

A Closed-Form, Consistent and Robust Solution to Uncalibrated Photometric Stereo Via Local Diffuse Reflectance Maxima

Thoma Papadhimetri · Paolo Favaro

Received: 30 November 2012 / Accepted: 7 October 2013 / Published online: 20 October 2013
© Springer Science+Business Media New York 2013

Abstract Images of an object under different illumination are known to provide strong cues about the object surface. A mathematical formalization of how to recover the normal map of such a surface leads to the so-called uncalibrated photometric stereo problem. In the simplest instance, this problem can be reduced to the task of identifying only three parameters: the so-called generalized bas-relief (GBR) ambiguity. The challenge is to find additional general assumptions about the object, that identify these parameters uniquely. Current approaches are not consistent, i.e., they provide different solutions when run multiple times on the same data. To address this limitation, we propose exploiting *local diffuse reflectance* (LDR) maxima, i.e., points in the scene where the normal vector is parallel to the illumination direction (see Fig. 1). We demonstrate several noteworthy properties of these maxima: a closed-form solution, computational efficiency and GBR consistency. An LDR maximum yields a simple closed-form solution corresponding to a semi-circle in the GBR parameters space (see Fig. 2); because as few as two diffuse maxima in different images identify a unique solution, the identification of the GBR parameters can be achieved very efficiently; finally, the algorithm is consistent as it always returns the same solution given the same data. Our algorithm is also remarkably robust: It can obtain an accurate estimate of the GBR parameters even with extremely high levels of outliers in the detected maxima (up to 80 % of the observations). The method is validated on real data and achieves state-of-the-art results.

Keywords Photometric stereo · Uncalibrated · Diffuse · Maxima · GBR · Ambiguity

1 Introduction

In this paper, we consider images captured under unknown distant point source illumination and satisfying the Lambertian model, i.e., where surface brightness looks the same from any viewing direction. We show that locations of maximal diffuse brightness in the captured images carry very useful geometrical information about shape (in the form of a unit normal field) and light. We demonstrate their potential in the case of *uncalibrated photometric stereo*, where no prior knowledge about the illumination, geometry, and reflectance albedo is available. In the most general formulation of this problem, the normal field of the object can be obtained from the Lambertian reflectance model only up to a nine-parameter linear ambiguity (Hayakawa 1994). This ambiguity can be further reduced to three parameters, the so-called *generalized bas-relief ambiguity* (GBR) (Belhumeur et al. 1999), via the integrability constraint, which enforces that a valid surface can be reconstructed from the estimated normal field.

Finding a way to fix the GBR parameters with assumptions as realistic as possible is still an open research problem. We make the assumption that one can detect¹ and exploit maxima of the *Lambertian diffuse reflectance* (LDR) component (as opposed to the specular reflectance component (Drbohlav and Chantler 2005; Drbohlav and Sara 2002; Lager and Fua 2008; Oren and Nayar 1997) or the temporal maxima (Koppal and Narasimhan 2006)). We show that each of these

T. Papadhimetri (✉) · P. Favaro
University of Bern, Bern, Switzerland
e-mail: thoma.papadhimetri@iam.unibe.ch

P. Favaro
e-mail: paolo.favaro@iam.unibe.ch

¹ Ideally, our detection would require a uniform albedo. However, because of its robustness, our algorithm can handle spatially varying albedos (see Sect. 6.4 and 6.5).

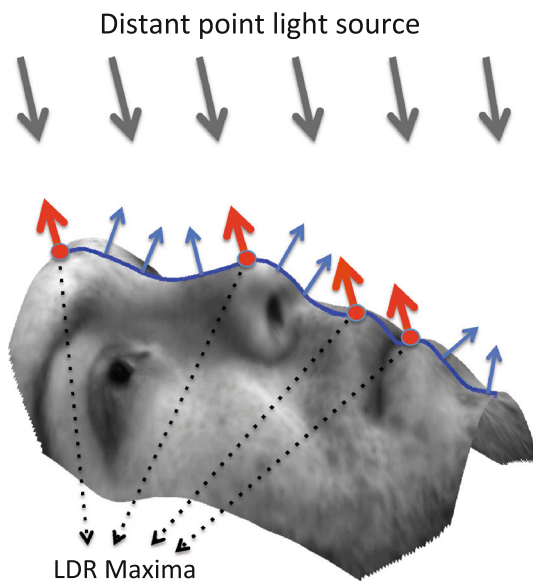


Fig. 1 Some Lambertian diffuse reflectance (LDR) maxima on the surface of a face in one input image. The lighting direction is shown at the very top with *gray arrows*. For illustration purposes, we look only at a 1D profile (*blue solid line*) of the 3D surface. Along this profile we show some of the normals to the surface as *blue arrows*. Those normals that are parallel to the illumination direction are shown as *red arrows* and their corresponding point (shown in *red*) on the surface is an LDR maximum (Color figure online)

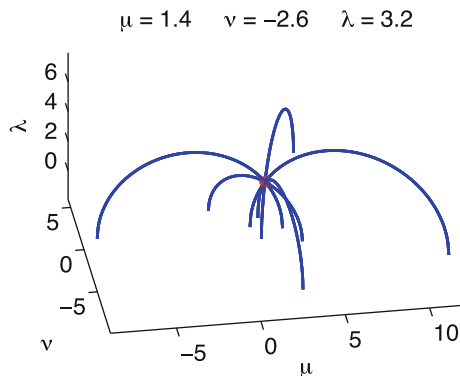


Fig. 2 GBR ambiguity solved via LDR constraints. The integrability constraint alone allows the GBR parameters (μ, ν, λ) to take any value in the 3D space. A single LDR maximum constrains the GBR parameters to a semi-circle (which can be written in closed-form). Two different LDR maxima are sufficient to generate a unique solution. Indeed, two maxima correspond to two semi-circles, and these intersect at one point (whose calculation is very efficient). All LDR maxima define other semi-circles all intersecting at the true GBR parameters (which shows the consistency of this constraint)

maxima restricts the space of GBR ambiguities to a semi-circle in 3-D space, and that as few as two maxima from differently illuminated images yield a unique solution for the GBR parameters (see Fig. 2). Moreover, the semi-circles can be efficiently computed in closed-form (see Sect. 6.1).

The most fundamental challenge that we address in this paper is the *consistency* of the assumptions, i.e., when the

assumptions uniquely identify the same solution. We show that there is a limited choice of constraints that satisfy the consistency property. As we argue in the next sections, our proposed solution is one such choice. Consistency is not only needed to guarantee that the solution is reliable, but also to make the algorithm fast. Indeed, current inconsistent state-of-the-art methods obtain a stable estimate by running their algorithms several times with different initial conditions and by averaging all the estimates, a process that may take several minutes and yield a solution with low confidence. In contrast, our algorithm needs to be run only once and yields a reliable estimate in about 14 s with non optimized code, as we show in Sect. 6.

2 Prior Work and Contributions

Photometric stereo (Woodham 1980) is a method for estimating shape and reflectance of an object, using three (Woodham 1980) or more (Wu and Tang 2005) images of a static object taken from a fixed viewpoint and under different lighting conditions. When the illumination directions and intensities are known, the problem can be solved via a linear system.²

When no prior knowledge about the illumination, geometry, and reflectance is available, the problem is called *uncalibrated photometric stereo* (UPS). UPS has been addressed with a variety of techniques that make explicit or implicit use of the GBR ambiguity (Belhumeur et al. 1999):

2.1 Explicit GBR Disambiguation Methods

In (Alldrin et al. 2007), the GBR ambiguity is solved by minimizing the entropy of the albedo distribution following the argument that incorrect GBR parameters result in spreading albedo values. This approach relies on the assumption that albedos are based on a few intensity values. A similar philosophy is to group the normal-albedo distribution based on intensity-color appearance (Shi et al. 2010). The use of color, when available, allows albedo clustering based on the chromaticity. Another example is to introduce additional constraints to determine the GBR ambiguity by exploiting specularities of glossy surfaces (Drbohlav and Chantler 2005; Drbohlav and Sara 2002; Tan et al. 2007). These methods rely on the ability to correctly detect specularities and on the assumption that objects in the scene have a non-negligible specular component. In other investigations, the GBR ambiguity has been eliminated by exploiting interreflections (Chandraker et al. 2005) or by considering the Torrance and Sparrow reflectance model (Georghiadis 2003).

² This result holds under the Lambertian image formation model, orthographic projection, by considering distant point light sources and when shadows and interreflections can be ignored.

Lagger and Fua (2008) notice that both specular and diffuse maxima provide useful information about the illumination direction. However, they discard diffuse maxima and focus only on specular ones.

2.2 Implicit GBR Disambiguation Methods

Other approaches instead exploit shadows (Okabe et al. 2009; Sunkavalli and Zickler 2010), dimensionality reduction (Sato et al. 2007), the bilateral symmetry of isotropic bidirectional reflectance distribution function (BRDF) (Alldrin and Kriegman 2007) or consider changing the viewpoint and exploit the Helmholtz reciprocity principle (Zickler et al. 2002). Hertzmann and Seitz (2005) use a reference object in the scene to generate intensity-normal look-up tables, and can handle fairly general reflectance functions (more precisely, all materials of a single object must be a linear combination of a given basis of materials). A recent work of Chandraker et al. (2011) exploits a particular image acquisition setup and spatial and temporal image derivatives to obtain an uncalibrated photometric invariant that can deal with a general isotropic BRDF. Other methods generalize the lighting environment and implicitly extend the Lambertian model so that different normals see a different portion of the illumination hemisphere (Basri and Jacobs 2001).

2.3 Contributions

We introduce LDR maxima (Favaro and Papadhimitri 2012) to solve the GBR ambiguity in UPS explicitly. Our main contributions are:

- A closed-form solution given an LDR maximum;
- A robust estimator to find the GBR parameters that tolerates up to 80% incorrect LDR maxima;
- A proof that the solution is GBR-consistent.

Finally, we experimentally demonstrate that our method handles a wide range of real-world scenes (with varying albedo distribution), is the most computationally efficient UPS algorithm, and achieves state-of-the-art results.

3 Photometric Stereo

Photometric stereo is the task of recovering the surface of an object with normal field \mathbf{n} , given K input images taken from a fixed viewpoint, and illuminated by known distant point light sources with direction \mathbf{l} . We denote the unit-normal to the surface at the pixel index p with $\mathbf{n}_p \in S^2$ for $p = 1, \dots, P$, where P is the total number of pixels in an image. Similarly, we denote with $\mathbf{l}_k \in S^2$ the unit-normal corresponding to the k -th illumination direction. Now, let us call $\mathbf{L}_k = e_k \mathbf{l}_k$

the *light vector*, where e_k denotes the light intensity. Also, let $\mathbf{N}_p = \rho_p \mathbf{n}_p$ represent the *generic normal* vector and ρ_p be the albedo coefficient (i.e., the intensity of the surface texture). Then, the intensity measured at a pixel index p with illumination k for the Lambertian case (and under the assumption of a linear response of the camera and orthographic projection) is:

$$I_{p,k} = \rho_p \langle \mathbf{n}_p, \mathbf{l}_k \rangle e_k = \mathbf{N}_p^\top \mathbf{L}_k \tag{1}$$

where $\langle \cdot, \cdot \rangle$ denotes inner product and $(\cdot)^\top$ the transpose of a vector (or of a matrix). By sorting the pixels in lexicographical order we can rearrange all the generic normal vectors in a matrix $\mathbf{N} = [\mathbf{N}_1 \dots \mathbf{N}_P] \in \mathbb{R}^{3 \times P}$, which we call *normal matrix*; similarly we can rearrange the light vectors into a matrix $\mathbf{L} = [\mathbf{L}_1 \dots \mathbf{L}_K] \in \mathbb{R}^{3 \times K}$, which we call *light matrix*. Then, we can write Eq. (1) in the following compact matrix form

$$\mathbf{I} = \mathbf{N}^\top \mathbf{L} \tag{2}$$

where $\{\mathbf{I}\}_{p,k} = I_{p,k}$. If light directions and intensities are given, solving photometric stereo is equivalent to solving the linear system (2) in the unknown normal matrix \mathbf{N} . Finally, extracting the albedo ρ_p from the generic normal vectors \mathbf{N}_p can be done via $\rho_p = \|\mathbf{N}_p\|_2$, where $\|\cdot\|_2$ denotes the ℓ_2 norm.

4 Uncalibrated Photometric Stereo

If the light matrix is unknown, \mathbf{N} and \mathbf{L} can be obtained up to a linear ambiguity $\mathbf{G} \in GL(3)$ (Hayakawa 1994), where:

$$\mathbf{I} = \mathbf{N}^\top \overbrace{\mathbf{G}^{-1} \mathbf{G}}^{\text{identity}} \mathbf{L}. \tag{3}$$

It can be shown (Belhumeur et al. 1999) that if we impose the *integrability constraint*, then the ambiguities take the form

$$\mathbf{G} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ \mu & \nu & \lambda \end{bmatrix} \tag{4}$$

where the three parameters $\lambda \neq 0, \mu, \nu \in \mathbb{R}$ represent the group of GBR transformations. In this formulation it is apparent that solving UPS amounts to fixing these three parameters. For simplicity, we also fix the sign of λ to be positive. Define $\hat{\mathbf{N}} \doteq \mathbf{G}^{-\top} \mathbf{N}$, the *pseudo-normals*, and $\hat{\mathbf{L}} \doteq \mathbf{G} \mathbf{L}$, the *pseudo-lights*. An initial pair $(\hat{\mathbf{N}}, \hat{\mathbf{L}})$ can be computed with the algorithm of Yuille and Snow (1997). Our task is to find $\hat{\lambda}, \hat{\mu},$ and $\hat{\nu}$ such that $\hat{\mathbf{G}}^\top \hat{\mathbf{N}} = \mathbf{N}$ and $\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}} = \mathbf{L}$ where

$$\hat{\mathbf{G}}^\top = \begin{bmatrix} 1 & 0 & \hat{\mu} \\ 0 & 1 & \hat{\nu} \\ 0 & 0 & \hat{\lambda} \end{bmatrix} \quad \text{and} \quad \hat{\mathbf{G}}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -\frac{\hat{\mu}}{\hat{\lambda}} & -\frac{\hat{\nu}}{\hat{\lambda}} & \frac{1}{\hat{\lambda}} \end{bmatrix}. \quad (5)$$

5 Reconstruction Consistency

As mentioned in the introduction and in the discussion on prior work, the main unsolved challenge in uncalibrated photometric stereo is that most proposed solutions do not yield a unique estimate for the normal matrix. In this section we define this issue more formally and show how it can be addressed.

Given the pseudo-normals $\hat{\mathbf{N}}$ and a GBR solution $\hat{\mu}$, $\hat{\nu}$ and $\hat{\lambda}$, the reconstructed normal matrix $\tilde{\mathbf{N}}$ can be written as follows

$$\tilde{\mathbf{N}} = \hat{\mathbf{G}}^\top \mathbf{G}^{-\top} \mathbf{N} = \begin{bmatrix} 1 & 0 & \hat{\mu} \\ 0 & 1 & \hat{\nu} \\ 0 & 0 & \hat{\lambda} \end{bmatrix} \begin{bmatrix} 1 & 0 & -\frac{\mu}{\lambda} \\ 0 & 1 & -\frac{\nu}{\lambda} \\ 0 & 0 & \frac{1}{\lambda} \end{bmatrix} \\ \mathbf{N} = \begin{bmatrix} 1 & 0 & \frac{\hat{\mu}-\mu}{\lambda} \\ 0 & 1 & \frac{\hat{\nu}-\nu}{\lambda} \\ 0 & 0 & \frac{\hat{\lambda}-\lambda}{\lambda} \end{bmatrix} \mathbf{N}. \quad (6)$$

When the estimated GBR parameters match the true ones, the vector $[\frac{\hat{\mu}-\mu}{\lambda} \quad \frac{\hat{\nu}-\nu}{\lambda} \quad \frac{\hat{\lambda}-\lambda}{\lambda}]^\top$ is equal to $[0 \ 0 \ 1]^\top$ and $\tilde{\mathbf{N}} = \mathbf{N}$. In general, however, we might have some error in our estimate of the parameters. Let us define such error as

$$\eta = \begin{bmatrix} \frac{\hat{\mu}-\mu}{\lambda} \\ \frac{\hat{\nu}-\nu}{\lambda} \\ \frac{\hat{\lambda}-\lambda}{\lambda} \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} = \frac{1}{\lambda} \begin{bmatrix} \hat{\mu} - \mu \\ \hat{\nu} - \nu \\ \hat{\lambda} - \lambda \end{bmatrix}. \quad (7)$$

Current algorithms return an error vector η that depends on the true GBR parameters μ , ν , and λ . Hence, for any pseudo-normal $\hat{\mathbf{N}}$ one obtains a different $\tilde{\mathbf{N}}$ estimate. This is a highly undesirable behavior of a UPS algorithm, and to obtain a stable result, the current strategy is to average several estimates obtained with different pseudo-normals. The implicit assumption is that the error η is zero-mean. Current published results on real data (Alldrin et al. 2007; Shi et al. 2010), however, show relatively large average angular errors, which indicate that these assumptions are unfounded.

We therefore argue that one should only introduce constraints that result in estimates that do not change with the GBR parameters. We formalize this in the next definition.

Definition 1 Given a pseudo-normals and pseudo-lights pair $(\hat{\mathbf{N}}, \hat{\mathbf{L}})$, we define a constraint on $\hat{\mu}$, $\hat{\nu}$ and $\hat{\lambda}$, to be *GBR-consistent* when the error η is independent of the true parameters μ , ν , and λ .

The next question is then: Given the pseudo-normals and the pseudo-lights, what type of GBR-consistent constraints in

the estimated GBR parameters can one devise? In general, an algorithm will introduce M constraints that can be written as

$$f_j(\hat{\mu}, \hat{\nu}, \hat{\lambda}, \hat{\mathbf{N}}, \hat{\mathbf{L}}) = 0 \quad \forall j = 1, \dots, M. \quad (8)$$

We can then see how the constraints depend on the true GBR parameters by explicitly writing the pseudo-normals at a pixel p and the pseudo-lights in the k -th image as

$$\begin{bmatrix} \hat{n}_{1,p} \\ \hat{n}_{2,p} \\ \hat{n}_{3,p} \end{bmatrix} = \begin{bmatrix} n_{1,p} - \frac{\mu}{\lambda} n_{3,p} \\ n_{2,p} - \frac{\nu}{\lambda} n_{3,p} \\ \frac{n_{3,p}}{\lambda} \end{bmatrix} \quad (9)$$

$$\begin{bmatrix} \hat{l}_{1,k} \\ \hat{l}_{2,k} \\ \hat{l}_{3,k} \end{bmatrix} = \begin{bmatrix} l_{1,k} \\ l_{2,k} \\ \mu l_{1,k} + \nu l_{2,k} + \lambda l_{3,k} \end{bmatrix}.$$

Then, according to the GBR-consistency definition, it must be possible to write the constraints such that the estimates $\hat{\mu}$, $\hat{\nu}$, and $\hat{\lambda}$ are in the form

$$\hat{\mu} = \mu + \lambda \eta_1 \quad (10)$$

$$\hat{\nu} = \nu + \lambda \eta_2 \quad (11)$$

$$\hat{\lambda} = \lambda(1 + \eta_3) \quad (12)$$

where $\eta = [\eta_1 \ \eta_2 \ \eta_3]^\top$ and $\nabla_{\mu,\nu,\lambda} \eta = \mathbf{0}$ (i.e., the error is independent of the true GBR parameters). Therefore, we can conclude that the constraints will be GBR-consistent if there exists an η such that for any μ , ν , and λ , and $\forall j = 1, \dots, M$

$$f_j \left(\mu + \lambda \eta_1, \nu + \lambda \eta_2, \lambda(1 + \eta_3), \left\{ n_{1,p} - \frac{\mu}{\lambda} n_{3,p}, n_{2,p} - \frac{\nu}{\lambda} n_{3,p}, \frac{n_{3,p}}{\lambda} \right\}_{p=1,\dots,P}, \left\{ l_{1,k}, l_{2,k}, \mu l_{1,k} + \nu l_{2,k} + \lambda l_{3,k} \right\}_{k=1,\dots,K} \right) = 0. \quad (13)$$

We will show in the next sections that the constraints based on LDR maxima are indeed GBR-consistent. However, as an illustration of the analysis presented so far, let us introduce a simple example of a GBR-consistent constraint. Suppose that a novel algorithm for UPS can reliably detect a pixel p where the true normal is $[0 \ 0 \ -1]^\top$. Then, one could devise the constraint

$$f(\hat{\mu}, \hat{\nu}, \hat{\lambda}, \hat{\mathbf{N}}, \hat{\mathbf{L}}) = \frac{\hat{\mu}}{\hat{\lambda}} - \hat{n}_{1,p} = 0 \quad (14)$$

i.e., that $\hat{\mu}$ and $\hat{\lambda}$ are linearly related through $\hat{n}_{1,p}$. By substituting the expression of $\hat{\mu}$ and $\hat{\lambda}$ and $\hat{n}_{1,p}$ above, one obtains

$$\frac{\mu}{\lambda(1 + \eta_3)} + \frac{\eta_1}{1 + \eta_3} - n_{1,p} + \frac{\mu}{\lambda} n_{3,p} = 0 \quad (15)$$

which is verified for any λ and μ with error $\eta_1 = -n_{1,p}/n_{3,p}$ and $\eta_3 = -1 - 1/n_{3,p}$. Finally, notice that the estimators

would achieve $\eta = \mathbf{0}$ because $\mathbf{n}_p = [0 \ 0 \ -1]^\top$. Of course, this example is not sufficient to uniquely determine all the GBR parameters and also it is not clear how one can detect such pixels directly from images. Indeed, this shows how the design of an algorithm that is feasible, robust, and satisfying the GBR-consistency, is challenging.

6 Lambertian Diffuse Reflectance Maxima

We assume that the scene contains curved objects such that there exists a nonempty set \mathcal{S}_k of spatial maxima (where “spatial” refers to the image domain) of the inner product $\langle \mathbf{n}_p, \mathbf{l}_k \rangle$, for a given k -th illumination. Geometrically, we assume that there exist points in the scene where both principal curvatures are positive. We call such points *Lambertian diffuse reflectance* maxima. Notice that we have a maximum only when \mathbf{n}_p and \mathbf{l}_k point in the same direction. Since both vectors are unit-normal, they must also coincide.

In the next sections we introduce the LDR maxima constraints. We first impose them by looking for the maximum of the inner product between normals and lights (see Sect. 6.1). Then, we show that this formulation fully represents all the constraints by carrying out analysis imposing that at LDR maxima normals and lights must be parallel (see Sect. 6.2). Our proof is based on the fact that both approaches yield the same set of solutions (i.e., semi-circles), albeit through different calculations.

6.1 A Closed-Form Solution

For each k , an LDR maximum at $p \in \mathcal{S}_k$ constrains the GBR ambiguity to the following set

$$(\tilde{\mu}, \tilde{\nu}, \tilde{\lambda}) \in \arg \max_{\tilde{\mu}, \tilde{\nu}, \tilde{\lambda}} \left\langle \frac{\hat{\mathbf{G}}^\top \hat{\mathbf{N}}_p}{\|\hat{\mathbf{G}}^\top \hat{\mathbf{N}}_p\|}, \frac{\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k}{\|\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k\|} \right\rangle. \quad (16)$$

Notice the use of the symbol \in , as the solution of the above maximization problem is not unique given a single LDR maximum. The entire set of solutions can be fully characterized in closed-form. First, by using $I_{p,k} = \langle \hat{\mathbf{G}}^\top \hat{\mathbf{N}}_p, \hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k \rangle$, let us rewrite the above problem as

$$(\tilde{\mu}, \tilde{\nu}, \tilde{\lambda}) \in \arg \min_{\tilde{\mu}, \tilde{\nu}, \tilde{\lambda}} \|\hat{\mathbf{G}}^\top \hat{\mathbf{N}}_p\| \|\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k\|. \quad (17)$$

Next, we are ready to state the main result in this manuscript:

Theorem 1 *The set of minima*

$$(\mu, \nu, \lambda) \in \arg \min_{\hat{\mu}, \hat{\nu}, \hat{\lambda}} \|\hat{\mathbf{G}}^\top \hat{\mathbf{N}}_p\| \|\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k\|, \quad \forall p \in \mathcal{S}_k \quad (18)$$

lies on a segment between the 2-D points $[\mu_0 \ v_0]^\top$ and $[\mu_1 \ v_1]^\top$, where

$$\begin{aligned} \mu_0 &= \mu_1 + \frac{\hat{l}_{1,k}}{\sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}} \theta & \mu_1 &= -\frac{\hat{n}_{1,p}}{\hat{n}_{3,p}} \\ \nu_0 &= \nu_1 + \frac{\hat{l}_{2,k}}{\sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}} \theta & \nu_1 &= -\frac{\hat{n}_{2,p}}{\hat{n}_{3,p}} \\ \theta &= \frac{\hat{\mathbf{N}}_p^\top \hat{\mathbf{L}}_k}{\hat{n}_{3,p} \sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}} \end{aligned} \quad (19)$$

and we defined $\hat{\mathbf{L}}_k = [\hat{l}_{1,k} \ \hat{l}_{2,k} \ \hat{l}_{3,k}]^\top$ and $\hat{\mathbf{N}}_p = [\hat{n}_{1,p} \ \hat{n}_{2,p} \ \hat{n}_{3,p}]^\top$. Given a 2-D point $[\hat{\mu} \ \hat{\nu}]^\top = [\mu_0 \ v_0]^\top + (1 - \alpha) ([\mu_1 \ v_1]^\top - [\mu_0 \ v_0]^\top)$ onto the segment, where $\alpha \in [0, 1]$, the third GBR parameter is uniquely determined via

$$\hat{\lambda} = \sqrt{\alpha(1 - \alpha)} |\theta| \quad (20)$$

and the trajectory of the 3-D point $[\hat{\mu} \ \hat{\nu} \ \hat{\lambda}]^\top$ in the parameter $\alpha \in [0, 1]$ forms a semi-circle of diameter $|\theta|$.

Proof See Appendix. □

The next result tells us how many LDR maxima we need to solve the GBR ambiguity.

Lemma 1 *If two Lambertian diffuse reflectance maxima correspond to two different pseudo-normals \hat{n}_{p_a} and \hat{n}_{p_b} with*

$$\frac{\hat{n}_{1,p_a}}{\hat{n}_{2,p_a}} \neq \frac{\hat{n}_{1,p_b}}{\hat{n}_{2,p_b}}. \quad (21)$$

then there exists a unique solution in the GBR parameters.

Proof See Appendix. □

To illustrate the results, we show in Figs. 2 and 3 how LDR maxima constrain the GBR parameters segments in the (μ, ν) plane and to semi-circles in the (μ, ν, λ) space. All segments or semi-circles intersect at the same point, which corresponds to the correct solution.

Next, we prove a remarkable property of LDR maxima: The GBR parameters obtained from the LDR maxima are GBR consistent. In practice, this means that our method yields the same solution every time it is run, unlike competing methods (see Sect. 7).

Proposition 1 *The constraints on the GBR parameters introduced by an LDR maximum are GBR-consistent.*

Proof See Appendix. □

6.2 Using Pseudo-Normals as Pseudo-Lights

In the previous section, given an LDR we have defined a single equation based on the inner product between a nor-

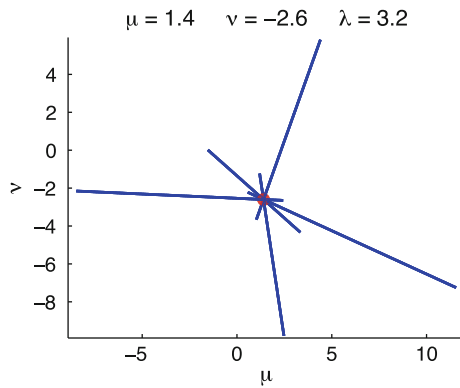


Fig. 3 GBR ambiguity solved via LDR constraints. Each segment is the set of solutions in the parameters $\tilde{\mu}$ and \tilde{v} for one LDR maximum. Segments must intersect at the correct GBR solution

mal vector and a light. However, from the definition of LDR one can derive up to 3 equations in 4 unknowns. The question is then of whether LDRs carry any extra constraint in the GBR solution beyond what has been already presented in Theorem 1. As we will find out, the answer is negative: All the constraints of an LDR point are fully captured in Theorem 1.

Suppose that the set \mathcal{S}_k is given and that one is also given $\hat{\mathbf{N}} = \mathbf{G}^{-T}\mathbf{N}$, the pseudo-normals, and $\hat{\mathbf{L}} = \mathbf{G}\mathbf{L}$, the pseudo-lights. We now impose that the LDR constraint corresponds to having \mathbf{N}_p parallel to \mathbf{L}_k when $p \in \mathcal{S}_k$:

$$\frac{\hat{\mathbf{G}}^T \hat{\mathbf{N}}_p}{\|\hat{\mathbf{G}}^T \hat{\mathbf{N}}_p\|} = \frac{\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k}{\|\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k\|}. \tag{22}$$

We can write the above equation as

$$c_{p,k} \hat{\mathbf{L}}_k = \hat{\mathbf{G}} \hat{\mathbf{G}}^T \hat{\mathbf{N}}_p \tag{23}$$

where $c_{p,k} = \frac{\|\hat{\mathbf{G}}^T \hat{\mathbf{N}}_p\|}{\|\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k\|}$. Notice that although the scalar $c_{p,k}$ depends on the GBR parameters, it also depends on the (unknown) albedo and (unknown) light intensity. Therefore, it can be considered as a free variable. Then, we can finally solve these equations in the 4 unknowns $\hat{\mu}$, \hat{v} , $\hat{\lambda}$, and $c_{p,k}$.

One can verify that these equations are satisfied by the solution in Theorem 1 for any $\alpha \in [0, 1]$ with

$$c_{p,k} = \alpha \frac{\theta n_3}{\sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}}. \tag{24}$$

Hence, Eq. (22) does not provide any additional constraints to the solution in Theorem 1.

6.3 Existence of LDR Maxima

One may wonder whether LDR maxima are commonly found in images. Indeed, surfaces that do not generate LDR maxima do exist. These are surfaces whose normals are never parallel to any light direction; for example, this can easily occur with slopes or cylinders. In these cases it is not possible to detect LDR maxima reliably. Similarly, more general surfaces can also have no LDR maxima. However, in practice this is rarely a problem for general surfaces of interest with a sufficiently uniform sampling of illumination directions. Indeed, we experimentally find an average of 251 LDR maxima per dataset, among those used in Sect. 7. These statistics are very favorable as only 2 LDR maxima are sufficient to solve UPS.

6.4 Detection and Localization of LDR Maxima

In this section we describe a procedure to detect LDR maxima. The main idea is that we can do so by detecting local intensity maxima. In general, however, intensity maxima do not correspond to LDR maxima. In fact, detection of the maxima of $I_{p,k}$ requires

$$\nabla I_{p,k} = \nabla \rho_p \mathbf{N}_p^T \mathbf{l}_k e_k + \rho_p \nabla \mathbf{N}_p^T \mathbf{l}_k e_k = 0, \tag{25}$$

while the LDR maxima are determined by

$$\nabla \mathbf{N}_p^T \mathbf{l}_k = 0. \tag{26}$$

Thus, even small variations of the albedo could result in a displacement of the detected maximum. However, we show that an LDR maximum can be found within a small disc around the detected intensity maximum, whose radius Δ depends on the albedo gradient magnitude and the curvature of the surface.

We approximate the surface of the object around an intensity maximum with a hemisphere. Let

$$\mathbf{P} \doteq [p_1 \ p_2 \ \sqrt{R^2 - p_1^2 - p_2^2}]^T \in \mathbb{R}^3, \tag{27}$$

with $p_1^2 + p_2^2 \leq R^2$, be a point on the hemisphere with radius $R \in \mathbb{R}_+$. If a pixel index p corresponds to the pixel coordinates $[p_1 \ p_2]^T$, then the normals to this surface region can be directly defined as $\mathbf{N}_p = \mathbf{P}/\|\mathbf{P}\|$. Notice that in this case the correct LDR maximum lies at the pixel coordinates $R[l_{1,k} \ l_{2,k}]^T$. If $[p_1 \ p_2]^T$ denotes the detected intensity maximum, we are interested in finding an estimate of the disc radius Δ via the distance

$$\Delta = \left\| R \begin{bmatrix} l_{1,k} \\ l_{2,k} \end{bmatrix} - \begin{bmatrix} p_1 \\ p_2 \end{bmatrix} \right\|. \tag{28}$$

Suppose that the surface is Lambertian and we capture an image with illumination direction \mathbf{l}_k , where $l_{3,k} > 0$, then the measured intensity is $I_{p,k} = \rho_p \frac{\mathbf{P}^\top \mathbf{l}_k}{\|\mathbf{P}\|} e_k$. Detection of the maxima of $I_{p,k}$ requires

$$\nabla I_{p,k} = \nabla \rho_p \frac{\mathbf{P}^\top \mathbf{l}_k}{\|\mathbf{P}\|} e_k + \rho_p \frac{\nabla \mathbf{P}^\top}{\|\mathbf{P}\|} \left(\mathbf{I}_d - \frac{\mathbf{P} \mathbf{P}^\top}{\|\mathbf{P}\|^2} \right) \mathbf{l}_k e_k = 0 \tag{29}$$

where \mathbf{I}_d denotes the identity matrix. Notice that since

$$\nabla \mathbf{P}^\top = \begin{bmatrix} 1 & 0 & -\frac{p_1}{\sqrt{R^2 - p_1^2 - p_2^2}} \\ 0 & 1 & -\frac{p_2}{\sqrt{R^2 - p_1^2 - p_2^2}} \end{bmatrix} \tag{30}$$

we have $\nabla \mathbf{P}^\top \mathbf{P} = \mathbf{0}$, which, combined with Eq. (25), yields the constraint

$$\nabla \rho_p \mathbf{P}^\top \mathbf{l}_k + \rho_p \nabla \mathbf{P}^\top \mathbf{l}_k = 0. \tag{31}$$

Assume that the ratio between the albedo gradient and the albedo magnitude is bound by a constant $B > 0$, i.e., $\left\| \frac{\nabla \rho_p}{\rho_p} \right\| \leq B$. Then, we can write

$$\left\| \nabla \mathbf{P}^\top \mathbf{l}_k \right\| = \left\| \frac{\nabla \rho_p}{\rho_p} \mathbf{P}^\top \mathbf{l}_k \right\| \leq \left\| R \frac{\nabla \rho_p}{\rho_p} \right\| \leq BR, \tag{32}$$

where we used the identity $\|\mathbf{P}\| = R$. Finally, we obtain

$$\left\| \begin{bmatrix} l_{1,k} \\ l_{2,k} \end{bmatrix} - \begin{bmatrix} p_1 \\ p_2 \end{bmatrix} \frac{l_{3,k}}{\sqrt{R^2 - p_1^2 - p_2^2}} \right\| \leq BR. \tag{33}$$

For a reliable detection and localization we consider LDR maxima that are not too far off the detected maxima. This means that locally, we can approximate \mathbf{P} with the hemisphere tangent plane at $R\mathbf{l}_k$. This means that

$$\sqrt{R^2 - p_1^2 - p_2^2} \simeq R l_{3,k} \tag{34}$$

and we obtain

$$\Delta = \left\| R \begin{bmatrix} l_{1,k} \\ l_{2,k} \end{bmatrix} - \begin{bmatrix} p_1 \\ p_2 \end{bmatrix} \right\| \leq BR^2. \tag{35}$$

Therefore, the correct localization of LDR maxima will be a tradeoff between the curvature of the local surface and the albedo gradient: Surfaces with high curvature can tolerate high albedo gradients and, vice versa, surfaces with small curvature can tolerate little albedo variation. We experimentally find that $\Delta = 1$ is a good tradeoff.

As done in [Lagger and Fua \(2008\)](#), we improve the selection of LDR maxima by discarding maxima that appear at

the same spatial location under different illuminations (two are sufficient to decide) as they are most likely generated by the albedo texture. Third, we discard maxima with low pixel intensities, e.g., maxima from the k -th input image with intensity less than $(\max_p I_{p,k} - \min_p I_{p,k})/2$. We experimentally found that such maxima tend to be less reliable than those with high intensity.

6.5 Dealing with Outliers

In practice, the sets $\{\mathcal{S}_k\}_{k=1,\dots,K}$ contain many outliers. We take outliers into account by using a robust estimation procedure.

To simplify the notation, let us define $\mathbf{y} \doteq [\mu \ v \ \lambda]^\top$, the vector of GBR parameters. Let S be the sum of the cardinalities of the sets $\{\mathcal{S}_k\}_{k=1,\dots,K}$. Also, let $\{\mathbf{y}_i\}_{i=1,\dots,S(S-1)/2}$ be the set of all the intersections between any two semi-circles defined by two detected maxima lying in one of the sets $\{\mathcal{S}_k\}_{k=1,\dots,K}$. We consider the intersections $\{\mathbf{y}_i\}_{i=1,\dots,S(S-1)/2}$ as samples of a random variable \mathbf{Y} defining the GBR parameters. We can obtain an approximation to the distribution $\mathbf{f}_\mathbf{Y}$ of the random variable \mathbf{Y} via the empirical distribution

$$\mathbf{f}_\mathbf{Y}(\mathbf{y}) \simeq \frac{2}{S(S-1)} \sum_{i=1}^{\frac{S(S-1)}{2}} \delta(\mathbf{y} - \mathbf{y}_i). \tag{36}$$

Then, the GBR parameters can be chosen as the median $\bar{\mathbf{y}}$ of the empirical distribution, i.e.,

$$(\tilde{\mu}, \tilde{v}, \tilde{\lambda}) \equiv \bar{\mathbf{y}} = \arg \min_{\hat{\mathbf{y}}} \mathbf{E}_{\mathbf{f}_\mathbf{Y}} [|\hat{\mathbf{y}} - \mathbf{y}|] \tag{37}$$

where $\mathbf{E}_{\mathbf{f}_\mathbf{Y}}$ denotes the expectation with respect to the distribution $\mathbf{f}_\mathbf{Y}$. This formalization makes it possible to further improve this outlier rejection method by using other approximations for $\mathbf{f}_\mathbf{Y}$ (for example, with a Kernel density estimator).

Remark 1 As we have seen in the previous section, the localization of an LDR maximum via a local brightness maximum is affected by the albedo of the surface. However, as long as the albedo introduces a small random error, the likelihood of distinguishing the correct LDRs via the median is high. That is because with high likelihood only the correct detected LDR maxima will agree on the same correct GBR solution. It is however possible to find albedos that introduce a non-random error leading to agreement on an incorrect GBR solution. Fortunately, although these cases are possible, they are extremely rare: Such albedos must be a function of the normal components so as to consistently introduce the same error at corresponding normals. One artificial example of such albedos can be built by taking the ratio between two images captured under different illumination (see right image

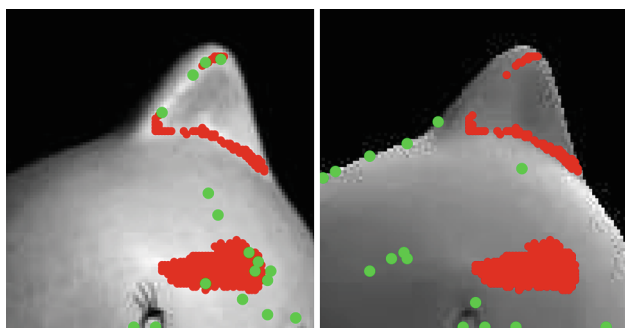


Fig. 4 Detected image brightness maxima are shown as *green dots*. The image regions where the normals are almost parallel to the light direction are shown in *red*. We show the detection of LDR maxima in the original image (*left*) and in the ratio image (*right*). Notice that as the albedo is no longer independent of the normal map, the ratio may cause the LDR maxima to vote with strong consensus for one or many incorrect GBR solutions. In this illustration the detected LDR maxima in the ratio image (*right*) are quite far from the correct ones (Color figure online)

in Fig. 4). In this case the albedo becomes the inverse of a linear combination of the normals. Let us briefly analyze the ratio of two images k and i . The detection of LDR maxima on these “image ratios” relies on the brightness maxima, which can be found at pixels p where

$$\nabla \left(\frac{I_{p,k}}{I_{p,i}} \right) = \nabla \left(\frac{\mathbf{N}_p^\top \mathbf{l}_k e_k}{\mathbf{N}_p^\top \mathbf{l}_i e_i} \right) = 0 \Rightarrow \frac{\nabla \mathbf{N}_p^\top \mathbf{l}_k}{\mathbf{N}_p^\top \mathbf{l}_k} = \frac{\nabla \mathbf{N}_p^\top \mathbf{l}_i}{\mathbf{N}_p^\top \mathbf{l}_i}. \tag{38}$$

Notice that LDR maxima would instead require $\nabla \mathbf{N}_p^\top \mathbf{l}_k = 0$, but the above equation may instead impose that $\nabla \mathbf{N}_p^\top \mathbf{l}_k \neq 0$ as LDR maxima are typically not shared across different input images.

6.6 Robustness

Because the proposed algorithm can compute the GBR parameters very efficiently, we can evaluate how well it can tolerate several levels of outliers and errors in the detected LDR maxima. We show that the algorithm can tolerate very high levels of outliers, while it is more sensitive to persistent misalignments between normals and lights. We synthetically generate a set of $K = 12$ light directions. Then, we generate $P = 500$ maxima by matching each normal to one of the light directions. With reference to Fig. 5, to simulate outliers in the detected LDR maxima, we corrupt outliers ratio normals every 100 normals. To simulate errors in the detected LDR maxima, we add to all the normals uniform noise in the range uniform noise ratio $\times [-0.1 + 0.1]$ (notice that the normals have magnitude equal to 1). Uniform noise corresponds to persistent misalignment. The performance is

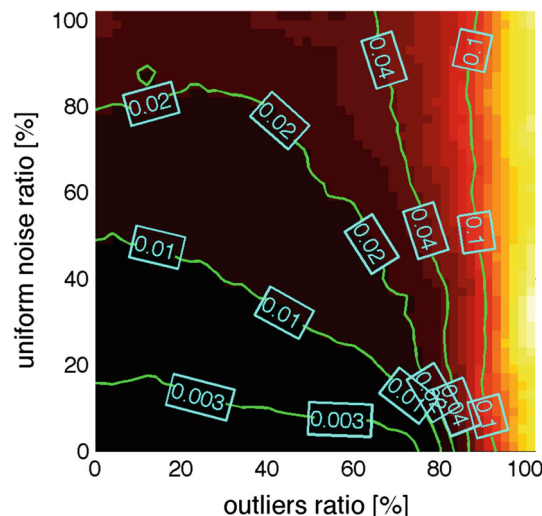


Fig. 5 Color coding of the estimation error against outliers and noise ratios. Error levels are shown with *solid green lines* and the corresponding values are indicated in the *cyan box* (Color figure online)

computed as the relative error $\epsilon = \frac{\|\mathbf{y} - \hat{\mathbf{y}}\|_2}{\|\mathbf{y}\|_2}$, where \mathbf{y} is the correct vector of GBR parameters and $\hat{\mathbf{y}}$ the estimated one. In Fig. 5 we show the results for several outlier and noise ratios. To ease visualization we smooth the error map and then quantize it. The outlier ratio is shown in the abscissa and ranges from 0 % (left) to 100 % (right). Similarly, the uniform noise ratio is shown in the ordinate axis and ranges from 0 % (bottom) to 100 % (top). Notice that the darkest region (bottom-left) denotes no more than 0.3 % relative error and it allows up to 75 % outliers even when the remaining 25 % of the maxima have around 10 % uniform additive noise.

7 Experiments

To validate our method we use nine real object datasets: **Redfish** and **Octopus** of 5 images each (courtesy of Neil Alldrin³); **Cat**, **Buddha**, **Owl**, **Horse** and **Rock** of 12 images each (courtesy of Dan Goldman and Steven Seitz⁴); **Puppet** of 15 images taken in our lab with a Canon 5D Mark II; **Face** of 64 images (Yale Face Database (Georghiades et al. 2001)). All the experiments were run with Matlab on a Mac (2.66 GHz CPU, Intel Core Duo) platform.

7.1 Image Acquisition of the Puppet

To take the photos of the Puppet, we used a Canon 5D Mark II camera and a 30-watt incandescent light source attached on a tripod. We placed the light source in 15 different positions

³ <http://neilalldrin.com/research>

⁴ <http://www.cs.washington.edu/education/courses/csep576/05wi/projects/project3/project3.htm>

(far enough from the object in order to satisfy the distant point light source assumption) and took three photos for each position. The three photos were then averaged to reduce noise. Ambient image subtraction was also performed. Because the light source has constant intensity, we scaled each image with a random scalar between 0.5 and 1.5 (this was done for the face dataset as well). An object mask was also created by simple thresholding.

7.2 Image Pre-processing

Equation (2) dictates that the intensity matrix \mathbf{I} should have rank 3. Because real world data are affected by noise, shadows, specularities and other non/Lambertian effects, the rank will be different, typically bigger, than 3. Hence, we perform a pre-processing step based on a recent convex optimization approach (Lin et al. 2009; Wu et al. 2011) that recovers the low-rank intensity matrix without having any information about what entries are corrupted:

$$\min_{\mathbf{A}, \mathbf{E}} \|\mathbf{A}\|_* + \gamma \|\mathbf{E}\|_1 \quad s.t. \quad \mathbf{I} = \mathbf{A} + \mathbf{E} \quad (39)$$

where \mathbf{A} is the noise-free low-rank data and \mathbf{E} is a sparse matrix with the outliers. $\|\cdot\|_*$ and $\|\cdot\|_1$ denote the nuclear and L^1 norm respectively. $\gamma > 0$ is a weighting parameter that defines the amount of outliers in the data. This parameter depends on the input images and, as suggested in Wright et al. (2009), we fix it to $\gamma = \frac{\kappa}{\sqrt{P}}$ where κ is a constant and P is the total number of pixels of a single input image. We fixed $\kappa = 1.7$ for all datasets having at least 12 images and $\kappa = 3$ for the others. Figure 6 shows how the pre-processing algorithm removes shadows and specularities in the case of the Face dataset.

7.3 Implementation Details

We impose the integrability constraint as described in Yuille and Snow (1997) and the surface pseudo-normals $\hat{\mathbf{N}}$ and pseudo-lights $\hat{\mathbf{L}}$ are obtained up to the GBR ambiguity. The initial group of all the LDR maxima is detected by using



Fig. 6 Removal of non Lambertian effects. On the *left* we show one of the input images. In the *middle* image we show the detected non Lambertian components (shadows and specularities) via the low-rank matrix recovery pre-processing step. Finally, on the *right* image we show the low-rank and noise-free data used in our algorithm

the built-in Matlab function “imregionalmax” after applying a small Gaussian smoothing to all the input images. This function gives as output a set of single pixels for each image corresponding to the intensity local maxima. As described in Sect. 6.4 we select as local maxima all pixels in a disc neighborhood of radius 1 pixel. Then, we discard all local maxima that appear (at least twice) at the same spatial location for different illumination conditions and those that have low intensity values.

7.4 Results

Because no ground truth data is available for these datasets, we consider as ground truth for the normal maps the ones obtained from the calibrated photometric stereo method. We compare our experimental results with the state-of-the-art in uncalibrated photometric stereo (Shi et al. 2010; Alldrin et al. 2007). For the accuracy evaluation we consider the

Table 1 Performance comparison with the entropy minimization (EM) method (Alldrin et al. 2007) and the self calibrating photometric stereo (SCPS) method (Shi et al. 2010)

$K = 12$	A/E	Rock	Buddha	Horse	Cat	Owl
EM	ϵ	22.16	15.05	20.65	15.39	18.48
	σ	1.88	2.19	3.85	3.78	5.58
SCPS	ϵ	24.88	13.58	21.01	6.15	10.47
	σ	7.42	4.93	9.57	2.83	4.75
LDR	ϵ	11.61	4.98	4.80	5.37	6.63
	σ	0	0	0	0	0
LDR with GT	ϵ	2.67	3.11	2.27	2.72	5.76
	σ	0	0	0	0	0
LDR without preprocessing	ϵ	14.30	5.97	17.03	10.16	6.91
	σ	0	0	0	0	0
$K = 5$	A/E	Octopus	Red fish	$K = 15$	Puppet	
EM	ϵ	9.03	8.63	EM	26.36	
	σ	0.76	1.14		2.39	
SCPS	ϵ	13.23	7.60	SCPS	NA	
	σ	9.85	4.32		NA	
LDR	ϵ	6.64	5.60	LDR	12.15	
	σ	0	0		0	
LDR with GT	ϵ	2.87	4.47	LDR with GT	5.63	
	σ	0	0		0	
LDR without preprocessing	ϵ	11.39	10.82	LDR with GT	20.15	
	σ	0	0		0	

We show the mean (ϵ) and standard deviation (σ) of the angular error of the estimated normal maps

mean angular error of the estimated normals with respect to the calibrated case. In Table 1 we show the performance comparison of all methods in the case of datasets with 12 input images (top) and in the case of datasets with 5 and 15 input images (bottom). Notice that $\sigma < 10^{-12}$ in all the reconstructions via LDR maxima as predicted by our analysis (see Proposition 1). Since our method provides a GBR-consistent solution, we always get the same estimate no matter what pseudo-normals and pseudo-lights we start from. This is a major advantage over prior work.

Figure 7 shows the comparison of the obtained normal and depth maps for the methods (Shi et al. 2010; Alldrin



Fig. 7 Comparison of reconstructed normal and depth maps. From left to right, each column corresponds to the first (x axis), second (y axis), and third (z axis) components of the normal map \mathbf{n} and the depth map, respectively. We show the results obtained from the calibrated photometric stereo method (top row), the results obtained from our method (second row), the results obtained from the entropy minimization method (third row) and the results obtained from the SCPS method (bottom row)

et al. 2007) in the case of the **Buddha** dataset. The depth map is obtained by integrating the normal map with a Poisson solver (Agrawal and Raskar 2006). Figure 8 shows all the depth maps obtained by our method across all the datasets (bottom row) and the depth maps obtained from the calibrated photometric stereo method (middle row).

All datasets yield visually reasonable surface reconstructions. However, in the **Owl** dataset one might notice an artifact (a spike) of the reconstructed surface on the left eye. This artifact is caused by specular highlights (probably because the eye is made of glass) that were not removed by the preprocessing step. Notice that this artifact is present also in the calibrated reconstruction. One possible explanation for the failure of the preprocessing algorithm (Lin et al. 2009; Wu et al. 2011) is that the outliers in this case have very low intensities that can be confused with small Gaussian noise. In Fig. 9 we show an enlargement of that region in one of the input images. To demonstrate that outliers are the cause of the artifact, we mask out these pixels (by using some simple thresholding) before reconstructing the surface. In Fig. 9 one can also see the comparison of the surface reconstruction around the eye with and without outliers. The reconstruction without outliers is much smoother and looks qualitatively correct.

Finally, the average time⁵ to run our method, implemented in non/optimized Matlab code, on all the datasets is 13.5 s (we also include the preprocessing step). Notice that the average time for Alldrin et al. (2007) is 62 s, while for Shi et al. (2010) is 10 min. Moreover, since our standard deviation is zero, we only need to run the algorithm once, while the other methods need to be run several times. To appreciate the fine details of the reconstructed surfaces, in Fig. 10 we show two views of the surfaces obtained for all datasets. The first and the third columns are the results obtained with photometric stereo. The other two columns are obtained with our method. We also tested our method on face reconstruction by using data from Yale's face database and the reconstructed depth map and surfaces are shown in Fig. 11.

7.5 Discussion

In the second row (from the bottom) of each table in Table 1, we also show the performance when the LDR maxima are guaranteed to be true LDR maxima (as they are selected from the ground truth normal map—obtained with calibrated photometric stereo). The performance based on the ground truth shows that the algorithm would not achieve 0 error even if the detected maxima were correct. As the theory is exact, this limitation can only be due to non Lambertian compo-

⁵ Notice that the running time of our algorithm depends mostly on the number of detected local maxima and much less on the size of the input images.

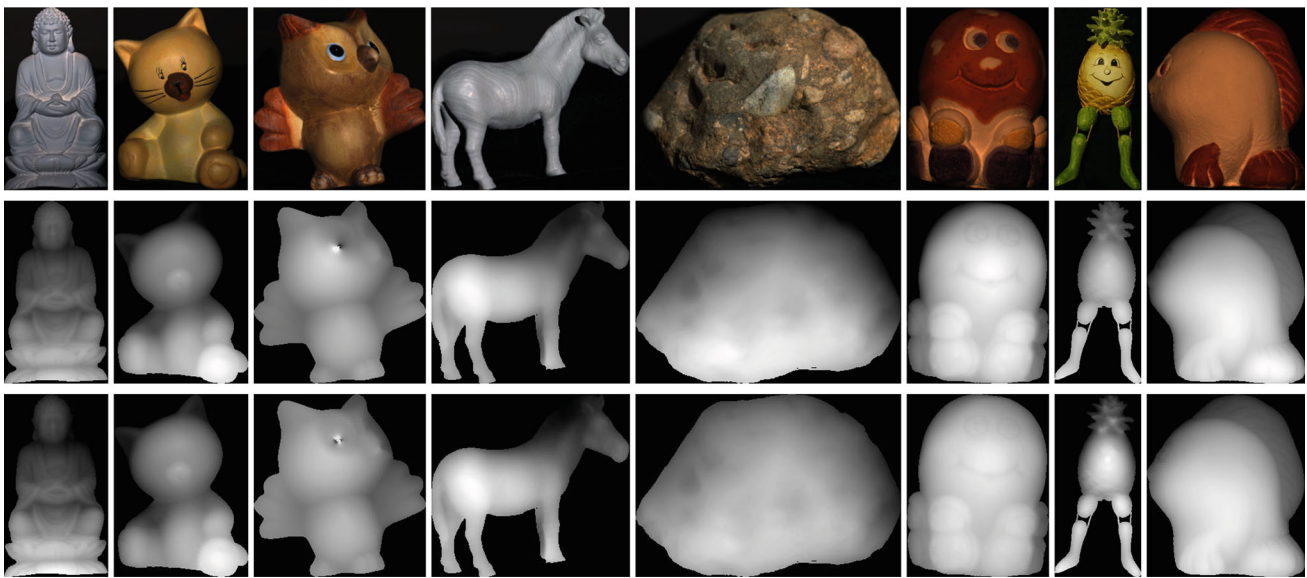


Fig. 8 Reconstructed depth maps. One of the input images (*top row*). Graylevel-coded depth map reconstructed with calibrated photometric stereo (*middle row*). Graylevel-coded depth map obtained with our method (*bottom row*)

nents still present in the model. Indeed, the pre-processing step (Lin et al. 2009; Wu et al. 2011) may not fully remove specularities and shadows as already observed in the case of the **Owl** dataset. Finally, in the last row we show the mean angular errors without the pre-processing step. Although we still achieve lower error rates compared to prior methods on many datasets, the reconstruction accuracy degrades. This is true especially for datasets (e.g., **Horse**) where there is a non-negligible specular component.

7.6 Limitations and Future Work

Our method is entirely based on the diffuse component of the images. Even though the detection of maxima is very robust, their estimation degrades in the presence of shadows, noise, interreflections and specularities. We found the latter ones particularly important. Thus, a possible suggestion of future work is to find a way to detect/model specularities and use them (at the moment we just discard them) together with the LDR maxima to better disambiguate the uncalibrated photometric stereo problem.

8 Conclusion

In this paper we presented a simple and fast method for solving the generalized bas relief (GBR) ambiguity of the uncalibrated photometric stereo problem. Our method makes no stringent assumptions about the distribution of lights, only that they are sufficiently different, and can handle a wide variety of albedos and surfaces. We introduce a novel Lam-

bertian constraint for the distribution of lights and normals, which can be computed in closed-form. We use this new constraint to design a very robust algorithm that achieves the best results to date and provides the same estimate regardless of the GBR ambiguity. Key to our approach is a reliable detection of Lambertian diffuse reflectance maxima in the input images and a robust formulation of the GBR estimation process.

Appendix

In this section we provide the detailed proofs for the LDR solutions, how a unique solution can be identified given 2 LDR maxima, and that these solutions satisfy the GBR consistency property. First, we need to introduce a simple technique to obtain analytical solutions in a minimization problem.

Lemma 2 *Let ϕ be a function of two variables and ϕ_{ξ_1} be its derivative with respect to the first argument ξ_1 . Then, if $\phi_{\xi_1}(\xi_1, \xi_2) = 0$ if and only if $\xi_1 = f(\xi_2)$ for any ξ_2 , the following two minimization problems yield the same extrema*

$$(P_1) \min_{\xi_1, \xi_2} \phi(\xi_1, \xi_2) \tag{40}$$

$$(P_2) \min_{\xi_2} \phi(f(\xi_2), \xi_2) \text{ with } \xi_1 = f(\xi_2).$$

Proof Problem P_1 gives the following conditions on the gradients of ϕ :

$$\phi_{\xi_1}(\xi_1, \xi_2) = 0 \quad \phi_{\xi_2}(\xi_1, \xi_2) = 0 \tag{41}$$

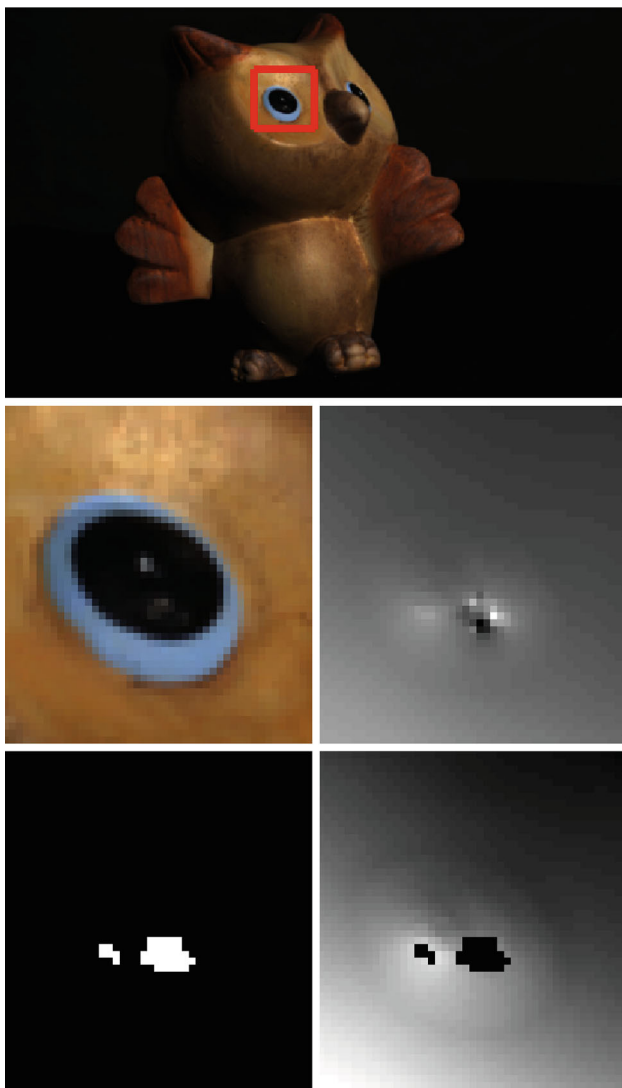


Fig. 9 On the *top* we show one of the input images of the **Owl** dataset. In *red* we show the selected patch. In the *second row—left* image we show the selected patch where some outliers are visible inside the pupil. In the *second row—right* image we show the reconstructed depth map without outlier rejection. Notice that the outliers cause sharp discontinuities in the depth map. In the *bottom—left* image we show the detected specular highlights (simple thresholding based on the standard deviation across different illumination settings). Finally in the *bottom—right* image we show the reconstructed depth map after rejection of the detected specular highlights. Notice that, unlike in the previous case, the recovered depth map does not have sharp discontinuities (Color figure online)

since $\phi_{\xi_1}(\xi_1, \xi_2) = 0$ implies that $\xi_1 = f(\xi_2)$, we must have

$$\phi_{\xi_1}(f(\xi_2), \xi_2) = 0 \quad \phi_{\xi_2}(f(\xi_2), \xi_2) = 0 \quad (42)$$

Problem P_2 gives the following condition on the gradients of ϕ :

$$\phi_{\xi_1}(f(\xi_2), \xi_2) \times f_{\xi_2}(\xi_2) + \phi_{\xi_2}(f(\xi_2), \xi_2) = 0. \quad (43)$$

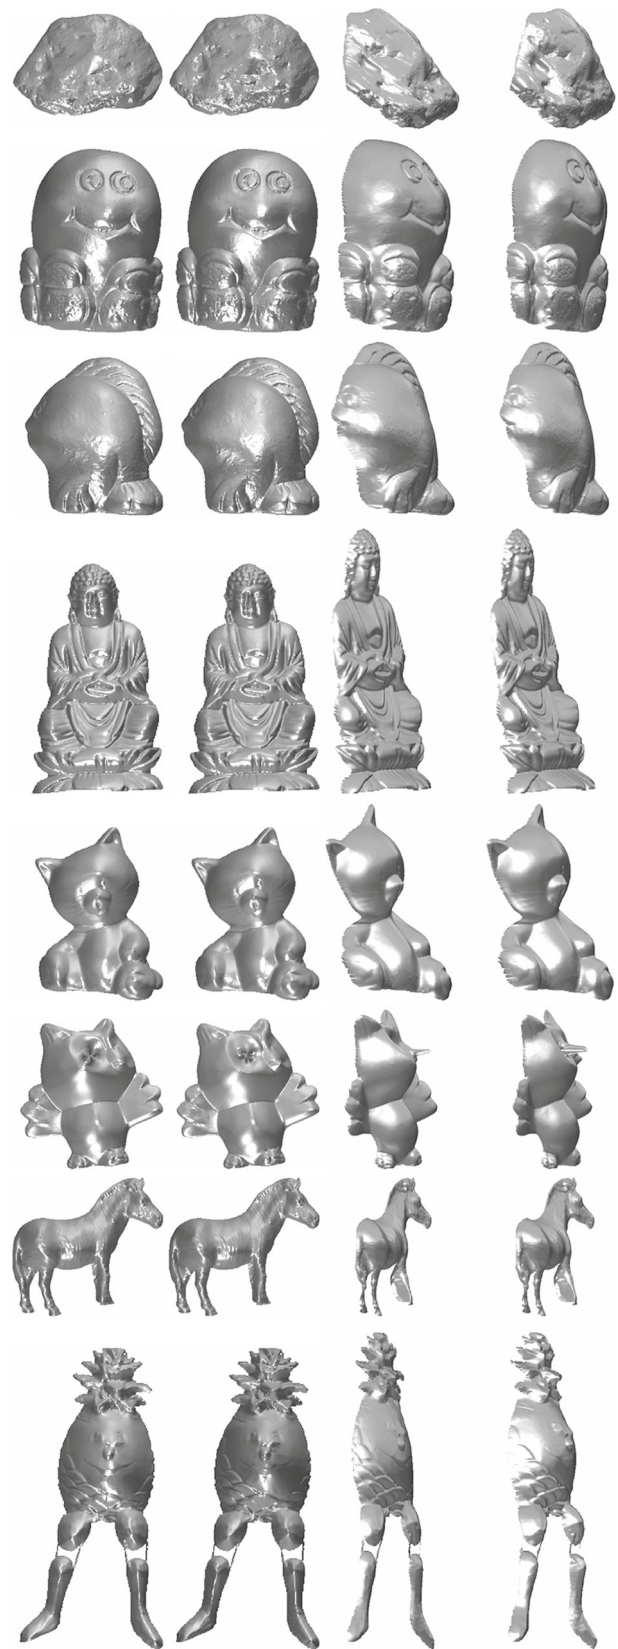


Fig. 10 Surface reconstructions. Rendered surfaces obtained with the calibrated photometric stereo method (*first and third column*). Rendered surfaces obtained with our method (*second and fourth column*)

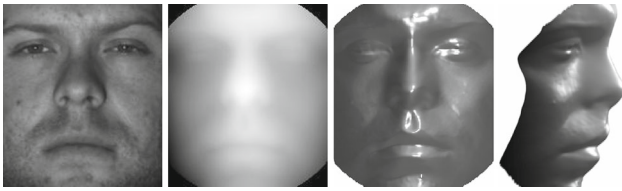


Fig. 11 Results on Yale’s Face database. One of the input images (*left*). Reconstructed depth map (*second from the left*). Two views of the reconstructed surface (*third and fourth from the left*)

Since $\phi_{\xi_1}(f(\xi_2), \xi_2) = 0$ by definition of f , the first term in the above equation disappears and we have exactly the same set of equations as problem P_1 . Notice that the same result can be easily extended to more than 2 variables. \square

We intend to use the above lemma in the Theorem 1 to derive an analytic solution.

Proof of Theorem 1 Firstly, rather than minimizing Eq. (17) we can equivalently minimize the cost function

$$\min_{\hat{\mu}, \hat{\nu}, \hat{\lambda}} \|\hat{\mathbf{G}}^T \hat{\mathbf{N}}_p\|^2 \|\hat{\mathbf{G}}^{-1} \hat{\mathbf{L}}_k\|^2, \quad \forall p \in \mathcal{S}_k. \quad (44)$$

Now, recall that

$$\hat{\mathbf{G}}^T = \begin{bmatrix} 1 & 0 & \hat{\mu} \\ 0 & 1 & \hat{\nu} \\ 0 & 0 & \hat{\lambda} \end{bmatrix} \quad \text{and} \quad \hat{\mathbf{G}}^{-1} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -\frac{\hat{\mu}}{\hat{\lambda}} & -\frac{\hat{\nu}}{\hat{\lambda}} & \frac{1}{\hat{\lambda}} \end{bmatrix}; \quad (45)$$

then, $\forall p \in \mathcal{S}_k$ the cost function can be explicitly written as

$$\begin{aligned} \min_{\hat{\mu}, \hat{\nu}, \hat{\lambda}} \frac{1}{\hat{\lambda}^2} & \left((\hat{n}_{1,p} + \hat{\mu}\hat{n}_{3,p})^2 + (\hat{n}_{2,p} + \hat{\nu}\hat{n}_{3,p})^2 + \hat{\lambda}^2 \hat{n}_{3,p}^2 \right) \\ & \times \left(\hat{\lambda}^2 (\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2) + (-\hat{\mu}\hat{l}_{1,k} - \hat{\nu}\hat{l}_{2,k} + \hat{l}_{3,k})^2 \right). \end{aligned} \quad (46)$$

Notice that the cost is a convex function in the parameter $\hat{\lambda} \in (0, \infty)$. Consider the case $\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2 = 0$, then, the minimization problem has infinite solutions with $\hat{\lambda} = \infty$ and any $\hat{\mu}, \hat{\nu} \in \mathbb{R}$. This is a degenerate case not only for the LDR constraint theory, but also for photometric stereo in general. Notice that the first two components of the normalized pseudo-lights always coincide with the first two components of the true lights. Then, in this case, the true light direction can be immediately obtained from the pseudo-light. Notice that with this type of illumination the normals can never be retrieved from the Lambertian model. Now, suppose that $\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2 \neq 0$ and the values of $\hat{\mu}$ and $\hat{\nu}$ are given; then, a necessary condition for a minimum is that the first order derivative of the cost with respect to $\hat{\lambda}$ is zero,

$$\begin{aligned} \hat{n}_{3,p}^2 \hat{\lambda}^4 (\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2) &= (\hat{n}_{1,p} + \hat{\mu}\hat{n}_{3,p})^2 + (\hat{n}_{2,p} + \hat{\nu}\hat{n}_{3,p})^2 \\ &\times (-\hat{\mu}\hat{l}_{1,k} - \hat{\nu}\hat{l}_{2,k} + \hat{l}_{3,k})^2 \end{aligned} \quad (47)$$

which immediately yields the unique solution (in $(0, \infty)$)

$$\hat{\lambda} = \sqrt[4]{\frac{[(\hat{n}_{1,p} + \hat{\mu}\hat{n}_{3,p})^2 + (\hat{n}_{2,p} + \hat{\nu}\hat{n}_{3,p})^2] (-\hat{\mu}\hat{l}_{1,k} - \hat{\nu}\hat{l}_{2,k} + \hat{l}_{3,k})^2}{\hat{n}_{3,p}^2 (\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2)}}. \quad (48)$$

Now, let us use Lemma 2, substitute the above solution in the cost (46) and solve the minimization in the two remaining unknowns $\hat{\mu}$ and $\hat{\nu}$; one obtains the equivalent problem

$$\begin{aligned} \arg \min_{\hat{\mu}, \hat{\nu}} & \sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2} \sqrt{(\hat{n}_{1,p} + \hat{\mu}\hat{n}_{3,p})^2 + (\hat{n}_{2,p} + \hat{\nu}\hat{n}_{3,p})^2} \\ & + |\hat{n}_{3,p}| \left| -\hat{\mu}\hat{l}_{1,k} - \hat{\nu}\hat{l}_{2,k} + \hat{l}_{3,k} \right| \end{aligned} \quad (49)$$

which is a convex function in both $\hat{\mu}$ and $\hat{\nu}$ (as each energy term is convex in the unknowns). Let us assume that $\hat{\mu}, \hat{\nu}$ are such that $-\hat{\mu}\hat{l}_{1,k} - \hat{\nu}\hat{l}_{2,k} + \hat{l}_{3,k} \neq 0$. Then, we can compute the first order derivatives of the cost in Eq. (49) in both $\hat{\mu}$ and $\hat{\nu}$; by equating them to zero we obtain

$$\hat{\nu} = \frac{\hat{l}_{2,k}}{\hat{l}_{1,k}\hat{n}_{3,p}} (\hat{n}_{1,p} + \hat{n}_{3,p}\hat{\mu}) - \frac{\hat{n}_{2,p}}{\hat{n}_{3,p}}. \quad (50)$$

Now suppose that $\hat{\mu}, \hat{\nu}$ are such that $-\hat{\mu}\hat{l}_{1,k} - \hat{\nu}\hat{l}_{2,k} + \hat{l}_{3,k} = 0$; then, we have

$$\hat{\nu} = \frac{\hat{l}_{3,k} - \hat{\mu}\hat{l}_{1,k}}{\hat{l}_{2,k}}. \quad (51)$$

To find the minimum we plug both solutions in the cost and look for $\hat{\mu}$ yielding the smallest value among the two cases. In the first case, we can simplify the minimization to

$$\min_{\hat{\mu}} \frac{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}{|\hat{l}_{1,k}|} \left[|\hat{n}_{1,p} + \hat{\mu}\hat{n}_{3,p}| + \left| \hat{n}_{1,p} + \hat{\mu}\hat{n}_{3,p} - \frac{\hat{l}_{1,k}\hat{\mathbf{N}}_p^T \hat{\mathbf{L}}_k}{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2} \right| \right] \quad (52)$$

and, in the second case, we have

$$\begin{aligned} \min_{\hat{\mu}} & \sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2} \\ & \times \sqrt{(\hat{n}_{1,p} + \hat{\mu}\hat{n}_{3,p})^2 + \left(\hat{n}_{2,p} + \frac{\hat{l}_{3,k} - \hat{\mu}\hat{l}_{1,k}}{\hat{l}_{2,k}} \hat{n}_{3,p} \right)^2} \end{aligned} \quad (53)$$

In the first case, the solutions in $\hat{\mu}$ are readily obtained as $\hat{\mu} \in [\min(\hat{\mu}_0, \hat{\mu}_1), \max(\hat{\mu}_0, \hat{\mu}_1)]$ where

$$\hat{\mu}_1 = -\frac{\hat{n}_{1,p}}{\hat{n}_{3,p}} \quad \hat{\mu}_0 = \frac{\hat{l}_{1,k}\hat{l}_{2,k}\hat{n}_{2,p} - \hat{l}_{2,k}^2\hat{n}_{1,p} + \hat{l}_{3,k}\hat{l}_{1,k}\hat{n}_{3,p}}{\hat{n}_{3,p}(\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2)}. \quad (54)$$

The gradient of the second cost gives

$$\hat{\mu}_* = \frac{-\hat{l}_{2,k}^2 \hat{n}_{1,p} + \hat{l}_{1,k} \hat{l}_{2,k} \hat{n}_{2,p} + \hat{l}_{1,k} \hat{l}_{3,k} \hat{n}_{3,p}}{\hat{n}_{3,p} (\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2)} \tag{55}$$

which is identical to $\hat{\mu}_0$. Hence, we only need to substitute $\hat{\mu}_0$ in both costs. A few calculations show that the value of both Eqs. (52) and (53) is $|\hat{\mathbf{N}}_p^\top \hat{\mathbf{L}}_k|$ and, therefore, the solutions are $\hat{\mu} \in [\min(\hat{\mu}_0, \hat{\mu}_1), \max(\hat{\mu}_0, \hat{\mu}_1)]$. Now, given all the solutions in $\hat{\mu}$ we can substitute in the previous expressions and obtain an explicit solution for \hat{v} and for $\hat{\lambda}$. By using the relationship in Eq. (50), we have $\hat{v} \in [\min(\hat{v}_0, \hat{v}_1), \max(\hat{v}_0, \hat{v}_1)]$ where

$$\hat{v}_1 = -\frac{\hat{n}_{2,p}}{\hat{n}_{3,p}} \quad \hat{v}_0 = \frac{\hat{l}_{1,k} \hat{l}_{2,k} \hat{n}_{1,p} - \hat{l}_{1,k}^2 \hat{n}_{2,p} + \hat{l}_{2,k} \hat{l}_{3,k} \hat{n}_{3,p}}{\hat{n}_{3,p} (\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2)} \tag{56}$$

Notice also that given $\theta = \frac{\hat{\mathbf{N}}_p^\top \hat{\mathbf{L}}_k}{\hat{n}_{3,p} \sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}}$, the following identities hold

$$\hat{\mu}_0 = \hat{\mu}_1 + \frac{\hat{l}_{1,k} \theta}{\sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}} \quad \hat{v}_0 = \hat{v}_1 + \frac{\hat{l}_{2,k} \theta}{\sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}} \tag{57}$$

Equation (20) can be readily obtained by direct substitution of the point $[\hat{\mu} \ \hat{v}]^\top = [\mu_0 \ v_0]^\top + (1 - \alpha)([\mu_1 \ v_1]^\top - [\mu_0 \ v_0]^\top)$ in Eq. (48). Finally, we show that the solution in α is a semi-circle by computing its center and diameter. Firstly, define $(\bar{\mu}, \bar{v}, 0)$ as the center of the circle. By using Eq. (57), we obtain

$$\bar{\mu} = \hat{\mu}_1 + \frac{\hat{l}_{1,k} \theta}{2\sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}} \quad \bar{v} = \hat{v}_1 + \frac{\hat{l}_{2,k} \theta}{2\sqrt{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2}} \tag{58}$$

and

$$\begin{aligned} (\hat{\mu} - \bar{\mu})^2 &= \left(\alpha - \frac{1}{2}\right)^2 \frac{\hat{l}_{1,k}^2 \theta^2}{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2} \\ (\hat{v} - \bar{v})^2 &= \left(\alpha - \frac{1}{2}\right)^2 \frac{\hat{l}_{2,k}^2 \theta^2}{\hat{l}_{1,k}^2 + \hat{l}_{2,k}^2} \\ \hat{\lambda}^2 &= \alpha(1 - \alpha)\theta^2. \end{aligned} \tag{59}$$

By summing the above equations and by taking the square root, one can immediately draw the conclusion that the diameter is the constant θ .

Proof of Lemma 1 If two Lambertian diffuse reflectance maxima correspond to two different pseudo-normals \hat{n}_{p_a} and \hat{n}_{p_b} with

$$\frac{\hat{n}_{1,p_a}}{\hat{n}_{2,p_a}} \neq \frac{\hat{n}_{1,p_b}}{\hat{n}_{2,p_b}} \tag{60}$$

then there is a unique solution in the GBR parameters.

Proof Suppose that two LDR maxima with pixel indices p_a and p_b , obtained with illumination k_a and k_b respectively, correspond to two pseudo-normals with different directions; then, by using the LDR maxima constraint, the pseudo-normals will have the same direction as the pseudo-light directions, and therefore the corresponding pseudo-lights will be different and satisfy an inequality corresponding to the one above. Now, to find the explicit intersection between the two segments identified by these two LDR maxima, we need to solve

$$\begin{aligned} \begin{bmatrix} \mu_{1,a} \\ \nu_{1,a} \end{bmatrix} + \alpha \left(\begin{bmatrix} \mu_{0,a} \\ \nu_{0,a} \end{bmatrix} - \begin{bmatrix} \mu_{1,a} \\ \nu_{1,a} \end{bmatrix} \right) \\ = \begin{bmatrix} \mu_{1,b} \\ \nu_{1,b} \end{bmatrix} + \beta \left(\begin{bmatrix} \mu_{0,b} \\ \nu_{0,b} \end{bmatrix} - \begin{bmatrix} \mu_{1,b} \\ \nu_{1,b} \end{bmatrix} \right) \end{aligned} \tag{61}$$

in the unknown α and β . Let us define

$$\theta_a \doteq \frac{\hat{\mathbf{N}}_{p_a}^\top \hat{\mathbf{L}}_{k_a}}{\hat{n}_{3,p_a} (\hat{l}_{1,k_a}^2 + \hat{l}_{2,k_a}^2)} \quad \theta_b \doteq \frac{\hat{\mathbf{N}}_{p_b}^\top \hat{\mathbf{L}}_{k_b}}{\hat{n}_{3,p_b} (\hat{l}_{1,k_b}^2 + \hat{l}_{2,k_b}^2)} \tag{62}$$

Then, the intersection can be computed by solving the following linear system

$$\begin{bmatrix} \hat{l}_{1,k_a} \theta_a & -\hat{l}_{1,k_b} \theta_b \\ \hat{l}_{2,k_a} \theta_a & -\hat{l}_{2,k_b} \theta_b \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} \mu_{1,b} - \mu_{1,a} \\ \nu_{1,b} - \nu_{1,a} \end{bmatrix} \tag{63}$$

To have a unique solution the determinant of the 2×2 matrix on the left hand side must be nonzero. Hence, we find

$$-\hat{l}_{1,k_a} \hat{l}_{2,k_b} \theta_a \theta_b + \hat{l}_{1,k_b} \hat{l}_{2,k_a} \theta_a \theta_b \neq 0 \tag{64}$$

which can be simplified as

$$\frac{\hat{l}_{1,k_a}}{\hat{l}_{2,k_a}} \neq \frac{\hat{l}_{1,k_b}}{\hat{l}_{2,k_b}} \tag{65}$$

The latter constraint is satisfied due to the inequality on the pseudo-normals in the assumptions. \square

Proof of Proposition 1 The transformation $\hat{\mathbf{G}}^\top \mathbf{G}^{-\top}$, where μ, ν , and λ have been estimated via the LDR constraints is GBR-consistent.

Proof When we estimate the unit normals from the pseudo-normals we compute

$$\tilde{\mathbf{N}} = \hat{\mathbf{G}}^\top \hat{\mathbf{N}} / \|\hat{\mathbf{G}}^\top \hat{\mathbf{N}}\| = \hat{\mathbf{G}}^\top \mathbf{G}^{-\top} \mathbf{N} / \|\hat{\mathbf{G}}^\top \mathbf{G}^{-\top} \mathbf{N}\| \tag{66}$$

Hence, if $\hat{\mathbf{G}}^\top \mathbf{G}^{-\top}$ is GBR consistent, then the recovered normals $\tilde{\mathbf{N}}$ do not depend on the GBR parameters. By computing $\hat{\mathbf{G}}^\top \mathbf{G}^{-\top}$ we obtain

$$\hat{\mathbf{G}}^\top \mathbf{G}^{-\top} = \begin{bmatrix} 1 & 0 & \frac{\hat{\mu} - \mu}{\lambda} \\ 0 & 1 & \frac{\hat{\nu} - \nu}{\lambda} \\ 0 & 0 & \frac{\hat{\lambda}}{\lambda} \end{bmatrix} \tag{67}$$

Thus, we need to show that $\frac{\hat{\mu}-\mu}{\lambda}$, $\frac{\hat{\nu}-\nu}{\lambda}$ and $\frac{\hat{\lambda}}{\lambda}$ do not depend on the GBR parameters μ , ν , and λ , when $\hat{\mu}$, $\hat{\nu}$, and $\hat{\lambda}$ are computed from the LDR constraints. By direct substitution and by using $\hat{\mathbf{L}} = \mathbf{GL}$ and $\hat{\mathbf{N}} = \mathbf{G}^{-1}\mathbf{N}$ we obtain the following equalities, which depend only on the data and α ,

$$\eta_1 = \frac{\hat{\mu} - \mu}{\lambda} = \alpha \frac{l_{1,k}I_{p,k}}{n_{3,p}(l_{1,k}^2 + l_{2,k}^2)} - \frac{n_{1,p}}{n_{3,p}} \quad (68)$$

$$\eta_2 = \frac{\hat{\nu} - \nu}{\lambda} = \alpha \frac{l_{2,k}I_{p,k}}{n_{3,p}(l_{1,k}^2 + l_{2,k}^2)} - \frac{n_{2,p}}{n_{3,p}} \quad (69)$$

$$\eta_3 = \frac{\hat{\lambda}}{\lambda} = \sqrt{\alpha(1 - \alpha)} \frac{I_{p,k}}{|n_{3,p}|(l_{1,k}^2 + l_{2,k}^2)} \quad (70)$$

To conclude the proof we need to show that α does not depend on the GBR parameters. To this purpose, consider the intersection between the segments generated by two LDR maxima as in Eq. (61). Then, by direct substitution we obtain

$$\begin{aligned} & \begin{bmatrix} -\frac{n_{1,pa}\lambda}{n_{3,pa}} + \mu \\ -\frac{n_{2,pa}\lambda}{n_{3,pa}} + \nu \end{bmatrix} + \alpha \frac{\lambda I_{pa,ka}}{|n_{3,pa}|(l_{1,ka}^2 + l_{2,ka}^2)} \begin{bmatrix} l_{1,ka} \\ l_{2,ka} \end{bmatrix} \\ &= \begin{bmatrix} -\frac{n_{1,pb}\lambda}{n_{3,pb}} + \mu \\ -\frac{n_{2,pb}\lambda}{n_{3,pb}} + \nu \end{bmatrix} + \beta \frac{\lambda I_{pb,kb}}{|n_{3,pb}|(l_{1,kb}^2 + l_{2,kb}^2)} \begin{bmatrix} l_{1,kb} \\ l_{2,kb} \end{bmatrix} \end{aligned} \quad (71)$$

this is equivalent to

$$\begin{aligned} & \begin{bmatrix} -\frac{n_{1,pa}}{n_{3,pa}} \\ -\frac{n_{2,pa}}{n_{3,pa}} \end{bmatrix} + \alpha \frac{I_{pa,ka}}{|n_{3,pa}|(l_{1,ka}^2 + l_{2,ka}^2)} \begin{bmatrix} l_{1,ka} \\ l_{2,ka} \end{bmatrix} \\ &= \begin{bmatrix} -\frac{n_{1,pb}}{n_{3,pb}} \\ -\frac{n_{2,pb}}{n_{3,pb}} \end{bmatrix} + \beta \frac{I_{pb,kb}}{|n_{3,pb}|(l_{1,kb}^2 + l_{2,kb}^2)} \begin{bmatrix} l_{1,kb} \\ l_{2,kb} \end{bmatrix} \end{aligned} \quad (72)$$

which is independent of the GBR parameters. Therefore, α and β are also independent of the GBR parameters. \square

References

Agrawal, A., & Raskar, R. (2006). What is the range of surface reconstructions from a gradient field. In *ECCV* (pp. 578–591). Berlin:Springer.

Alldrin, N., & Kriegman, D. (2007). Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. *International Journal of Computer Vision*, 21, 1–8.

Alldrin, N., Mallick, S., & Kriegman, D. (2007). Resolving the generalized bas-relief ambiguity by entropy minimization. In *Computer Vision and Pattern Recognition* (pp. 1–7, 17–22).

Basri, R., & Jacobs, D. (2001). Photometric stereo with general, unknown lighting. In *IEEE Conference on Computer Vision and Pattern Recognition* (pp. 374–381).

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

Chandraker, M., Bai, J., & Ramamoorthi, R. (2011). A theory of photometric reconstruction for unknown isotropic reflectances. In *IEEE Conference on Computer Vision and Pattern Recognition*.

Chandraker, M. K., Kahl, F., & Kriegman, D. J. (2005). Reflections on the generalized bas-relief ambiguity. In *Conference on Computer Vision and, Pattern Recognition* (pp. 788–795).

Drbohlav, O., & Chantler, M. (2005). Can two specular pixels calibrate photometric stereo?. In *Proceedings of International Conference on Computer Vision* (pp. 1850–1857).

Drbohlav, O., & Sara, R. (2002). Specularities reduce ambiguity of uncalibrated photometric stereo. In *European Conference on Computer Vision* (pp. 46–62).

Favaro, P., & Papadhimetri, T. (2012). A closed-form solution to uncalibrated photometric stereo via diffuse maxima. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 821–828).

Georghades, A. (2003). Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *Proceedings of International Conference on Computer Vision* (pp. 816–823).

Georghades, A. S., Belhumeur, P. N., & Kriegman, D. J. (2001). From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23, 643–660.

Hayakawa, H. (1994). Photometric stereo under a light-source with arbitrary motion. *Journal of the Optical Society of America*, 11(11), 3079–3089.

Hertzmann, A., & Seitz, S. (2005). Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *Pattern Analysis and Machine Intelligence*, 27(8), 1254–1264.

Koppal, S. J., & Narasimhan, S. G. (2006). Clustering appearance for scene analysis. *IEEE Conference on Computer Vision and Pattern Recognition*, 2, 1323–1330.

Lagger, P., & Fua, P. (2008). Retrieving multiple light sources in the presence of specular reflections and texture. *Computer Vision and Image Understanding*, 111, 207–218.

Lin, Z., Chen, M., & Ma, Y. (2009). *The augmented lagrange multiplier method for exact recovery of corrupted low-rank matrices*. UIUC Technical, Report UILU-ENG-09-2215.

Okabe, T., Sato, I., & Sato, Y. (2009). Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions. In *Proceedings of International Conference on Computer Vision* (pp. 1693–1700).

Oren, M., & Nayar, S. K. (1997). A theory of specular surface geometry. *International Journal of Computer Vision*, 24, 105–124.

Sato, I., Okabe, T., Yu, Q., & Sato, Y. (2007). Shape reconstruction based on similarity in radiance changes under varying illumination. In *Proceedings of International Conference on Computer Vision* (pp. 1–8, 14–21).

Shi, B., Matsushita, Y., Wei, Y., Xu, C., & Tan, P. (2010). Self-calibrating photometric stereo. In *IEEE Conference on Computer Vision and, Pattern Recognition* (pp. 1118–1125).

Sunkavalli, H. P. K., & Zickler, T. (2010). Visibility subspaces: Uncalibrated photometric stereo with shadows. In *European Conference on Computer Vision, Part II* (pp. 251–264).

Tan, P., Mallick, S., Quan, L., Kriegman, D., & Zickler, T. (2007). Isotropy, reciprocity and the generalized bas-relief ambiguity. *Computer Vision and Pattern Recognition Conference*, 1(8), 17–22.

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

Wright, J., Ganesh, A., Rao, S., Peng, Y., & Ma, Y. (2009). Robust principal component analysis: Exact recovery of corrupted low-rank matrices by convex optimization. In *Proceedings of Neural Information Processing Systems (NIPS)*.

Wu, L., Ganesh, A., Shi, B., Matsushita, Y., Wang, Y., & Ma, Y. (2011). Robust photometric stereo via low-rank matrix completion

- and recovery. In *Asian Conference on Computer Vision* (pp. 703–717).
- Wu, T., & Tang, C. (2005). Dense photometric stereo using a mirror sphere and graph cut. *Computer Vision and, Pattern Recognition, 1*, 140–147.
- Yuille, A., & Snow, D. (1997). Shape and albedo from multiple images using integrability. In *Computer Vision and, Pattern Recognition* (pp. 158–164).
- Zickler, T., Belhumeur, P., & Kriegman, D. (2002). Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision, 49*(2–3), 215–227.