

Isotropy, Reciprocity and the Generalized Bas-Relief Ambiguity

Ping Tan^{1*} Satya P. Mallick² Long Quan¹ David J. Kriegman² Todd Zickler³

¹ The Hong Kong University of Science and Technology

² University of California, San Diego

³ Harvard School of Engineering and Applied Sciences

Abstract

A set of images of a Lambertian surface under varying lighting directions defines its shape up to a three-parameter Generalized Bas-Relief (GBR) ambiguity. In this paper, we examine this ambiguity in the context of surfaces having an additive non-Lambertian reflectance component, and we show that the GBR ambiguity is resolved by any non-Lambertian reflectance function that is isotropic and spatially invariant. The key observation is that each point on a curved surface under directional illumination is a member of a family of points that are in isotropic or reciprocal configurations. We show that the GBR can be resolved in closed form by identifying members of these families in two or more images. Based on this idea, we present an algorithm for recovering full Euclidean geometry from a set of uncalibrated photometric stereo images, and we evaluate it empirically on a number of examples.

1. Introduction

Most problems in computer vision are simplified in the presence of perfectly diffuse, or Lambertian, surfaces. According to the Lambertian model, the bidirectional reflectance distribution function (BRDF) is a constant function of the viewing and illumination directions. By assuming that surfaces are well-represented by this model, one can build powerful tools for stereo reconstruction, shape from shading, motion estimation, segmentation, photometric stereo, and a variety of other visual tasks.

Most surfaces are not Lambertian, however, so we often seek ways of generalizing these powerful Lambertian-based tools. One common approach is to assume that non-Lambertian phenomena occur only in small regions of an image, and to treat these regions as outliers or ‘missing data’. Another approach is to model these phenomena using parametric representations of reflectance that are more com-

plex than the Lambertian model. The latter approach has the important advantage of using all of the available image data, but it also has a significant limitation. Even relatively simple reflectance models (such as the Phong or Cook-Torrance models) severely complicate the image analysis problem, and since they are only applicable for limited classes of surfaces, this approach generally requires new and complex analysis for each application and each material class.

Recently, we have witnessed acceleration in the development of a third approach to handling non-Lambertian scenes—one that is based on exploiting more general properties of surface reflectance. This approach stems from the observation that even though there is a wide variety of materials in the world, there are common reflectance phenomena that are exhibited by broad classes of these materials. By building tools that exploit these properties, one can build vision systems that are more likely to succeed in real-world, non-Lambertian environments. One early example of this approach is Shafer’s development of the dichromatic model, which exploits the fact that additive diffuse and specular components of reflectance often differ in color [15].

Two important reflectance phenomena are isotropy and Helmholtz reciprocity. On a small surface patch, the BRDF is defined as the ratio of the reflected radiance in direction (θ_o, ϕ_o) to the received irradiance from direction (θ_i, ϕ_i) . It is typically denoted $f(\theta_i, \phi_i, \theta_o, \phi_o)$, where the parameters are spherical coordinates in the local coordinate system of the patch. Helmholtz reciprocity tells us that the BRDF is symmetric in its incoming and outgoing directions ($f(\theta_i, \phi_i, \theta_o, \phi_o) = f(\theta_o, \phi_o, \theta_i, \phi_i)$), and isotropy implies that there is no preferred azimuthal orientation or ‘grain’ to the surface ($f(\theta_i, \phi_i, \theta_o, \phi_o) = f(\theta_o, \theta_i, |\phi_o - \phi_i|)$). In computer vision, these properties have been exploited for surface reconstruction [11, 18], and since they effectively reduce the BRDF domain, they have also been used extensively for image-based rendering in computer graphics (e.g., [7, 14]).

In this paper, we seek to exploit isotropy and reciprocity more broadly. We show that an image of a curved surface (convex or not) under parallel projection and distant illumina-

*Much of this work was completed while P. Tan was visiting Harvard University.

nation contains observations of distinct surface points that have equivalent local view and illumination geometry under isotropy and reciprocity. By studying the structure of these equivalence classes, we derive intensity-based constraints on the field of surface normals. As an application, we show that these constraints are sufficient to resolve the generalized bas-relief (GBR) ambiguity that is inherent to uncalibrated photometric stereo.

1.1. The GBR Ambiguity

It is well established that a set of images of a Lambertian surface under varying, distant lighting do not completely determine its Euclidean shape. Given any such set of images, the surface can only be recovered up to a three-parameter ambiguity—the GBR ambiguity [1, 10]. Significant effort has been devoted to understanding when and how this ambiguity can be resolved. It is known, for example, that when a surface is Lambertian, the GBR ambiguity can be resolved in the presence of interreflections [2], or when relative albedo values and/or source strengths are known [1, 8]. It is also known that the GBR can be resolved when surface reflectance can be represented using one of two specific non-Lambertian reflectance models: the Torrance-Sparrow model [6] or the ‘Lambertian plus specular spike’ model [4, 5].

In this paper, we investigate the relationship between the GBR and surfaces with more general non-Lambertian reflectance. We study reflectance that is a linear combination of a Lambertian diffuse components and an isotropic¹ (but otherwise arbitrary) specular component:

$$f(\mathbf{x}, \theta_i, \phi_i, \theta_o, \phi_o) = \rho(\mathbf{x}) + f_s(\theta_i, \theta_o, |\phi_i - \phi_o|). \quad (1)$$

Here, \mathbf{x} denotes a point on the surface, so that the diffuse component varies spatially (i.e., the surface has ‘texture’), while the specular component is spatially invariant. This model is quite generic, and it generalizes all existing analysis of the GBR ambiguity in the context of non-Lambertian reflectance [4, 5, 6], since all of these consider special cases of Eq. 1. Given a surface with reflectance of this form, one can obtain a reconstruction of the surface (up to an unknown GBR transformation) using existing techniques for diffuse/specular image separation (e.g., [16, 13]) and by applying uncalibrated Lambertian photometric stereo [8, 17] to the diffuse component. The specular component then provides additional information that can be used to resolve the GBR ambiguity.

One of the key results of this paper is that two images (with sources in general positions) are sufficient to resolve

¹There seems to be some confusion in the use of the term *isotropy* in the vision and graphics communities. In some cases (e.g. [9]) it implies dependence on the absolute difference of azimuthal angles $|\phi_i - \phi_o|$, but in others it only implies dependence on the signed difference $(\phi_i - \phi_o)$, with the additional absolute value being a separate property termed ‘bilateral symmetry’ (e.g. [12]). In this paper, we use the former interpretation.




Figure 1. On the visible hemisphere of Gauss sphere, two normals form an *isotropic pair* if they lie at the intersections of two circles centered at source direction \mathbf{s} and view direction \mathbf{v} . If the BRDF is isotropic (but otherwise arbitrary) the observed intensity at these points will be equal.

the GBR for any surface with reflectance as shown in Eq. 1. Isotropy and reciprocity are sufficient to resolve the GBR, and no additional assumptions are required.

The remainder of the paper is organized as follows. Sect. 2 describes the general structure that isotropy and reciprocity induce on the Gauss sphere for surfaces with isotropic, spatially-invariant reflectance. Sect. 3 discusses how these structures are affected by a GBR transformation and how they can be exploited to resolve the GBR ambiguity. Finally, Sect. 4 describes a complete system for photometric stereo auto-calibration and evaluates it empirically.

2. Isotropic and Reciprocal Image Structure

We begin by exploring the geometric structure of surface points that are in isotropic and reciprocal configurations under a distant point light source and an orthographic view. We assume that the Gauss map is known (e.g., we are given the output of a photometric stereo algorithm), so that all analysis can be performed on the Gauss sphere. A fixed orthographic observer is located in direction \mathbf{v} , and a global coordinate system is chosen so that the z -axis is aligned with this direction. A distant point source is located in direction \mathbf{s} . Both \mathbf{v} and \mathbf{s} are vectors of unit length, as are all vectors in the remainder of this paper unless stated otherwise. We use the phrase *principal meridian* to refer to the great circle through \mathbf{v} and \mathbf{s} .

Definition 1. Two surface normals \mathbf{n} and \mathbf{n}' form an *isotropic pair* under source \mathbf{s} if they satisfy

$$\mathbf{n}'^T \mathbf{s} = \mathbf{n}^T \mathbf{s} \text{ and } \mathbf{n}'^T \mathbf{v} = \mathbf{n}^T \mathbf{v}.$$

A geometric interpretation is shown in Fig. 1. Two normals form an isotropic pair if they lie at the intersections of two circles centered at \mathbf{v} and \mathbf{s} on the Gauss sphere. If we consider the local coordinate system for the BRDF domain at each of two such surface normals, it is clear that the incoming and outgoing elevation angles (θ_i and θ_o) are the same in both cases. This is because the tetrahedron formed by unit vectors $(\mathbf{n}, \mathbf{s}, \mathbf{v})$ is equivalent under reflection to that formed by $(\mathbf{n}', \mathbf{s}, \mathbf{v})$. Also, if we were to project \mathbf{v} and \mathbf{s} onto




Figure 2. On the visible hemisphere of Gauss sphere, two normals form a *reciprocal pair* if they lie at the intersections of circles that are centered at \mathbf{s} and \mathbf{v} and whose radii are exchanged. In this example, there are four reciprocal pairs: $\mathbf{m} \leftrightarrow \mathbf{n}$, $\mathbf{m}' \leftrightarrow \mathbf{n}'$, $\mathbf{m} \leftrightarrow \mathbf{n}'$, and $\mathbf{m}' \leftrightarrow \mathbf{n}$. If the BRDF is isotropic and reciprocal, the observed BRDF value (but not necessarily the radiance) at these points will be equal.

the plane orthogonal to each normal, it is clear that the magnitude of the angular difference between these projections (i.e., $|\phi_i - \phi_o|$) would also be the same. This remark follows.

Remark 1. Assuming an isotropic and spatially invariant BRDF, the emitted radiance from two points whose surface normals form an isotropic pair will be equal.

Definition 2. Two surface normals \mathbf{n} and \mathbf{m} form a *reciprocal pair* under light source \mathbf{s} if they satisfy

$$\mathbf{m}^\top \mathbf{s} = \mathbf{n}^\top \mathbf{v} \text{ and } \mathbf{m}^\top \mathbf{v} = \mathbf{n}^\top \mathbf{s}.$$

This condition can similarly be interpreted in terms of the intersections of circles centered at \mathbf{s} and \mathbf{v} . The important difference is that the radii of the two circles are swapped for the two normals. As depicted in Fig. 2, there are four possible intersections derived from all combinations of circles with centers at \mathbf{s} or \mathbf{v} and having two different radii. This family of four normals comprises two isotropic pairs ($\mathbf{m} \leftrightarrow \mathbf{m}'$ and $\mathbf{n} \leftrightarrow \mathbf{n}'$) and four reciprocal pairs ($\mathbf{m} \leftrightarrow \mathbf{n}$, $\mathbf{m}' \leftrightarrow \mathbf{n}'$, $\mathbf{m}' \leftrightarrow \mathbf{n}$ and $\mathbf{m} \leftrightarrow \mathbf{n}'$). Using an argument similar to that above, it is clear that the local light and view directions at any two normals of a reciprocal pair are such that the following is true.

Remark 2. Assuming an isotropic and reciprocal BRDF, the observed BRDF value (but not the emitted radiance, which is the BRDF multiplied by $\mathbf{n}^\top \mathbf{s}$) at two points whose surface normals form a reciprocal pair will be equal.

The conditions in Def. 1 induce an equivalence relation (i.e., they satisfy reflexivity, transitivity and symmetry) so we can say that two normals are ‘isotropically equivalent’ if they form an isotropic pair. The conditions for a reciprocal pair do not induce an equivalence relation, however, as a normal is generally not reciprocal to itself. To form an equivalence relation that exploits reciprocity, the conditions for isotropic and reciprocal pairs must be combined; we can say that two normals are ‘equivalent under combined isotropy and reciprocity’ if they form either an isotropic pair or a reciprocal pair. These equivalence relations partition

the Gauss sphere into equivalence classes of normals. To further investigate the geometric structure of these equivalence classes, it will be helpful to define two families of curves on the Gauss sphere that, among other things, have the nice property of being unions of equivalence classes.

Definition 3. An *isotropic curve* under \mathbf{s} is a great circle on the Gauss sphere that satisfies

$$(\mathbf{x}^\top \mathbf{s}) + \alpha(\mathbf{x}^\top \mathbf{v}) = 0$$

where x is a point on the Gauss sphere and α is a constant.

An isotropic curve has the following properties. First, if a normal is on an isotropic curve, its isotropic corresponding normal is also on the curve. (This can be easily verified by substituting Def. 1 into the isotropic curve equation.) In other words, an isotropic curve is a union of isotropic equivalence classes. Second, the one-parameter family of isotropic curves partitions the Gauss sphere. Third and finally, the emitted radiance along the curve is symmetric due to the symmetric arrangement of isotropic pairs along it (e.g., Fig. 3). This last property plays an important role in subsequent analysis, so we present it as a formal Remark.

Remark 3. For any isotropic BRDF, the emitted radiance along an isotropic curve is a symmetric function when the curve is parameterized by the signed angle (ψ , say) to the principal meridian.

Definition 4. A *reciprocal curve* under \mathbf{s} is a conical curve on the Gauss sphere that satisfies

$$(\mathbf{x}^\top \mathbf{v})(\mathbf{x}^\top \mathbf{s}) + \alpha(\mathbf{v}, \mathbf{s}, \mathbf{x})^2 = 0$$

where x is a point on the Gauss sphere and α is a constant.

Here $(\mathbf{x}, \mathbf{y}, \mathbf{z}) = (\mathbf{x} \times \mathbf{y})^\top \mathbf{z}$ denotes the scalar triple product. A reciprocal curve has properties that are analogous to its isotropic counterpart. If a normal is on a given reciprocal curve, its isotropic/reciprocal corresponding normals are on the same curve. Thus, it is a union of ‘combined isotropic/reciprocal equivalence classes’. Also, like isotropic curves, the one-parameter family of reciprocal curves (parameterized by α) partitions the Gauss sphere minus points on the principal meridian. Finally, the symmetric arrangement of reciprocal and isotropic pairs within a reciprocal curve induces symmetry in the BRDF along it.

Remark 4. For any isotropic and reciprocal BRDF, we can choose a parameterization of a reciprocal curve such that the BRDF value (emitted radiance divided by $\mathbf{n}^\top \mathbf{s}$) along the curve is a symmetric function. One such parameterization is given by the azimuthal angle between a normal on the curve and the plane that contains the bisector of \mathbf{s} and \mathbf{v} and is perpendicular to the principal meridian.

An example of a reciprocal curve is shown in Fig. 4. To provide a geometric interpretation of isotropic and reciprocal curves, it is useful to interpret the Gauss sphere

as being mapped to a plane. An elliptic plane is obtained by a gnomonic (or central) projection that maps a point of the sphere from the center of the sphere onto the tangent plane at \mathbf{v} . It maps a point to a point, a great circle to a line, and a conical curve to a conic. The elliptic plane is a real projective plane with an elliptic metric [3]. The three great circles $c_1 = \{\mathbf{x} : \mathbf{v}^\top \mathbf{x} = 0\}$, $c_2 = \{\mathbf{x} : \mathbf{s}^\top \mathbf{x} = 0\}$, and $c_3 = \{\mathbf{x} : (\mathbf{v} \times \mathbf{s})^\top \mathbf{x} = 0\}$ are mapped into three lines in the elliptic plane: \mathbf{l}_1 is the line at infinity, $\mathbf{l}_2 = \mathbf{s}$, and $\mathbf{l}_3 = \mathbf{v} \times \mathbf{s}$.

On the elliptic plane, isotropic curves form a pencil (linear family) of lines by \mathbf{l}_1 and \mathbf{l}_2 . Reciprocal curves form a pencil of conics going through the double points $\mathbf{p} = \mathbf{l}_1 \times \mathbf{l}_3$ (the intersection of \mathbf{l}_1 and \mathbf{l}_2) and $\mathbf{q} = \mathbf{l}_2 \times \mathbf{l}_3$ (the intersection of \mathbf{l}_1 and \mathbf{l}_2). The conic pencil is a linear family of parabola directed by the light source direction \mathbf{s} , and it touches the point at infinity \mathbf{q} .

3. Behavior under GBR Transformations

The isotropic and reciprocal structures described in the previous section exist whenever a curved surface is illuminated from a distant point-source and viewed orthographically. They provide constraints between the intensities observed at distinct surface points and the orientation of the surface normals at these points. Since these constraints are valid for any isotropic and reciprocal BRDF, they may find use for a variety of tasks.

In this section, we analyze the behavior of these structures when a GBR transformation is applied. We show that, somewhat surprisingly, a GBR transformation generally maps isotropic/reciprocal curves ‘as sets’ to other isotropic/reciprocal curves. At the same time, normals within each curve generally move relative to one another, thereby breaking the symmetry in the radiance functions along the curve (see Figs. 3 and 4). As a result, given an initial reconstruction up to an arbitrary GBR ambiguity, we can establish the Euclidean reconstruction by finding the GBR transformation that restores this symmetry structure.

3.1. GBR Transformations

Given three or more uncalibrated images of a Lambertian surface, the field of surface normals $\mathbf{n}(\mathbf{x})$ (defined over the image plane parameterized by $\mathbf{x} \in \mathbb{R}^2$) and the source directions \mathbf{s}_i can only be recovered up to an invertible linear transformation of \mathbb{R}^3 [8]. It has been shown that by imposing integrability on the surface, this general linear transformation is restricted to lie in the group of GBR transformations, which are 3×3 matrices of the form [1]

$$\mathbf{G} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{pmatrix},$$

with $\mu, \nu, \lambda \in \mathbb{R}$. A GBR transformation affects the normal field and source directions according to

$$\bar{\mathbf{n}} = \mathbf{G}^{-\top} \mathbf{n} / \|\mathbf{G}^{-\top} \mathbf{n}\|, \quad \bar{\mathbf{s}} = \mathbf{G} \mathbf{s} / \|\mathbf{G} \mathbf{s}\|, \quad (2)$$

and

$$\mathbf{n} = \mathbf{G}^\top \bar{\mathbf{n}} / \|\mathbf{G}^\top \bar{\mathbf{n}}\|, \quad \mathbf{s} = \mathbf{G}^{-1} \bar{\mathbf{s}} / \|\mathbf{G}^{-1} \bar{\mathbf{s}}\| \quad (3)$$

is the effect of the inverse transformation.

It is easy to verify that a GBR transformation affects neither view direction \mathbf{v} nor the principal meridian. Isotropy and reciprocity, however, are in general destroyed by the GBR since $\mathbf{n}^\top \mathbf{s} \neq \bar{\mathbf{n}}^\top \bar{\mathbf{s}}$. As a result, if we are given a reconstruction up to an unknown GBR transformation and we are able to find the pairs of transformed isotropic and reciprocal normals, we expect that the GBR can be solved.

A comment on terminology: in the subsequent discussion, we will be interested in describing the manner in which two normals $\mathbf{n} \leftrightarrow \mathbf{m}$ of an isotropic or reciprocal pair under \mathbf{s} are affected by a GBR transformation. Since a GBR transformation does not preserve isotropy and reciprocity, we are generally uninterested in normals that form isotropic/reciprocal pairs under $\bar{\mathbf{s}}$ in the sense of Defs. 1 and 2. Instead, we say that transformed normal $\bar{\mathbf{n}}$ *corresponds* to transformed normal $\bar{\mathbf{m}}$ if the pre-images of these normals \mathbf{n} and \mathbf{m} form an isotropic/reciprocal pair under \mathbf{s} . The story is different, however, for isotropic and reciprocal curves. Since these curves are preserved as sets under a GBR transformation, it *does* make sense to consider isotropic and reciprocal curves under $\bar{\mathbf{s}}$ in the sense of Defs. 3 and 4.

3.2. Isotropy and GBR Transformations

We first look at the action of a GBR transformation on isotropic pairs and isotropic curves.

Proposition 1. A GBR transformation maps each isotropic curve under \mathbf{s} ‘as a set’ to an isotropic curve under $\bar{\mathbf{s}}$.

To prove this, we first look at how a pair of isotropic normals is transformed by \mathbf{G} . Suppose $\bar{\mathbf{n}}$ is a GBR-transformed normal and $\bar{\mathbf{x}}$ is its unknown isotropic correspondence. Since the pre-images of $\bar{\mathbf{n}}$ and $\bar{\mathbf{x}}$ form an isotropic pair under \mathbf{s} , we know $\mathbf{x}^\top \mathbf{v} = \mathbf{n}^\top \mathbf{v}$ and $\mathbf{x}^\top \mathbf{s} = \mathbf{n}^\top \mathbf{s}$. By substituting Eq. 3 into these equations, we obtain the following linear constraint for the position of $\bar{\mathbf{x}}$ corresponding to $\bar{\mathbf{n}}$:

$$(\bar{\mathbf{x}}^\top \bar{\mathbf{s}}) + \alpha(\bar{\mathbf{x}}^\top \bar{\mathbf{v}}) = 0,$$

with

$$\alpha = -(\bar{\mathbf{n}}^\top \bar{\mathbf{s}}) / (\bar{\mathbf{n}}^\top \bar{\mathbf{v}}).$$

Thus, if \mathbf{n} and \mathbf{x} form an isotropic pair under \mathbf{s} , then $\bar{\mathbf{n}}$ and $\bar{\mathbf{x}}$ lie on an isotropic curve under $\bar{\mathbf{s}}$.

To complete the proof, we can explicitly derive a mapping between the pre-GBR and post-GBR isotropic curves by substituting the inverse transformation equations (Eq. 3) into the expression above. This yields the following equation for the pre-image of the isotropic curve:

$$(\mathbf{x}^\top \mathbf{s}) + \beta(\mathbf{x}^\top \mathbf{v}) = 0,$$

where $\beta = \alpha \|\mathbf{G} \mathbf{s}\| / \lambda$. This is also an isotropic curve, which proves Proposition 1.




Figure 3. (Left) Top view of the Gauss sphere. Green points represent isotropic pairs, which are arranged symmetrically about the principal meridian on an isotropic curve prior to a GBR transformation. These points are mapped to the yellow points by a GBR transformation, and are no longer symmetrically arranged within the curve. (Right) Emitted radiance as a function of signed angle from the principal meridian along the isotropic curves before and after the GBR. Since isotropy is not preserved by the GBR, the symmetry in the radiance function is lost.

A GBR transformation maps isotropic curves to isotropic curves, but as demonstrated in Fig. 3, since it does not preserve isotropy, it destroys the symmetry of the intensity function along the curve. Thus, assuming that we can identify corresponding normals $\bar{\mathbf{n}}$ and $\bar{\mathbf{n}}'$ on an isotropic curve in the GBR-transformed system, we can obtain a constraint on the unknown GBR transformation by imposing the condition that their pre-images \mathbf{n} and \mathbf{n}' form an isotropic pair under \mathbf{s} . To formulate this constraint, we substitute Eq. 3 into the equations in Def. 1. These equations are homogeneous in λ , which is therefore eliminated, and they provide a linear equation in μ and ν :

$$\bar{s}_y\mu - \bar{s}_x\nu + c = 0, \quad (4)$$

where $c = (\bar{\mathbf{v}}, \bar{\mathbf{s}}, \bar{\mathbf{n}} + \bar{\mathbf{n}}') / (\bar{\mathbf{n}}^\top \bar{\mathbf{v}} + \bar{\mathbf{n}}'^\top \bar{\mathbf{v}})$.

There are two important observations to make regarding this equation. First, it is independent of λ . This means that isotropy is preserved by a classical bas-relief transformation ($\mathbf{G} = \text{diag}(1, 1, \lambda)$), and λ cannot be resolved using isotropic constraints alone. Second, the coefficients of μ and ν are independent of $\bar{\mathbf{n}}$ and $\bar{\mathbf{n}}'$ and depend only on the source direction $\bar{\mathbf{s}}$. As a result, each image provides only one independent linear constraint on μ and ν . This constraint can be geometrically interpreted as a requirement for symmetry of isotropic pairs $\mathbf{n} \leftrightarrow \mathbf{n}'$ about the principal meridian characterized by $(\mathbf{v}, \mathbf{s}, \mathbf{n} + \mathbf{n}') = 0$. To recover μ and ν , we need at least one additional independent equation, which requires a second image under source $\bar{\mathbf{s}}'$ whose projection onto the xy -plane is different from that of $\bar{\mathbf{s}}$. Since the GBR transformation does not affect the xy -plane projection of a source direction (nor the principal meridian as mentioned earlier), this leads to the following proposition.

Proposition 2. For any isotropic, spatially-invariant BRDF, the GBR parameters μ and ν can be uniquely determined from any two images under source directions that are not coplanar with the view direction.

3.3. Reciprocity and Bas-relief Transformations

Once μ, ν are known, we can apply the appropriate inverse transform that corrects for the additive plane and reduces the GBR ambiguity to a classic bas-relief ambiguity, which is represented by the diagonal matrix $\mathbf{G}' = \text{diag}(1, 1, \lambda)$. Here, we examine the effects of a bas-relief transformation on reciprocal pairs and reciprocal curves.

Proposition 3. A bas-relief transformation maps each reciprocal curve under \mathbf{s} ‘as a set’ to a reciprocal curve under $\bar{\mathbf{s}}$.

The proof is similar to the isotropic case, consider $\bar{\mathbf{n}} \leftrightarrow \bar{\mathbf{x}}$ whose pre-images form a reciprocal pair under \mathbf{s} . The definition for a reciprocal pair states that $\mathbf{x}^\top \mathbf{v} = \mathbf{n}^\top \mathbf{s}$ and $\mathbf{x}^\top \mathbf{s} = \mathbf{n}^\top \mathbf{v}$. Substituting the expressions from Eq. 3 into these equations (with \mathbf{G} reduced to a bas-relief transformation) and eliminating variable λ , one obtains

$$(\bar{\mathbf{x}}^\top \mathbf{v})(\bar{\mathbf{x}}^\top \bar{\mathbf{s}}) + \alpha(\mathbf{v}, \bar{\mathbf{s}}, \bar{\mathbf{x}})^2 = 0,$$

with

$$\alpha = -(\bar{\mathbf{n}}^\top \mathbf{v})(\bar{\mathbf{n}}^\top \bar{\mathbf{s}}) / (\mathbf{v}, \bar{\mathbf{s}}, \bar{\mathbf{n}})^2.$$

Thus, if \mathbf{n} and \mathbf{x} form a reciprocal pair under \mathbf{s} , then $\bar{\mathbf{n}}$ and $\bar{\mathbf{x}}$ lie on a reciprocal curve under $\bar{\mathbf{s}}$.

Next, to derive the relationship between reciprocal curves related by bas-relief transformations, we apply the inverse transformation to the reciprocal curve equation above to obtain the pre-image of this curve:

$$(\mathbf{x}^\top \mathbf{v})(\mathbf{x}^\top \mathbf{s}) + \beta(\mathbf{v}, \mathbf{s}, \mathbf{x})^2 = 0,$$

where $\beta = \alpha\lambda / \|\mathbf{G}\mathbf{s}\|$. This is also a reciprocal curve, which completes the proof of Proposition 3.

Since reciprocity is not preserved by a bas-relief transformation, two corresponding normals $\bar{\mathbf{n}}$ and $\bar{\mathbf{m}}$ in the transformed system can be used to obtain a constraint on λ by imposing the condition that the pre-images of these normals form a reciprocal pair under source \mathbf{s} . Substituting Eq. 3 (again with $\mathbf{G} = \text{diag}(1, 1, \lambda)$) into the conditions for a reciprocal curve and solving for λ , we obtain:

$$(1 - s_z^2)\lambda^2 = (\bar{\mathbf{m}}^\top \bar{\mathbf{s}})(\bar{\mathbf{n}}^\top \bar{\mathbf{s}}) / (\bar{\mathbf{m}}^\top \bar{\mathbf{v}})(\bar{\mathbf{n}}^\top \bar{\mathbf{v}}) - s_z^2. \quad (5)$$

Geometrically, this equation for λ derives from the symmetry of reciprocal pairs $\mathbf{m} \leftrightarrow \mathbf{n}$ under \mathbf{s} . The bisector of these normals must lie in the plane containing the bisector of \mathbf{v} and \mathbf{s} (characterized by $(\mathbf{s} + \mathbf{v}, \mathbf{s} \times \mathbf{v}, \mathbf{m} + \mathbf{n}) = 0$) that is perpendicular to the principal meridian. See Fig. 4.

Equation 5 provides the magnitude of the parameter λ , and its sign can be determined as that of $\bar{\mathbf{m}}^\top \bar{\mathbf{s}} / \bar{\mathbf{n}}^\top \mathbf{v}$ or $\bar{\mathbf{m}}^\top \mathbf{v} / \bar{\mathbf{n}}^\top \bar{\mathbf{s}}$. Thus we have:

Proposition 4. For a spatially-invariant, isotropic, and reciprocal BRDF, the bas-relief parameter λ can be uniquely determined from any image whose source direction is not collinear with the view direction.




Figure 4. (Left) Green points represent reciprocal and isotropic pairs, which are symmetrically arranged on a reciprocal curve prior to a bas-relief transformation. These are mapped to the yellow points by a bas-relief transform and are no longer symmetrically located. (Right) The BRDF as a function of position along the reciprocal curves before and after the transformation. This function is initially symmetric, but the symmetry is destroyed by the bas-relief transformation.

4. Application

The results of the previous section tell us that the auto-calibration of photometric stereo can be performed from as few as two images by exploiting isotropy and reciprocity. Given an initial reconstruction of normals and sources, the GBR ambiguity can be recovered from Eqs. 4 and 5. All that is required is to identify at least one isotropic pair in each of two images (with sources in general position), and at least one reciprocal pair in any single image.

In order to test these results, we conducted experiments with both real and synthetic data. Each set of images was captured/rendered using directional illumination and an orthographic camera. In the case of real data, the camera was radiometrically calibrated so that image intensity could be directly related to emitted scene radiance. Diffuse and specular reflection components of the input images were separated using a color-based technique [13], and an initial reconstruction was obtained by applying an existing method for uncalibrated Lambertian photometric stereo [17] to the diffuse images. Then, by identifying reciprocal and isotropic pairs in the specular images we were able to resolve the GBR ambiguity.

4.1. Isotropic and Reciprocal Correspondence

To obtain isotropic correspondence for a given normal $\bar{\mathbf{n}}$, we search along the isotropic curve containing $\bar{\mathbf{n}}$ until we locate another normal with the same intensity. By establishing at least one such correspondence in two or more images, we can recover GBR parameters μ and ν using Eq. 4 and reduce the GBR ambiguity to a classic bas-relief ambiguity.

Once we correct for μ and ν , we use a similar procedure to establish reciprocal correspondence. The reciprocal correspondence for a given normal $\bar{\mathbf{n}}$ can be located by searching along the reciprocal curve until we locate another normal with the same BRDF value. The process is complicated by the fact that we do not have direct access to the BRDF value along the curve. We only know the radiance (image




Figure 5. The accuracy of an estimate of the the bas-relief parameter λ can be assessed by searching a reciprocal curve for reciprocal correspondence and computing an error measure based on this correspondence (see text for details). Here, the error reaches its global minimum at the ground truth value of $\lambda = 1.3$.

intensity) $I(\bar{\mathbf{n}}) = f(\bar{\mathbf{n}})(\bar{\mathbf{n}} \cdot \mathbf{s})$. We note that

$$I(\bar{\mathbf{m}}) = I(\bar{\mathbf{n}})(\bar{\mathbf{m}}^T \mathbf{s}) / (\bar{\mathbf{n}}^T \mathbf{s}) = I(\bar{\mathbf{n}}) \lambda \| \mathbf{G}^{-1}(\lambda) \bar{\mathbf{s}} \| (\bar{\mathbf{n}}^T \mathbf{v}) / (\bar{\mathbf{n}}^T \bar{\mathbf{s}}),$$

which allows us to assess the accuracy of any estimate λ_1 by: 1) using this equation to estimate the radiance $I(\bar{\mathbf{m}})$; 2) searching the reciprocal curve for a correspondence $\bar{\mathbf{m}}$ that has this intensity; 3) using $\bar{\mathbf{n}}$ and this estimate of $\bar{\mathbf{m}}$ to compute λ_2 using the Eq. 5; and 4) computing the error $\|\lambda_1 - \lambda_2\|^2$. Based on this procedure, we do an exhaustive one-dimensional search over a wide interval $\lambda \in [-10, 10]$. Figure 5 shows a typical profile of the error computed over this interval, which generally reaches a global minimum at the ground truth value of λ .

As an aside, we note that previous work using the ‘specular spike’ model of reflectance [4, 5] can be derived as a special case of this method. In that case, the correspondence problem is trivial because each specular normal is its own isotropic and reciprocal corresponding normal. The result is the same, however, and the GBR ambiguity can be fully resolved from two images [4].

4.2. Results

Figure 6 shows results on a synthetic example consisting of thirty-six rendered input images of a Venus model using a Cook-Torrance BRDF. The top row of this figure shows a linearly coded normal map, where the RGB channels represent x, y, and z components, respectively. The second row shows the corresponding surface height fields obtained by integrating the normal fields. From left to right, the columns show results of calibrated photometric stereo (i.e., ‘ground truth’); uncalibrated photometric stereo [17]; and our auto-calibrated photometric stereo. These results demonstrate that the proposed approach successfully resolves the GBR ambiguity that remains after the traditional uncalibrated result [17], and gives results that are very close to the calibrated case. The GBR parameters recovered by our method are $(\mu, \nu, \lambda) = (1.2, 0.9, 1.3)$.

Figures 7 and 8 show results for experiments on real images. In these figures, the top row shows one of the input images along with the separated diffuse and specular components. The middle and bottom rows show normal maps




Figure 6. Results from thirty-six synthetic images rendered with a Cook-Torrance BRDF. Top row: linearly coded normal maps, where r, g, b channels represent x, y, z components. Bottom row: surface height fields. Results in columns from left to right: calibrated lighting directions; traditional uncalibrated photometric stereo [17]; our auto-calibrated photometric stereo algorithm, which successfully resolves the GBR ambiguity and obtains results comparable to the calibrated case.

Ground truth s	Results from [17]	Auto-calibrated s	Error
(0.34, 0.26, 0.90)	(0.13, 0.20, -0.97)	(0.33, 0.52, 0.79)	16
(0.35, -0.30, 0.89)	(0.20, -0.07, -0.98)	(0.39, -0.15, 0.90)	8
(-0.27, -0.34, 0.90)	(-0.02, 0.05, -0.99)	(-0.07, 0.19, 0.98)	33
(-0.23, 0.11, 0.97)	(-0.02, 0.03, -0.99)	(-0.05, 0.11, 0.99)	10

Table 1. Accuracy of recovered source directions for four images of the pear dataset. Right-most column shows angular error.

and surface height fields as before, with columns from left to right representing: 1) calibrated Lambertian photometric stereo applied to the diffuse images; 2) uncalibrated Lambertian photometric stereo [17] applied to the diffuse images; and 3) our auto-calibrated results that use the specular images to resolve the GBR ambiguity. For the calibrated case, we use source directions that were measured from mirrored spheres during acquisition. The GBR parameters (μ, ν, λ) recovered by our method are $(2.1, -1.2, -3.3)$ and $(-2, -1.2, 3.1)$ for the pear and fish, respectively.

Table 1 and Figure 9 provide quantitative evaluation results. For the pear data, Table 1 compares the source directions recovered by our method to the ground-truth source directions measured during acquisition. Here, the right-most column provides angular error in degrees. Figure 9 shows average angular errors between the normals recovered using our approach and those recovered using calibrated photometric stereo with the measured source directions. For the Venus, pear and fish examples, our method achieves average angular errors of 2.8, 4 and 6.7 degrees and maximum angular errors of 44, 32 and 26.3 degrees, respectively. The sources of error include: 1) inaccuracies in isotropic and reciprocal correspondence; 2) imaging noise; and 3) inaccurate diffuse/specular separation.




Figure 7. Results from four input images of a pear. Top row: one input image with separated diffuse and specular components. Middle row: linearly encoded normal map. Bottom row: surface height fields. Columns from left to right show photometric stereo results using: calibrated lighting directions; an uncalibrated approach [17]; and our auto-calibrated approach that resolves the GBR and provides results comparable to the calibrated case.




Figure 8. Results from seven input images of a plastic toy fish. Top row: one input image with separated diffuse and specular components. Middle row: linearly encoded normal map. Bottom row: surface height fields. Columns from left to right show photometric stereo results using: calibrated lighting directions; an uncalibrated approach [17]; and our auto-calibrated approach that resolves the GBR and provides results comparable to the calibrated case.

Finally, Fig. 10 shows surfaces obtained by integrating the recovered normal fields from the calibrated, uncalibrated, and auto-calibrated methods. In each case, our auto-calibrated procedure significantly improves the uncalibrated results by resolving the GBR ambiguity.




Figure 9. Surface normal angular error between the results of our auto-calibrated method and calibrated photometric stereo results.




Figure 10. Surfaces recovered from integrating estimated normal fields. Rendering of the recovered surface from a novel view point. Columns from left to right show photometric stereo results using: calibrated lighting directions; an uncalibrated approach [17]; and our auto-calibrated approach.

5. Conclusion

This paper demonstrates that the generalized bas-relief ambiguity can be resolved for any surface that has an additive specular reflectance component that is spatially-invariant, isotropic and reciprocal. It shows that two images are sufficient to resolve the GBR, and presents a practical algorithm for doing so. The result is an auto-calibrating system for photometric stereo that can be applied to a very wide variety of surfaces.

More broadly, this paper demonstrates the utility of two very general reflectance properties: isotropy and reciprocity. It shows that any image of a surface (convex or not) under directional illumination and orthographic view contains observations of distinct surface points that are in isotropic and reciprocal configurations. By analyzing these equivalence classes, it reveals patterns of intensity on the Gauss sphere that can be used as constraints on surface geometry. In the future, these constraints could potentially be used in other ways for the analysis of scenes with complex, non-Lambertian surfaces.

6. Acknowledgements

P. Tan and L. Quan were supported by Hong Kong RGC project 619005, 619006 and RGC/NSFC project N_HKUST602/05. T. Zickler was supported by NSF CAREER Award IIS-0546408. D. Kriegman and S. Mallick were partially supported under NSF grants IIS-0308185 and NSF/EIA-0303622.

References

- [1] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. *IJCV*, 35(1):33–44, 1999.
- [2] M. K. Chandraker, F. Kahl, and D. Kriegman. Reflections on the generalized bas-relief ambiguity. In *CVPR*, 2005.
- [3] H. S. M. Coxeter. *Introduction to Geometry*, 2nd Edition. Wiley, 1989.
- [4] O. Drbohlav and M. Chantler. Can two specular pixels calibrate photometric stereo? In *ICCV*, volume 2, pages 850–1857, 2005.
- [5] O. Drbohlav and R. Sara. Specularities reduce ambiguity of uncalibrated photometric stereo. In *ECCV*, volume 2, pages 46–60, 2002.
- [6] A. S. Georgiades. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In *ICCV*, pages 816–823, 2003.
- [7] T. Hawkins and P. E. P. Debevec. A dual light stage. In *Rendering Techniques 2005 (Proc. Eurographics Symposium on Rendering)*, 2005.
- [8] K. Hayakawa. Photometric stereo under a light source with arbitrary motion. *J. Opt Soc. Am.*, 11(11), 1994.
- [9] J. J. Koenderink and A. J. van Doorn. Phenomenological description of bidirectional surface reflection. *JOSA A*, 15:2903–2912, 1998.
- [10] D. J. Kriegman and P. N. Belhumeur. What shadows reveal about object structure. *JOSA A*, 18, 2001.
- [11] J. Lu and J. Little. Reflectance and shape from images using a collinear light source. *IJCV*, 32(3):1–28, 1999.
- [12] S. R. Marschner. *Inverse rendering for computer graphics*. PhD thesis, Cornell University, 1998.
- [13] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. *Journal of Optical Society of America A*, 11(11):2990 – 3002, November 1994.
- [14] P. Sen, B. Chen, G. Garg, S. Marschner, M. Horowitz, M. Levoy, and H. P. A. Lensch. Dual photography. In *ACM Transactions on Graphics (ACM SIGGRAPH)*, 2005.
- [15] S. Shafer. Using color to separate reflection components. *COLOR research and applications*, 10(4):210–218, 1985.
- [16] R. T. Tan and K. Ikeuchi. Separating reflection components of textured surfaces using a single image. *PAMI*, 27:178–193, 2005.
- [17] A. L. Yuille and D. Snow. Shape and albedo from multiple images using integrability. In *CVPR*, pages 158–164, 1997.
- [18] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *IJCV*, 49(2/3):215–227, September 2002.