O. Drbohlav and M. Chantler: Can Two Specular Pixels Calibrate Photometric Stereo?
In *ICCV 2005: Proceedings of the 10th IEEE International Conference on Computer Vision*, vol. II, pp. 1850–1857, Beijing, China, 2005.

Can Two Specular Pixels Calibrate Photometric Stereo?

Ondrej Drbohlav and Mike Chantler
Texture Lab, School of Mathematical and Computer Sciences
Heriot-Watt University
Edinburgh, UK

Abstract

Lambertian photometric stereo with unknown light source parameters is ambiguous. Provided that the object imaged constitutes a surface, the ambiguity is represented by the group of Generalised Bas-Relief (GBR) transformations. We show that this ambiguity is resolved when specular reflection is present in two *images taken under two different light source directions. We identify all configurations of the two directional lights which are singular and show that they can easily be tested for. While previous work used optimisation algorithms to apply the constraints implied by the specular reflectance component, we have developed a* linear *algorithm to achieve this goal. Our theory can be utilised to construct fast algorithms for automatic reconstruction of smooth glossy surfaces.*

1. Introduction

Photometric stereo [9] is an established method for acquiring reflectance parameters and surface normals of opaque surfaces. The reflectance and normals are recovered for every pixel from a collection of images taken by a static camera under different illumination conditions. Sufficiently accurate results are often obtained under an assumption that the surface reflectance behaves according to Lambert's model [8] which represents the reflectance at a point by a single parameter called *albedo*. The basic entity of Lambertian photometric stereo is a *scaled normal vector*. The scaled normal's direction is given by the surface normal, while its magnitude is equal to the albedo, at a given point.

If the light source directions and intensities utilised for obtaining the input data are not known, Lambertian photometric stereo reconstructs the scaled normal vectors only up to a global, essentially affine ambiguity [7]. Yet explicit calibration of the light sources involves measuring their intensity and position in space, or having a calibration object in the scene, which is both limiting and non-trivial. Any possibility of reducing the ambiguity without *explicit* calibration of light sources carries great practical potential. This is why the problem of photometric stereo auto-calibration has received considerable interest in recent years [6, 10, 1, 2, 3, 5].

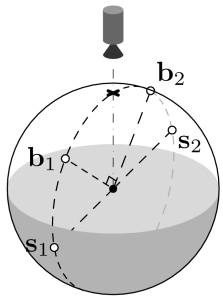

Figure 1: Specularities in two images are sufficient to calibrate photometric stereo unless the two lights $\mathbf{s}_1$ and $\mathbf{s}_2$ which produced them are in the configuration depicted. Singular configurations occur when the light sources are antipodal—or equivalently, when the two normals $\mathbf{b}_1$ and $\mathbf{b}_2$ accommodating the specularities for respective images are perpendicular and in one plane with the viewing vector.

The following methods have been developed for reducing the ambiguity of photometric stereo.

When at least six light sources are of equal intensity (or when albedo is uniform for at least six normals at a curved surface), the original ambiguity represented by the $GL(3)$ group (the group of all invertible 3×3 matrices) reduces to the group of orthogonal transformations $O(3)$ [6, 10, 1].

Another possibility is to apply the integrability constraint that requires the normals recovered by photometric stereo to correspond to a continuous surface [1, 4]. As shown by Belhumeur et al. [1], in this case the original ambiguity is reduced to an ambiguity represented by the group of generalised bas-relief (GBR) transformations. Importantly, the integrability and equal intensity constraints combined reduce the ambiguity to binary convex/concave ambiguity.

Drbohlav and Šára [3] pointed out that specular (mirror-like) reflections in images, which of course do not follow the Lambertian law, should not be just discarded in a preprocessing step, but instead they should be employed to reveal further information about the surface geometry. This resulted in the *consistent viewpoint* constraint which requires that light directions flipped along the respective specular normals all give the viewing vector. The constraint can be constructed from four specularities produced under four different lights in general configuration. It reduces the ambiguity to the 1dof set of rotations/reflections about the viewing direction. When combined with the integrability constraint, the ambiguity is again reduced to a convex/concave one [3].

Georghiades developed a similar approach [5] and presented an optimisation algorithm for solving uncalibrated photometric stereo in the case that the reflectance is a composition of Lambertian and specular (modelled as Torrance-Sparrow) components. He used a hard integrability constraint (no depth discontinuities were allowed) and claimed that three distinct light sources should be sufficient to resolve the ambiguity up to a finite number of solutions at worst, and with four light sources one of these solutions is selected uniquely.

In this paper, we present novel theory which shows that, contrary to what might be induced at a first sight from previous work [5, 3], only *two* specular pixels observed under two different light directions are needed in order to remove the ambiguity for an object with integrable surface. The necessary assumption is that the specular reflectance lobe is of limited width, such that i) specular pixels can be detected as outliers to Lambert's model and the surface can be reconstructed from the remaining Lambertian data, and ii) normals observed to reflect in a specular manner satisfy the consistent viewpoint constraint.

We derive all configurations of two lights which are *singular*, and show that they can easily be tested for.

Finally, we present a *linear* algorithm for applying the consistent viewpoint constraint. This allows the development of fast and robust methods for automatic reconstruction of smooth glossy surfaces.

2. Notation and Concepts

Surface reflection. Within this paper we adopt the following assumptions:

i) A single distant point light source is used to illuminate the object. Multiple sources are excluded.
ii) The diffuse reflectance is Lambertian with spatial albedo varying arbitrarily. The surface is *not* required to be of constant albedo.
iii) The specular component is not modelled by a parametric model; only a geometrical constraint binding the orientation of the specular normal and the viewing and light directions is used in the development of the theory.

Diffuse component. The intensity i observed at an illuminated non-specular pixel is, according to the Lambertian model,

$$i = \rho\sigma\cos\theta\,, \tag{1}$$

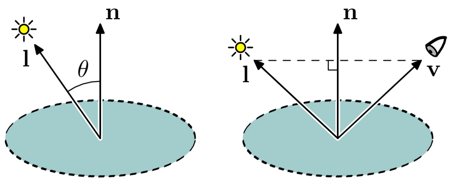

(a) (b)

Figure 2: (a) Reflection geometry for the Lambertian model of reflectance. The intensity observed is dependent on the cosine of the angle of incidence θ between the light source vector $\mathbf{l}$ and the normal vector $\mathbf{n}$. (b) At a specular point, the surface normal $\mathbf{n}$ is the bisector between the light vector $\mathbf{l}$ and the viewing vector $\mathbf{v}$.

where ρ is the albedo which represents how much light is reflected back into the air in the form of diffuse reflection, σ is the light source intensity, and θ is the angle of incidence (see Fig. 2(a)). This can be rewritten as

$$i = (\rho\mathbf{n})^\top(\sigma\mathbf{l}) = \mathbf{b}^\top\mathbf{s}\,, \tag{2}$$

where $\mathbf{n}$ is the unit surface normal vector and $\mathbf{l}$ is the unit light source vector. The right-most part of the equation expresses the Lambertian reflectance in a compact way, using $\mathbf{b} = \rho\mathbf{n}$ (the normal vector scaled by albedo), and $\mathbf{s} = \sigma\mathbf{l}$ (the light vector scaled by the light intensity); the two vectors $\mathbf{b}$ and $\mathbf{s}$ will be termed *scaled normal vector* and *scaled light vector* throughout this paper.

Specular component. The specular component is represented by a geometrical constraint stating that the specular normal $\mathbf{n}$ is a bisector between the light source direction $\mathbf{l}$ and the viewing direction $\mathbf{v}$ (see Fig. 2(b)), meaning that

$$\mathbf{v} = 2(\mathbf{n}^\top\mathbf{l})\mathbf{n} - \mathbf{l}\,. \tag{3}$$

Coordinate system. The z-axis points towards the viewpoint, and axes x–y span the image plane. Hence the viewing vector $\mathbf{v}$ has coordinates $\mathbf{v} = [0, 0, 1]^\top$.

Gauss sphere and normal sets. The normal sets are defined on a Gauss sphere as depicted in Fig. 3: Having the viewing direction $\mathbf{v}$, the occluding boundary set $\mathcal{O}$ is defined as a set of normals perpendicular to $\mathbf{v}$. The visible set $\mathcal{V}$ is defined as the set of normals which make strictly acute angles with $\mathbf{v}$.

Photometric stereo and ambiguity. Having N pixels and $M \geq 3$ illuminations, the extension of equation (2) can be written as $i_{kl} = {\mathbf{b}_k}^\top\mathbf{s}_l$ where $\mathbf{b}_k$ is the scaled normal at the k-th pixel, $\mathbf{s}_l$ is the l-th light, and i_{kl} denotes the intensity of the k-th pixel under the l-th illumination. In a matrix form, this can be written as

$$\mathtt{I} = [\mathbf{b}_1, \mathbf{b}_2, \ldots, \mathbf{b}_N]^\top[\mathbf{s}_1, \mathbf{s}_2, \ldots, \mathbf{s}_M] = \mathtt{B}^\top\mathtt{S}\,, \tag{4}$$

where matrix I stores all pixel intensity measurements, matrix B collects the $\mathbf{b}_k$'s and matrix S collects the $\mathbf{s}_l$'s. If the light sources are calibrated and therefore matrix S is known, the normals B can easily be computed from the system of equations (4). If the light matrix S is unknown, the problem of factorising the data matrix I into B and S is a classical bilinear calibration-estimation problem [7]. The solution is not unique because if $\{\mathtt{B}, \mathtt{S}\}$ is a solution then $\{\mathtt{XB}, \mathtt{X}^{-\top}\mathtt{S}\}$ where $\mathtt{X} \in GL(3)$ is an equally valid solution, as $(\mathtt{XB})^\top \mathtt{X}^{-\top}\mathtt{S} = \mathtt{B}^\top\mathtt{X}^\top\mathtt{X}^{-\top}\mathtt{S} = \mathtt{B}^\top\mathtt{S}$. However, requiring the surface normals to be integrable reduces the ambiguity from $\mathtt{X} \in GL(3)$ to $\mathtt{X} \in S_{gbr}$ where the group of generalised bas-relief (GBR) transformations S_{gbr} is the set whose elements are the following matrices [1]:

$$\mathtt{X} = \begin{bmatrix} \lambda & 0 & \mu \\ 0 & \lambda & \nu \\ 0 & 0 & \tau \end{bmatrix}, \qquad \begin{array}{l} \lambda \neq 0, \tau \neq 0\,, \\ \mu, \nu \in \mathbb{R}\,. \end{array} \tag{5}$$

Specular pairs. Let $\mathbf{s}$ be a scaled light, and $\mathbf{b}$ be a scaled normal at a pixel which is observed to reflect in a specular manner under the light $\mathbf{s}$. Then $\mathbf{b}$ and $\mathbf{s}$ are called a *specular pair* and are denoted $\mathbf{b} \odot \mathbf{s}$.

Consistent viewpoint constraint. If the light of direction $\mathbf{l}$ and normal of direction $\mathbf{n}$ satisfy Equation (3) for the viewing direction $\mathbf{v}$ then they are said to obey the *consistent viewpoint* constraint.

Note on scale. The photometric stereo ambiguity includes the ambiguity of a *global* scale in recovered scaled normals B and scaled lights S because scaling them by κ and $1/\kappa$ ($\kappa \neq 0$), respectively, does not change $\mathtt{B}^\top\mathtt{S}$. The global scale affects only the overall scaling of recovered albedos and light intensities, and is very difficult to determine without sophisticated radiometric measurements. Fortunately, the global scale is hardly ever even needed in practise; in applications such as texture mapping, albedos are commonly normalised before being used.

To avoid unnecessarily complicated presentation, we adopt the convention that we omit the words 'up to a scale' whenever there is no risk of confusion. This applies to the whole article.

3. Theory

This Section presents the principal results of this paper. Firstly, Section 3.1 shows that there exists a linear (SVD-based) algorithm for enforcing the consistent viewpoint constraint. Secondly, Section 3.2 provides a proof that two specular pairs are sufficient for enforcing the constraint, and all singular configurations are identified as well. The theory is developed under an assumption that the normals and lights have been determined up to a GBR ambiguity, i.e. that the integrability constraint has already been enforced.

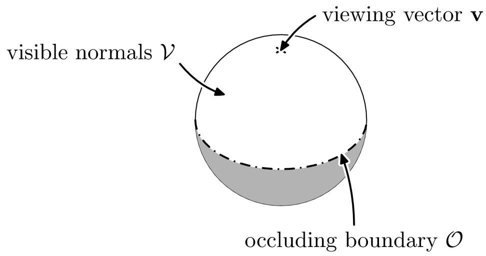

Figure 3: Gauss sphere and definition of sets on it: the viewing vector $\mathbf{v}$, the occluding boundary set $\mathcal{O}$ consisting of directions perpendicular to $\mathbf{v}$, and the set of visible normals $\mathcal{V}$ which make *strictly* acute angle with $\mathbf{v}$.

3.1. Linear algorithm

This Section presents theory leading to a linear algorithm enforcing the consistent viewpoint constraint on integrable normals. The algorithm is outlined in Table 1 (Steps 3–4).

Let $\mathbf{b} \odot \mathbf{s}$ be a specular pair, and let A denote the GBR transformation which makes the data compatible with the consistent viewpoint constraint (3). Substituting the transformed unit normal $\mathbf{n} = \mathtt{A}\mathbf{b}/\|\mathtt{A}\mathbf{b}\|$ and the transformed unit light $\mathbf{l} = \mathtt{A}^{-\top}\mathbf{s}/\|\mathtt{A}^{-\top}\mathbf{s}\|$ into (3), there must hold

$$\mathbf{v} = \frac{2\left[(\mathtt{A}\mathbf{b})^\top(\mathtt{A}^{-\top}\mathbf{s})\right]\mathtt{A}\mathbf{b}}{\|\mathtt{A}\mathbf{b}\|^2\|\mathtt{A}^{-\top}\mathbf{s}\|} - \frac{\mathtt{A}^{-\top}\mathbf{s}}{\|\mathtt{A}^{-\top}\mathbf{s}\|}\,. \tag{6}$$

Left-multiplying both sides of the equation by matrix $\mathtt{A}^\top$ and by a scalar $\|\mathtt{A}\mathbf{b}\|^2\|\mathtt{A}^{-\top}\mathbf{s}\|$ we obtain

$$\begin{aligned} &\|\mathtt{A}\mathbf{b}\|^2\|\mathtt{A}^{-\top}\mathbf{s}\|\mathtt{A}^\top\mathbf{v} = \\ &= 2[(\mathtt{A}\mathbf{b})^\top(\mathtt{A}^{-\top}\mathbf{s})]\mathtt{A}^\top\mathtt{A}\mathbf{b} - \|\mathtt{A}\mathbf{b}\|^2\mathbf{s}\,. \end{aligned} \tag{7}$$

First, as A is a GBR transformation taking the form of (5), expression $\mathtt{A}^\top\mathbf{v}$ on the left-hand side of the equation can be rewritten as $\mathtt{A}^\top\mathbf{v} = \tau\mathbf{v}$. Hence, with substitutions

$$\alpha = \|\mathtt{A}\mathbf{b}\|^2\|\mathtt{A}^{-\top}\mathbf{s}\|\tau\,, \tag{8}$$

$$\mathtt{P} = \mathtt{A}^\top\mathtt{A}\,, \tag{9}$$

Equation (7) is rewritten as

$$\alpha\mathbf{v} = 2(\mathbf{b}^\top\mathbf{s})\mathtt{P}\mathbf{b} - (\mathbf{b}^\top\mathtt{P}\mathbf{b})\mathbf{s}\,, \tag{10}$$

where we applied the fact that $(\mathtt{A}\mathbf{b})^\top(\mathtt{A}^{-\top}\mathbf{s}) = \mathbf{b}^\top\mathbf{s}$, and rewrote the factor $\|\mathtt{A}\mathbf{b}\|^2$ as $\|\mathtt{A}\mathbf{b}\|^2 = (\mathbf{b}^\top\mathtt{A}^\top\mathtt{A}\mathbf{b}) = (\mathbf{b}^\top\mathtt{P}\mathbf{b})$. Now, left-multiplying this equation by $\mathbf{b}^\top$ we obtain

$$\alpha\mathbf{b}^\top\mathbf{v} = (\mathbf{b}^\top\mathtt{P}\mathbf{b})(\mathbf{b}^\top\mathbf{s})\,, \tag{11}$$

and, therefore, an alternative to (8) for expressing α is (provided that $\mathbf{b}^\top\mathbf{v} \neq 0$) is

$$\alpha = \frac{(\mathbf{b}^\top\mathtt{P}\mathbf{b})(\mathbf{b}^\top\mathbf{s})}{\mathbf{b}^\top\mathbf{v}}\,. \tag{12}$$

Putting this expression for α back into (10) gives

$$\frac{(\mathbf{b}^\top \mathtt{P}\mathbf{b})(\mathbf{b}^\top \mathbf{s})}{\mathbf{b}^\top \mathbf{v}}\mathbf{v} = 2(\mathbf{b}^\top \mathbf{s})\mathtt{P}\mathbf{b} - (\mathbf{b}^\top \mathtt{P}\mathbf{b})\mathbf{s} \tag{13}$$

which is finally rewritten as (still requiring $\mathbf{b}^\top \mathbf{v} \neq 0$)

$$(\mathbf{b}^\top \mathtt{P}\mathbf{b})(\mathbf{b}^\top \mathbf{s})\mathbf{v} = 2(\mathbf{b}^\top \mathbf{s})(\mathbf{b}^\top \mathbf{v})\mathtt{P}\mathbf{b} - (\mathbf{b}^\top \mathtt{P}\mathbf{b})(\mathbf{b}^\top \mathbf{v})\mathbf{s}\,. \tag{14}$$

This equation is clearly homogeneous and linear in the elements of matrix $\mathtt{P}$, and all the other entries in the equation are known. The matrix $\mathtt{P}$, being $\mathtt{A}^\top\mathtt{A}$ where $\mathtt{A}$ is a GBR (5), is

$$\mathtt{P} = \begin{bmatrix} \lambda^2 & 0 & \mu\lambda \\ 0 & \lambda^2 & \nu\lambda \\ \mu\lambda & \nu\lambda & \tau^2+\mu^2+\nu^2 \end{bmatrix} \stackrel{\text{def}}{=} \begin{bmatrix} p_1 & 0 & p_3 \\ 0 & p_1 & p_4 \\ p_3 & p_4 & p_2 \end{bmatrix}. \tag{15}$$

Hence there are four unknowns p_1, p_2, p_3, p_4 to be found. Each specular pair $\mathbf{b} \odot \mathbf{s}$ produces one vector equation (14). With sufficient number of specular pairs, the matrix $\mathtt{P}$ will be determined up to a scale but otherwise uniquely, by solving the system of linear equations

$$\mathtt{M}\mathbf{p} = \mathbf{0}\,, \tag{16}$$

where $\mathtt{M}$ is constructed from the known entries in (14) and $\mathbf{p} = [p_1, p_2, p_3, p_4]^\top$. Finally, matrix $\mathtt{A}$ can be found by factorising the matrix $\mathtt{P} = \mathtt{A}^\top\mathtt{A}$ as follows: $\lambda = \pm\sqrt{p_1}$ is computed first, $\mu = p_3/\lambda$ and $\nu = p_4/\lambda$ follows. The last unknown τ is then computed as $\tau = \pm\sqrt{p_2 - \mu^2 - \nu^2}$. In summary, the matrix $\mathtt{A}$ is computed from $\mathtt{P}$ as

$$\mathtt{A} \;=\; \begin{bmatrix} s^\pm\sqrt{p_1} & 0 & s^\pm\frac{p_3}{\sqrt{p_1}} \\ 0 & s^\pm\sqrt{p_1} & s^\pm\frac{p_4}{\sqrt{p_1}} \\ 0 & 0 & t^\pm\sqrt{p_2 - \left(\frac{p_3^2}{p_1} + \frac{p_4^2}{p_1}\right)} \end{bmatrix}$$

$$s^\pm = \pm 1\,, t^\pm = \pm 1\,. \tag{17}$$

The sign $s^\pm$ reflects the surface around the plane perpendicular to the viewing direction, and corresponds to the convex-concave ambiguity. The sign $t^\pm$ reflects the third normal component and should always be chosen such that the normals are inclined towards the viewer.

Note that Equations (12–14) are all subject to the constraint that $\mathbf{b}^\top\mathbf{v} \neq 0$. Under such a constraint, Equation (14) is *equivalent* with the consistent viewpoint constraint (6) because operations used in between the two involve only multiplication by a nonzero constant, or by an invertible matrix. What is left to be analysed is whether there is any discrepancy between Equations (14) and (6) in case that $\mathbf{b}^\top\mathbf{v} = 0$. First, from (11) it follows that $(\mathbf{b}^\top\mathbf{v} = 0) \Rightarrow (\mathbf{b}^\top\mathbf{s} = 0)$ because $\alpha \neq 0$ (due to (8) and regularity of $\mathtt{A}$) and $\mathbf{b}^\top\mathtt{P}\mathbf{b} \neq 0$ (due to $\mathtt{P}$ being positive definite). Thus $\mathbf{b}^\top\mathbf{v} = 0$ implies vanishing of (14). On the other hand, Lemma 2 in the Appendix says that when $\mathbf{b}^\top\mathbf{v} = 0$ then i) the true normal $\mathtt{A}\mathbf{b}$ is on an occluding boundary, and ii) for specular pairs with a normal on an occluding boundary, the consistent viewpoint constraint is valid *always*, under *arbitrary* GBR transformations. In summary, therefore, Equation (14) is equivalent to (6) if $\mathbf{b}^\top\mathbf{v} \neq 0$; and if $\mathbf{b}^\top\mathbf{v} = 0$ then the whole equation (14) vanishes while the original constraint represented by (6) is valid under all GBR transformations. Equation (14) thus exhaustively represents the consistent viewpoint constraint.

Having developed the linear algorithm, the question now is how many specular pairs are needed in order that matrix $\mathtt{M}$ in (16) is of rank 3 and the solution $\mathbf{p}$ is therefore fully constrained.

1. Take the input data $\mathtt{I}$ and factorise it into the normals $\mathtt{B}$ and the lights $\mathtt{S}$ according to (4).
2. Enforce the integrability of normals $\mathtt{B}$ by computing the matrix $\mathtt{A}^{\mathtt{int}}$ using the algorithm [10]. Then make the data consistent with the integrability constraint by $\mathbf{b} \hookleftarrow \mathtt{A}^{\mathtt{int}}\mathbf{b}$ for all normals, and $\mathbf{s} \hookleftarrow (\mathtt{A}^{\mathtt{int}})^{-\top}\mathbf{s}$ for all lights.
3. Identify images in which specularity is present. For each of them, take the scaled normal $\mathbf{b}$ at a specular pixel and the scaled light $\mathbf{s}$ corresponding to that image, and form a linear system of equations for $\mathbf{p}$, as described in Equations (14–16).
4. Solve for $\mathbf{p}$ and compute the matrix $\mathtt{A}$ as described by Equation (17). Apply the transformation to normals and lights obtained in Step 2.

Table 1: The overview of the algorithm for calibrating photometric stereo.

3.2. Number of specular pairs needed

The problem of what constitutes a sufficient number of specular pairs is analysed in a natural basis in which the normals and lights are the true (unambiguous) ones. Note that this choice of basis does not affect the generality of the results. With a sufficient number of specular pairs, $\mathtt{P}$ will be constrained to be an identity matrix and the vector $\mathbf{p}$ thus should be $\mathbf{p} = [1, 1, 0, 0]^\top$ in such case.

Our analysis is led as follows. Having a true specular pair $\mathbf{b} \odot \mathbf{s}$, we find that Equation (14) can only be valid if the normal $\mathbf{b}$ is an eigenvector of $\mathtt{P}$. Subsequently, we analyse what eigenvectors the matrix $\mathtt{P}$ has depending on the values of p_1, p_2, p_3, p_4, after which we will be ready to state the Theorem concerning the sufficiency of two specular pairs for disambiguating the photometric stereo. The Theorem concludes this Section.

Lemma 1 (Specular normal is an eigenvector of P) *Let* $\mathbf{b} \odot \mathbf{s}$ *be a true specular pair. Any solution* P *to (14) must be such that the specular normal* $\mathbf{b}$ *is its eigenvector.*

Proof 1 Having the true scaled normal $\mathbf{b}$ which is specular, the corresponding light source can be constructed as*

$$\mathbf{s} = 2(\mathbf{b}^\top \mathbf{v})\mathbf{b} - (\mathbf{b}^\top \mathbf{b})\mathbf{v}\,. \tag{18}$$

From this equation it follows by left-multiplying by $\mathbf{b}^\top$ that

$$\mathbf{b}^\top \mathbf{s} = (\mathbf{b}^\top \mathbf{b})(\mathbf{b}^\top \mathbf{v})\,. \tag{19}$$

Utilising (19) and (18) in (14), we obtain

$$\begin{aligned}(\mathbf{b}^\top \mathtt{P}\mathbf{b})(\mathbf{b}^\top \mathbf{b})\cancel{(\mathbf{b}^\top \mathbf{v})}\mathbf{v} &= 2(\mathbf{b}^\top \mathbf{b})(\mathbf{b}^\top \mathbf{v})^{\cancel{2}}\mathtt{P}\mathbf{b} \\ &-(\mathbf{b}^\top \mathtt{P}\mathbf{b})\cancel{(\mathbf{b}^\top \mathbf{v})}\left\{2(\mathbf{b}^\top \mathbf{v})\mathbf{b} - (\mathbf{b}^\top \mathbf{b})\mathbf{v}\right\}\end{aligned} \tag{20}$$

and after expanding the second line we find that the left-hand term $(\mathbf{b}^\top \mathtt{P}\mathbf{b})(\mathbf{b}^\top \mathbf{b})\mathbf{v}$ cancels out with the identical term on the right hand side. This leaves

$$(\mathbf{b}^\top \mathbf{b})\cancel{(\mathbf{b}^\top \mathbf{v})}\mathtt{P}\mathbf{b} = (\mathbf{b}^\top \mathtt{P}\mathbf{b})\cancel{(\mathbf{b}^\top \mathbf{v})}\mathbf{b}\,. \tag{21}$$

This equation is obviously of the form of a characteristic equation for P. □

Following this, P is analysed as for its spectrum and eigenvectors. These results are derived in full detail in the Appendix (Lemma 3 and 4), and are summarised in Table 2 which shows the three basic cases which together give all possibilities for the values of $\mathbf{p} = [p_1, p_2, p_3, p_4]^\top$.

Case (i) Matrix P is a scaled identity matrix, having $p_1 = p_2$ on its diagonal. In that case, p_1 is a three-fold eigenvalue and the set of eigenvectors $\mathbb{E}_{p_1}$ is equal to the sphere $\mathcal{S}$.

Case (ii) Matrix P is a diagonal matrix with its first two diagonal elements equal to p_1 and the third one equal to $p_2 \neq p_1$. Then the eigenvector corresponding to p_2 is the viewing vector ($\mathbb{E}_{p_2} = \mathbf{v}$) and the set of eigenvectors corresponding to the two-fold eigenvalue p_1 is the occluding boundary $\mathbb{E}_{p_1} = \mathcal{O}$.

Case (iii) Matrix P has off-diagonal elements p_3 and p_4 which are not jointly zero. In that case, the spectrum is non-degenerate and one of the eigenvalues is always p_1. Its corresponding eigenvector is on the occluding boundary ($\mathbb{E}_{p_1} \in \mathcal{O}$). The other two eigenvectors are perpendicular to it, as well as to each other.

	Case (i)	Case (ii)	Case (iii)
Matrix el.	$p_1 = p_2$ $p_3 = 0$ $p_4 = 0$	$p_1 \neq p_2$ $p_3 = 0$ $p_4 = 0$	$p_3 \neq 0$ $\vee\, p_4 \neq 0$
Eigvals ϵ	$\{p_1, p_1, p_1\}$	$\{p_1, p_1, p_2\}$	$\{p_1, x, y\}$ all distinct
Eigenvectors $\mathbb{E}$	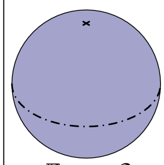 $\mathbb{E}_{p_1} = \mathcal{S}$	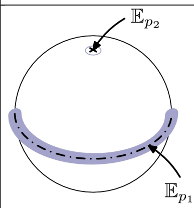 $\mathbb{E}_{p_2}$ $\mathbb{E}_{p_1}$	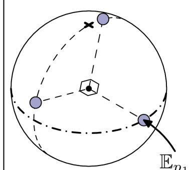 $\mathbb{E}_{p_1}$

Table 2: Summarising the characteristics of P depending on values of its elements.

With these results, we can now prove the principal theorem of this paper (the Theorem is illustrated in Fig. 1).

Theorem [Two specular pairs are sufficient] *Let there be two true specular pairs* $\mathbf{b}_1 \odot \mathbf{s}_1$ *and* $\mathbf{b}_2 \odot \mathbf{s}_2$. *If i) both* $\mathbf{b}_1$ *and* $\mathbf{b}_2$ *are strictly visible, and ii) the two light sources* $\mathbf{s}_1$ *and* $\mathbf{s}_2$ *are linearly independent (i.e. not of the same nor opposite directions) then the two pairs constrain the matrix* P *to be an identity matrix.*

Proof. The first requirement ($\mathbf{b}_1, \mathbf{b}_2 \in \mathcal{V}$) is needed because, as we discussed previously, the normals on an occluding boundary $\mathcal{O}$ are non-informative (cf. Lemma 2 in the Appendix). Next, $\mathbf{s}_1$ and $\mathbf{s}_2$ clearly must not be of the same direction, as $\mathbf{b}_1$ and $\mathbf{b}_2$ would then be the same as well, and that would correspond to employing just one specular pair. Having two distinct lights then, the corresponding specular normals will be distinct as well. Considering these $\mathbf{b}_1$ and $\mathbf{b}_2$, they can only be eigenvectors of Case (i) or Case (iii) because Case (ii) can not accommodate two distinct eigenvectors (the only allowable eigenvector would be $\mathbf{v}$, cf. Table 2). They thus favour Case (i) unless they are perpendicular and in one plane with $\mathbf{v}$ in which a case they would favour Case (iii) as well. But this can happen only when the corresponding lights are antipodal. □

4. Experiment

To show the performance of the theory developed in the previous Section, we conducted an experiment with a smooth glossy object: a glazed china teapot (see Fig. 4). The images were acquired by a 12 bit cooled camera (COOL-1300 by

*Such light vector s possibly differs in intensity from the true light vector. But this does not affect the validity of the proof.

Vosskühler) under tungsten illumination (150W, stabilised direct current). The light was moved by hand around the object. No information about lights has been measured nor recorded. The data were processed as outlined by the algorithm in Table 1. In a more detail,

1. Input data I were factorised into B and S using robust factorisation as in [3].
2. Integrability was enforced as in [10].
3. Specular points in two images (see Fig. 4) were marked manually or, identified by a robust algorithm.
4. The consistent viewpoint constraint was enforced as described in Section 3.1.

The results are illustrated in Fig. 5. The ambiguity in recovered scaled normals makes it seem as if the teapot has been illuminated from an alternative direction from that used to create the image. This is obvious in the first column where the scaled normals obtained by factorisation are re-illuminated from the right, top and camera directions. When the integrability constraint is applied (see the second column) then this is still true, although there is one image (the one illuminated from the camera direction) which looks dim but otherwise as it is expected. This is in agreement with the form of GBR transformations (5) whose effect on the z-component of recovered normals is a pure scaling. It is only when both the integrability and the consistent viewpoint constraints are enforced (see the third column) that the ambiguity is removed†. This fact is further demonstrated on recovered albedos (see last two rows in Fig. 5) and integrated shape (see Fig. 6).

We also made a prototype of an algorithm which can work fully automatically. The only difference with respect to the basic algorithm is that the specularities are not marked manually. Instead, a threshold is set which identifies the *candidates* for specular points. RANSAC is then used to identify the consistent set of specularities. The results obtained were very similar to the outcome of the basic algorithm. Full implementation of the algorithm is the subject of future work.

5. Summary and Conclusions

As a principal result of this paper, we have shown that the answer to the question posed in the title is positive. If the object imaged constitutes a surface then photometric stereo can be calibrated from as few as two specular pixels, provided that these are produced by lights which are not antipodal.

This has several practical implications. Firstly, the theory presented in this paper provides a simple guideline for input data acquisition: the only fact to keep in mind is to

†The remaining convex-concave ambiguity was resolved by hand.

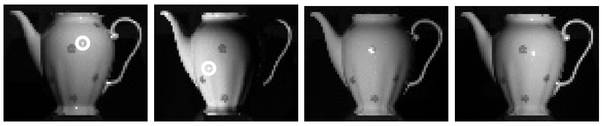

Figure 4: Four out of twenty images used for photometric stereo reconstruction. The two specularities identified in the first two images by a circle were used for applying the consistent viewpoint constraint in the basic algorithm. In a robust extension of the algorithm, manual marking of specularities is no longer needed.

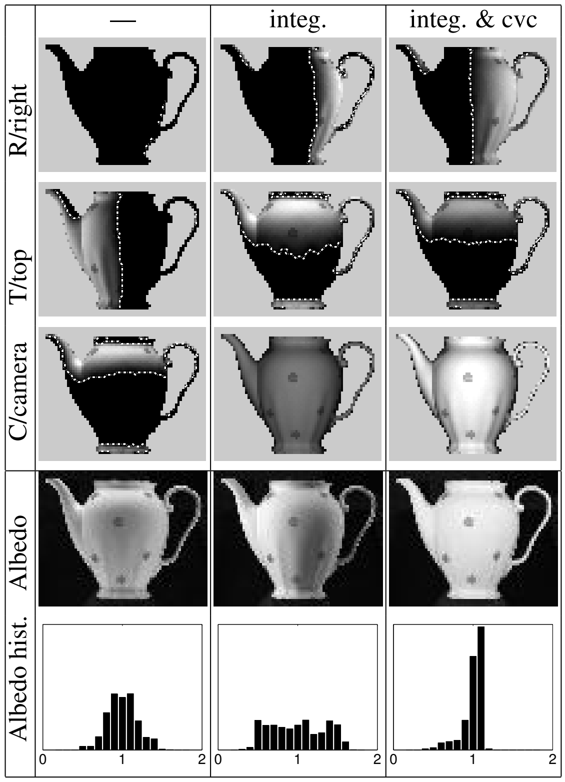

Figure 5: Removing the ambiguity: Scaled normals re-illuminated from the right (R), from the top (T), from the camera (C), and the recovered albedo, with constraints applied as indicated above the table (—: no constraint, integ: integrability constraint, cvc: consistent viewpoint constraint). It is only after both integrability and consistent viewpoint constraints are applied (last column) that the shape and albedo are recovered unambiguously. A light source of unit intensity was used for the rendering. The white lines mark the shadow boundaries which are again symmetric as expected in the last column. Also note that the histogram of albedo (last row) is sharp for the last case, as should be for a uniform object (for sake of fair comparison, the albedo is mean-normalised to 1 and the scales of both axes are equal in all three histogram plots).

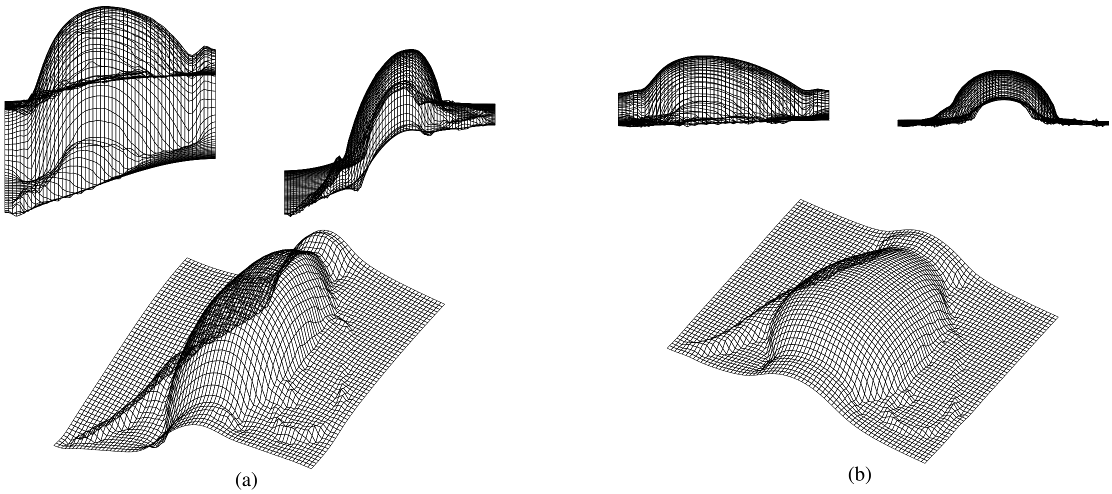

Figure 6: The result of integration before (a) and after (b) enforcing the consistent viewpoint constraint.

produce two specularities under lights which are not antipodal. In practise, it is often the case that the lights are not inclined from the viewing direction more than, say, 70 degrees, in which case antipodal lighting cannot occur. Secondly, for RANSAC-based algorithms which can separate 'proper' specular reflections from inter-reflections, it is important that the number of points to be sampled from the specularity candidate set is as few as *two* because it means that such algorithms can find the solution quickly even in presence of a high number of outliers.

The result of this paper is directly applicable for surfaces with a very sharp specular lobe. As the specular pixels on such surfaces are saturated in ordinary cameras, attempts to match the input data with intensities suggested by a parametric model are difficult.

This paper also impacts on the case where the specular lobe is broader, and thus more accurate parametric optimisation methods can be used. In that case, the linear algorithm presented in this article represents a way to obtain an initial guess for starting the optimisation algorithm.

Acknowledgements

O. Drbohlav acknowledges the generous support by the Marie Curie Intra-European Fellowship No. 506053 (PhoCal) within the 6th European Community Framework Programme. The authors wish to thank Radim Šára from the Center for Machine Perception (CMP, Czech Technical University in Prague) for fruitful discussions and comments on the manuscript; and the anonymous reviewers for their suggestions regarding the paper content and future work. The data used in this article were acquired at the Czech Technical University when O. Drbohlav was with the CMP, Prague.

References

[1] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. *International Journal of Computer Vision*, 35(1):33–44, 1999.

[2] O. Drbohlav and R. Šára. Unambiguous determination of shape from photometric stereo with unknown light sources. In *Proc. IEEE International Conference on Computer Vision*, volume 1, pages 581–586, 2001.

[3] O. Drbohlav and R. Šára. Specularities reduce ambiguity of uncalibrated photometric stereo. In *Proc. European Conference on Computer Vision*, volume 2, pages 46–60, 2002.

[4] J. Fan and L. B. Wolff. Surface curvature and shape reconstruction from unknown multiple illumination and integrability. *Computer Vision and Image Understanding*, 65(2):347–359, 1997.

[5] A. Georghiades. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In *Proc. IEEE International Conference on Computer Vision*, volume 2, pages 816–823, 2003.

[6] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *Journal of Optical Society of America A*, 11(11):3079–3089, 1994.

[7] J. J. Koenderink and A. J. van Doorn. The generic bilinear calibration-estimation problem. *International Journal of Computer Vision*, 23(3):217–234, 1997.

[8] J. H. Lambert. *Photometria sive de mensura et gradibus luminis, colorum et umbra*. Augustae Vindelicorum, Basel, 1760.

[9] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.

[10] A. Yuille and D. Snow. Shape and albedo from multiple images using integrability. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pages 158–164, 1997.

Appendix

Lemma 2 (Specularities on the occluding boundary) *Let* $\mathbf{b} \odot \mathbf{s}$ *be a specular pair and let* $\mathbf{b}^\top \mathbf{v} = 0$*. Such a pair represents a specular reflection occurring on an occluding boundary, and is* non-informative *in the sense that it does not constrain the GBR ambiguity in any way.*

Proof 2 If $\mathbf{b}^\top \mathbf{v} = 0$ then the third component of $\mathbf{b}$ is zero, and clearly will stay zero under any GBR transformation (5) (i.e. $(\mathtt{X}\mathbf{b})^\top \mathbf{v} = 0$). As the matrix $\mathtt{A}$ which disambiguates the data is a GBR, we conclude that the true normal $\mathtt{A}\mathbf{b}$ must be on an occluding boundary.

Next, if the true normal is on the occluding boundary and is specular then the corresponding light source $\mathbf{s}$ has to be directly opposite the viewing direction, $\mathbf{s} = -\sigma \mathbf{v}$ (recall that $\sigma > 0$ denotes light intensity). Under a GBR (5), the light $\mathbf{s}$ is transformed by $\mathtt{X}^{-\top}$ which takes the form [1]

$$\mathtt{X}^{-\top} = \begin{bmatrix} \frac{1}{\lambda} & 0 & 0 \\ 0 & \frac{1}{\lambda} & 0 \\ -\frac{\mu}{\lambda\tau} & -\frac{\nu}{\lambda\tau} & \frac{1}{\tau} \end{bmatrix}, \qquad \begin{array}{l} \lambda \neq 0, \tau \neq 0\,, \\ \mu, \nu \in \mathbb{R}\,. \end{array} \tag{22}$$

which implies $\mathtt{X}^{-\top}\mathbf{s} = 1/\tau \mathbf{s}$, and thus the direction of the light source is preserved. In addition, $\mathtt{X}\mathbf{b}$ stays on the occluding boundary. It thus follows that a specular pair $\mathbf{b} \odot \mathbf{s}$ with $\mathbf{b}^\top \mathbf{v} = 0$ will *always* be in accordance with the consistent viewpoint constraint (3) under *any* GBR transformation. □

Lemma 3 (Spectrum of $\mathtt{P}$) *Depending on the values of* p_1, p_2, p_3, p_4*, the spectrum* Σ *of matrix* $\mathtt{P}$ *from Eq. (15) is:*

(i) If $p_1 = p_2 \wedge p_3 = p_4 = 0$ *then* $\Sigma = \{p_1, p_1, p_1\}$*,*

(ii) If $p_1 \neq p_2 \wedge p_3 = p_4 = 0$ *then* $\Sigma = \{p_1, p_1, p_2\}$*, and*

(iii) If $p_3 \neq 0 \vee p_4 \neq 0$ *then* $\Sigma = \{p_1, x, y\}$ *with* p_1, x, y *being distinct.*

Proof 3 The first two cases (i) and (ii) are obvious as, under the conditions specified, matrix $\mathtt{P}$ is diagonal. To show the non-degeneracy of eigenvalues under the condition $p_3 \neq 0 \vee p_4 \neq 0$ needed for Case (iii) requires some analysis. We will prove the claim by contradiction, showing that the degeneracy of eigenvalues requires $p_3 = 0 \wedge p_4 = 0$. Denoting the eigenvalue by ϵ, the characteristic polynomial of the matrix $\mathtt{P}$ is

$$(p_1 - \epsilon)^2 (p_2 - \epsilon) - (p_1 - \epsilon)(p_3^2 + p_4^2)\,, \tag{23}$$

which shows that one of the eigenvalues is always equal to $\epsilon = p_1$. The other two eigenvalues are computed by solving the quadratic equation in ϵ

$$\begin{aligned} (p_1 - \epsilon)(p_2 - \epsilon) - (p_3^2 + p_4^2) &= \\ \epsilon^2 - \epsilon(p_1 + p_2) + p_1 p_2 - (p_3^2 + p_4^2) &= 0\,. \end{aligned} \tag{24}$$

The spectrum will be degenerate only if equal to $\Sigma = \{p_1, p_1, x\}$, or to $\Sigma = \{p_1, x, x\}$, with $x > 0$. As for the first possibility, $\Sigma = \{p_1, p_1, x\}$, setting $\epsilon = p_1$ in (24) implies $p_3^2 + p_4^2 = 0$ and, therefore, $p_3 = p_4 = 0$ which is in contradiction with the assumption that at least one of these two is nonzero. The second possibility $\Sigma = \{p_1, x, x\}$ would require the discriminant of the quadratic equation (24) to vanish. This would mean

$$\begin{aligned} 0 = (p_1 + p_2)^2 - 4(p_1 p_2 - p_3^2 - p_4^2) &= \\ (p_1 - p_2)^2 + 4(p_3^2 + p_4^2)\,. \end{aligned} \tag{25}$$

Therefore, besides $p_1 = p_2$ this requires $p_3 = p_4 = 0$, which is a contradiction again. The last case (iii) is thus proved.

Finally, note that the three cases listed cover all possibilities for the values of p_1, p_2, p_3, p_4 and are mutually exclusive. Hence the lemma addresses the spectrum of all matrices $\mathtt{P}$. □

Lemma 4 (Eigenvectors of $\mathtt{P}$) *Depending on the values of* p_1, p_2, p_3, p_4*, the eigenvectors* $\mathbb{E}$ *of the matrix* $\mathtt{P}$ *are as follows:*

(i) If $p_1 = p_2 \wedge p_3 = p_4 = 0$ *then all vectors from the unit sphere are eigenvectors:* $\mathbb{E}_{p_1} = \mathcal{S}$*.*

(ii) If $p_1 \neq p_2 \wedge p_3 = p_4 = 0$ *then* $\mathbb{E}_{p_1} = \mathcal{O}$ *and* $\mathbb{E}_{p_2} = \mathbf{v}$*.*

(iii) If $p_3 \neq 0 \vee p_4 \neq 0$ *then there are exactly three eigenvectors. They form an orthogonal basis and the one corresponding to* $\epsilon = p_1$ *is on an occluding boundary,* $\mathbb{E}_{p_1} \in \mathcal{O}$*.*

Proof 4 The claims (i) and (ii) are again clearly true (the matrix $\mathtt{P}$ is under their conditions diagonal and of spectrum identified in Lemma 3). Claim (iii) again requires some analysis. First, matrix $\mathtt{P}$ is symmetric and positive-definite and hence the eigenvectors corresponding to different eigenvalues are orthogonal. Additionally, it was shown in Lemma 3 that the eigenvalues are distinct and, therefore, it follows that the eigenvectors are exactly three and form an orthogonal basis. Thus, it is left to be proved that $\mathbb{E}_{p_1}$ is on an occluding boundary, which is easy. For $\epsilon = p_1$, the corresponding singular matrix is

$$\mathtt{P} - p_1 \mathtt{1} = \begin{bmatrix} 0 & 0 & p_3 \\ 0 & 0 & p_4 \\ p_3 & p_4 & p_2 - p_1 \end{bmatrix}. \tag{26}$$

As p_3 and p_4 are not jointly zero, the corresponding eigenvector $\mathbb{E}_{p_1}$ is then clearly $[p_4, -p_3, 0]^\top$. The third component of it is zero, and therefore it is indeed on an occluding boundary, $\mathbb{E}_{p_1} \in \mathcal{O}$. □