

Unambiguous determination of shape from photometric stereo with unknown light sources

Ondřej Drbohlav and Radim Šára
Center for Machine Perception
Czech Technical University in Prague
{drbohlav,sara}@cmp.felk.cvut.cz

Abstract

Photometric stereo with uncalibrated lights determines surface orientations ambiguously up to any regular transformation. If the surface reflectance model is separable with respect to the illumination and viewing directions, its inherent symmetries enable to design two previously unrecognized constraints on normals that reduce this ambiguity. The two constraints represent projections of normals onto planes perpendicular to the viewing and illumination directions, respectively. We identify the classes of transformations that leave each constraint invariant. We construct the constraints using polarization measurement under the assumption of separable reflectance model for smooth dielectrics. We verify that applying the first constraint together with the integrability constraint results in bas-relief ambiguity, while application of the second constraint on integrable normals reduces the ambiguity to convex/concave ambiguity. Importantly, the latter result is also obtained when the first and second constraints alone are combined.

1. Introduction

Photometric stereo [10] determines local surface orientations from images of an object viewed from a fixed position, but illuminated from at least three non-collinear directions. Most often, it is assumed that the surface reflectance is Lambertian, which implies that the pixel brightness $I_{i,j}$ of the i -th pixel in an image of the object's appearance under illumination by the j -th light is:

$$I_{i,j} = (S_j \mathbf{l}_j)^T (\rho_i \mathbf{n}_i) = \mathbf{s}_j^T \mathbf{b}_i, \quad (1)$$

where \mathbf{n}_i is the outward-pointing normal of a surface patch projected onto the i -th pixel, ρ_i is the diffuse albedo of the patch, and \mathbf{l}_j is the unit vector in the direction of the j -th light source whose intensity is S_j . Vector \mathbf{s}_j groups S_j and \mathbf{l}_j together, as does \mathbf{b}_i for ρ_i and \mathbf{n}_i . The equation (1) is

written in a matrix form as $\mathbf{I} = \mathbf{s}^T \mathbf{b}$. In classical photometric stereo [10], an object is illuminated by three lights of known direction and intensity (thus \mathbf{s} , called *the lights* for brief, is known). The matrix \mathbf{b} (the normals multiplied by corresponding albedos, called *the normals* for brief) is then obtained by inverting (1): $\mathbf{b} = (\mathbf{s}^T)^{inv} \mathbf{I}$. In this paper, however, we are interested in the case when the lights \mathbf{s} are uncalibrated. It is known [3] that the matrix \mathbf{I} can be then factorized into two matrices $\bar{\mathbf{s}}$ and $\bar{\mathbf{b}}$ such that $\mathbf{I} = \bar{\mathbf{s}}^T \bar{\mathbf{b}}$. The factorization accommodates inherent ambiguity [5] represented by a group of invertible matrices $\mathbf{GL}(3)$ whose elements \mathbf{A} transform the particular factorization solution $\{\bar{\mathbf{s}}, \bar{\mathbf{b}}\}$ into all other solutions $\{\bar{\mathbf{s}} \mapsto (\mathbf{A}^T)^{inv} \bar{\mathbf{s}}, \bar{\mathbf{b}} \mapsto \mathbf{A} \bar{\mathbf{b}}\}$. Among these, indeed, there is exactly one $\tilde{\mathbf{A}} \in \mathbf{GL}(3)$ that transforms $\bar{\mathbf{s}}$ and $\bar{\mathbf{b}}$ into the true lights \mathbf{s} and the true normals \mathbf{b} ; all other transformations lead to solutions that are called *the pseudolights* and *the pseudonormals* throughout this paper. If we want to determine this $\tilde{\mathbf{A}}$, additional information about normals and/or lights must be employed [10, 5]. Such information takes on the form of additional constraints that must hold, besides $\mathbf{I} = \mathbf{s}^T \mathbf{b}$. One fairly general condition is the integrability constraint on normals that represents the fact that recovered normals must form a continuous surface. Application of this constraint reduces the ambiguity from $\mathbf{GL}(3)$ to the group of generalized bas-relief transformations [1]. While the integrability condition can usually be applied, this is not true of other attempts to reduce the photometric stereo ambiguity that were suggested in the literature. One such possibility is to estimate normal orientations at several points by an independent method and using them to disambiguate the photometric stereo [3]. Another proposal, if occluding boundary can be identified, two of the three needed non-collinear normals are uniquely defined on the occluding boundary [4]. Another possibility is to assume that all light sources are of equal or known intensity. Yet another one is that there is a curved object (or its known part) in the scene which is of constant albedo. Such possibilities were employed and/or discussed in [3, 11, 12].

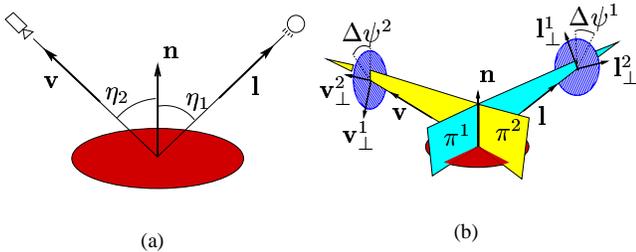


Fig. 1 The geometry of illumination and viewing.

Integrability requires the assumption that the observed entity is a continuous surface. The knowledge of normals at several surface points is not generally available. Next, the occluding boundary is hard to detect reliably. As for the constant-albedo assumption, it necessarily restricts the class of recoverable surfaces. Hence, most promising so far, in the sense of applicability, could be the combination of equal-intensity light source assumption with the integrability assumption, that taken together result in the convex/concave ambiguity [11].

In this paper, we offer an alternative way to disambiguate the photometric stereo by designing previously unrecognized constraints on normals. Under the conditions of distant point light source illumination and orthographic imaging, we show that it is possible to determine, up to a scale, two entities: the projection of each surface normal onto the plane orthogonal to the viewing direction \mathbf{v} , and the projection of each normal onto the plane orthogonal to the illumination direction \mathbf{l} . The former entity is called *the v-tilt*, since it corresponds to common meaning of *tilt*, being the orientation of a surface normal projected onto the retinal plane. The latter entity is then called *the l-tilt*, by analogy with the former one. We then identify in this paper the classes of linear transformations that leave the \mathbf{v} -tilt and the \mathbf{l} -tilt of surface normals unchanged. Subsequently, we make three key observations: First, constraining the normals to be consistent with the \mathbf{v} -tilt and \mathbf{l} -tilt reduces the photometric stereo ambiguity into convex/concave ambiguity. Second, the \mathbf{v} -tilt constraint combined with the integrability constraint results in bas-relief ambiguity (group of transformations that scale the depth of the recovered surface). Third, making the normals consistent with the \mathbf{l} -tilt and integrability constraints results in reducing the ambiguity into the convex/concave ambiguity.

Our approach relies on the assumption that the surface reflectance at any surface point is separable with respect to illumination and viewing directions. Such reflectance model implies two fundamental symmetries. The first symmetry follows from that the radiance does not change when the illumination (no matter if point or extended) is mirrored by any plane containing a given surface normal. The second

symmetry follows from that the radiance does not change when the detector is mirrored by any plane containing a surface normal. The first symmetry enables the construction of the \mathbf{l} -tilt, with the plane of incidence (denoted π^1 in Fig. 1(b)) taken as the symmetry plane. The second symmetry implies the possibility to construct the \mathbf{v} -tilt with the emittance plane (denoted π^2 in Fig. 1(b)) playing the role of the symmetry plane.

The paper is organized as follows. In Section 2 we present an example of experimental procedure that determines the \mathbf{v} -tilt and \mathbf{l} -tilt employing just discussed symmetries. With working procedure at hand, we return to the problem of disambiguating the photometric stereo in Section 3 where we identify transformations under which each of the tilt constraints is preserved. The analytical derivation is followed by experimental demonstration in Section 4.

2. Tilt constraints construction

In this section we show how the tilt constraints may be constructed thanks to underlying symmetries of the model of reflectance. We demonstrate that on an example of reflection on smooth dielectric surface. Reflectance model for diffuse reflection on such surfaces [8] possesses the essential property of being separable with respect to the illumination and viewing directions. The separability follows from the assumption that light which penetrates under the smooth optical boundary becomes completely depolarized by internal scattering, as well as its distribution becomes azimuth-independent with respect to the surface normal before it leaves the sub-surface back out into the air [9]. The model is constructed by enriching the Lambert's model by two transmissivity terms $\mathcal{T}(\eta_1, n)$, $\mathcal{T}(\eta_2, n)$ as follows:

$$I = S\rho \mathcal{T}(\eta_1, n)\mathcal{T}(\eta_2, n) \cos \eta_1, \quad (2)$$

where n is the index of refraction, the angles η_1 and η_2 are respectively the incidence angle and the emittance angle (see Fig. 1(a)) and other variables are as in (1). The two transmissivity terms correspond to the reality of incident light refracting into the sub-surface (term $\mathcal{T}(\eta_1, n)$) and, after internal scattering, part of it refracting back out into the air (term $\mathcal{T}(\eta_2, n)$). Having a separable reflectance model, symmetries discussed at the end of the previous section imply that if we mount a linear polarizing filter in front of the light source, the radiance of a surface patch projected onto a particular pixel will not change while the light source is mirrored by the plane of incidence. The pixel brightness will hence be the same for orientations $\{\Delta\psi^1, -\Delta\psi^1\}$ of the polarizer ($\Delta\psi^1$ is the inclination of the polarizing direction from the plane of incidence, see Fig. 1(b)). If the dependence of radiance on the polarizer orientation is enough discriminative, it will be possible to identify the plane of

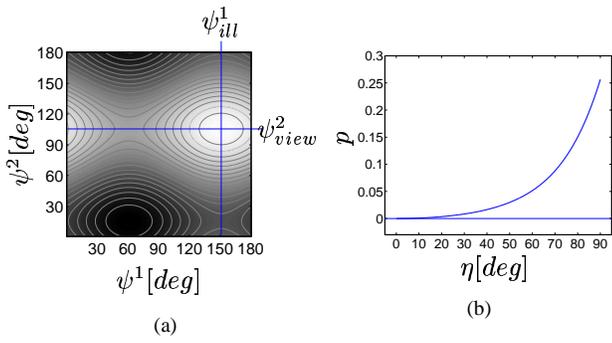


Fig. 2 (a) Diffuse radiance as a function of orientations ψ_1 and ψ_2 of the two polarizers. (b) Dependence of the partial polarization on the angle of incidence/emittance, displayed for refraction index $n = 1.3$.

symmetry and hence to get the **I**-tilt of the normal. From an elementary optics it follows that this dependence is described by the following substitution in (2):

$$\mathcal{T}(\eta_1, n) \rightarrow \frac{1}{2} \{ \mathcal{T}_{\parallel}(\eta_1, n) \cos^2 \Delta\psi^1 + \mathcal{T}_{\perp}(\eta_1, n) \sin^2 \Delta\psi^1 \}, \quad (3)$$

where the terms \mathcal{T}_{\perp} and \mathcal{T}_{\parallel} are the transmissivities [2] for light polarized perpendicularly and parallelly to illumination plane. The pixel brightness varies harmonically with the polarizer orientation ψ^1 which is measured as ordinary polar angle within the polarizer coordinate system $\{\mathbf{I}_{\perp}^1, \mathbf{I}_{\parallel}^1\}$. It is therefore sufficient to capture images for several orientations of the polarizer and then apply a harmonic fit in each pixel. The maximum is indeed found at the incidence plane orientation ψ_{ill}^1 , for which $\Delta\psi^1 = \psi^1 - \psi_{ill}^1 = 0$ since $\mathcal{T}_{\perp}(\eta, n) < \mathcal{T}_{\parallel}(\eta, n)$ for all $\eta > 0$. The reliability of the symmetry plane determination depends on the relative magnitude of the harmonic variation:

$$p_{\mathbf{I}} = \frac{\mathcal{T}_{\parallel}(\eta_1, n) - \mathcal{T}_{\perp}(\eta_1, n)}{\mathcal{T}_{\parallel}(\eta_1, n) + \mathcal{T}_{\perp}(\eta_1, n)}. \quad (4)$$

We call $p_{\mathbf{I}}$ the **I**-partial polarization. Fig. 2(b) shows the dependence of the partial polarization on the angle of incidence.

The **v**-tilt is then constructed analogously. If a linear polarizing filter is mounted in front of the camera, then the other symmetry that we identified ensures that image intensities measured under polarizer orientations $\{\Delta\psi^2, -\Delta\psi^2\}$ are the same ($\Delta\psi^2$ is the inclination of the polarizer orientation from the plane of emittance π^2 , see Fig. 1(b)). Dependence of the pixel brightness on the camera polarizer orientation can be described by the substitution [9]

$$\mathcal{T}(\eta_2, n) \rightarrow \frac{1}{2} \{ \mathcal{T}_{\parallel}(\eta_2, n) \cos^2 \Delta\psi^2 + \mathcal{T}_{\perp}(\eta_2, n) \sin^2 \Delta\psi^2 \}. \quad (5)$$

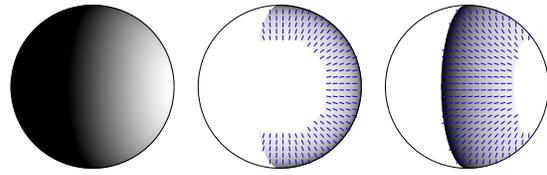


Fig. 3 A sphere illuminated from the right (left image), the **v**-tilt (middle) and the **I**-tilt (right) shown as needle maps. The tilts are superimposed on images of the **v**-partial polarization and **I**-partial polarization, respectively. Dark points correspond to high partial polarization values.

The analogous form of (3) and (5) enables us to quickly come to the result that the **v**-tilt can be determined as the orientation of the camera polarizer under which the maximum brightness is perceived. Reliability of the tilt evaluation raises with the angle of emittance and is linked to the **v**-partial polarization $p_{\mathbf{v}}$ that is defined in analogy with (4). To give this section a summary, the term for reflectance is separable with respect to orientation of the light polarizer and the camera polarizer (see Fig. 2(a)). It is possible to determine the **I**-tilt constraint and then the **v**-tilt constraint on each normal by first rotating the polarizer placed in front of the light source and then in front of the camera. The **I**-partial polarization and **v**-partial polarization serve as factors that express the reliability of the tilt constraint determination. This is illustrated in Fig. 3 where both the **v**-tilts and **I**-tilts are displayed where the **v**-partial polarization and **I**-partial polarization, respectively, exceed 3%.

3. Transformations that preserve tilts

We discussed in the introduction section that photometric stereo can determine the normals and the lights only up to an arbitrary linear transformation $\mathbf{A} \in \mathbf{GL}(3)$ that can not be resolved unless further information about normals and/or lights is known. This information takes on the form of additional conditions that must hold besides (1) and generally reduces the ambiguity. The tilt constraints on normals are examples of such conditions and in this section we identify the classes of transformations that leave each tilt type invariant.

Theorem 1. The only transformations the **v**-tilt of the normals is invariant to are given by:

$$\mathbf{A} = [\mathbf{v}_{\perp} | \mathbf{v}] \begin{bmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ \mu & \nu & \tau \end{bmatrix} [\mathbf{v}_{\perp} | \mathbf{v}]^T, \quad \lambda, \tau \neq 0, \quad (6)$$

where $[\mathbf{v}_{\perp} | \mathbf{v}]$ is the orthonormal triple with first two vectors inherited from the retinal plane and the third vector equal to the direction of viewing (see Fig. 1(b)).

Proof. Let \mathcal{S}_v be the set of indices such that the v -tilt of $\{\mathbf{b}_i; i \in \mathcal{S}_v\}$ is defined. The invariance of the v -tilt with respect to transformation of the true normals \mathbf{b} into pseudonormals $\bar{\mathbf{b}} = \mathbf{A}\mathbf{b}$ is written as $\mathbf{v}_\perp^T \bar{\mathbf{b}}_i = \lambda_i \mathbf{v}_\perp^T \mathbf{b}_i$, or more precisely and with substitution $\bar{\mathbf{b}} = \mathbf{A}\mathbf{b}$:

$$\forall i \in \mathcal{S}_v \exists \lambda_i \neq 0 : \mathbf{v}_\perp^T \mathbf{A}\mathbf{b}_i = \lambda_i \mathbf{v}_\perp^T \mathbf{b}_i. \quad (7)$$

If the normals $\{\mathbf{b}_i; i \in \mathcal{S}_v\}$ are enough diverse, then it must hold $\forall i \in \mathcal{S}_v$ that $\lambda_i = \lambda$. Consequently, it implies that $\mathbf{v}_\perp^T \mathbf{A} = \lambda \mathbf{v}_\perp^T$, which in turn means that (6) must hold. Parameter λ must be non-zero from (7) and τ must be non-zero for the matrix \mathbf{A} to be invertible. \square

Theorem 2. Let \mathbf{l} be the illumination direction and \mathbf{l}_\perp be the light polarizer coordinate system that makes an orthonormal basis with \mathbf{l} (see Fig. 1(b)). The only transformations the \mathbf{l} -tilt of normals is invariant to are given by

$$\mathbf{A} = [\bar{\mathbf{l}}_\perp | \bar{\mathbf{l}}] \begin{bmatrix} \alpha & 0 & 0 \\ 0 & \alpha & 0 \\ 0 & 0 & \beta \end{bmatrix} [\mathbf{l}_\perp | \mathbf{l}]^T \quad \alpha, \beta \neq 0. \quad (8)$$

Besides α and β , the free parameters are represented by a vector triple $[\bar{\mathbf{l}}_\perp | \bar{\mathbf{l}}]$ that is constrained to be orthonormal and of the same orientation as $[\mathbf{l}_\perp | \mathbf{l}]$.

Proof. Let \mathcal{S}_l be the set of indices such that the \mathbf{l} -tilt of $\{\mathbf{b}_i; i \in \mathcal{S}_l\}$ is defined. The condition that the \mathbf{l} -tilt of the normals preserves under transformation of the true normals onto the pseudonormals $\bar{\mathbf{b}} = \mathbf{A}\mathbf{b}$ means that

$$\forall i \in \mathcal{S}_l \exists \alpha_i \neq 0 : \bar{\mathbf{l}}_\perp^T \bar{\mathbf{b}}_i = \alpha_i \mathbf{l}_\perp^T \mathbf{b}_i, \quad (9)$$

where $\bar{\mathbf{l}}_\perp$ is a pair of vectors that represents a new plane of projection for evaluation of the \mathbf{l} -tilt and is therefore constrained to form an orthonormal basis with the transformed light direction $\bar{\mathbf{l}}$, and in addition, this basis has to be of the same orientation as $[\mathbf{l}_\perp | \mathbf{l}]$. New light direction is given by $\bar{\mathbf{l}} = \beta (\mathbf{A}^T)^{inv} \mathbf{l}$, $\beta = S/\bar{S}$ being the ratio of the light intensity before and after the transformation. If the normals $\{\mathbf{b}_i; i \in \mathcal{S}_l\}$ are enough diverse, then it must hold $\forall i \in \mathcal{S}_l$ that $\alpha_i = \alpha$. Taking this into account and writing the equivalent of the light transformation formula gives:

$$\bar{\mathbf{l}}_\perp^T \mathbf{A} = \alpha \mathbf{l}_\perp^T, \quad (10)$$

$$\bar{\mathbf{l}}^T \mathbf{A} = \beta \mathbf{l}^T, \quad (11)$$

which implies (8) where α must be non-zero from (9) and β must be non-zero for the matrix \mathbf{A} to be invertible. \square

No of the v -tilt and \mathbf{l} -tilt constraints alone removes the photometric stereo ambiguity. Combining one with the other, or one of them with an independent constraint will result in various degree of the ambiguity reduction. Our intention is to work primarily with constraints that can be widely applied and we therefore restrict our attention to three of them: the v -tilt, the \mathbf{l} -tilt, and the integrability constraint.

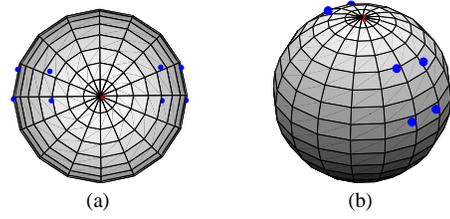


Fig. 4 (a) The directions of eight light sources, displayed on a sphere from the camera point of view. The viewing direction shown as the central dot. (b) The same, rotated for better illustration.

It is known that the only transformations the integrability constraint is invariant to are [1]:

$$\mathbf{A} = [\mathbf{v}_\perp | \mathbf{v}] \begin{bmatrix} \lambda' & 0 & \mu' \\ 0 & \lambda' & \nu' \\ 0 & 0 & \tau' \end{bmatrix} [\mathbf{v}_\perp | \mathbf{v}]^T \quad \lambda', \tau' \neq 0. \quad (12)$$

What follows is the identification of classes of transformations that represent the ambiguity that remains after pairwise applications of the three constraints. While Point 3 below is obvious, Points 1 and 2 can be proven by straightforward analysis which is omitted here for lack of space.

1. Application of the v -tilt and \mathbf{l} -tilt constraints reduces the ambiguity into the convex/concave case because the only common elements of transformation classes (6) and (8) are given by $\mathbf{A} = [\mathbf{v}_\perp | \mathbf{v}] \text{diag} [\pm\lambda, \pm\lambda, \lambda] [\mathbf{v}_\perp | \mathbf{v}]^T$ with $\lambda \neq 0$.
2. Application of the integrability and \mathbf{l} -tilt constraints results in convex/concave ambiguity as in Point 1.
3. Application of the v -tilt and integrability constraints results in bas-relief ambiguity, because the only transformations belonging to classes (6) as well as (12) are $\mathbf{A} = [\mathbf{v}_\perp | \mathbf{v}] \text{diag} [\lambda, \lambda, \tau] [\mathbf{v}_\perp | \mathbf{v}]^T$ with $\lambda, \tau \neq 0$.

4. Experiment

The object we used to test our approach was a glazed china teapot (see Fig. 5(a)). The teapot was illuminated from eight directions in turn (see Fig. 4; the ground truth positions of light sources were measured by a geometric stereo-based method). For each light source position, we grabbed two sets of images by a 14-bit camera. The first and second sets were formed by images taken for six orientations of the camera polarizer and for six orientations of the light polarizer, respectively. Both the camera polarizer

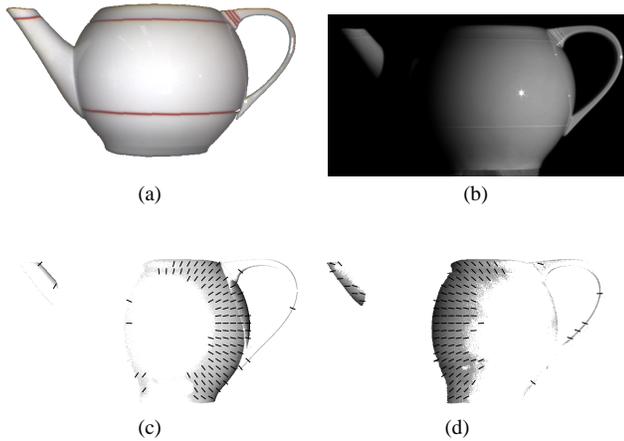


Fig. 5 The test object (a) illuminated from the right (b), the \mathbf{v} -tilt (c) and the \mathbf{l} -tilt (d) showed as needle maps superimposed on the \mathbf{v} -partial polarization and \mathbf{l} -partial polarization values, respectively (dark points correspond to high partial polarization values).

and the light polarizer were rotated in 30deg steps. For each light we computed the \mathbf{v} -tilt and \mathbf{l} -tilt of every illuminated surface patch by applying the harmonic fit to the first and second set of images, respectively, in a pixel-wise manner. An example of the \mathbf{v} -tilt and the \mathbf{l} -tilt determined for the teapot illuminated from the right is shown in Fig. 5(b-d). Note that the tilts look fairly as expected (cf. Fig. 3). Only tilts whose partial polarization exceeded 3% were trusted. Besides that, tilts of specularly reflecting surface patches were marked as invalid. Specularities, as well as shadows, were identified by simple thresholding.

Lambert’s reflectance model was used for factorization of photometric data into pseudonormals and pseudolights, because it is close to (2) for most incidence and emittance angles [8]. The pseudonormals and pseudolights were computed by an iterative factorization algorithm similar to the one published in [11]. Images of a teapot illuminated from a given direction were computed by averaging all images grabbed under that illumination. Shadows and specularities were not included into the factorization process. Subsequently, different sets of constraints were applied on pseudonormals in order to reduce the photometric stereo ambiguity. The results are shown in Fig. 6. The integrability constraint was employed first, using the algorithm published in [11] (see the first image row). We observe that the pseudonormals illuminated by the light of constant intensity from three directions (the retinal plane axes directions and the camera direction) do not correspond to the appearance of a teapot illuminated from the left, bottom and the viewpoint direction. This is consistent with the remaining ambiguity (12). We subsequently let $\lambda' = 1$ and found μ', ν'

and τ' such that the integrable pseudonormals transformed by (12) were consistent with the \mathbf{v} -tilt constraint. This was done by least squares with the optimization function defined as squared sum of angular differences between measured \mathbf{v} -tilt and the \mathbf{v} -tilt of transformed pseudonormals. If a given surface patch was illuminated by more than one light during the measurement, there were multiple evaluations of \mathbf{v} -tilt available from which the one with highest magnitude of harmonic brightness variation was selected. The results are shown in the second row. The images have correct shading, except for the last one which is dark. This is consistent with our expectation about the resulting ambiguity (scaling of recovered surface depth). The \mathbf{l} -tilt constraints were imposed on integrable normals in a similar way. The goal of optimization was to reach consistency of the transformed integrable normals with the \mathbf{l} -tilt. To maintain good stability of the result, the optimization was done with respect to \mathbf{l} -tilts associated with all eight lights simultaneously. Illuminated normals are shown in the third row. The appearance is consistent with the expected convex/concave ambiguity (that was resolved by hand). Finally, we applied the \mathbf{v} -tilt constraint on the (unintegrable) pseudonormals. The least squares optimization now searched for two vectors to which the pseudonormals were projected and this projection was compared with the evaluated \mathbf{v} -tilt. The result is in the fourth row and is consistent with the group of transformations (6). Subsequent optimization of the third row of the transformation matrix in (6) with respect to consistency with the \mathbf{l} -tilt constraint resulted in pseudonormals shown in the last row of Fig. 6. The ambiguity was reduced to the convex/concave case and was subsequently resolved by hand.

To summarize, only the \mathbf{v} -tilt and \mathbf{l} -tilt constraints combined, or the integrability and \mathbf{l} -tilt constraints combined can recognize the normals *and* the lights correctly, as expected. In Fig. 7 we show directions of the eight light sources estimated by the five methods together with the ground truth.

5. Conclusions

In this paper we have shown that if we are able to measure surface normal tilts with respect to camera and light directions the ambiguity of uncalibrated photometric stereo reduces to only convex/concave case. To recover tilts we employed partial polarization of light reflected on smooth dielectric surface, although this is not the only possibility. The tilt with respect to the light direction alone may be combined with the integrability constraint to obtain the convex/concave ambiguity as well but this requires an explicit assumption about surface continuity.

Previous work on shape from polarization vision used only the fact that the *specular* reflection component is po-

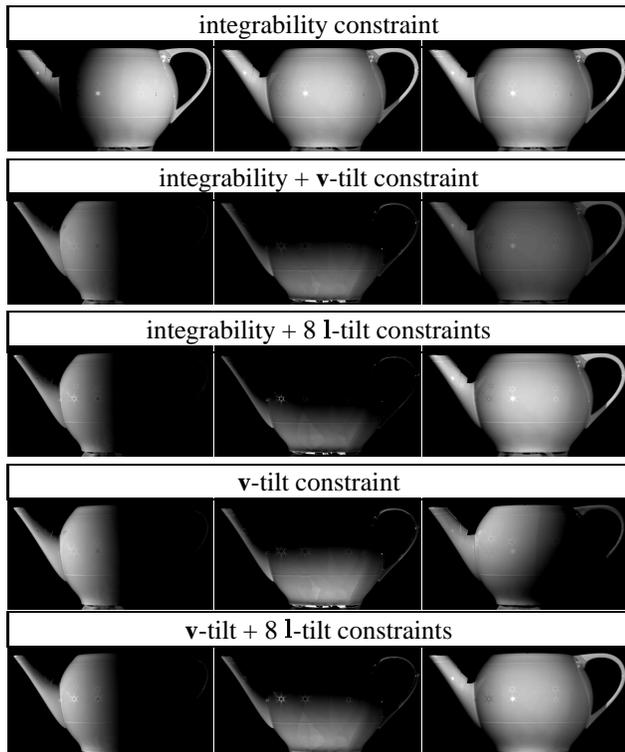


Fig. 6 The results of pseudonormal transformations conforming to the constraints that are written above each image row. The images show the computed normals illuminated by the light source of constant intensity from directions of the first and second retinal plane axes and the viewing direction, respectively.

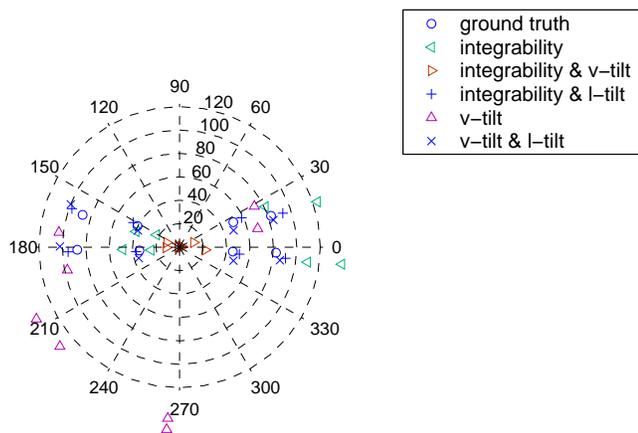


Fig. 7 Illumination directions estimated under different sets of constraints imposed on normals. Directions are displayed in spherical coordinates. Only the combination of the v -tilt and l -tilt constraints, as well as of the integrability and l -tilt constraints determines the illumination directions correctly.

larized [6, 7]. The polarization of the diffusely reflected light is too weak to recover useful dense information about the object shape but, as we showed, may be used to provide disambiguation constraints for uncalibrated photometric stereo.

Acknowledgement

This research was supported by the Grant Agency of the Czech Republic under Project GACR 102/01/1371 and in part by the Czech Ministry of Education under Research Programme J04/98:212300013, as well as by the EU Fifth Framework Programme project 1999-29017.

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] M. Born and E. Wolf. *Principles of Optics. Electromagnetic Theory of Propagation, Interference and Diffraction of Light*. Cambridge University Press, 1980.
- [3] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *Journal of Optical Society of America A*, 11(11):3079–3089, 1994.
- [4] J. J. Koenderink. What does the occluding contour tell us about solid shape? *Perception*, 13(3):321–330, 1984.
- [5] J. J. Koenderink and A. J. van Doorn. The generic bilinear calibration-estimation problem. *International Journal of Computer Vision*, 23(3):217–234, 1997.
- [6] M. Saito, Y. Sato, K. Ikeuchi, and H. Kashiwagi. Measurement of surface orientations of transparent objects using polarization in highlight. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp. 381–386, 1999.
- [7] L. B. Wolff. Surface orientation from two camera stereo with polarizers. In *Proc. Conference of Optics, Illumination and Image Sensing for Machine Vision IV*, vol. 1194 of *Proc. of SPIE - The International Society for Optical Engineering*, pp. 287–297, 1990.
- [8] L. B. Wolff. Diffuse-reflectance model for smooth dielectric surfaces. *Journal of Optical Society of America A*, 11(11):2956–2968, 1994.
- [9] L. B. Wolff and T. E. Boulton. Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7):635–657, 1991.
- [10] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [11] A. Yuille and D. Snow. Shape and albedo from multiple images using integrability. In *Proc. IEEE Conference on Computer Vision and Pattern Recognition*, pp. 158–164, 1997.
- [12] A. L. Yuille, D. Snow, R. Epstein, and P. N. Belhumeur. Determining generative models of objects under varying illumination: Shape and albedo from multiple images using SVD and integrability. *International Journal of Computer Vision*, 35(3):203–222, 1999.