathbf{s} as fixed or known quantities. Now, suppose there exists another surface-albedo pair (\bar{S}, \bar{a}) and we will show that this leads to an inconsistency.

Let \mathbf{F} denote the forward operator which, given light sources \mathbf{s} , produces images L , cf. (8). Hence, we have:

$$L = \bar{\mathbf{F}} \bar{\mathbf{s}} = \mathbf{F} \mathbf{s}. \quad (15)$$

The light sources \mathbf{s} and $\bar{\mathbf{s}}$ must be related by a GBR transformation, i.e., $\bar{\mathbf{s}} = \mathbf{G} \mathbf{s}$, as surface points where no interreflections are present must obey the traditional GBR. Inserting this into (15) and realizing that it must hold for *any* light source \mathbf{s} results in

$$\bar{\mathbf{F}} \mathbf{G} = \mathbf{F}.$$

The above equation is equivalent to:

$$\bar{a}\bar{\mathbf{n}}^\top \mathbf{G} = (\mathbf{I} - \bar{a}\bar{\mathbf{K}})\mathbf{F}, \quad (16)$$

where \mathbf{I} and $\bar{\mathbf{K}}$ should be interpreted as operators (the identity operator and the interreflection kernel operator, respectively). Factorizing out \bar{a} , we get $\bar{a}(\bar{\mathbf{n}}^\top \mathbf{G} + \bar{\mathbf{K}}\mathbf{F}) = \mathbf{F}$, or equivalently,

$$\bar{\mathbf{n}}^\top \mathbf{G} + \bar{\mathbf{K}}\mathbf{F} \sim \mathbf{F}, \quad (17)$$

where \sim denotes equality up to scale. As \mathbf{F} is considered known, the above equation can be regarded as an implicit equation in the unknown depths \bar{z} (and \mathbf{G}). For each surface point $\bar{\mathbf{x}}$, (17) says that there are two constraints for every depth $\bar{z}(\bar{\mathbf{x}})$ - three equations minus one scale factor. Thus, this is an overconstrained system with independent equations¹ and for general \mathbf{F} there is no solution for \bar{z} . The proposition is proved. \square

It turns out to be fruitful to look at the problem in a lower dimension as well, that is, the (orthographic) projection of a curve in the two-dimensional plane to a one-dimensional line. We will refer to this as the 2D case.

Proposition 3. *In the 2D case, there exists a family of transformations of a Lambertian curve-albedo pair (\mathcal{S}, a) , with light sources correspondingly transformed, that results in the same shading configuration under an orthographic (1D) camera when interreflections are present.*

Proof. Suppose that there exists a curve-albedo pair $(\bar{\mathcal{S}}, \bar{a})$ different from (\mathcal{S}, a) . Following the same derivations as in Proposition 2 all the way to (17), one obtains that the hypothesized curve $\bar{\mathcal{S}}$ must satisfy,

$$\bar{\mathbf{n}}^\top \mathbf{G} + \bar{\mathbf{K}}\mathbf{F} \sim \mathbf{F}.$$

Contrary to the 3D case, there is now only one constraint for every depth $\bar{z}(x)$ - two equations minus one scale factor. Another way to look at the constraint, is by eliminating the scale factor, and regard the resulting system as an operator equation:

$$\mathbf{E}_{\mathbf{G}}(\bar{z}) = 0,$$

which depends on \mathbf{G} . The true surface with $\mathbf{G} = \mathbf{I}$ satisfies of course the equation $\mathbf{E}_{\mathbf{G}}(z) = 0$. And from functional analysis it is known that the solution for an integral operator with a smooth integral kernel generally depends continuously on \mathbf{G} , and hence we will generally be able to find a solution \bar{z} with $\mathbf{G} \neq \mathbf{I}$. The corresponding albedo \bar{a} is found through (16). Since \mathbf{G} is a 2×2 matrix in the 2D case, there is a family of solutions. \square

¹The independence follows from the general surface assumption (similar to [12]), but it can also be verified with random, synthetic data.

While the above proposition explains the observed multiplicity in reconstructions of 2D curves (see experimental section), we need to delve deeper into the interreflection kernel's elements to explain the ambiguity in the case of 3D surfaces with a translational symmetry.

Definition 1. *A surface \mathcal{S} has a translational symmetry about a direction with direction cosines (α, β, γ) in some Cartesian coordinate system when $[x, y, z]^\top$ lies on \mathcal{S} if and only if $[x + t\alpha, y + t\beta, z + t\gamma]^\top$ lies on $\mathcal{S} \ \forall t \in \mathbb{R}$.*

Proposition 4. *For a surface \mathcal{S} with translational symmetry about some direction \mathbf{d} , there exists a family of transformations of the Lambertian surface-albedo pair (\mathcal{S}, a) , with light sources correspondingly transformed, that results in the same shading configuration under an orthographic camera when interreflections are present.*

Proof. In a 2D world, if $\tilde{\mathbf{r}}$ is the distance to a point light source from any point on the plane, light fall-off over the plane will vary as $1/\tilde{\mathbf{r}}$ instead of the $1/\tilde{\mathbf{r}}^2$ fall-off we commonly encounter in a 3D world. Thus, the form of the interreflection kernel for a 2D curve is given by

$$K(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}') = \frac{(\mathbf{n}^\top(-\tilde{\mathbf{r}}))(\mathbf{n}'^\top \tilde{\mathbf{r}})}{(\tilde{\mathbf{r}}^\top \tilde{\mathbf{r}})^{3/2}} \quad (18)$$

where $\tilde{\mathbf{x}} = [x(s), y(s)]^\top$ is a point on the 2D curve parametrized by arc length and $\tilde{\mathbf{r}} = \tilde{\mathbf{x}} - \tilde{\mathbf{x}}'$.

Now consider a 3D surface with a translational symmetry. Without loss of generality, let the axis of symmetry be the z-axis, that is, the surface is an infinite 3D extrusion of a 2D curve on the xy -plane. Therefore, any point $[x, y, z]^\top$ that lies on the surface can be parametrized as $[\tilde{\mathbf{x}}(s), z]^\top$ where s is the arc length parameter of the projection of the surface on the xy -plane.

For a point $[\tilde{\mathbf{x}}(s), 0]^\top$, averaging out the contribution from all points $[\tilde{\mathbf{x}}(s'), z]^\top$, $z \in \mathbb{R}$, we get the value of the kernel at (s, s') . As the component of the surface normal at any point along the direction of symmetry is zero, the product $\mathbf{n}^\top \mathbf{r}$ remains the same along the direction of symmetry. Here $\mathbf{r} = [\tilde{\mathbf{x}}(s), 0]^\top - [\tilde{\mathbf{x}}(s'), z]^\top = [\tilde{\mathbf{x}}(s) - \tilde{\mathbf{x}}(s'), -z]^\top$. As in [4], the 3D-kernel can be computed as

$$\begin{aligned} K(s, s') &= \int_{-\infty}^{\infty} \frac{(\mathbf{n}^\top(-\mathbf{r}))(\mathbf{n}'^\top \mathbf{r})}{(\mathbf{r}^\top \mathbf{r})^2} dz \\ &= (\mathbf{n}^\top(-\mathbf{r}))(\mathbf{n}'^\top \mathbf{r}) \int_{-\infty}^{\infty} \frac{1}{(\tilde{\mathbf{r}}^\top \tilde{\mathbf{r}} + z^2)^2} dz \\ &= \frac{\pi}{2} \frac{(\mathbf{n}^\top(-\mathbf{r}))(\mathbf{n}'^\top \mathbf{r})}{(\tilde{\mathbf{r}}^\top \tilde{\mathbf{r}})^{3/2}} \end{aligned} \quad (19)$$

which is the same as the interreflection kernel for a 2D curve up to scale. Hence, the problem can be expressed with a 2D-kernel, which is ambiguous (Proposition 3). \square

The implication of this result is that the multiplicity of solutions that we postulated would occur for 2D curves will also be observed for an infinitely long 3D surface with translational symmetry. Approximately similar effects would be observed for surfaces that are finite, but long in the direction of translational symmetry. Moreover, since estimation procedures are sensitive to noise, a 3D surface which is not symmetrical, but close to symmetrical, is also expected to yield an ambiguous reconstruction.

4. Experiments

4.1. Surface Estimation

The surface recovery is through a gradient descent minimization over the unknown depth, albedo and light sources of the least-squares errors between the measured intensity L_m and that obtained from the reprojection, using the discrete forward solution (8). Iteratively, the objective $\|L - L_m\|$ as well as the Jacobian with respect to the unknowns are computed and the current estimate is updated using the Levenberg-Marquardt formula. Convergence is assumed when the relative magnitude of updates falls below a reasonable threshold.

The initial estimate used for all experiments (if not otherwise stated) was obtained from the (pseudo) normal field recovered from traditional uncalibrated photometric stereo ignoring interreflections, disambiguated up to an unknown GBR using the technique in [19]. For noisy measurements, one can of course never be sure to reach the global minimum. The optimization was usually found to converge to the expected 3D reconstruction, though we have encountered local minima for noise-free cases.

4.2. Simulations

Consider a (synthetic) 2D curve with both concave and convex regions (Figure 3). Using noise-free (1D) images and a 2D line as initialization, a pseudo-curve and a 2D curve with interreflections are reconstructed. Both these curves reproduce exactly the same 1D images and hence the reprojection error is zero as expected.

Among all GBR-transformed pseudo-curves, the one with the same light sources directions as the original one is plotted. The pseudo-curve is less concave than the original in regions where interreflections are present, cf. [14]. The recovered 2D curve (whose set of images under all possible lighting sources is the same as the original curve) is some non-linear transformation of the original, supporting our theory in Proposition 3.

Next, consider a 3D surface with (approximate) translational symmetry, such as two inclined, intersecting planes, see Figure 4(a). Uncalibrated photometric stereo where mutual illumination effects are ignored gives a "pseudo-surface" that is much less concave the original, Figure 4(b).




Figure 3. Ambiguity in 2D shape recovery from uncalibrated photometric stereo with interreflections.

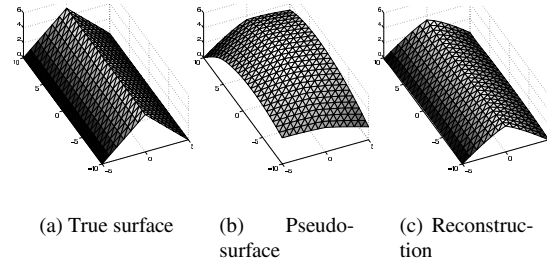


Figure 4. A simulation to illustrate the reconstruction ambiguity we predicted for the case of translational symmetry. (a) The input geometry. (b) The "pseudo-surface" obtained by uncalibrated photometric stereo without consideration of interreflections. (c) The recovered (ambiguous) surface from our algorithm. The optimization converges to a local minimum that produces nearly the same images as the true surface. Note the difference in height and curvature between the reconstructed and original shapes.

The relative error is 0.6%. Note that there may not exist a pseudo-surface which exactly reproduces the pseudo-normal field as a pseudo-normal field can not be expected to be integrable in general. From Proposition 4, we expect an ambiguity in shape recovery with uncalibrated stereo reconstruction (even though the surface is not an infinitely extruded one). And indeed, starting from a flat plane as initialization, the optimization algorithm converges to the local minimum in Figure 4(c). Relative error between images produced by the recovered surface and light sources and the original surface and light sources is less than 0.1%.





Figure 5. Recovered shape for the composed sine and cosine surface from uncalibrated photometric stereo using interreflections. Initialization was a flat planar surface. The optimization converges to a minimum that is indistinguishable from the original.

Finally, a (more) general 3D surface was simulated (Fig-

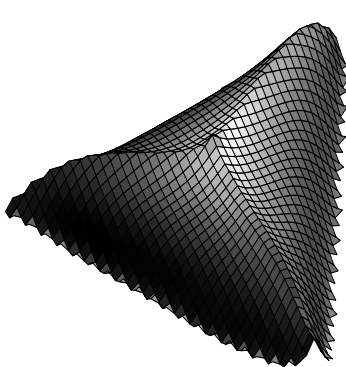


(a) Input image 1


(b) Input image 2

(c) Input image 3


(d) Input image 4



(e) Outside view of reconstruction from uncalibrated photometric stereo



(f) Outside view of reconstruction from calibrated photometric stereo with interreflection effects removed



(g) Outside view of reconstruction from our uncalibrated photometric stereo disambiguated by interreflections

Figure 6. *Comparison of reconstruction approaches* (a)-(d) The input images. (e) "Pseudo-surface" from uncalibrated photometric stereo without considering interreflections disambiguated up to a GBR. (f) Reconstruction from calibrated photometric stereo that compensates the effect of interreflections [14]. (g) Surface recovered by our uncalibrated photometric stereo taking interreflections into account.

ure 5(a)). With just a flat plane as initialization, the reconstruction algorithm converges to a surface very close to original, cf. Figure 5(b) with negligible relative error. This is in agreement with Proposition 2 that the GBR ambiguity is resolved for a general 3D surface with interreflections.

4.3. Experiments on Real Data

The paper surface used in the experiments below was verified to be Lambertian up to 98.1% in the following way. The paper was wrapped around a cylinder and imaged under nearly frontal lighting. The intensity profile across a horizontal cross-section (averaged across the vertical dimension of the cylinder) was fitted to a cosine curve and the goodness of fit gave an estimate of the degree of Lambertianity.

We demonstrate Euclidean reconstruction for general 3D surfaces in consonance with the theory presented in Proposition 2. In the process, we provide a comparison to uncalibrated photometric stereo assuming absence of interreflections as well as illustrate the similarity of reconstruction between calibrated and uncalibrated stereo when interreflections are considered.

The input (Figures 6(a) - 6(d)) consisted of images of the interior of a tetrahedron that displayed strong interreflections. Uncalibrated photometric stereo ignoring mutual illumination yields a surface that differs from the "pseudo-surface" by a 3×3 transformation, which is further reduced to a GBR-transformation by the method in [19]. The relative error is 16%. An exterior view of this reconstruction is presented in Figure 6(e) and it is obvious that the recovered "pseudo-surface" greatly differs from the true surface.

An implementation of the calibrated photometric stereo technique in [14] which iteratively compensates for interreflections yields the surface in Figure 6(f). We compare this with the reconstruction obtained by our uncalibrated approach with interreflections (Figure 6(g)). The two appear nearly the same with straight edges and planar faces. The errors are 1.3% and 0.6% for the calibrated and uncalibrated cases, respectively, and the resolution of the quadrilateral mesh is 1039 patches.

The differences and similarities between the various reconstructions can be better appreciated by a comparison of the height profiles across a horizontal cross-section of

the three reconstructions (Figure 7). We would expect the profile of an object containing flat planes to be composed of nearly straight lines, while the effects of interreflection would cause the pseudo-surface recovered by traditional calibrated photometric stereo to be noticeably curved. The angles between the four manually calibrated light source directions and the estimated ones differ by an average of 10° .




Figure 7. Comparison of shape recovery across a cross-section of the tetrahedron.

As another demonstration of our algorithm, we recover the surface of an open book in which interreflections occur between the two sides (Figure 8). The relative error is 1.1% and a mesh of 1256 triangles is used.




Figure 8. The first column shows two out of four input images of the open book. The second column shows the reconstructed surface of the book.

Finally, we tested our algorithm on more complex inputs, namely faces. The face input images in Figure 1 were synthetically generated under various point light sources from a real 3D head scan. The reconstructed face is compared against the structures obtained by traditional photometric stereo methods. On the one hand, there is relatively little interreflection in this data set, but on the other hand, the images are noise-free. The reprojection error is negligible.

5. Discussions

We have shown that it is possible to obtain a Euclidean reconstruction of a surface with uncalibrated photometric stereo in the presence of interreflections. The theoretical findings have been supported by several experiments. However, from a practical point of view, further work is needed. For example, it is unclear exactly how sensitive the method is to noise and under what circumstances would modelling interreflections be advantageous. A more efficient solution

to the forward problem (for example, using a Gauss-Seidel technique) will expedite the optimization and allow for higher resolutions in the reconstruction.

Acknowledgments Support for this work was provided by NSF under IIS-0308185 and EIA-0224431, the U. C. MICRO program and the Swedish Research Council. Many thanks to Sameer Agarwal for several fruitful discussions and insightful comments on the nature of interreflections.

References

- [1] P. Belhumeur, D. Kriegman, and A. Yuille. The bas-relief ambiguity. *IJCV*, 35(1):33–44, November 1999.
- [2] M. Cohen and D. Greenberg. The hemi-cube: a radiosity solution for complex environments. *ACM Computer Graphics*, 19(3):31–40, July 1985.
- [3] O. Drbohlav and R. Sara. Specularities reduce ambiguity of uncalibrated photometric stereo. In *ECCV*, pages 46–62, 2002.
- [4] D. Forsyth and A. Zisserman. Mutual illumination. In *CVPR*, pages 466–473, 1989.
- [5] D. Forsyth and A. Zisserman. Reflections on shading. *PAMI*, 13(7):671–679, July 1991.
- [6] A. Georghiades. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In *ICCV*, pages 816–823, 2003.
- [7] A. Gilchrist. The perception of surface blacks and whites. *Scientific American*, 240:112–123, 1979.
- [8] H. Hayakawa. Photometric stereo under a light-source with arbitrary motion. *JOSA-A*, 11(11):3079–3089, 1994.
- [9] B. Horn and M. Brooks, editors. *Shape from Shading*. MIT Press, 1989.
- [10] J. Kajiya. The rendering equation. In *SIGGRAPH*, pages 143–150, 1986.
- [11] J. Koenderink and A. van Doorn. Geometrical modes as a general method to treat diffuse interreflections in radiometry. *JOSA*, 73(6):843–850, June 1983.
- [12] D. Kriegman and P. Belhumeur. What shadows reveal about object structure. *JOSA-A*, 18(8):1804–1813, 2001.
- [13] P. Moon. On interreflections. *JOSA*, 30(5):195–205, 1940.
- [14] S. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. In *ICCV*, pages 2–11, 1990.
- [15] F. Tricomi. *Integral Equations*. Dover Publications, 1957.
- [16] T. Wada, H. Ukida, and T. Matsuyama. Shape from shading with interreflections under proximal light source: 3D shape reconstruction of unfolded book surface from a scanner image. In *ICCV*, pages 66–71, 1995.
- [17] R. Woodham. Photometric method for determining surface orientation from multiple images. *OptEng*, 19(1):139–144, January 1980.
- [18] A. Yuille, J. Coughlan, and S. Konishi. The KGBR viewpoint-lighting ambiguity. *JOSA-A*, 20(1):24–31, January 2003.
- [19] A. Yuille and D. Snow. Shape and albedo from multiple images using integrability. In *CVPR*, pages 158–164, 1997.