

A Projective Framework for Radiometric Image Analysis

Ping Tan¹ Todd Zickler²

¹ Department of Electrical & Computer Engineering, National University of Singapore

² Harvard School of Engineering and Applied Sciences

Abstract

Different materials reflect light in different ways, and reflectance interacts with shape, lighting, and viewpoint to determine an object's image. Common materials exhibit diverse reflectance effects, and this is a significant source of difficulty for radiometric image analysis. One strategy for dealing with this diversity is to build computational tools that exploit reflectance symmetries, such as reciprocity and isotropy, that are exhibited by broad classes of materials. In this paper, we advocate the real projective plane as a tool for representing and exploiting these symmetries. In this approach, each point in the plane represents a surface normal that is visible from a fixed viewpoint, and reflectance symmetries are analyzed in terms of the geometric structures that they induce. We provide an overview of these structures and explore applications to both calibrated and uncalibrated photometric stereo.

1. Introduction

Radiometric methods for image analysis, including photometric shape analysis and reflectometry, seek to infer scene information from measurements of image irradiance. The difficulty of this task stems from two facts: 1) surface reflectance varies dramatically between materials; and 2) it is tightly coupled with shape, illumination, and viewpoint in image measurements. One way to deal with these complexities is to exploit reflectance symmetries that are common to many materials. By taking advantage of isotropy, bilateral symmetry, reciprocity, and other symmetries, one can develop radiometric techniques that succeed for very broad classes of materials.

This paper presents a geometric framework for describing the joint constraints on shape, viewpoint and illumination that are induced by symmetries in surface reflectance. We represent the hemisphere of surface normals that are visible from a fixed orthographic view by the real projective plane (Fig. 1), and we analyze reflectance symmetries in terms of the geometric structures that they induce in this plane. This representation provides concise and intuitive descriptions of the symmetries, and it provides a convenient




Figure 1. The hemisphere of surface normals visible from direction v is represented by a plane obtained by gnomonic projection. This is the real projective plane, where great circles map to lines and the equator maps to a line at infinity. Reciprocity, isotropy, and other reflectance symmetries can be studied in terms of the geometric structures that they induce in this plane.

tool for developing new techniques for image analysis.

To demonstrate the utility of this framework we use it to develop new techniques for both calibrated and uncalibrated photometric stereo. In uncalibrated photometric stereo, we build upon existing work to show that constraints induced by isotropy and reciprocity in a single image are sufficient to resolve the generalized bas-relief ambiguity. This builds upon the work of Tan et al. [10], who require combined constraints from *two* images to produce the equivalent result. In the calibrated case, we show that isotropy and reciprocity constraints can be used to recover full Euclidean structure from images captured under a known, view-centered cone of light sources. This is achieved by completing the partial reconstruction provided by the method of Alldrin and Kriegman [1].

2. Background and related work

At an appropriate scale, the reflectance of opaque materials is described by the bi-directional reflectance distribution function, or BRDF. The BRDF describes the manner in which incident radiant flux is modulated by a uniform surface patch. It is a positive function of four angular dimensions and is written $f(\omega_i, \omega_o)$, where ω_i and ω_o are unit vectors on the hemisphere centered about the patch normal. These are the directions of incident and reflected flux, respectively, and they are often expressed in spherical coordinates: (θ_i, ϕ_i) and (θ_o, ϕ_o) .

Many materials exhibit reflectance symmetries. Reciprocity, for example, guarantees that the BRDF is symmetric about the input and output directions: $f(\omega_i, \omega_o) = f(\omega_o, \omega_i)$. In many cases, the BRDF domain $(\theta_i, \phi_i, \theta_o, \phi_o)$ can be further ‘projected’ onto the 3D domain $(\theta_i, \theta_o, \phi_i - \phi_o)$ and then folded onto $(\theta_i, \theta_o, |\phi_i - \phi_o|)$. The projection is acceptable whenever a BRDF is unchanged by rotations of the input and output directions (as a fixed pair) about the surface normal; and additional folding is acceptable whenever it is unchanged by reflecting the output direction about the incident plane. Materials that satisfy these two criteria are said to be *isotropic* and *bilaterally-symmetric*, respectively. (It is also common to use *isotropic* to mean both, and we will do so here.) Finally, in a number of cases [6, 8], the BRDF domain can be further reduced because it is unchanged when the input and output directions are rotated (again, as a fixed pair) about the *half-vector*, or the vector that bisects them. Materials that exhibit this last symmetry might be described as being *half-vector symmetric*.

When one or more of these symmetries is satisfied, radiance measurements that are captured at symmetrically-equivalent local view and illumination directions must be equal, and this induces joint constraints on shape, viewpoint and illumination. This has been exploited, for example, for surface reconstruction using isotropy [5], bilateral symmetry [1, 10], reciprocity [12], and half-vector symmetry [2]. The advantage of such symmetry-based approaches is that they avoid the use of low-parameter BRDF models (Lambertian, Ward, Cook-Torrance, etc.), which are often inaccurate and/or introduce inconvenient non-linearities.

The main goal of this paper is to provide an intuitive framework for analyzing the constraints induced by reflectance symmetries in the hopes of facilitating the development of new symmetry-based methods for image analysis as well as a better understanding of those that already exist.

3. Reflectance symmetries on the plane

The real projective plane provides an effective tool for analyzing reflectance symmetries, of which we focus on reciprocity and isotropy in this paper. (We use the term *isotropy* to include bilateral-symmetry, as discussed above.) As shown in Fig. 1, we represent each visible normal $\mathbf{n} = (n_1, n_2, n_3)^\top$ as a point in the 2D projective plane created by a gnomonic (or central) projection of the unit hemisphere onto the tangent plane passing through the viewing direction $\mathbf{v} = (0, 0, 1)^\top$. In this mapping, great circles map to lines, and the equator maps to a line at infinity. Both points and lines are represented using homogenous three-vectors, for which we use the notation $x = (x_1, x_2, x_3)$, while noting that $x \simeq \alpha x$ represent the same point for any $\alpha \in \mathbb{R}/\{0\}$. (We use \simeq to indicate equality up to scale.) The plane is equipped with an elliptic metric: the distance

between any two points is given by the angular difference between the corresponding rays in \mathbb{R}^3 , and the angle between lines is given by the angle between corresponding planes in \mathbb{R}^3 . (It has thus been termed the *elliptic plane* [4].) In this paper we use bold font to represent normalized vectors, i.e., $\mathbf{x} \triangleq x/|x|$.

We are interested in geometric structures induced by reciprocity and isotropy for images captured with directional lighting. Let \mathbf{v} be the direction of an orthographic observer, and let \mathbf{s} be the direction of the source. (For illustrative purposes, we assume \mathbf{s} lies in the upper hemisphere and can be associated with a point on the projective plane. This assumption can be relaxed.) We refer to the line connecting \mathbf{v} and \mathbf{s} as the *principal meridian* and use the notation $vs \triangleq \mathbf{v} \times \mathbf{s}$. We also refer to the midpoint on the line segment vs as the *half-vector* and use the notation $h \triangleq \mathbf{v} + \mathbf{s}$.

3.1. Isotropic pair

Following the definition of Tan et al. [10], a pair of normals \mathbf{n} and \mathbf{n}' is defined as an *isotropic pair* if and only if their local view and lighting conditions are equivalent up to isotropy. If the surface reflectance is the same at two surface points whose normals form an isotropic pair, the radiant exitance (BRDF multiplied by the cosine foreshortening factor) will be equal at these two points. For a scene illuminated from direction \mathbf{s} and observed from direction \mathbf{v} necessary and sufficient conditions for two normals to form an isotropic pair with respect to source \mathbf{s} are [10]:

$$\mathbf{s}^\top \mathbf{n} = \mathbf{s}^\top \mathbf{n}', \quad \mathbf{v}^\top \mathbf{n} = \mathbf{v}^\top \mathbf{n}'.$$

In the sequel, we find it useful to provide the following alternative set of conditions which, as shown in Proposition 6 of Appendix A, are equivalent to those above:

$$(\mathbf{v} \times \mathbf{s})^\top (\mathbf{n} + \mathbf{n}') = 0, \tag{1}$$

$$\frac{\mathbf{s}^\top \mathbf{n}}{\mathbf{v}^\top \mathbf{n}} = \frac{\mathbf{s}^\top \mathbf{n}'}{\mathbf{v}^\top \mathbf{n}'}. \tag{2}$$

Equations (1) and (2) can be interpreted geometrically as saying that the principal meridian vs is the perpendicular bisector of the line segment nn' . Specifically, Eq. 1 says that the mid-point of the line segment nn' (given by $\mathbf{n} + \mathbf{n}'$) lies on the principal meridian vs , and Eq. 2 says nn' is perpendicular to vs (see Proposition 8 of Appendix B). In Fig. 2, both $(\mathbf{n}, \mathbf{n}')$ and $(\mathbf{m}, \mathbf{m}')$ are isotropic pairs.

3.2. Reciprocal pair

Following the definition of Tan et al. [10], a pair of normal directions \mathbf{n} and \mathbf{m} is defined as a *reciprocal pair*, if and only if their local view and lighting conditions are equivalent up to isotropy and reciprocity. If the surface reflectance is the same at two surface points whose normals form a reciprocal pair, the BRDF value (radiant exitance divided by the cosine foreshortening factor) that is sampled at those two points will be equal. Necessary and sufficient conditions for two normals to form a reciprocal pair with respect to source \mathbf{s} are [10]:

$$\mathbf{s}^\top \mathbf{n} = \mathbf{v}^\top \mathbf{m}, \quad \mathbf{v}^\top \mathbf{n} = \mathbf{s}^\top \mathbf{m}.$$

As for isotropic pairs, it is helpful to consider an alternative, equivalent (see Proposition 7 of Appendix A) set of conditions:

$$\mathbf{s}^\top (\mathbf{n} + \mathbf{m}) = \mathbf{v}^\top (\mathbf{n} + \mathbf{m}), \quad (3)$$

$$\mathbf{s}^\top (\mathbf{n} \times \mathbf{m}) = \mathbf{v}^\top (\mathbf{n} \times \mathbf{m}), \quad (4)$$

that can also be interpreted geometrically. Equation 3 says the line connecting point $\mathbf{n} + \mathbf{m}$ and the halfway vector \mathbf{h} is perpendicular to the principal meridian (multiply the equation by $\mathbf{v}^\top (\mathbf{v} + \mathbf{s}) = \mathbf{s}^\top (\mathbf{v} + \mathbf{s})$ and apply Proposition 8 of Appendix B). Equation 4 says the line nm intersects the principal meridian at the point $(\mathbf{v} + \mathbf{s}) \times (\mathbf{v} \times \mathbf{s}) = \mathbf{h} \times (\mathbf{v} \times \mathbf{s})$ (multiply the equation by $\mathbf{v}^\top (\mathbf{v} + \mathbf{s}) = \mathbf{s}^\top (\mathbf{v} + \mathbf{s})$ and apply Proposition 9 in Appendix B). Notably, this intersection point is independent of \mathbf{m} and \mathbf{n} . In Fig. 2, both (\mathbf{n}, \mathbf{m}) and $(\mathbf{n}', \mathbf{m}')$ are reciprocal pairs.

Given any point \mathbf{n} , its reciprocal correspondence \mathbf{m} can be determined as follows. First, find the point $\mathbf{n} + \mathbf{m}$ as the intersection of two lines: 1) the join of points \mathbf{n} and $\mathbf{h} \times (\mathbf{v} \times \mathbf{s})$; and 2) the line through \mathbf{h} that is perpendicular to the principal meridian. Then, using the elliptic metric, \mathbf{m} is uniquely determined by \mathbf{n} and $\mathbf{n} + \mathbf{m}$.

3.3. Isotropic-reciprocal quadrilateral

As shown in Fig. 2, four normals $(\mathbf{n}, \mathbf{m}, \mathbf{m}', \mathbf{n}')$ form an *isotropic-reciprocal quadrilateral* with respect to \mathbf{s} if and only if both (\mathbf{n}, \mathbf{m}) and $(\mathbf{n}', \mathbf{m}')$ are reciprocal pairs and both $(\mathbf{n}, \mathbf{n}')$ and $(\mathbf{m}, \mathbf{m}')$ are isotropic pairs. An isotropic-reciprocal quadrilateral consists of four normals with local lighting and view directions that are equivalent up to isotropy and reciprocity. Thus, for a surface with uniform reflectance (or one with a uniform separable component) that is isotropic and reciprocal, the (uniform component of the) BRDF values observed at four surface points whose normals form such a quadrilateral will be equal.

According to the discussion above, both nn' and mm' are perpendicular to the principal meridian, and both nm and $n'm'$ intersect the principal meridian at $\mathbf{h} \times (\mathbf{v} \times \mathbf{s})$. So the quadrilateral is an isosceles trapezoid. Further, the center of the quadrilateral is the half-vector \mathbf{h} because the diagonals of the quadrilateral intersect there.

4. Transformations

The geometric structures described in previous section exist whenever an isotropic surface is observed orthographically under directional illumination, and they may be exploited for compression, reflectometry, reconstruction, or any other visual task. In the following, we focus on reconstruction via photometric stereo and show how these structures can be used to resolve shape ambiguities in both uncalibrated and calibrated cases. These applications often involve linear transformations of a surface's normal field. In uncalibrated Lambertian photometric stereo, for example, a linear




Figure 2. In projective plane, an isotropic-reciprocal quadruplet is an isosceles trapezoid centered at \mathbf{h} with two sides perpendicular to vs and two sides intersect at $\mathbf{h} \times (\mathbf{v} \times \mathbf{s})$.

algorithm yields the normal field up to an arbitrary linear transformation, and for differentiable surfaces, this ambiguity can be further reduced to a three-parameter generalized bas-relief (GBR) transformation [3].

Linear transformations of the normal field correspond to projective transformations of the plane, so this section discusses the behavior of our symmetry-induced structures under projective transformations. We show that only a uniform scaling of the plane (i.e., a classic bas-relief transformation) and rotation about the view direction preserve isotropic pairs, and that only an identity transformation preserves both isotropic and reciprocal pairs.

Proposition 1. A rotation (about the origin) and a uniform scaling are the only linear transformations that preserve isotropic pairs with respect to two or more lighting directions that are non-coplanar with the view direction.

Proof. Start with an arbitrary projective transformation. Two isotropic pairs $(\mathbf{n}_1, \mathbf{n}'_1)$ and $(\mathbf{n}_2, \mathbf{n}'_2)$ with respect to the same lighting direction define a point on the line at infinity $n_1 n'_1 \times n_2 n'_2$. Thus, from two lighting directions non-coplanar with the view direction, two distinct points can be located on the line at infinity, and the transformation can be reduced to an affinity by restoring the line at infinity. Affine transformations preserve the mid-point of a line segment. The join of the mid-point of $n_1 n'_1$ with the mid-point of $n_2 n'_2$ is a line l perpendicular to $n_1 n'_1$. This provides a pair of perpendicular directions, i.e. l and $n_1 n'_1$. From two lighting directions non-coplanar with the view direction, two such perpendicular pairs can be identified, which reduces the affine transformation to a similarity. Similarities preserve the perpendicular bisector of a line segment. As the original is on the perpendicular bisector of $n_1 n'_1$, it can be determined by intersecting two such perpendicular bisectors from two lighting directions non-coplanar with the view direction. Hence, the translation can be resolved and the similarity is reduced to a rotation about the origin and uniform scaling. Finally, it is easy to verify that isotropic pairs are preserved under a rotation and uniform scaling. \square This proof provides an 8-normal algorithm (two isotropic

pairs in each of two images) for reducing an arbitrary projective ambiguity to a scale and rotation.

Proposition 2. If the principal meridian vs is known, a classic bas-relief transformation is the only linear transformation that preserves isotropic pairs with respect to two or more sources that are non-collinear with the view direction.

Proof. In the projective plane, the classical bas-relief transformation is a uniform scaling. Given Proposition 1, we need only prove that the principal meridian vs fixes the rotation. This is straight-forward, as the rotation angle is determined by rotating the perpendicular bisector of any isotropic pair to align it with vs . \square

Proposition 3. If the lighting directions are known, the identity transformation is the only linear transformation that preserves isotropic-reciprocal quadrilaterals with respect to two or more lighting directions that are non-collinear with the view direction.

Proof. Since \mathbf{s} is known, vs can be determined by its join with \mathbf{v} , the origin. Given Proposition 2, the only remaining transformation is a uniform scaling. If \mathbf{s} is known, the scaling factor can be determined by requiring the center of isosceles trapezoid locating at $\mathbf{h} = \mathbf{v} + \mathbf{s}$. \square

5. Applications

In photometric stereo, we seek 3D shape from multiple images recorded at a fixed viewpoint with variable illumination. Methods are referred to as being ‘calibrated’ or ‘uncalibrated’ depending on whether the lighting is known or unknown *a priori*. We address both cases in this section.

5.1. Uncalibrated photometric stereo

In uncalibrated photometric stereo, the shape of a smooth Lambertian surface can only be recovered up to a three-parameter family of surfaces, all of which are related by so-called generalized bas-relief (GBR) transformations [3]. Tan et al. [10] demonstrated that the constraints induced by isotropy and reciprocity are sufficient to resolve the GBR ambiguity when one considers the combined constraints from two different images; and they presented an auto-calibrating method that produces a Euclidean reconstruction from images of any smooth surface having a spatially-varying BRDF with a uniform separable component (e.g., of the form $f(x, y, \omega_i, \omega_o) = f_1(x, y) + f_2(\theta_i, \theta_o, |\phi_i - \phi_o|)$). In this section, we use the insights from the previous section to prove that constraints from a *single* image are sufficient to obtain the same result.

Following Belhumeur et al. [3], a GBR transformation affects the normals and source directions according to

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

with




Figure 3. In the projective plane, the original isotropic-reciprocal quadruplet is transformed by a scaling followed with a translation. The illumination direction is transformed separately.

$$\mathbf{G} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{pmatrix}, \quad \mathbf{G}^{-\top} \simeq \begin{pmatrix} \lambda & 0 & -\mu \\ 0 & \lambda & -\nu \\ 0 & 0 & 1 \end{pmatrix}.$$

In the projective plane, this simplifies to $\bar{\mathbf{n}} \simeq \mathbf{G}^{-\top} \mathbf{n}$, meaning that a GBR consists of a uniform scaling of the surface normals by λ and a translation by $(-\mu, -\nu)$, as shown in Fig. 3. The view direction \mathbf{v} is unchanged by the transformation, and the light direction transforms as $\bar{\mathbf{s}} \simeq \mathbf{G} \mathbf{s}$. Furthermore, since $vs = \mathbf{v} \times \mathbf{s} = \mathbf{v} \times \bar{\mathbf{s}}$, the principal meridian is also unchanged by a GBR. It follows immediately from Proposition 2, that isotropy alone is sufficient to reduce the GBR to a classic bas-relief ambiguity from two images.

In order to completely resolve the GBR ambiguity from a single image, we consider one isotropic-reciprocal quadrilateral. Prior to a GBR transformation, the quadrilateral $nmm'n'$ is an isosceles trapezoid centered at \mathbf{h} and symmetric about the vs . A GBR transformation scales and translates the quadrilateral to $\bar{n}\bar{m}\bar{n}'\bar{n}'$ and moves the source to a different point $\bar{\mathbf{s}}$ on the principal meridian (see Fig. 3). To resolve the GBR ambiguity, we must find the transformation that maps $\bar{n}\bar{m}\bar{n}'\bar{n}'$ back to its canonical position.

The unknown transformation can be determined from the isotropic and reciprocal constraints described in the previous section because the required geometric structures are destroyed by a GBR. While the transformed line segment $\bar{n}\bar{n}'$ is still perpendicular to the principal meridian $v\bar{\mathbf{s}}$, its midpoint $\bar{\mathbf{n}} + \bar{\mathbf{n}}'$ no longer lies upon it. Similarly, the join of point $\bar{\mathbf{n}} + \bar{\mathbf{m}}$ and the transformed half-vector $\bar{\mathbf{h}} = \mathbf{v} + \bar{\mathbf{s}}$ is no longer perpendicular to the principal meridian $v\bar{\mathbf{s}}$, and the intersection of lines $\bar{m}\bar{n}$ and $\bar{m}'\bar{n}'$ is no longer at $\mathbf{h} \times (\mathbf{v} \times \bar{\mathbf{s}})$.

Mathematically, this means that while Eq. (2) is invariant under a GBR transformation, Eqs. (1,3,4) are not. There are three equalities destroyed by the GBR, and each provides a constraint on the unknown transformation. Since this transformation has exactly three degrees of freedom, the GBR can be uniquely determined using a single isotropic-reciprocal quadrilateral that is detected in a single image.

Proposition 4. The GBR ambiguity is resolved by the

isotropy and reciprocity constraints in a single image.

Proof. By isotropy and reciprocity, an isotropic-reciprocal quadrilateral $(\mathbf{n}, \mathbf{m}, \mathbf{n}', \mathbf{m}')$ with respect to \mathbf{s} , satisfies:

$$\begin{aligned} (\mathbf{v} \times \mathbf{s})^\top (\mathbf{n} + \mathbf{n}') &= 0, \\ \mathbf{s}^\top (\mathbf{n} + \mathbf{m}) &= \mathbf{v}^\top (\mathbf{n} + \mathbf{m}), \\ \mathbf{s}^\top (\mathbf{n} \times \mathbf{m}) &= \mathbf{v}^\top (\mathbf{n} \times \mathbf{m}). \end{aligned}$$

Substituting $\mathbf{n} = \mathbf{G}^\top \bar{\mathbf{n}} / \|\mathbf{G}^\top \bar{\mathbf{n}}\|$, $\mathbf{s} = \mathbf{G}^{-1} \bar{\mathbf{s}} / \|\mathbf{G}^{-1} \bar{\mathbf{s}}\|$, we obtain three equations in the unknown parameters μ, ν, λ ¹:

$$(\mathbf{v}, \bar{\mathbf{s}}, \mathbf{g}) + \frac{(\mathbf{v}, \bar{\mathbf{s}}, \bar{\mathbf{n}} + \bar{\mathbf{n}}')}{\mathbf{v}^\top (\bar{\mathbf{n}} + \bar{\mathbf{n}}')} = 0, \quad (5)$$

$$\bar{\mathbf{s}}^\top (\text{diag}(\lambda^2, \lambda^2, 0) + \mathbf{g}\mathbf{g}^\top) \bar{\mathbf{s}} = \left(\frac{\bar{\mathbf{s}}^\top (\bar{\mathbf{n}} + \bar{\mathbf{m}})}{\mathbf{v}^\top (\bar{\mathbf{n}} + \bar{\mathbf{m}})} \right)^2, \quad (6)$$

$$\bar{\mathbf{s}}^\top (\text{diag}(\lambda^2, \lambda^2, 0) + \mathbf{g}\mathbf{g}^\top) (\bar{\mathbf{n}} \times \bar{\mathbf{m}}) = \lambda (\mathbf{g}, \bar{\mathbf{n}}, \bar{\mathbf{m}}). \quad (7)$$

Here, $\text{diag}(\cdot)$ represents a 3×3 diagonal matrix, (\cdot, \cdot, \cdot) is the triple scalar product of its arguments, and \mathbf{g} is the translation vector $[-\mu, -\nu, 1]^\top$. These equations represent three independent constraints on the GBR parameters. \square

According to this proposition, the GBR ambiguity is resolved once we have identified a quadrilateral $\bar{n}\bar{m}\bar{n}'\bar{m}'$ whose pre-image $n\bar{m}n'm'$ is an isotropic-reciprocal quadrilateral with respect to the pre-image of $\bar{\mathbf{s}}$. It still remains to discuss how to identify such a quadrilateral in an input image. Given an arbitrary normal $\bar{\mathbf{n}}$, we must locate the appropriate corresponding normals $\bar{\mathbf{n}}'$ and $\bar{\mathbf{m}}$ (computationally, $\bar{\mathbf{m}}'$ is not required.) These two normals can be located sequentially by making use of the fact that the BRDF (or intensity) must be equal to that at $\bar{\mathbf{n}}$.

Since the GBR includes only translation and scale, the isotropic match $\bar{\mathbf{n}}'$ must lie on the (known) line that passes through $\bar{\mathbf{n}}$ and is orthogonal to the principal meridian (this line is the ‘‘isotropic curve’’ described in [10]). The match can be obtained as the intersection of this line and the iso-intensity contour on the plane that passes through $\bar{\mathbf{n}}$. Once $\bar{\mathbf{n}}'$ is determined, the symmetry axis of the sought trapezoid is determined as the bisector of $\bar{n}\bar{n}'$, and by finding the translation that takes this to the principal meridian, we resolve one degree of freedom in the GBR ambiguity. Specifically, if we re-parameterize Eqs. (5–7) using $\mu' = \bar{s}_2\mu - \bar{s}_1\nu$ and $\nu' = \bar{s}_1\mu + \bar{s}_2\nu$, the first equation depends only on μ' , and it is this translational degree of freedom that is resolved.

Next, we locate the reciprocal match, $\bar{\mathbf{m}}$, and resolve the remaining parameters ν' and λ using the re-parameterized Eqs. (6) and (7). This is achieved via exhaustive 2D search as follows. A hypothesis (ν'_1, λ_1) yields hypotheses for the point $\mathbf{h} \times (\mathbf{v} \times \mathbf{s})$ and BRDF value at each point. These in turn induce a hypothesis for the reciprocal match $\bar{\mathbf{m}}$ as the

¹Since \mathbf{G}^\top is an affine, the midpoint of ab is mapped to the midpoint of $\bar{a}\bar{b}$. Thus, we can use $\mathbf{G}^\top \mathbf{a} / \|\mathbf{G}^\top \mathbf{a}\| + \mathbf{G}^\top \mathbf{b} / \|\mathbf{G}^\top \mathbf{b}\| = \mathbf{G}^\top (\bar{\mathbf{a}} + \bar{\mathbf{b}}) / \|\mathbf{G}^\top (\bar{\mathbf{a}} + \bar{\mathbf{b}})\|$. We also use $\mathbf{G}^\top \mathbf{a} \times \mathbf{G}^\top \mathbf{b} / \|\mathbf{G}^\top \mathbf{a} \times \mathbf{G}^\top \mathbf{b}\| = \mathbf{G}^{-1} (\mathbf{a} \times \mathbf{b}) / \|\mathbf{G}^{-1} (\mathbf{a} \times \mathbf{b})\|$.

intersection of the iso-BRDF curve and the join of \mathbf{n} and the hypothesized $\mathbf{h} \times (\mathbf{v} \times \mathbf{s})$. Along with the known values of $\bar{\mathbf{n}}$ and $\bar{\mathbf{n}}'$, this hypothesis for $\bar{\mathbf{m}}$ induces its own estimate (ν'_2, λ_2) of the unknown GBR parameters through Eqs. (6) and (7). In this way, the expression $(\nu'_1 - \nu'_2)^2 + (\lambda_1 - \lambda_2)^2$ provides a measure of inconsistency, and the exhaustive 2D search is used to minimize this inconsistency.

Relation to previous work. This analysis sheds new light on the method of Tan et al. [10]. First, Prop. 2 shows that their two-image method for reducing the ambiguity to a classic bas-relief ambiguity using isotropy alone is minimal in the sense of requiring the fewest possible images and matched normals. But Prop. 4 shows that a method with significantly fewer requirements can be obtained by considering isotropy and reciprocity *jointly*. To obtain a Euclidean reconstruction, [10] requires five normals in two images (one isotropic pair in each image, and one reciprocal match in either). In contrast, the proposed method requires only three normals in one image. Thus, like any method that requires minimal data, the proposed method will have particular advantages when used in conjunction with robust estimation techniques such as RANSAC.

Experimental Results. Results of the proposed procedure are in Fig. 4. Following [10], each input image was decomposed into diffuse and specular components [7], and diffuse images were used to obtain a surface up to the GBR ambiguity [11]. Then, the specular component of a *single image* (top) was assumed to be spatially homogeneous and was mapped onto the projective plane via the Gaussian sphere. The GBR ambiguity was resolved by locating a single quadrilateral in this plane as described above.

The second row of Fig. 4 shows linearly-coded normal fields, where horizontal, vertical, and out-of-page components are represented by red, green, and blue, respectively. The third row shows the surface depth computed by integrating the recovered normals, and the bottom row shows the recovered surfaces. Columns from left to right show results of: 1) calibrated photometric stereo (with known lighting directions); 2) uncalibrated reconstruction up to the GBR; and 3) our auto-calibrated result. The proposed method resolves the GBR and yields a reconstruction that is very close to the calibrated result. (The median angular error in the normal field is 7.95 degrees, which is comparable to the results obtained in [10] using two images.)

5.2. Calibrated Photometric Stereo

The auto-calibrated approach described in the previous section can be applied whenever the BRDF (or a separable component of it) is isotropic and spatially uniform on the surface. More general reflectance can be accommodated when more information is available about the light sources. A very general method is that of Alldrin and Kriegman [1],




Figure 4. Resolving the GBR from a single image. *Top*: One input image and its diffuse and specular components. *Second row*: recovered normal fields, linearly-coded in RGB. *Third row*: depth maps from integrated normal fields. *Fourth row*: Reconstructed surfaces. Columns from left to right correspond to: “ground truth” obtained using calibrated sources; uncalibrated reconstruction up to a GBR; and our auto-calibrated result.

which provides a partial reconstruction for surfaces with isotropic reflectance that varies *arbitrarily* between surface points. Given a set of images $I(x, y, t)$ captured using a cone of known source directions $\mathbf{s}(t), t \in [0, 2\pi)$ centered about view direction \mathbf{v} , this method yields one component of the normal at every image point (x, y) . Specifically, for each pixel it provides the plane spanned by the unknown surface normal and the view direction, but the remaining degree of freedom in each normal cannot be recovered without additional information. In other words, if the surface is differentiable, the surface gradient direction can be recovered at each point, but the gradient magnitude is unknown. This means that one can recover the ‘iso-depth contours’ of the surface, but that these curves cannot be ordered [1].

In this section, we use reflectance symmetries on the projective plane to complete this partial reconstruction. Consider a surface $\mathcal{S} = \{x, y, z(x, y)\}$ that is described by a height field $z(x, y)$ on the image plane. A surface point with gradient z_x, z_y is mapped via the Gaussian sphere to point $\mathbf{n} \simeq (z_x, z_y, -1)$ in the projective plane, and the ambiguity in gradient magnitude from [1] corresponds to a transformation of normal field $\bar{\mathbf{n}}(x, y) \simeq \text{diag}(1, 1, \lambda(x, y))\mathbf{n}(x, y)$, where the per-pixel scaling $\lambda(x, y)$ is unknown. As depicted




Figure 5. In calibrated photometric stereo, one component of each surface normal is recovered by exploiting isotropy at each point [1]. Determining the remaining degree of freedom at each point can be interpreted as finding an unknown translation of the normal along the line through \mathbf{v} on the projective plane.

in Fig. 5, this can be interpreted as a per-pixel bas-relief transformation, where each normal \mathbf{n} is translated arbitrarily along the line vn . Now, an isotropic pair $(\mathbf{n}, \mathbf{n}')$ has two properties: 1) \mathbf{n}, \mathbf{n}' are equally distant from the origin; and 2) lines vn and vn' are symmetric across the principle meridian. The per-pixel transformation destroys the first property but preserves the second. In what follows, we seek to resolve the shape ambiguity by restoring the former.

Having equal distance to \mathbf{v} , isotropic pairs lie on circles centered about the view direction. Any two normals on such a circle is necessarily an isotropic pair with respect to one of the sources in $\{\mathbf{s}(t)\}$. Thus, for any normal \mathbf{n} , the view-centered circle on which it lies can be interpreted as the union of its isotropic matches under the set of sources $\{\mathbf{s}(t)\}$. Now, a view-centered circle is also the Gaussian-image of surface points having equal gradient magnitude: $\|\nabla z\| = \text{constant}$. Thus, if we could locate surface points with normals on such circle, we would recover a surface curve of constant gradient magnitude—a curve we will refer to as an ‘iso-slope contour’. To get a sense of how this would constrain the surface, consider that when only the iso-depth contours are known, the surface can be recovered—at best—up to a differentiable function [1]. This is because any two differentiable height fields $z_1(x, y)$ and $z_2 = h(z_1)$ will have the same set of iso-depth contours for any differentiable function $h(\cdot)$. Here we show that if they also possess the same iso-slope contours, then this arbitrary differentiable function is reduced to a classic bas-relief transformation (i.e., a linear scaling of depth).

Proposition 5. In the general case, if differentiable height fields $z_1(x, y)$ and $z_2 = h(z_1)$ are related by a differentiable function h and possess equivalent sets of iso-slope contours, the function h is linear.

Proof. From the functional relationship between z_1, z_2 :

$$\|\nabla z_2\|^2 = (\partial h / \partial z_1)^2 \|\nabla z_1\|^2.$$

If z_1 and z_2 possess the same set of iso-slope contours, either $(\partial h/\partial z_1)^2$ is constant or $\|\nabla_{z_1}\|^2 = \|\nabla_{z_2}\|^2 = 0$ along each contour. Since sets of iso-depth and iso-slope contours are generically distinct, this implies that $\partial h/\partial z_1$ is constant and h is linear. \square

Once the surface is known up to a scale factor, by Prop. 3 this factor can be determined using the known source direction for any one input images. Also, it can be shown that the ‘accidental’ case in which the sets of iso-depth and iso-slope contours are equivalent corresponds to a surface-of-revolution with the view direction \mathbf{v} as the symmetry axis.

In light of Prop. 5, it is desirable to be able to identify the iso-slope contour passing through a given image point. With the partial reconstruction, a normal $\bar{\mathbf{n}}$ is recovered up to a per-pixel bas-relief transformation. This transformation does not change the line $v\bar{\mathbf{n}}$, hence, the isotropic match $\bar{\mathbf{n}}'$ must lie on the line symmetric to $v\bar{\mathbf{n}}$ across the principal meridian. If the surface has uniform reflectance (or has a uniform separable component), the match $\bar{\mathbf{n}}'$ can be located by intersecting this line with the iso-intensity contour passing through $\bar{\mathbf{n}}$. Such isotropic matches $\bar{\mathbf{n}}'(t)$ under all light directions $\mathbf{s}(t), t \in [0, 2\pi)$ define the iso-slope contour.

This simple method for identifying iso-slope contours can be extended to surfaces having spatially-varying isotropic reflectance of the form:

$$f(x, y, \omega_i, \omega_o) = f_1(x, y) + f_2(x, y)f_3(\theta_i, \theta_o, |\phi_i - \phi_o|). \quad (8)$$

Let $I(x, y, t)$ be the recorded radiance, and at each image point (x, y) , shift and normalize these observations as

$$I_n(x, y, t) \triangleq \frac{I(x, y, t - \phi_n) - \min_t I(x, y, t)}{(\max_t I(x, y, t) - \min_t I(x, y, t))}, \quad (9)$$

where $\phi_n \in [0, 2\pi)$ is the azimuthal component of the surface normal as recovered by [1], and t is extended periodically: $t \rightarrow t + 2k\pi$ for integer k . Then, if the spatially-varying BRDF is of the form in Eq. (8), a necessary condition for two points (x_1, y_1) and (x_2, y_2) to have normal directions forming an isotropic pair is $I_n(x_1, y_1, t) = I_n(x_2, y_2, t) \forall t \in [0, 2\pi)$. This is because normalizing the temporal radiance at each pixel to $[0, 1]$ removes the effects of the spatially-varying reflectance terms f_1 and f_2 . This constraint can be used in the matching procedure above by using it in place of the iso-intensity contours.

Experimental Results. Figure 6 shows some iso-slope contours (in red) recovered using Eq. (9) to identify isotropic matches $\bar{\mathbf{n}}'(t)$ for the points marked in yellow. For comparison, the (generally distinct) iso-depth contours (recovered as in [1]) are shown in blue. The result of our procedure is shown on the right, and as a form of ‘ground truth’, we show the corresponding contours taken from the complete reconstructions obtained by Alldrin et al. [2]². Despite the differences between these two reconstruction procedures and




Figure 6. Recovered iso-slope and iso-depth curves. Red and blue curves are the iso-slope and iso-depth curves respectively. These two set of curves often do not coincide. Our result is shown on the right. For comparison, on the left is result computed from the reconstructions of [2].

their underlying BRDF models, the results are quite consistent.

6. Conclusion

We advocate the real projective plane as a tool for radiometric image analysis based on reflectance symmetries. We study the geometric structure induced in this plane by reciprocity and isotropy and show how this structure can be exploited to resolve or reduce shape ambiguities in both uncalibrated and calibrated photometric stereo.

While we restrict our attention to reciprocity and isotropy in this paper, the same framework can be used to study half-vector symmetry [6, 8], and it is likely that this symmetry provides much stronger joint constraints on shape, lighting and viewpoint. This is especially interesting in the context of calibrated photometric stereo, where one might be able to derive conditions for uniqueness that supplement recent empirical results [2]. Conditions for uniqueness based on parametric BRDF models exist [9], but for the most part, conditions that avoid the restrictions of low-parameter BRDF models have yet to be discovered.

Finally, we have exploited these symmetries for photometric stereo-based reconstruction, but one could also explore other applications, including reflectometry, illuminant estimation, and compression of appearance data.

7. Acknowledgements

We would like to thank Long Quan for the insightful discussion in the early stage of the project and Neil Alldrin for sharing his code and data. P. Tan and was supported by

²See additional results at <http://vision.ucsd.edu/~nalldrin/research/>

Singapore FRC Grant R-263-000-477-112. T. Zickler was supported by NSF Career Award IIS-0546408 and a fellowship from the Alfred P. Sloan Foundation.

References

- [1] N. Alldrin and D. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In *Proc. CVPR*, 2007.
- [2] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *Proc. CVPR*, 2008.
- [3] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. *IJCV*, 35(1):33–44, 1999.
- [4] H. S. M. Coxeter. *Introduction to Geometry, 2nd Edition*. Wiley, 1989.
- [5] J. Lu and J. Little. Reflectance and Shape from Images Using a Collinear Light Source. *IJCV*, 32(3):213–240, 1999.
- [6] F. Romeiro, Y. Vasilyev, and T. Zickler. Passive reflectometry. In *Proc. ECCV*, 2008.
- [7] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. *JOSA A*, 11:2990–3002, 1994.
- [8] M. Stark, J. Arvo, and B. Smits. Barycentric parameterizations for isotropic BRDFs. *IEEE T-VCG*, 11(2):126–138, 2005.
- [9] H. D. Tagare and R. J. P. deFigueiredo. A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *IEEE T-PAMI*, 13(2):133–152, 1991.
- [10] P. Tan, S. P. Mallick, L. Quan, D. Kriegman, and T. Zickler. Isotropy, reciprocity and the generalized bas-relief ambiguity. In *Proc. CVPR*, 2007.
- [11] A. Yuille and D. Snow. Shape and Albedo from Multiple Images Using Integrability. In *Proc. CVPR*, 1997.
- [12] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In *Proc. ECCV*, 2002.

A. Isotropic and reciprocal pairs

Proposition 6. For $\mathbf{v} = (0, 0, 1)'$ and any other three unit vectors $\mathbf{s}, \mathbf{n}, \mathbf{n}'$ where $\mathbf{n} \neq \pm \mathbf{n}'$, $\mathbf{s}^\top \mathbf{n} = \mathbf{s}^\top \mathbf{n}'$ and $\mathbf{v}^\top \mathbf{n} = \mathbf{v}^\top \mathbf{n}'$, if and only if $(\mathbf{v} \times \mathbf{s})^\top (\mathbf{n} + \mathbf{n}') = 0$ and $\frac{\mathbf{s}^\top \mathbf{n}}{\mathbf{v}^\top \mathbf{n}} = \frac{\mathbf{s}^\top \mathbf{n}'}{\mathbf{v}^\top \mathbf{n}'}$.

Proof. We first prove the forward direction. The result $\frac{\mathbf{s}^\top \mathbf{n}}{\mathbf{v}^\top \mathbf{n}} = \frac{\mathbf{s}^\top \mathbf{n}'}{\mathbf{v}^\top \mathbf{n}'}$ is trivial since it is simply the quotient of the two equations. Next, $(\mathbf{v} \times \mathbf{s})^\top (\mathbf{n} + \mathbf{n}') = 0 \Leftrightarrow -s_2(n_1 + n'_1) + s_1(n_2 + n'_2) = 0 \Leftrightarrow \frac{s_1}{s_2} = \frac{n_1 + n'_1}{n_2 + n'_2}$. On the other hand, since $n_3 = \mathbf{v}^\top \mathbf{n} = \mathbf{v}^\top \mathbf{n}' = n'_3$, we have $\mathbf{s}^\top \mathbf{n} = \mathbf{s}^\top \mathbf{n}' \Leftrightarrow$

$n_1 s_1 + n_2 s_2 = n'_1 s_1 + n'_2 s_2 \Leftrightarrow \frac{s_1}{s_2} = -\frac{n_2 - n'_2}{n_1 - n'_1}$. Hence, $(\mathbf{v} \times \mathbf{s})^\top (\mathbf{n} + \mathbf{n}') = 0 \Leftrightarrow \frac{n_1 + n'_1}{n_2 + n'_2} = -\frac{n_2 - n'_2}{n_1 - n'_1} \Leftrightarrow n_1^2 + n_2^2 = n_1'^2 + n_2'^2 \Leftrightarrow n_3 = n'_3$, which is true.

Now, we prove the reverse direction. $\frac{\mathbf{s}^\top \mathbf{n}}{\mathbf{v}^\top \mathbf{n}} = \frac{\mathbf{s}^\top \mathbf{n}'}{\mathbf{v}^\top \mathbf{n}'} \Leftrightarrow \frac{s_1}{s_2} = -\frac{n_2 n'_3 - n'_2 n_3}{n_1 n'_3 - n'_1 n_3}$ and $\frac{\mathbf{s}^\top \mathbf{n}}{\mathbf{v}^\top \mathbf{n}} = \frac{\mathbf{s}^\top \mathbf{n}'}{\mathbf{v}^\top \mathbf{n}'} \Leftrightarrow \frac{s_1}{s_2} = \frac{n_1 + n'_1}{n_2 + n'_2}$. Hence, from these two equations we have $-\frac{n_2 n'_3 - n'_2 n_3}{n_1 n'_3 - n'_1 n_3} = \frac{n_1 + n'_1}{n_2 + n'_2} \Leftrightarrow (n_1 n'_3 - n'_1 n_3)(n_1 + n'_1) + (n_2 n'_3 - n'_2 n_3)(n_2 + n'_2) = 0 \Leftrightarrow (1 - n_3^2)n'_3 - (1 - n_3'^2)n_3 + (n_1 n'_1 + n_2 n'_2)(n'_3 - n_3) = 0 \Leftrightarrow (n'_3 - n_3)(1 + \mathbf{n}^\top \mathbf{n}') = 0$. Since $\mathbf{n} \neq -\mathbf{n}'$, we have $\mathbf{v}^\top \mathbf{n}' = n'_3 = n_3 = \mathbf{v}^\top \mathbf{n}$. Then from $\frac{\mathbf{s}^\top \mathbf{n}}{\mathbf{v}^\top \mathbf{n}} = \frac{\mathbf{s}^\top \mathbf{n}'}{\mathbf{v}^\top \mathbf{n}'}$ we immediately obtain $\mathbf{s}^\top \mathbf{n}' = \mathbf{s}^\top \mathbf{n}$. \square

Proposition 7. Four unit vectors $\mathbf{v}, \mathbf{s}, \mathbf{n}, \mathbf{m}$ where \mathbf{n}, \mathbf{m} non-collinear satisfy $\mathbf{s}^\top \mathbf{n} = \mathbf{v}^\top \mathbf{m}$ and $\mathbf{v}^\top \mathbf{n} = \mathbf{s}^\top \mathbf{m}$, if and only if $\mathbf{s}^\top (\mathbf{n} + \mathbf{m}) = \mathbf{v}^\top (\mathbf{n} + \mathbf{m})$ and $\mathbf{s}^\top (\mathbf{n} \times \mathbf{m}) = \mathbf{v}^\top (\mathbf{n} \times \mathbf{m})$.

Proof. As \mathbf{n}, \mathbf{m} are distinct, they define a coordinate system in \mathbb{R}^3 along with vector $\mathbf{n} \times \mathbf{m}$. The result follows directly from representing \mathbf{v} and \mathbf{s} in this coordinate system. \square

B. Identities in the projective plane

Proposition 8. Given that $\mathbf{v} = (0, 0, 1)$ is the origin in the 2D projective plane, for any other three points $\mathbf{a}, \mathbf{b}, \mathbf{s}$ on the plane, the line ab is perpendicular to the line vs if and only if $(\mathbf{s}^\top \mathbf{a})/(\mathbf{v}^\top \mathbf{a}) = (\mathbf{s}^\top \mathbf{b})/(\mathbf{v}^\top \mathbf{b})$.

Proof. The line ab and vs are $\mathbf{a} \times \mathbf{b}$ and $\mathbf{v} \times \mathbf{s}$ respectively. Let the intersection of ab and vs with the line at infinity $\mathbf{l}_\infty = (0, 0, 1)$ be $\mathbf{A} \triangleq ab \times \mathbf{l}_\infty$ and $\mathbf{B} \triangleq vs \times \mathbf{l}_\infty$. Then $ab \perp vs$ if and only if \mathbf{A}, \mathbf{B} are harmonic conjugate points with respect to the circular points $\mathbf{I} = (1, i, 0)'$ and $\mathbf{J} = (1, -i, 0)'$, i.e., the cross ratio $cr(\mathbf{A}, \mathbf{B}, \mathbf{I}, \mathbf{J}) = -1$. Now, $cr(\mathbf{A}, \mathbf{B}, \mathbf{I}, \mathbf{J}) = -1 \Leftrightarrow A_1 B_1 + A_2 B_2 = 0$, and this reduces to $(s_1 a_1 + s_2 a_2) b_3 = (s_1 b_1 + s_2 b_2) a_3 \Leftrightarrow (\mathbf{s}^\top \mathbf{a})(\mathbf{v}^\top \mathbf{b}) = (\mathbf{s}^\top \mathbf{b})(\mathbf{v}^\top \mathbf{a}) \Leftrightarrow (\mathbf{s}^\top \mathbf{a})/(\mathbf{v}^\top \mathbf{a}) = (\mathbf{s}^\top \mathbf{b})/(\mathbf{v}^\top \mathbf{b})$. \square

Proposition 9. For any four points $\mathbf{a}, \mathbf{b}, \mathbf{s}$ and \mathbf{v} , the line ab intersects sv at the point $(\mathbf{v} + \mathbf{s}) \times (\mathbf{v} \times \mathbf{s})$ if and only if $(\mathbf{s}^\top (\mathbf{a} \times \mathbf{b})) / (\mathbf{v}^\top (\mathbf{a} \times \mathbf{b})) = (\mathbf{s}^\top (\mathbf{v} + \mathbf{s})) / (\mathbf{v}^\top (\mathbf{v} + \mathbf{s}))$ ³

Proof. The line ab and vs are $\mathbf{a} \times \mathbf{b}$ and $\mathbf{v} \times \mathbf{s}$ respectively. It is clear that $(\mathbf{v} + \mathbf{s}) \times (\mathbf{v} \times \mathbf{s})$ is a point on the line vs . This point is the intersection of ab and vs if it is also a point on ab , i.e., $(\mathbf{a} \times \mathbf{b})^\top ((\mathbf{v} + \mathbf{s}) \times \mathbf{s} \times \mathbf{v}) = 0$. On the other hand, $(\mathbf{v} + \mathbf{s}) \times \mathbf{s} \times \mathbf{v} = \mathbf{s}(\mathbf{v}^\top (\mathbf{v} + \mathbf{s})) - \mathbf{v}(\mathbf{s}^\top (\mathbf{v} + \mathbf{s}))$. Therefore, $(\mathbf{a} \times \mathbf{b})^\top ((\mathbf{v} + \mathbf{s}) \times \mathbf{s} \times \mathbf{v}) = 0 \Leftrightarrow (\mathbf{a} \times \mathbf{b})^\top \mathbf{s}(\mathbf{v}^\top (\mathbf{v} + \mathbf{s})) = (\mathbf{a} \times \mathbf{b})^\top \mathbf{v}(\mathbf{s}^\top (\mathbf{v} + \mathbf{s})) \Leftrightarrow (\mathbf{s}^\top (\mathbf{a} \times \mathbf{b})) / (\mathbf{v}^\top (\mathbf{a} \times \mathbf{b})) = (\mathbf{s}^\top (\mathbf{v} + \mathbf{s})) / (\mathbf{v}^\top (\mathbf{v} + \mathbf{s}))$. \square

³ According to Prop. 8, this is equivalent to the join of $(\mathbf{v} + \mathbf{s})$ and $(\mathbf{a} \times \mathbf{b})$ being perpendicular to vs .