

On the well-posedness of uncalibrated photometric stereo under general lighting

Mohammed Brahim, Yvain Quéau, Bjoern Haefner and Daniel Cremers

Abstract Uncalibrated photometric stereo aims at estimating the 3D-shape of a surface, given a set of images captured from the same viewing angle, but under unknown, varying illumination. While the theoretical foundations of this inverse problem under directional lighting are well-established, there is a lack of mathematical evidence for the uniqueness of a solution under general lighting. On the other hand, stable and accurate heuristical solutions of uncalibrated photometric stereo under such general lighting have recently been proposed. The quality of the results demonstrated therein tends to indicate that the problem may actually be well-posed, but this still has to be established. The present paper addresses this theoretical issue, considering first-order spherical harmonics approximation of general lighting. Two important theoretical results are established. First, the orthographic integrability constraint ensures uniqueness of a solution up to a global concave-convex ambiguity, which had already been conjectured, yet not proven. Second, the perspective integrability constraint makes the problem well-posed, which generalizes a previous result limited to directional lighting. Eventually, a closed-form expression for the unique least-squares solution of the problem under perspective projection is provided, allowing numerical simulations on synthetic data to empirically validate our findings.

Mohammed Brahim
Technical University of Munich, e-mail: brahimi@in.tum.de

Yvain Quéau
Normandie Univ., UNICAEN, ENSICAEN, CNRS, e-mail: yvain.queau@ensicaen.fr

Bjoern Haefner
Technical University of Munich, e-mail: bjoern.haefner@tum.de

Daniel Cremers
Technical University of Munich, e-mail: cremers@tum.de

1 Introduction

Among the many photographic techniques which can be considered for the 3D-reconstruction of a still surface, photometric stereo [43] is often considered as a first choice when it comes to the recovery of very thin geometric structures. Nevertheless, the classic formulation of photometric stereo requires illumination to be highly controlled: each image must be captured under a single collimated light source at infinity, and the direction and relative intensity of each source must be calibrated beforehand. In practice, this restricts possible applications of the technique to laboratory setups where collimation of light can be ensured and a (possibly tedious) calibration procedure can be carried out.

Considering uncalibrated general lighting i.e., lighting induced by unknown, non-collimated sources and in the presence of ambient lighting, would both drastically simplify the 3D-scanning process for non-experts, and allow to bring photometric stereo outside of the lab [37]. The theoretical foundations of the problem under uncalibrated directional lighting are well-understood: the solution can be recovered only up to a linear transformation [12]. When integrability is enforced, this linear ambiguity reduces to the generalized bas-relief one under orthographic projection [44], and vanishes under perspective projection [25]. This work rather focuses on uncalibrated general lighting represented using first-order spherical harmonics [2, 35], in which case the solution can be recovered only up to a Lorentz transformation [3] and it has been conjectured - but not proven yet, that additional constraints such as integrability may reduce this ambiguity. One reason for thinking that this conjecture might hold is that stable numerical implementations of uncalibrated photometric stereo under general illumination have been proposed recently, under both orthographic [1, 23] and perspective [11] projections. Despite having no theoretical foundation, the results provided therein do not exhibit a significant low-frequency bias which would reveal an underlying ambiguity: empirically, the problem seems well-posed.

The objective of this paper is thus to establish the uniqueness of a solution to the problem of uncalibrated photometric stereo under general illumination, represented by first-order spherical harmonics. After discussing the classic case of directional lighting in Section 2, we characterize in Section 3 the ambiguities arising in uncalibrated photometric stereo under first-order spherical harmonics lighting. Then, we show in Section 4 that imposing integrability of the sought normal field resolves such ambiguities. In the orthographic case, only a global concave-convex ambiguity remains, hence the ambiguity is characterized by a single binary degree of freedom (in the directional case, there are three real degrees of freedom characterizing the generalized bas-relief ambiguity [4]). Under perspective projection, the problem becomes completely well-posed, which generalizes the result of [25] to more general lighting. In this case the solution can even be determined in closed-form, as shown in Section 5. Section 6 eventually recalls our findings, and suggests future research directions.

2 Preliminaries: photometric stereo under directional lighting

Assuming a Lambertian surface is observed from a still camera under $m \geq 1$ different directional lighting indexed by $i \in \{1, \dots, m\}$, the graylevel in the i -th image can be modeled as follows:

$$I^i(x) = \rho(x) \mathbf{n}(x)^\top \mathbf{l}^i, \quad \forall x \in \Omega, \quad (1)$$

where $\Omega \subset \mathbb{R}^2$ is the reconstruction domain (projection of the 3D-surface onto the image plane), $\rho(x) > 0$ is the albedo at the surface point conjugate to pixel x , $\mathbf{n}(x) \in \mathbb{S}^2 \subset \mathbb{R}^3$ is the unit-length outward normal at this point, and $\mathbf{l}^i \in \mathbb{R}^3$ is a vector oriented towards the light source whose norm represents the relative intensity of the source. Photometric stereo consists, given a set of graylevel observations $I^i, i \in \{1, \dots, m\}$, in estimating the shape (represented by the surface normal \mathbf{n}) and the reflectance (represented by the surface albedo ρ). Depending whether the lighting vectors \mathbf{l}^i are known or not, the problem is called calibrated or uncalibrated.

2.1 Calibrated photometric stereo under directional lighting

Woodham showed in the late 70s [42] that $m \geq 3$ images captured under non-coplanar, known lighting vectors were sufficient to solve this problem. Indeed, defining for every $x \in \Omega$ the following observation vector $\mathbf{i}(x) \in \mathbb{R}^m$, lighting matrix $\mathbf{L} \in \mathbb{R}^{m \times 3}$ and surface vector $\mathbf{m}(x) \in \mathbb{R}^3$:

$$\mathbf{i}(x) = \begin{bmatrix} I^1(x) \\ \vdots \\ I^m(x) \end{bmatrix}, \quad \mathbf{L} = \begin{bmatrix} \mathbf{l}^{1\top} \\ \vdots \\ \mathbf{l}^{m\top} \end{bmatrix}, \quad \mathbf{m}(x) = \rho(x) \mathbf{n}(x), \quad (2)$$

the set of equations (1) can be rewritten as a linear system in $\mathbf{m}(x)$:

$$\mathbf{i}(x) = \mathbf{L} \mathbf{m}(x), \quad \forall x \in \Omega. \quad (3)$$

Provided that \mathbf{L} is of rank three, (3) admits a unique least-squares solution in $\mathbf{m}(x)$, from which the normal and albedo can be extracted according to:

$$\rho(x) = |\mathbf{m}(x)|, \quad \mathbf{n}(x) = \frac{\mathbf{m}(x)}{|\mathbf{m}(x)|}. \quad (4)$$

Such a simple least-squares approach may be replaced by robust variational or learning-based strategies to ensure robustness [13, 17, 33, 34]. There also exist numerical solutions for handling non-Lambertian reflectance models [7, 15, 21, 40], non-distant light sources [18, 22, 29, 32], or the ill-posed cases where $m = 2$ [16, 20, 24, 31] or $m = 1$ [5, 8, 39, 45]. Such issues are not addressed in the present paper which rather focuses on the theoretical foundations of uncalibrated photometric stereo.

2.2 Uncalibrated photometric stereo under directional lighting

The previous strategy relies on the knowledge of the lighting matrix \mathbf{L} , and it is not straightforward to extend it to unknown lighting. Let us illustrate this in the discrete setting, denoting the pixels as x_j , $j \in \{1, \dots, n\}$ where $n = |\Omega|$ is the number of pixels, and stack all the observations in an observation matrix $\mathbf{I} \in \mathbb{R}^{m \times n}$ and all the surface vectors in a surface matrix $\mathbf{M} \in \mathbb{R}^{3 \times n}$:

$$\mathbf{I} = [\mathbf{i}(x_1), \dots, \mathbf{i}(x_n)], \quad \mathbf{M} = [\mathbf{m}(x_1), \dots, \mathbf{m}(x_n)]. \quad (5)$$

Now, the set of n linear systems (3) can be represented compactly as:

$$\mathbf{I} = \mathbf{L}\mathbf{M} \quad (6)$$

where both the lighting matrix \mathbf{L} and the surface matrix \mathbf{M} are unknown. Since we know that \mathbf{L} should be of rank three, a joint least-squares solution in (\mathbf{L}, \mathbf{M}) can be computed using truncated singular value decomposition [12]. Nevertheless, such a solution is not unique, since given a possible solution (\mathbf{L}, \mathbf{M}) , any couple $(\mathbf{L}\mathbf{A}^{-1}, \mathbf{A}\mathbf{M})$ with $\mathbf{A} \in GL(3, \mathbb{R})$ is another solution:

$$\mathbf{I} = \mathbf{L}\mathbf{M} = (\mathbf{L}\mathbf{A}^{-1})(\mathbf{A}\mathbf{M}), \quad \forall \mathbf{A} \in GL(3, \mathbb{R}), \quad (7)$$

or equivalently, in the continuous setting:

$$\mathbf{i}(x) = \mathbf{L}\mathbf{m}(x) = (\mathbf{L}\mathbf{A}^{-1})(\mathbf{A}\mathbf{m}(x)), \quad \forall (x, \mathbf{A}) \in \Omega \times GL(3, \mathbb{R}). \quad (8)$$

However, not any surface matrix \mathbf{M} (or \mathbf{m} -field, in the continuous setting) is acceptable as a solution. Indeed, this encodes the geometry of the surface, through its normals. Assuming that the surface is regular, its normals should satisfy the so-called integrability (or zero-curl) constraint. Let us assume orthographic projection and denote $\mathbf{n}(x) := [n_1(x), n_2(x), n_3(x)]^T$ the surface normal at 3D-point conjugate to pixel x . Let us further represent the surface as a Monge patch $(x, z(x))$ where $z : \mathbb{R}^2 \rightarrow \mathbb{R}^+$ is the depth map. Let us assume this map z is twice differentiable, and let $\nabla z(x) = [z_u(x), z_v(x)]^T \in \mathbb{R}^2$ be its gradient in some orthonormal basis (u, v) of the image plane. The integrability constraint is essentially a particular form of Schwarz' theorem, which implies that

$$z_{uv}(x) = z_{vu}(x), \quad \forall x \in \Omega. \quad (9)$$

The surface normal being defined as $\mathbf{n}(x) = \frac{[z_u(x), z_v(x), -1]^T}{\sqrt{z_u(x)^2 + z_v(x)^2 + 1}}$, Eq. (9) can be

rewritten as $\begin{pmatrix} n_1 \\ n_3 \end{pmatrix}_v(x) = \begin{pmatrix} n_2 \\ n_3 \end{pmatrix}_u(x)$, $\forall x \in \Omega$. Since $\mathbf{m}(x) = \rho(x)\mathbf{n}(x)$, $\forall x \in \Omega$, this is equivalent to:

$$\begin{pmatrix} m_1 \\ m_3 \end{pmatrix}_v(x) = \begin{pmatrix} m_2 \\ m_3 \end{pmatrix}_u(x), \quad \forall x \in \Omega. \quad (10)$$

Now, let us assume that one has found an \mathbf{m} -field solution of the lhs of (8), which further satisfies the integrability constraint (10) (in the discrete setting, this can be achieved using matrix factorization [44] or convex optimization techniques [36]). It can be shown that not all transformations \mathbf{A} in the rhs of (8) preserve this constraint. Indeed, the only ones which are acceptable are those of the generalized bas-relief (GBR) group. Such matrices define a subgroup of $GL(3, \mathbb{R})$ under the matrix product, and have the following form [4]:

$$\mathbf{G} = \begin{pmatrix} \lambda & 0 & -\mu \\ 0 & \lambda & -\nu \\ 0 & 0 & 1 \end{pmatrix}, \quad \mathbf{G}^{-1} = \begin{pmatrix} 1 & 0 & \mu/\lambda \\ 0 & 1 & \nu/\lambda \\ 0 & 0 & 1 \end{pmatrix}, \quad (\lambda, \mu, \nu) \in \mathbb{R}^3 \text{ and } \lambda \neq 0. \quad (11)$$

The three parameters μ, ν and λ characterize the GBR ambiguity inherent to uncalibrated photometric stereo under directional illumination and orthographic viewing. Intuitively, they can be understood as follows: any set of photometric stereo images can be reproduced by scaling the surface shape (this is the role of λ), adding a plane to it (this is the role of μ and ν), and moving the lighting vectors accordingly. If one is given a prior on the distribution of albedo values, on that of lighting vectors, or on the surface shape, then the three parameters can be estimated i.e., the ambiguity can be resolved. The literature on that particular topic is extremely dense, see e.g. [38] for an overview, [6] for a modern numerical solution based on deep learning, and [27] for an application to RGB-D sensing. As we shall prove later in Section 4.1, in the case of non-directional lighting represented using first-order spherical harmonics, the ambiguity is much simpler since it comes down to a global concave/convex one.

To terminate this discussion on uncalibrated photometric stereo under directional lighting, let us mention another important result. The previous discussion assumed that the camera is orthographic. Instead, let us now assume perspective projection, denote $x = (u, v)$ the pixel coordinates with respect to the principal point and $f > 0$ the focal length. The surface is now represented as the set of 3D-points $z(x) [u/f, v/f, 1]^T$, and the surface normal is now defined as (see, e.g., [30]):

$$\mathbf{n}(x) = \frac{1}{\sqrt{f^2 |\nabla z(x)|^2 + (-z(x) - [u, v]^T \nabla z(x))^2}} \begin{bmatrix} f \nabla z(x) \\ -z(x) - [u, v]^T \nabla z(x) \end{bmatrix}. \quad (12)$$

Using this new expression, writing Schwarz' theorem (8) yields an integrability constraint which slightly differs from (10) (see Corollary 2 in Appendix 3). Yet, this slight difference is of major importance, because the set of linear transformations which preserve this condition is restricted to the identity matrix. This means, under perspective projection and directional lighting the uncalibrated photometric stereo problem is well-posed. This result was established in [25]. As we shall prove later in Section 4.2, such a result can actually be extended to more general lighting represented using first-order spherical harmonics. Let us now elaborate on such a modeling of general lighting, and characterize the ambiguities therein.

3 Characterizing the ambiguities in uncalibrated photometric stereo under general lighting

The image formation model (1) is a simplified model, corresponding to the presence of a single light source located at infinity. However, this assumption is difficult to ensure in real-world experiments, and it would be more convenient to have at hand an image formation model accounting for general lighting (to handle multiple light sources, ambient lighting, etc.).

3.1 Spherical harmonics approximation of general lighting

The most general image formation model for Lambertian surfaces would integrate the incident lighting received from all directions $\mathbf{u}_l \in \mathbb{S}^2$:

$$I^i(x) = \rho(x) \int_{\mathbb{S}^2} s^i(x, \mathbf{u}_l) k(\mathbf{n}(x), \mathbf{u}_l) d\mathbf{u}_l, \quad \forall x \in \Omega, \quad (13)$$

where we denote $s^i(x, \mathbf{u}_l) \in \mathbb{R}$ the intensity of the light source in direction $\mathbf{u}_l \in \mathbb{S}^2$ at the surface point conjugate to pixel x in the i -th image, and $k(\mathbf{n}(x), \mathbf{u}_l) = \max\{\mathbf{n}(x)^\top \mathbf{u}_l, 0\}$ is the irradiance at this point (the max operator encodes self-shadows).

Assuming a single light source illuminates the scene in the i -th image, and neglecting self-shadows, then Eq. (13) obviously comes down to the simplified model (1). However, there exist other simplifications of the integral model (13), which allow to handle more general illumination. Namely, the spherical harmonics approximation which were introduced simultaneously in [2] and [35]. In the present work we focus on first-order spherical harmonics approximation, which is known to capture approximately 87% of general lighting [10]. Using this approximation, (13) simplifies to (see the aforementioned papers for technical details):

$$I^i(x) = \rho(x) \begin{bmatrix} 1 \\ \mathbf{n}(x) \end{bmatrix}^\top \mathbf{l}^i, \quad \forall x \in \Omega, \quad (14)$$

with $\mathbf{l}^i \in \mathbb{R}^4$ a vector representing the general illumination in the i -th image. Denoting $\mathbf{L} = [\mathbf{l}^1, \dots, \mathbf{l}^m]^\top \in \mathbb{R}^{m \times 4}$ the general lighting matrix, System (14) can be rewritten in the same form as the directional one (6):

$$\mathbf{i}(x) = \mathbf{L}\mathbf{m}(x), \quad \forall x \in \Omega, \quad (15)$$

$$\text{with } \mathbf{m}(x) = \rho(x) \begin{pmatrix} 1 \\ \mathbf{n}(x) \end{pmatrix}. \quad (16)$$

3.2 Uncalibrated photometric stereo under first-order spherical harmonics lighting

Uncalibrated photometric stereo under first-order spherical harmonics lighting comes down to solving the set of linear systems (15) in terms of both the general lighting matrix \mathbf{L} and the \mathbf{m} -field (which encodes albedo and surface normals). In the directional case discussed previously, this was possible only up to an invertible linear transformation, as shown by (8). Despite appearing more complicated at first glance, the case of first-order spherical harmonics is actually slightly more favorable than the directional one: not all such linear transformations are acceptable, because they have to preserve the particular form of the \mathbf{m} -field, given in Eq. (16). That is to say, given one \mathbf{m} -field solution and another one \mathbf{Am} obtained by applying an invertible linear transformation $\mathbf{A} \in GL(4, \mathbb{R})$, the entries c_1, c_2, c_3, c_4 of \mathbf{Am} should respect the constraint $c_1^2 = c_2^2 + c_3^2 + c_4^2$ over Ω (cf. Eq. (16), remembering that each surface normal has unit length).

As discussed in [3], this means that ambiguities in uncalibrated photometric stereo under first-order spherical harmonics are characterized as follows:

$$\mathbf{i}(x) = \mathbf{Lm}(x) = (\mathbf{LA}^{-1})(\mathbf{Am}(x)), \quad \forall (x, \mathbf{A}) \in \Omega \times L_s, \quad (17)$$

where L_s is the space of scaled Lorentz transformations defined by

$$L_s = \{s\mathbf{A} \mid s \in \mathbb{R} \setminus \{0\} \text{ and } \mathbf{A} \in L\}, \quad (18)$$

with L the Lorentz group [28] arising in Einstein's theory of special relativity [9]:

$$L = \{\mathbf{A} \in GL(4, \mathbb{R}) \mid \forall \mathbf{x} \in \mathbb{R}^4, l(\mathbf{Ax}) = l(\mathbf{x})\}, \quad (19)$$

$$\text{with } l : (t, x, y, z) \mapsto x^2 + y^2 + z^2 - t^2. \quad (20)$$

In spite of the presence of the scaled Lorentz ambiguity in Eq. (17), several heuristical approaches to solve uncalibrated photometric stereo under general lighting have been proposed lately. Let us mention the approaches based on hemispherical embedding [1] and on equivalent directional lighting [23], which both deal with the case of orthographic projection, and the variational approach in [11] for that of perspective projection. The empirically observed stability of such implementations tends to indicate that the problem might be better-posed than it seems, as already conjectured in [3]. In order to prove this conjecture, we will show in Section 4 that not all scaled Lorentz transformations preserve the integrability of surface normals. To this end, we need to characterize algebraically a scaled Lorentz transformation.

3.3 Characterization of the scaled Lorentz transformation

We propose to characterize any ambiguity matrix $\mathbf{A} \in L_s$ in (17) by means of a scale factor $s \neq 0$ (one degree of freedom), a vector inside the unit \mathbb{R}^3 -ball $\mathbf{v} \in B(\mathbf{0}, 1)$ (three degrees of freedom, where $B(\mathbf{0}, 1) = \{\mathbf{x} \in \mathbb{R}^3, |\mathbf{x}| < 1\}$) and a 3D-rotation matrix $\mathbf{O} \in SO(3, \mathbb{R})$ (three degrees of freedom, hence a total of seven). More explicitly, any scaled Lorentz transformation can be characterized algebraically as follows:

Theorem 1 $\forall \mathbf{A} \in L_s, \exists! (s, \mathbf{v}, \mathbf{O}) \in \mathbb{R} \setminus \{0\} \times B(\mathbf{0}, 1) \times SO(3, \mathbb{R})$,

$$\mathbf{A} = s \begin{pmatrix} \epsilon_1(\mathbf{A}) \gamma & \epsilon_1(\mathbf{A}) \gamma \mathbf{v}^\top \mathbf{O} \\ \epsilon_2(\mathbf{A}) \gamma \mathbf{v} & \epsilon_2(\mathbf{A}) (\mathbf{I}_3 + \frac{\gamma^2}{1+\gamma} \mathbf{v} \mathbf{v}^\top) \mathbf{O} \end{pmatrix}, \quad (21)$$

with

$$\gamma = \frac{1}{\sqrt{1 - |\mathbf{v}|^2}}, \quad (22)$$

$$\epsilon_1(\mathbf{A}) = \begin{cases} 1 & \text{if } P_o(\mathbf{A}), \\ -1 & \text{else,} \end{cases} \quad (23)$$

$$\epsilon_2(\mathbf{A}) = \begin{cases} -1 & \text{if } (P_p(\mathbf{A}) \wedge \overline{P_o(\mathbf{A})}) \vee (\overline{P_p(\mathbf{A})} \wedge P_o(\mathbf{A})), \\ 1 & \text{else,} \end{cases} \quad (24)$$

and $P_p(\mathbf{A})$ stands for “ \mathbf{A} is proper”, $P_o(\mathbf{A})$ for “ \mathbf{A} is orthochronous”,

where we recall that a Lorentz matrix \mathbf{A} is “proper” iff it preserves the orientation of the Minkowski spacetime, and it is “orthochronous” iff it preserves the direction of the time, i.e.:

$$\mathbf{A} \in L \text{ is proper} \iff \det(\mathbf{A}) > 0, \quad (25)$$

$$\begin{aligned} \mathbf{A} \in L \text{ is orthochronous} &\iff \forall \mathbf{x} = [t, x, y, z]^\top \in \mathbb{R}^4, \text{sign}(t) = \text{sign}(t'), \\ &\text{where } \mathbf{A}\mathbf{x} = [t', x', y', z']^\top. \end{aligned} \quad (26)$$

The opposites are improper and non-orthochronous, and we note L_o^p, L_o^i, L_n^p and L_n^i the sets of Lorentz transformations which are respectively proper and orthochronous, improper and orthochronous, proper and non-orthochronous, and improper and non-orthochronous. The Lorentz group is the union of all these spaces, i.e. $L = L_o^p \cup L_o^i \cup L_n^p \cup L_n^i$.

Using Theorem 1 (whose proof can be found in Appendix 1) to characterize the underlying ambiguity of uncalibrated photometric stereo under general lighting, we are ready to prove that imposing integrability disambiguates the problem.

4 Integrability disambiguates uncalibrated photometric stereo under general lighting

As we have seen in the previous section, uncalibrated photometric stereo under general lighting is ill-posed without further constraints, since it is prone to a scaled Lorentz ambiguity, cf. Eq. (17). In this section, we prove that not all scaled Lorentz transformations preserve the integrability of the underlying normal field: in the case of orthographic projection, the only acceptable one is the one which globally exchanges concavities and convexities, while it is the identity in the perspective case. That is to say, the orthographic case suffers only from a global concave/convex ambiguity, while the perspective one is well-posed.

4.1 Orthographic case

First, let us prove that under orthographic projection and first-order spherical harmonics lighting, there are only two integrable solutions to uncalibrated photometric stereo, and they differ by a global concave/convex transformation. To this end, we consider the genuine solution $\mathbf{m}(x)$ of (15) corresponding to a normal field $\mathbf{n}(x)$ and albedo map $\rho(x)$, and another possible solution $\mathbf{m}^*(x) = \mathbf{A}\mathbf{m}(x)$, $\mathbf{A} \in L_s$, with $(\rho^*(x), \mathbf{n}^*(x))$ the corresponding albedo map and surface normals. If the pictured surface is regular (no kink or depth discontinuity), then the genuine normal field \mathbf{n} is integrable by construction. We establish in this subsection that if the other candidate normal field \mathbf{n}^* is assumed integrable as well, then both the genuine and the alternative solutions differ according to:

$$\begin{cases} \rho^*(x) = \alpha \rho(x) \\ n_1^*(x) = \lambda n_1(x) \\ n_2^*(x) = \lambda n_2(x) \\ n_3^*(x) = n_3(x) \end{cases}, \quad \forall x \in \Omega, \quad (27)$$

where $\alpha > 0$ scales all albedo values simultaneously (which is non-important, since the estimated albedo values can be scaled back as a post-processing step), and $\lambda \in \{-1, 1\}$ simultaneously inverts the sign of the first two components of all normal vectors i.e., it turns concavities into convexities and vice-versa.

More formally, this result can be stated as the following theorem:

Theorem 2 *Under orthographic projection, the only scaled Lorentz transformation $\mathbf{A} \in L_s$ which preserves integrability of normals is the following one, where $\alpha > 0$ and $\lambda \in \{-1, 1\}$:*

$$\mathbf{A} = \alpha \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & \lambda & 0 & 0 \\ 0 & 0 & \lambda & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}. \quad (28)$$

Proof Let $\mathbf{m} : \Omega \rightarrow \mathbb{R}^4$ a field with the form of Equation (16), and let ρ and \mathbf{n} the corresponding albedo map and normal field, assumed integrable. The normal field \mathbf{n} being integrable, $p = -\frac{n_1}{n_3}$ and $q = -\frac{n_2}{n_3}$ satisfy the integrability constraint $p_v = q_u$ over Ω . Denoting by (c_1, c_2, c_3, c_4) the four components of the field \mathbf{m} , and using the expression (16) of \mathbf{m} , this implies:

$$\begin{aligned} & \left(\frac{c_2}{c_4} \right)_v = \left(\frac{c_3}{c_4} \right)_u \quad \text{over } \Omega, \\ \iff & \frac{c_{2v}c_4 - c_2c_{4v}}{c_4^2} = \frac{c_{3u}c_4 - c_3c_{4u}}{c_4^2} \quad \text{over } \Omega, \\ \iff & (c_{2v} - c_{3u})c_4 + c_{4u}c_3 - c_{4v}c_2 = 0 \quad \text{over } \Omega. \end{aligned} \quad (29)$$

Let $\mathbf{m}^* = \mathbf{A}\mathbf{m}$, with \mathbf{A} a scaled Lorentz transformation having the form given by Theorem 1, and let ρ^* and \mathbf{n}^* the corresponding albedo map and normal field, assumed integrable. The same rationale as above on the alternative normal field \mathbf{n}^* yields:

$$(c_{2v}^* - c_{3u}^*)c_4^* + c_{4u}^*c_3^* - c_{4v}^*c_2^* = 0 \quad \text{over } \Omega. \quad (30)$$

Since $\mathbf{m}^* = \mathbf{A}\mathbf{m}$, (30) writes as:

$$\begin{aligned} & (A_{21}c_{1v} + A_{22}c_{2v} + A_{23}c_{3v} + A_{24}c_{4v} - A_{31}c_{1u} - A_{32}c_{2u} \\ & \quad - A_{33}c_{3u} - A_{34}c_{4u})(A_{41}c_1 + A_{42}c_2 + A_{43}c_3 + A_{44}c_4) \\ & \quad + (A_{41}c_{1u} + A_{42}c_{2u} + A_{43}c_{3u} + A_{44}c_{4u})(A_{31}c_1 + A_{32}c_2 + A_{33}c_3 + A_{34}c_4) \\ & \quad - (A_{41}c_{1v} + A_{42}c_{2v} + A_{43}c_{3v} + A_{44}c_{4v})(A_{21}c_1 + A_{22}c_2 + A_{23}c_3 + A_{24}c_4) \\ & \quad = 0 \quad \text{over } \Omega. \end{aligned} \quad (31)$$

Let us introduce the following notation, $1 \leq i < j \leq 4$, and $k \in \{u, v\}$:

$$c_k^{i,j}(x) = c_j(x)c_{ik}(x) - c_i(x)c_{jk}(x), \quad \forall x \in \Omega, \quad (32)$$

and denote as follows the minors of size two of matrix \mathbf{A} :

$$A_{k,l}^{i,j} = A_{ij}A_{kl} - A_{kj}A_{il}, \quad 1 \leq i < k \leq 4, \quad 1 \leq j < l \leq 4. \quad (33)$$

Then, factoring (31) firstly by the coefficients A_{ij} and after by $c_u^{i,j}$ and $c_v^{i,j}$ for every $(i, j) \in \{1, 2, 3, 4\}$ with $i < j$, we get:

$$\begin{aligned} & c_v^{1,2}A_{4,2}^{2,1} + c_v^{1,3}A_{4,3}^{2,1} + c_v^{1,4}A_{4,4}^{2,1} + c_v^{2,3}A_{4,3}^{2,2} + c_v^{2,4}A_{4,4}^{2,2} + c_v^{3,4}A_{4,4}^{2,3} \\ & - c_u^{1,2}A_{4,2}^{3,1} - c_u^{1,3}A_{4,3}^{3,1} - c_u^{1,4}A_{4,4}^{3,1} - c_u^{2,3}A_{4,3}^{3,2} - c_u^{2,4}A_{4,4}^{3,2} - c_u^{3,4}A_{4,4}^{3,3} = 0 \quad \text{over } \Omega. \end{aligned} \quad (34)$$

In addition, (29) also writes as:

$$c_v^{2,4} = c_u^{3,4} \quad \text{over } \Omega. \quad (35)$$

Thus, substituting $c_v^{2,4}$ by $c_u^{3,4}$, Eq. (34) can be rewritten as:

$$\mathbf{i}_o(x)^\top \mathbf{a} = 0, \quad \forall x \in \Omega, \quad (36)$$

where $\mathbf{i}_o(x) \in \mathbb{R}^{11}$ is the ‘‘orthographic integrability vector’’ containing factors $c_u^{i,j}(x)$ and $c_v^{i,j}(x)$, and $\mathbf{a} \in \mathbb{R}^{11}$ contain the minors $A_{k,l}^{i,j}$ of \mathbf{A} appearing in (34).

Let us assume that there exist at least 11 points $x \in \Omega$ such that the matrix formed by concatenating row-wise the vectors $\mathbf{i}_o(x)^\top$ is full-rank (in practice, this property is satisfied as long as the surface is not trivial). Then, the only solution to (36) is $\mathbf{a} = \mathbf{0}$, which is equivalent to the following equations:

$$\begin{cases} A_{4,3}^{3,2} = A_{4,4}^{3,2} = A_{4,4}^{2,3} = A_{4,3}^{2,2} = 0, \\ A_{4,4}^{3,3} = A_{4,4}^{2,2}, \\ A_{4,2}^{2,1} = A_{4,2}^{3,1} = A_{4,3}^{2,1} = A_{4,3}^{3,1} = A_{4,4}^{2,1} = A_{4,4}^{3,1} = 0. \end{cases} \quad (37)$$

According to Corollary 1 provided in Appendix 2, this implies that the submatrix of \mathbf{A} formed by its last three rows and columns is a scaled generalized bas-relief transformation, i.e.: $\exists! (\lambda, \mu, \nu, \beta) \in \mathbb{R}^4$ with $\lambda \neq 0, \beta \neq 0$, such that:

$$\mathbf{A} = \begin{pmatrix} A_{11} & A_{12} & A_{13} & A_{14} \\ A_{21} & \beta\lambda & 0 & -\beta\mu \\ A_{31} & 0 & \beta\lambda & -\beta\nu \\ A_{41} & 0 & 0 & \beta \end{pmatrix}. \quad (38)$$

By taking into account the last equation of System (37), we get :

$$\mathbf{A} = \begin{pmatrix} A_{11} & A_{12} & A_{13} & A_{14} \\ 0 & \beta\lambda & 0 & -\beta\mu \\ 0 & 0 & \beta\lambda & -\beta\nu \\ 0 & 0 & 0 & \beta \end{pmatrix}. \quad (39)$$

Identifying (39) with the expression in Theorem 1, $\mathbf{v} = \mathbf{0}$, $\gamma = 1$ and $s_{\mathbf{e}_2}(\mathbf{A})\mathbf{O} = \beta \begin{pmatrix} \lambda & 0 & -\mu \\ 0 & \lambda & -\nu \\ 0 & 0 & 1 \end{pmatrix}$. In addition, $\mathbf{O} \in SO(3, \mathbb{R})$, which implies $\mathbf{O}^\top \mathbf{O} = \mathbf{I}_3$. Thus, since $(\mathbf{e}_2(\mathbf{A}))^2 = 1$, we have: $(s_{\mathbf{e}_2}(\mathbf{A})\mathbf{O})^\top (s_{\mathbf{e}_2}(\mathbf{A})\mathbf{O}) = s^2 \mathbf{I}_3$. Equivalently:

$$\begin{aligned} & \beta \begin{pmatrix} \lambda & 0 & 0 \\ 0 & \lambda & 0 \\ -\mu & -\nu & 1 \end{pmatrix} \beta \begin{pmatrix} \lambda & 0 & -\mu \\ 0 & \lambda & -\nu \\ 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} s^2 & 0 & 0 \\ 0 & s^2 & 0 \\ 0 & 0 & s^2 \end{pmatrix}, \\ \iff & \beta^2 \begin{pmatrix} \lambda^2 & 0 & -\mu\lambda \\ 0 & \lambda^2 & -\nu\lambda \\ -\mu\lambda & -\nu\lambda & \mu^2 + \nu^2 + 1 \end{pmatrix} = \begin{pmatrix} s^2 & 0 & 0 \\ 0 & s^2 & 0 \\ 0 & 0 & s^2 \end{pmatrix}, \end{aligned} \quad (40)$$

which implies $\lambda^2 = 1, \mu = 0, \nu = 0, \beta^2 = s^2$.

Finally, $\det(s\epsilon_2(\mathbf{A})\mathbf{O}) = \beta\lambda^2\beta = \epsilon_2(\mathbf{A})s$, thus according to (39):

$$\mathbf{A} = s \begin{pmatrix} \epsilon_1(\mathbf{A}) & 0 & 0 & 0 \\ 0 & \epsilon_2(\mathbf{A})\lambda & 0 & 0 \\ 0 & 0 & \epsilon_2(\mathbf{A})\lambda & 0 \\ 0 & 0 & 0 & \epsilon_2(\mathbf{A}) \end{pmatrix}. \quad (41)$$

Plugging (41) into $\mathbf{m}^* = \mathbf{A}\mathbf{m}$, we obtain:

$$\begin{pmatrix} \rho^*(x) \\ \rho^*(x)n_1^*(x) \\ \rho^*(x)n_2^*(x) \\ \rho^*(x)n_3^*(x) \end{pmatrix} = s \begin{pmatrix} \epsilon_1(\mathbf{A})\rho(x) \\ \epsilon_2(\mathbf{A})\lambda\rho(x)n_1(x) \\ \epsilon_2(\mathbf{A})\lambda\rho(x)n_2(x) \\ \epsilon_2(\mathbf{A})\rho(x)n_3(x) \end{pmatrix}, \quad \forall x \in \Omega. \quad (42)$$

Now, knowing that albedos $\rho, \rho^* \geq 0$ (they represent the proportion of light which is reflected by the surface), and that the last component of normals $n_3, n_3^* \leq 0$ (the normals point toward the camera), Eq. (42) implies that $\epsilon_1(\mathbf{A})$ and $\epsilon_2(\mathbf{A})$ have exactly the same sign as s .

Two cases must be considered. If $s > 0$, then $\epsilon_1(\mathbf{A}) = \epsilon_2(\mathbf{A}) = 1$, and plugging these values into (41) we obtain the expression provided in Theorem 2. If $s < 0$, then $\epsilon_1(\mathbf{A}) = \epsilon_2(\mathbf{A}) = -1$, and we again get the expression provided in Theorem 2. \square

Using the expression of Theorem 2 to modify the genuine \mathbf{m} -field, and decomposing the latter field into albedo and normals according to (16), we observe that the transformation has the effect depicted in (27). That is to say, albedo values are globally scaled, while all concavities and convexities are simultaneously inverted. The global scale on the albedo should not be considered as a real issue, since such values are relative to the camera response function and the intensities of the light sources, and they can be manually scaled back in a post-processing step if needed. However, the residual global concave/convex ambiguity shows that uncalibrated photometric stereo under general lighting remains ill-posed in the orthographic case. Still, the ill-posedness is characterized by a single binary degree of freedom, which is to be compared with the three real degrees of freedom characterizing the GBR ambiguity arising in the case of directional lighting [4].

From a practical point of view, once an integrable normal field candidate has been found heuristically, using e.g. hemispherical embedding [1] or an equivalent directional lighting model [23], the residual ambiguity can easily be resolved manually. Such a manual final disambiguation is for instance advocated in [23]. As we shall see now, in the case of perspective projection the problem becomes even completely well-posed, which circumvents the need for any manual intervention.

4.2 Perspective case

Now we will prove that uncalibrated photometric stereo under first-order spherical harmonics lighting and perspective projection is well-posed. This means, imposing integrability restricts the admissible ambiguity matrices \mathbf{A} in (17) to the identity matrix (up to a factor scaling all albedo values without affecting the geometry):

Theorem 3 *Under perspective projection, the only scaled Lorentz transformation $\mathbf{A} \in L_s$ which preserves integrability of normals is the identity matrix, up to a scale factor $\alpha > 0$:*

$$\mathbf{A} = \alpha \mathbf{I}_4 \quad (43)$$

Proof Let $\mathbf{m} : \Omega \rightarrow \mathbb{R}^4$ a field with the form of Equation (16), whose normal map is integrable. Let $\mathbf{m}^* = \mathbf{A}\mathbf{m}$ another such field whose normal map is integrable, with $\mathbf{A} \in L_s$ a scaled Lorentz transformation having the form given by Theorem 1.

Denoting by (c_1, c_2, c_3, c_4) the four components of the field \mathbf{m} , and $(c_1^*, c_2^*, c_3^*, c_4^*)$ those of \mathbf{m}^* , the integrability constraint of the normal field associated to \mathbf{m}^* can be written as follows: (see Corollary 2 in Appendix 3)

$$u(c_u^*)_{u}^{2,3} + v(c_v^*)_{v}^{2,3} + f(c_v^*)_{v}^{2,4} - f(c_u^*)_{u}^{3,4} = 0 \text{ over } \Omega. \quad (44)$$

with the same notations as in (32).

As in the previous proof of Theorem 2, we substitute in the integrability constraint (44) the entries of $\mathbf{m}^* = \mathbf{A}\mathbf{m}$ with their expressions in terms of entries of \mathbf{A} and \mathbf{m} . Then, by factoring firstly by the coefficients A_{ij} and then by $c_u^{i,j}$ and $c_v^{i,j}$ for every $(i, j) \in \{1, 2, 3, 4\}$ with $i < j$, we get:

$$\begin{aligned} & \left(uc_u^{1,2} + vc_v^{1,2}\right) A_{3,2}^{2,1} + \left(uc_u^{1,3} + vc_v^{1,3}\right) A_{3,3}^{2,1} \\ & + \left(uc_u^{1,4} + vc_v^{1,4}\right) A_{3,4}^{2,1} + \left(uc_u^{2,3} + vc_v^{2,3}\right) A_{3,3}^{2,2} \\ & + \left(uc_u^{2,4} + vc_v^{2,4}\right) A_{3,4}^{2,2} + \left(uc_u^{3,4} + vc_v^{3,4}\right) A_{3,4}^{2,3} \\ & + f \left(c_v^{1,2} A_{4,2}^{2,1} + c_v^{1,3} A_{4,3}^{2,1} + c_v^{1,4} A_{4,4}^{2,1} + c_v^{2,3} A_{4,3}^{2,2} + c_v^{2,4} A_{4,4}^{2,2} + c_v^{3,4} A_{4,4}^{2,3} \right) \\ & - f \left(c_u^{1,2} A_{4,2}^{3,1} + c_u^{1,3} A_{4,3}^{3,1} + c_u^{1,4} A_{4,4}^{3,1} + c_u^{2,3} A_{4,3}^{3,2} + c_u^{2,4} A_{4,4}^{3,2} + c_u^{3,4} A_{4,4}^{3,3} \right) \\ & = 0 \text{ over } \Omega. \end{aligned} \quad (45)$$

By concatenating equations (45) for all pixels $x \in \Omega$, we get the following set of linear systems:

$$\mathbf{i}_p(x)^\top \mathbf{w} = 0, \quad \forall x \in \Omega, \quad (46)$$

where $\mathbf{w} \in \mathbb{R}^{18}$ contains all the minors $A_{k,l}^{i,j}$ of the ambiguity matrix \mathbf{A} in equation (45), and the ‘‘perspective integrability’’ vector $\mathbf{i}_p(x)$ depends only on u, v, f and $c_k^{i,j}$ i.e., known quantities. We will see later in Section 5, that numerically solving the set of equations (46) provides a simple way to numerically solve uncalibrated perspective photometric stereo under first-order spherical harmonics lighting.

If in addition we use the fact that \mathbf{m} fulfills the integrability constraint (44), we can substitute $(uc_u^{2,3} + vc_v^{2,3})$ by $(-fc_v^{2,4} + fc_u^{3,4})$ in the equation (45), and we get 17 summands instead of 18, turning (46) as follows

$$\mathbf{c}(x)^\top \mathbf{a} = 0 \text{ over } \Omega, \quad (47)$$

where $\mathbf{c}(x), \mathbf{a} \in \mathbb{R}^{17}$.

If we assume that there exist at least 17 pixels x such that the matrix formed by row-wise concatenation of the $\mathbf{c}(x)^\top$, $x \in \Omega$, is full-rank (this is the case in real-world scenarios if the surface is nontrivial), then $\mathbf{a} = 0$ and we get the following equations:

$$\begin{cases} A_{4,3}^{3,2} = A_{4,4}^{3,2} = A_{4,4}^{2,3} = A_{4,3}^{2,2} = 0, \\ A_{2,2}^{3,3} = A_{4,4}^{3,3}, \\ A_{3,3}^{4,4} = A_{4,4}^{4,4}, \\ A_{2,2}^{4,4} = A_{4,4}^{2,2}, \\ A_{3,3}^{4,4} = A_{4,4}^{3,3}, \\ A_{2,1}^{3,2} = A_{3,3}^{2,1} = A_{3,4}^{2,1} = A_{3,4}^{2,2} = A_{3,4}^{2,3} = 0, \\ A_{2,1}^{3,2} = A_{2,1}^{3,3} = A_{2,1}^{3,4} = A_{2,2}^{3,2} = A_{2,3}^{3,3} = 0, \\ A_{4,2}^{4,3} = A_{4,4}^{4,2} = A_{4,3}^{4,1} = A_{4,3}^{3,1} = A_{4,4}^{3,1} = 0. \end{cases} \quad (48)$$

According to the three first equations of System (48), the submatrix of \mathbf{A} formed by the last three rows and columns is a scaled generalized bas-relief transformation (see Corollary 1 in Appendix 2). That is to say, $\exists! (\lambda, \mu, \nu, \alpha) \in \mathbb{R}^4$ with $\lambda \neq 0, \alpha \neq 0$, such that:

$$\mathbf{A} = \begin{pmatrix} A_{11} & A_{12} & A_{13} & A_{14} \\ A_{21} & \alpha\lambda & 0 & -\alpha\mu \\ A_{31} & 0 & \alpha\lambda & -\alpha\nu \\ A_{41} & 0 & 0 & \alpha \end{pmatrix}. \quad (49)$$

Taking into account the other equations of system (48), we get $\lambda = 1, \mu = \nu = A_{21} = A_{31} = A_{41} = 0$. Then, the same arguments as those used around Eq. (41) yield $A_{12} = A_{13} = A_{14} = 0$, and the form (16) of $\mathbf{A}\mathbf{m}$ implies $A_{11} = \alpha$, which concludes the proof. \square

Let us remark that such a particular form of a scaled Lorentz transformation only scales all albedo values, leaving the geometry unchanged. From a practical point of view, this means that once an integrable candidate has been found, it corresponds to the genuine surface and there is no need to manually solve any ambiguity, unlike in the orthographic case. In the next section, we show that such a candidate can be estimated in closed-form in the discrete setting. This will allow us to empirically verify the validity of our theoretical results, through numerical experiments on simulated images.

5 Numerical solving of the perspective case

In this section, we derive a practical algorithm for solving perspective uncalibrated photometric stereo under first-order spherical lighting. More specifically, we provide a closed-form solution for an integrable normal field satisfying the image formation model (15), provided that the perspective camera is calibrated (i.e., its focal length and principal point are known).

5.1 Discrete formulation

First, let us reformulate the problem in the discrete setting. Let us stack all the observations $I^i(x)$, $i \in \{1, \dots, m\}$, $x \in \Omega$, in a matrix $\mathbf{I} \in \mathbb{R}^{m \times n}$, with $n = |\Omega|$ the number of pixels. Similarly to the directional lighting case represented by (6), the set of linear systems (15) can be rewritten in matrix form as:

$$\mathbf{I} = \mathbf{L}\mathbf{M}, \quad (50)$$

where $\mathbf{L} \in \mathbb{R}^{m \times 4}$ is the general lighting matrix, and $\mathbf{M} \in \mathbb{R}^{4 \times n}$ stacks all the unknown $\mathbf{m}(x)$ -vectors columnwise (each column $\mathbf{m}_j = \mathbf{m}(x_j)$ has thus the form given in Eq. (16)).

As shown in [3], a least-squares solution $(\mathbf{L}_1, \mathbf{M}_1)$ of (50) satisfying the constraint (16) can be obtained by singular value decomposition of \mathbf{I} . Since we know that any other \mathbf{M} -matrix solution of (50) differ from \mathbf{M}_1 according to a scaled Lorentz transform, the genuine solution $\mathbf{M}^* \in \mathbb{R}^{4 \times n}$ is given by

$$\mathbf{M}^* = \mathbf{A}\mathbf{M}_1, \quad (51)$$

with $\mathbf{A} \in L_s$ an unknown scaled Lorentz transformation.

In the last section we have seen that there exists a unique \mathbf{m} -field which both satisfies the image formation model and is integrable. This means, that if the pictured surface is regular (there is no kink or depth discontinuity), then the matrix \mathbf{A} in (51) is unique (up to scale). In fact, we only need the last three rows of matrix \mathbf{A} : left-multiplying the last three rows of the initial guess \mathbf{M}_1 by this submatrix, we obtain a matrix of size $3 \times n$ where the norm of the j -th column is the albedo at the surface point conjugate to pixel x_j , and normalizing each column yields the surface normal at this point.

The problem thus comes down to estimating the last three rows of matrix \mathbf{A} . According to Proposition 7 in Appendix 2, these rows can be written in the form $(\mathbf{v} | \mathbf{Q}) \in \mathbb{R}^{3 \times 4}$, where $\mathbf{v} \in \mathbb{R}^3$ and $\mathbf{Q} \in GL(3, \mathbb{R})$. Next we show how to estimate \mathbf{v} and \mathbf{Q} in closed-form, using a discrete analogous of the perspective integrability constraint.

5.2 Closed-form solution through discrete integrability enforcement

During the proof of Theorem 3, we showed that the integrability constraint yields the set of linear systems (46) over Ω . In the discrete setting, this set of equations can be written compactly as

$$\mathbf{I}_p \mathbf{w} = 0, \quad (52)$$

where $\mathbf{w} \in \mathbb{R}^{18}$ contains several minors of size 2 denoted by $(A_{k,l}^{i,j})$, and $\mathbf{I}_p \in \mathbb{R}^{n \times 18}$ is a ‘‘perspective integrability matrix’’ depending only upon the known camera parameters and entries of \mathbf{M}_1 .

Matrix \mathbf{I}_p is in general full-rank. Thus, the least-squares solution (up to scale) of (52) in terms of vector \mathbf{w} can be determined by singular value decomposition of \mathbf{I}_p : denoting by $\mathbf{I}_p = \mathbf{U}\Sigma\mathbf{V}^\top$ this decomposition, the solution \mathbf{w} is the last column of \mathbf{V} . We denote by $(\tilde{A}_{k,l}^{i,j}) = (\lambda A_{k,l}^{i,j})$ its entries, where $\lambda \neq 0$ denotes the unknown scale factor.

Now, recall that matrix $\mathbf{Q} \in \mathbb{R}^{3 \times 3}$ to be determined is the sub-matrix formed by the last three rows and columns of \mathbf{A} . It relates to the aforementioned minors according to

$$\mathbf{Q}^{-1} = \frac{1}{\det(\mathbf{Q})} \text{com}(\mathbf{Q})^\top = \frac{1}{\det(\mathbf{Q})} \begin{pmatrix} A_{4,4}^{3,3} & -A_{4,4}^{2,3} & A_{3,4}^{2,3} \\ -A_{4,4}^{3,2} & A_{4,4}^{2,2} & -A_{3,4}^{2,2} \\ A_{4,3}^{3,2} & -A_{4,3}^{2,2} & A_{3,3}^{2,2} \end{pmatrix}, \quad (53)$$

where $\text{com}(\mathbf{Q})$ is the comatrix of \mathbf{Q} . Thus:

$$\lambda \mathbf{Q}^{-1} = \frac{1}{\det(\mathbf{Q})} \mathbf{\Delta}^{-1}, \text{ where } \mathbf{\Delta} = \begin{pmatrix} \tilde{A}_{4,4}^{3,3} & -\tilde{A}_{4,4}^{2,3} & \tilde{A}_{3,4}^{2,3} \\ -\tilde{A}_{4,4}^{3,2} & \tilde{A}_{4,4}^{2,2} & -\tilde{A}_{3,4}^{2,2} \\ \tilde{A}_{4,3}^{3,2} & -\tilde{A}_{4,3}^{2,2} & \tilde{A}_{3,3}^{2,2} \end{pmatrix}^{-1}. \quad (54)$$

Hence, we can determine \mathbf{Q} up to scale:

$$\mathbf{Q} = (\lambda \det \mathbf{Q}) \mathbf{\Delta}. \quad (55)$$

Next, we turn our attention to the estimation of vector $\mathbf{v} \in \mathbb{R}^3$ (recall that this vector is formed by the first column and last three rows of \mathbf{A}), up to scale. To this end, we consider the last nine minors. For example, considering $\tilde{A}_{3,2}^{2,1}$:

$$\tilde{A}_{3,2}^{2,1} = \lambda (A_{21} A_{32} - A_{31} A_{22}) \quad (56)$$

$$= \lambda (A_{21} Q_{21} - A_{31} Q_{11}) \quad (57)$$

$$\stackrel{(55)}{=} \left(\lambda^2 \det(\mathbf{Q}) A_{21} \right) \Delta_{21} - \left(\lambda^2 \det(\mathbf{Q}) A_{31} \right) \Delta_{11}. \quad (58)$$

Let $\hat{\mathbf{v}} = \begin{pmatrix} \hat{v}_1 \\ \hat{v}_2 \\ \hat{v}_3 \end{pmatrix} = \lambda^2 \det(\mathbf{Q}) \begin{pmatrix} A_{21} \\ A_{31} \\ A_{41} \end{pmatrix} = (\lambda^2 \det(\mathbf{Q})) \mathbf{v}$. Eq. (58) can be written as:

$$\Delta_{21} \hat{v}_1 - \Delta_{11} \hat{v}_2 = \hat{A}_{3,2}^{2,1}. \quad (59)$$

In the same manner, by using all the other minors which involve A_{21} , A_{31} or A_{41} , we get the following over-constrained linear system:

$$\mathbf{S} \hat{\mathbf{v}} = \mathbf{b}, \quad (60)$$

where $\mathbf{S} \in \mathbb{R}^{9 \times 3}$ and $\mathbf{b} \in \mathbb{R}^9$. A least-squares solution for $\hat{\mathbf{v}}$ can be found using, e.g., the pseudo inverse:

$$\hat{\mathbf{v}} = \mathbf{S}^\dagger \mathbf{b}. \quad (61)$$

Besides,

$$\lambda^2 \det(\mathbf{Q})(\mathbf{v} | \mathbf{Q}) = (\lambda^2 \det(\mathbf{Q}) \mathbf{v} | \lambda^2 \det(\mathbf{Q}) \mathbf{Q}) \quad (62)$$

$$= (\hat{\mathbf{v}} | \lambda^3 \det(\mathbf{Q})^2 \mathbf{\Delta}), \quad (63)$$

and applying the determinant to both sides of Eq. (55) yields:

$$\lambda^3 \det(\mathbf{Q})^2 = \frac{1}{\det(\mathbf{\Delta})}. \quad (64)$$

Plugging (64) into (63), we eventually obtain the following closed-form expression for $(\mathbf{v} | \mathbf{Q})$:

$$(\mathbf{v} | \mathbf{Q}) = \frac{1}{\lambda^2 \det(\mathbf{Q})} \left(\hat{\mathbf{v}} | \frac{1}{\det(\mathbf{\Delta})} \mathbf{\Delta} \right). \quad (65)$$

Since λ and $\det(\mathbf{Q})$ in (65) are unknown, the solution $(\mathbf{v} | \mathbf{Q})$ is known only up to scale. As already stated, the actual value of this scale factor is non-important, since it only scales all albedo values simultaneously without affecting the geometry. Let us denote by $\tilde{\mathbf{M}}_1$ the submatrix formed by the last three rows of the initial guess \mathbf{M}_1 . Then, matrix $\tilde{\mathbf{M}}_2 = (\mathbf{v} | \mathbf{Q}) \tilde{\mathbf{M}}_1$ is a $3 \times n$ matrix where each column corresponds to one surface normal, scaled by the albedo. The norm of each column of $\tilde{\mathbf{M}}_2$ thus provides the sought albedo (up to scale), and normalizing each column provides the sought surface normal.

Therefore, we now have at hand a practical way to find an integrable normal field solving uncalibrated photometric stereo under general lighting. In the next subsection, we show on simulated data that such a solution indeed corresponds to the genuine surface, which provides an empirical evidence for the theoretical analysis conducted in the previous section.

5.3 Experiments

To empirically validate the well-posedness of perspective uncalibrated photometric stereo under general lighting, we implemented the previous algorithm in Matlab, and evaluated it against 16 synthetic datasets. These datasets were created by considering four 3D-shapes (“Armadillo”, “Bunny”, “Joyful Yell” and “Thai Statue”¹) and four different albedo maps (“White”, “Bars”, “Ebsd” and “Voronoi”). For each of the 16 combinations of 3D-shape and albedo, $m = 21$ images were simulated under varying first-order spherical harmonics illumination, and then converted to grayscale, as illustrated in Figure 1. Each image is of size 1600×1200 , and comes along with the ground-truth surface normals, reconstruction domain Ω , and intrinsic camera parameters (the focal length f , and the principal point used as reference for pixel coordinates). For the evaluation, we measured the mean angular error (MAE, in degrees) between the estimated normals and the ground-truth ones.

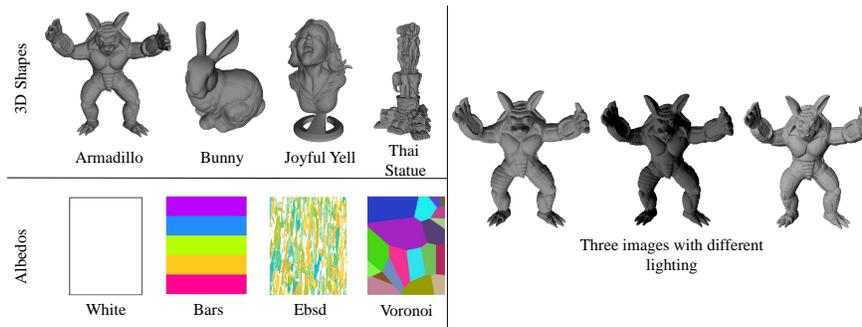


Fig. 1: The four 3D-shapes and four albedo maps used to create 16 (3D-shape, albedo) datasets. For each dataset, $m = 21$ images were rendered under varying first-order spherical harmonics lighting. On the right, we show three images of the (“Armadillo”, “White”) dataset.

As can be seen in Table 1, the mean angular error on the estimated normals is very low (less than 10 degrees for all datasets). This confirms that the geometry of the scene is unambiguously estimated. The images being synthesized without any additional noise or outlier to the Lambertian model (e.g., shadows or specularities), one may however be surprised that the mean angular error is non-zero. As suggested in [25], the observed residual errors may be due to the finite differences approximation of partial derivatives arising in the perspective integrability matrix (matrix \mathbf{I}_p in (52), which contains the partial derivatives of the entries of the initial \mathbf{m} -field, cf. (45)). In our implementation, we considered first-order finite differences: other choices of finite differences might reduce the error, yet we leave this as a perspective.

¹ Joyful Yell: <https://www.thingiverse.com/thing:897412>; other datasets: <http://www-graphics.stanford.edu/data/3dscanrep>

3D-shape	Albedo			
	White	Bars	Ebsd	Voronoi
Armadillo	2.01	1.81	2.13	2.03
Bunny	1.42	1.38	1.63	1.42
Joyful Yell	5.13	5.35	5.46	5.28
Thai Statue	6.33	6.40	6.46	7.62

Table 1: Mean angular error (in degrees), for each (3D-shape, albedo) combination. The error remains below 10 degrees for each dataset. This indicates that the genuine geometry is recovered, and empirically confirms the well-posedness of perspective uncalibrated photometric stereo under first-order spherical harmonics lighting.

Next, we evaluated the robustness of the proposed approach to an increasing amount of zero-mean, Gaussian noise added to the images of the (Armadillo, White) dataset. As can be seen in Table 2, the proposed method dramatically fails as soon as the noise becomes really perceptible (here, failure is observed when standard deviation $\sigma > 0.5\%$). For comparison, we also provide the results obtained with the state-of-the-art method [11], which is based on heuristical shape initialization followed by regularized nonconvex refinement. The heuristical nature of the initialization induces a non-negligible bias in shape estimation, which is clearly visible on noise-free data. However, this alternative is much more robust to noise.

Method	Standard deviation σ (in percents of the maximum intensity)								
	0.00	0.01	0.02	0.04	0.1	0.2	0.3	0.4	0.5
[11]	18.19	18.19	18.19	18.19	18.19	18.19	18.19	18.20	18.20
Ours	2.01	2.07	2.12	2.33	2.90	4.43	6.56	9.14	113.38

Table 2: Mean angular error (in degrees) on the (Armadillo, White) dataset, with increasing amount of zero-mean Gaussian noise added to the input images. When noise is negligible, the proposed method largely outperforms the state-of-the-art method from [11]. However, it should be discarded in the presence of strong noise.

This is not really surprising, since the proposed method is spectral, and the alternative one is based on evolved nonconvex optimization. In general, the former is faster (in our implementation on a recent computer, our results were obtained in less than 10 seconds, and the alternative ones in around 30 minutes), but the latter is more robust. Similar observations have been made in other computer vision communities, e.g. pose estimation: the 8-point algorithm [19] is usually replaced by bundle adjustment techniques [41] in order to handle the unavoidable noise arising in real-world data. Overall, the proposed algorithm should be considered only as a way to empirically confirm the well-posedness of the problem, yet on real-world data the existing numerical implementations should be preferred.

6 Conclusion and perspectives

We have studied the well-posedness of uncalibrated photometric stereo under general illumination, represented by first-order spherical harmonics. We have established that integrability reduces the scaled Lorentz ambiguity to a global concave/convex ambiguity in the orthographic case, and resolves it in the perspective one. As Table 3 summarizes, this generalizes previous results which were restricted to the directional lighting model. Still, open questions remain concerning further generalization of these results to even more evolved lighting models. For instance, future research on the topic may consider the case of unknown second-order spherical harmonics [3], or that of unknown nearby point light sources [26]. Such generalizations would be of interest from a practical perspective, because the former represents natural illumination very accurately [10], and the latter allows using inexpensive light sources such as LEDs [29].

Lighting model	Underlying ambiguity	Effect of imposing integrability	
		Orthographic	Perspective
Directional	9-dof (linear) [12]	3-dof (GBR) [44]	Well-posed [25]
SH ₁	6-dof (scaled Lorentz) [3]	1-dof (concave/convex)	Well-posed
SH ₂	9-dof (linear) [3]	?	?
Nearby point	4-dof (rotation and scale) [26]	?	?

Table 3: State-of-the-art of theoretical results concerning the well-posedness of uncalibrated photometric stereo under different lighting models (directional, spherical harmonics of order 1 and 2, or nearby point sources). We indicate the number of degrees of freedoms (dof) of the underlying ambiguity, and how imposing integrability reduces this number under both orthographic and perspective projection. The bold results refer to the findings in the present paper, and the question marks to remaining open problems.

References

1. Ofer Bartal, Nati Ofir, Yaron Lipman, and Ronen Basri. Photometric stereo by hemispherical metric embedding. *Journal of Mathematical Imaging and Vision*, 60(2):148–162, 2018.
2. Ronen Basri and David W Jacobs. Lambertian reflectances and linear subspaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(2):218–233, 2003.
3. Ronen Basri, David W Jacobs, and Ira Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of Computer Vision*, 72(3):239–257, 2007.
4. Peter N Belhumeur, David J Kriegman, and Alan L Yuille. The bas-relief ambiguity. *International journal of computer vision*, 35(1):33–44, 1999.
5. Michael Breuß, Emiliano Cristiani, Jean-Denis Durou, Maurizio Falcone, and Oliver Vogel. Perspective shape from shading: Ambiguity analysis and numerical approximations. *SIAM Journal on Imaging Sciences*, 5(1):311–342, 2012.

6. Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee K Wong. Self-calibrating deep photometric stereo networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8739–8747, 2019.
7. Lixiong Chen, Yinqiang Zheng, Boxin Shi, Art Subpa-asa, and Imari Sato. A microfacet-based model for photometric stereo with general isotropic reflectance. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (in press), 2019.
8. Jean-Denis Durou, Maurizio Falcone, and Manuela Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding*, 109(1):22–43, 2008.
9. Albert Einstein. Zur elektrodynamik bewegter körper. *Annalen der physik*, 322(10):891–921, 1905.
10. Darya Frolova, Denis Simakov, and Ronen Basri. Accuracy of spherical harmonic approximations for images of Lambertian objects under far and near lighting. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 574–587, 2004.
11. Bjoern Haefner, Zhenzhang Ye, Maolin Gao, Tao Wu, Yvain Quéau, and Daniel Cremers. Variational uncalibrated photometric stereo under general lighting. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2019.
12. Hideki Hayakawa. Photometric stereo under a light source with arbitrary motion. *Journal of the Optical Society of America A*, 11(11):3079–3089, 1994.
13. Satoshi Ikehata. CNN-PS: CNN-based photometric stereo for general non-convex surfaces. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
14. Arthur Jaffe. Lorentz transformations, rotations, and boosts, 2013. Online notes, accessed Sept. 2019. Url: http://home.ku.edu.tr/~amostafazadeh/phys517_518/phys517_2016f/Handouts/A_Jaffi_Lorentz_Group.pdf.
15. Maryam Khanian, Ali Sharifi Boroujerdi, and Michael Breuß. Photometric stereo for strong specular highlights. *Computational Visual Media*, 4(1):83–102, 2018.
16. Ryszard Kozera and Alexander Prokopenya. Second-order algebraic surfaces and two image photometric stereo. In *International Conference on Computer Vision and Graphics*, pages 234–247, 2018.
17. Junxuan Li, Antonio Robles-Kelly, Shaodi You, and Yasuyuki Matsushita. Learning to minify photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7568–7576, 2019.
18. Fotios Logothetis, Roberto Mecca, Yvain Qu' eau, and Roberto Cipolla. Near-field photometric stereo in ambient light. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2016.
19. H.C. Longuet-Higgins. A computer algorithm for reconstructing a scene from two projections. In Martin A. Fischler and Oscar Firschein, editors, *Readings in Computer Vision*, pages 61–62. Morgan Kaufmann, 1987.
20. Roberto Mecca and Maurizio Falcone. Uniqueness and approximation of a photometric shape-from-shading model. *SIAM Journal on Imaging Sciences*, 6(1):616–659, 2013.
21. Roberto Mecca, Yvain Quéau, Fotios Logothetis, and Roberto Cipolla. A single-lobe photometric stereo approach for heterogeneous material. *SIAM Journal on Imaging Sciences*, 9(4):1858–1888, 2016.
22. Roberto Mecca, Aaron Wetzler, Alfred M Bruckstein, and Ron Kimmel. Near field photometric stereo with point light sources. *SIAM Journal on Imaging Sciences*, 7(4):2732–2770, 2014.
23. Zhipeng Mo, Boxin Shi, Feng Lu, Sai-Kit Yeung, and Yasuyuki Matsushita. Uncalibrated photometric stereo under natural illumination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2936–2945, 2018.
24. Ruth Onn and Alfred Bruckstein. Integrability disambiguates surface recovery in two-image photometric stereo. *International Journal of Computer Vision*, 5(1):105–113, 1990.
25. Thoma Papadhimetri and Paolo Favaro. A new perspective on uncalibrated photometric stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1474–1481, 2013.
26. Thoma Papadhimetri and Paolo Favaro. Uncalibrated near-light photometric stereo. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2014.

27. Songyou Peng, Bjoern Haefner, Yvain Quéau, and Daniel Cremers. Depth super-resolution meets uncalibrated photometric stereo. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV) Workshops*, pages 2961–2968, 2017.
28. Henri Poincaré. Sur la dynamique de l'électron. *Comptes Rendus de l'Académie des Sciences*, 140:1504–1508, 1905.
29. Yvain Quéau, Bastien Durix, Tao Wu, Daniel Cremers, François Lauze, and Jean-Denis Durou. LED-based photometric stereo: Modeling, calibration and numerical solution. *Journal of Mathematical Imaging and Vision*, 60(3):313–340, 2018.
30. Yvain Quéau, Jean-Denis Durou, and Jean-François Aujol. Normal integration: a survey. *Journal of Mathematical Imaging and Vision*, 60(4):576–593, 2018.
31. Yvain Quéau, Roberto Mecca, Jean-Denis Durou, and Xavier Descombes. Photometric stereo with only two images: A theoretical study and numerical resolution. *Image and Vision Computing*, 57:175–191, 2017.
32. Yvain Quéau, Tao Wu, and Daniel Cremers. Semi-calibrated near-light photometric stereo. In *International Conference on Scale Space and Variational Methods in Computer Vision (SSVM)*, pages 656–668, 2017.
33. Yvain Quéau, Tao Wu, François Lauze, Jean-Denis Durou, and Daniel Cremers. A non-convex variational approach to photometric stereo under inaccurate lighting. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 99–108, 2017.
34. Georg Radow, Laurent Hoeltgen, Yvain Quéau, and Michael Breuß. Optimisation of classic photometric stereo by non-convex variational minimisation. *Journal of Mathematical Imaging and Vision*, 61(1):84–105, 2019.
35. Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*, pages 497–500, 2001.
36. Soumyadip Sengupta, Hao Zhou, Walter Forkel, Ronen Basri, Tom Goldstein, and David Jacobs. Solving uncalibrated photometric stereo using fewer images by jointly optimizing low-rank matrix completion and integrability. *Journal of Mathematical Imaging and Vision*, 60(4):563–575, 2018.
37. Boxin Shi, Kenji Inose, Yasuyuki Matsushita, Ping Tan, Sai-Kit Yeung, and Katsushi Ikeuchi. Photometric stereo using internet images. In *Proceedings of the International Conference on 3D Vision (3DV)*, volume 1, pages 361–368, 2014.
38. Boxin Shi, Zhe Wu, Zhipeng Mo, Dinglong Duan, Sai-Kit Yeung, and Ping Tan. A benchmark dataset and evaluation for non-Lambertian and uncalibrated photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(2):271–284, 2019.
39. Silvia Tozza and Maurizio Falcone. Analysis and approximation of some shape-from-shading models for non-Lambertian surfaces. *Journal of Mathematical Imaging and Vision*, 55(2):153–178, 2016.
40. Silvia Tozza, Roberto Mecca, Marti Duocastella, and Alessio Del Bue. Direct differential photometric stereo shape recovery of diffuse and specular surfaces. *Journal of Mathematical Imaging and Vision*, 56(1):57–76, 2016.
41. Bill Triggs, Philip F McLauchlan, Richard I Hartley, and Andrew W Fitzgibbon. Bundle adjustment – a modern synthesis. In *International Workshop on Vision Algorithms (ICCV Workshops)*, pages 298–372, 1999.
42. Robert J Woodham. Reflectance map techniques for analyzing surface defects in metal castings. Technical Report AITR-457, MIT, 1978.
43. Robert J Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
44. Alan Yuille and Daniel Snow. Shape and albedo from multiple images using integrability. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 158–164, 1997.
45. Ruo Zhang, Ping-Sing Tsai, James Edwin Cryer, and Mubarak Shah. Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999.

Appendix 1: Proof of Theorem 1

Theorem 1 characterizes scaled Lorentz transformations. Its proof relies on Propositions 1, 2 and 3 from Lorentz' group theory (proofs of these propositions can be found in [14]).

Proposition 1 $\forall \mathbf{A} \in L_o^P, \exists! (\mathbf{v}, \mathbf{O}) \in B(\mathbf{0}, 1) \times SO(3, \mathbb{R}),$

$$\mathbf{A} = \mathbf{S}(\mathbf{v}) \mathbf{R}(\mathbf{O}) = \begin{pmatrix} \gamma & \gamma \mathbf{v}^\top \mathbf{O} \\ \gamma \mathbf{v} & (\mathbf{I}_3 + \frac{\gamma^2}{1+\gamma} \mathbf{v} \mathbf{v}^\top) \mathbf{O} \end{pmatrix}, \quad (66)$$

where $(\mathbf{v}, \mathbf{O}) \in B(\mathbf{0}, 1) \times SO(3, \mathbb{R}), \gamma = \frac{1}{\sqrt{1-\|\mathbf{v}\|^2}},$

$$\mathbf{S}(\mathbf{v}) = \begin{pmatrix} \gamma & \gamma \mathbf{v}^\top \\ \gamma \mathbf{v} & \mathbf{I}_3 + \frac{\gamma^2}{1+\gamma} \mathbf{v} \mathbf{v}^\top \end{pmatrix}, \quad (67)$$

$$\mathbf{R}(\mathbf{O}) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & & & \\ 0 & \mathbf{O} & & \\ 0 & & & \end{pmatrix}. \quad (68)$$

Proposition 2 *The product of two proper/improper transformations is a proper one, and the product of a proper and improper transformations is an improper one. The same for the orthochronous property.*

Proposition 3 $\mathbf{T} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \in L_n^i,$ and $\mathbf{P} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \in L_o^i.$

Using these already known results, we propose the following characterization of Lorentz transformations:

Proposition 4 $\forall \mathbf{A} \in L, \exists! (\mathbf{v}, \mathbf{O}) \in B(\mathbf{0}, 1) \times SO(3, \mathbb{R}),$

$$\mathbf{A} = \begin{pmatrix} \epsilon_1(\mathbf{A}) \gamma & \epsilon_1(\mathbf{A}) \gamma \mathbf{v}^\top \mathbf{O} \\ \epsilon_2(\mathbf{A}) \gamma \mathbf{v} & \epsilon_2(\mathbf{A}) (\mathbf{I}_3 + \frac{\gamma^2}{1+\gamma} \mathbf{v} \mathbf{v}^\top) \mathbf{O} \end{pmatrix}. \quad (69)$$

Proof According to Propositions 2 and 3: $\mathbf{T} \in L_n^i$, $\mathbf{P} \in L_o^i \implies \mathbf{PT} \in L_n^p$.

Let $\mathbf{A} \in L_n^i$. Using Proposition 2, $\mathbf{TA} \in L_o^p$.

Therefore, according to Proposition 1, $\exists! (\mathbf{v}, \mathbf{O}) \in B(\mathbf{0}, 1) \times SO(3, \mathbb{R})$, $\mathbf{TA} = \mathbf{S}(\mathbf{v})\mathbf{R}(\mathbf{O})$. Since $\mathbf{TT} = \mathbf{I}_4$, this implies: $\exists! (\mathbf{v}, \mathbf{O}) \in B(\mathbf{0}, 1) \times SO(3, \mathbb{R})$, $\mathbf{A} = \mathbf{TS}(\mathbf{v})\mathbf{R}(\mathbf{O})$. In addition, $\epsilon_1(\mathbf{A}) = -1$ and $\epsilon_2(\mathbf{A}) = 1$, hence:

$$\mathbf{A} = \begin{pmatrix} \epsilon_1(\mathbf{A})\gamma & \epsilon_1(\mathbf{A})\gamma\mathbf{v}^\top\mathbf{O} \\ \epsilon_2(\mathbf{A})\gamma\mathbf{v} & \epsilon_2(\mathbf{A})(\mathbf{I}_3 + \frac{\gamma^2}{1+\gamma}\mathbf{v}\mathbf{v}^\top)\mathbf{O} \end{pmatrix}. \quad (70)$$

With the same reasoning, we get the result for all the other transformations. \square

Combining Proposition 4 and the definition (18) of scaled Lorentz transformations, we get Theorem 1.

7 Appendix 2: Some useful results on GBR and Lorentz matrices, and Corollary 1

The aim of this section is to prove Corollary 1, which was used in the proofs of Theorems 2 and 3. Its proof relies on a few results on GBR and Lorentz matrices, which we provide in the following.

Let us denote by G the group of generalized bas-relief (GBR) transformations, and by G_s that of scaled GBR transformations defined by:

$$G_s = \{s\mathbf{A} \mid s \in \mathbb{R} \setminus \{0\} \text{ and } \mathbf{A} \in G\} \quad (71)$$

Both are subgroups of $GL(3, \mathbb{R})$ under the matrix product. For all $\mathbf{B} = s\mathbf{A} \in G_s$, we call s the scale part of \mathbf{B} , and \mathbf{A} its GBR part.

Let $\mathbf{C} \in \mathbb{R}^{n \times n}$ with $n > 1$ and C_{ij} its entries. We will use the following notation for a minor of size two:

$$C_{k,l}^{i,j} = C_{ij}C_{kl} - C_{kj}C_{il} \quad (72)$$

where $1 \leq i < k \leq n$ and $1 \leq j < l \leq n$.

Such minors allow to characterize scaled GBR matrices:

Proposition 5 Let $\mathbf{A} \in \mathbb{R}^{3 \times 3}$, A_{ij} the entries of \mathbf{A} , then $\mathbf{A} \in G_s \iff \mathbf{A}$ is invertible and fulfills the following equations:

$$\begin{cases} A_{3,2}^{2,1} = A_{3,3}^{2,1} = A_{3,3}^{1,2} = A_{3,2}^{1,1} = 0, \\ A_{3,3}^{2,2} = A_{3,3}^{1,1}. \end{cases} \quad (73)$$

Proof See [4]. \square

Proposition 6 Let $\mathbf{v} \in B_0(1)$, $\gamma = \frac{1}{\sqrt{1-\|\mathbf{v}\|^2}}$, then $\mathbf{C} = \mathbf{I}_3 + \frac{\gamma^2}{1+\gamma} \mathbf{v}\mathbf{v}^\top$ is positive definite.

Proof Let $\mathbf{B} = \frac{\gamma^2}{1+\gamma} \mathbf{v}\mathbf{v}^\top$. We note $E_\lambda(\mathbf{B})$ the eigenspace associated to the eigenvalue λ of \mathbf{B} . \mathbf{B} is symmetric, thus according to the spectral theorem, all the eigenvalues of \mathbf{B} are real, and $\mathbb{R}^3 = \bigoplus_{i=1}^r E_{\lambda_i}(\mathbf{B})$ with $r \leq 3$ the number of eigenvalues, and $\{\lambda_i\}_{i=1..r}$ the eigenvalues of \mathbf{B} . Hence: $\dim(\mathbb{R}^3) = \sum_{i=1}^r \dim(E_{\lambda_i}(\mathbf{B}))$.

According to the rank-nullity theorem, $\dim(\text{Ker}(\mathbf{B})) + \text{rank}(\mathbf{B}) = 3$, and by definition $\text{rank}(\mathbf{B}) = 1$, thus $\dim(\text{Ker}(\mathbf{B})) = \dim(E_0(\mathbf{B})) = 2$. We deduce that $\exists! \lambda \in \mathbb{R} \setminus \{0\}$, $\mathbb{R}^3 = E_0(\mathbf{B}) \oplus E_\lambda(\mathbf{B})$ with $\dim(E_\lambda(\mathbf{B})) = 1$.

Let $\Pi_{\mathbf{v}}$ the orthogonal projection onto $\text{span}\{\mathbf{v}\}$, and let $\mathbf{x} \in \mathbb{R}^3 : \Pi_{\mathbf{v}}(\mathbf{x}) = \frac{\mathbf{v}\mathbf{v}^\top}{\|\mathbf{v}\|^2} \mathbf{x}$ and $\Pi_{\mathbf{v}}(\mathbf{v}) = \mathbf{v}$. We have: $\mathbf{B} = \frac{\gamma^2}{1+\gamma} \|\mathbf{v}\|^2 \Pi_{\mathbf{v}}$ and $\mathbf{B}\mathbf{v} = (\frac{\gamma^2}{1+\gamma} \|\mathbf{v}\|^2) \mathbf{v}$. Thus, $\frac{\gamma^2}{1+\gamma} \|\mathbf{v}\|^2$ is an eigenvalue of \mathbf{B} and $\lambda = \frac{\gamma^2}{1+\gamma} \|\mathbf{v}\|^2$.

Besides, $\frac{\lambda}{\gamma-1} = \frac{\gamma^2}{(\gamma-1)(\gamma+1)} \|\mathbf{v}\|^2 = \frac{\gamma^2}{\gamma^2-1} \|\mathbf{v}\|^2 = \frac{1}{1-\frac{1}{\gamma^2}} \|\mathbf{v}\|^2 = \frac{1}{1-(1-\|\mathbf{v}\|^2)} \|\mathbf{v}\|^2 = 1$.

Therefore, $\lambda = \gamma - 1$, hence the eigenvalues of \mathbf{B} are 0 and $(\gamma - 1)$.

Let $\alpha \in \{0, \gamma - 1\}$ and $\mathbf{u} \in E_\alpha(\mathbf{B})$. We have: $\mathbf{B}\mathbf{u} = \alpha\mathbf{u} \iff \mathbf{u} + \mathbf{B}\mathbf{u} = \mathbf{u} + \alpha\mathbf{u} \iff (\mathbf{I}_3 + \mathbf{B})\mathbf{u} = (\alpha + 1)\mathbf{u} \iff \mathbf{C}\mathbf{u} = (\alpha + 1)\mathbf{u}$. Thus, $1 > 0$ and $\gamma > 0$ are the eigenvalues of \mathbf{C} with $E_1(\mathbf{C}) = E_0(\mathbf{B})$ and $E_\gamma(\mathbf{C}) = E_{\gamma-1}(\mathbf{B})$.

Consequently, \mathbf{C} is positive definite. \square

Proposition 7 $\forall \mathbf{A}_s \in L_s$, the submatrix \mathbf{B} formed by the last 3 rows and 3 columns of \mathbf{A}_s is invertible.

Proof $\exists! (s, \tilde{\mathbf{A}}) \in \mathbb{R} \setminus \{0\} \times L$, $\mathbf{A}_s = s\tilde{\mathbf{A}}$. Therefore, using Proposition 4: $\exists! (\mathbf{v}, \mathbf{O}) \in B(\mathbf{0}, 1) \times SO(3, \mathbb{R})$, $\mathbf{B} = s e_2(\tilde{\mathbf{A}})(\mathbf{I}_3 + \frac{\gamma^2}{1+\gamma} \mathbf{v}\mathbf{v}^\top)\mathbf{O}$.

Since $\mathbf{O} \in SO(3, \mathbb{R})$, $\det(\mathbf{O}) = 1$. In addition, Proposition 6 implies that $\det(\mathbf{I}_3 + \frac{\gamma^2}{1+\gamma} \mathbf{v}\mathbf{v}^\top) > 0$. Therefore, $\det(\mathbf{B}) \neq 0$. \square

Corollary 1 Let $\mathbf{A}_s \in L_s$. If its entries A_{ij} fulfill the following equations:

$$\begin{cases} A_{4,3}^{3,2} = A_{4,4}^{3,2} = A_{4,4}^{2,3} = A_{4,3}^{2,2} = 0, \\ A_{4,4}^{3,3} = A_{4,4}^{2,2}, \end{cases} \quad (74)$$

then the submatrix \mathbf{B} of \mathbf{A}_s formed by the last 3 rows and 3 columns is a scaled GBR, i.e.:

$$\exists! (\lambda, \mu, \nu, \beta) \in \mathbb{R}^4 \text{ with } \lambda \neq 0, \beta \neq 0, \mathbf{A}_s = \begin{pmatrix} A_{11} & A_{12} & A_{13} & A_{14} \\ A_{21} & \beta\lambda & 0 & -\beta\mu \\ A_{31} & 0 & \beta\lambda & -\beta\nu \\ A_{41} & 0 & 0 & \beta \end{pmatrix}. \quad (75)$$

Proof According to Proposition 7, \mathbf{B} is invertible. Besides, \mathbf{A}_s fulfill equations (74) $\iff \mathbf{B}$ fulfill equations (73), thus according to Proposition 5, $\mathbf{B} \in G_s$. \square

8 Appendix 3: Perspective integrability constraint, and Corollary 2

The proof of Theorem 3 uses the expression of the integrability constraint under perspective projection. This expression is given by Corollary 2, which we prove in this section.

Let $z : \Omega \rightarrow \mathbb{R}$ a twice differentiable depth map of an object, $\mathbf{n} := [n_1, n_2, n_3]^\top : \Omega \rightarrow \mathbb{S}^2 \subset \mathbb{R}^3$ its normal field, and let us define the log depth map as:

$$\tilde{z} = \log(z), \quad (76)$$

and denote:

$$p = -\frac{n_1}{n_3}, \quad q = -\frac{n_2}{n_3}. \quad (77)$$

Proposition 8 *The gradient of the log-depth map $\tilde{z} = \log z$ is given by:*

$$\nabla \tilde{z} = [\hat{p}, \hat{q}]^\top, \quad \text{with} \quad (78)$$

$$\hat{p} = \frac{p}{f - up - vq} \quad \text{and} \quad \hat{q} = \frac{q}{f - up - vq}. \quad (79)$$

Proof By definition of \hat{p} and using the relationship (12) between the normal and depth maps:

$$\hat{p} = \frac{p}{f - up - vq} = \frac{-n_1}{fn_3 + un_1 + vn_2} = \frac{z_u}{z} = \tilde{z}_u. \quad (80)$$

Similarly, we show that $\hat{q} = \tilde{z}_v$. \square

Proposition 9 *Integrability of the depth map and integrability of the log-depth map are equivalent, i.e.:*

$$z_{uv} = z_{vu} \iff \tilde{z}_{uv} = \tilde{z}_{vu}. \quad (81)$$

Proof Assuming $z > 0$:

$$z_{uv} = z_{vu} \iff \frac{z_{uv}z - z_u z_v}{z^2} = \frac{z_{vu}z - z_u z_v}{z^2}, \quad (82)$$

$$\iff \left(\frac{z_u}{z} \right)_v = \left(\frac{z_v}{z} \right)_u, \quad (83)$$

$$\iff \tilde{z}_{uv} = \tilde{z}_{vu}. \quad (84)$$

Corollary 2 *Let $\mathbf{m} = [c_1, c_2, c_3, c_4]^\top : \Omega \rightarrow \mathbb{R}^4$ a field with the form of Equation (16). The underlying normal field is integrable iff the following relationship holds over Ω :*

$$u(c_2c_{3u} - c_{2u}c_3) + v(c_2c_{3v} - c_{2v}c_3) + f(c_2c_{4v} - c_{2v}c_4) + f(c_{3u}c_4 - c_3c_{4u}) = 0. \quad (85)$$

Proof The following proof is largely inspired by [25].

According to Propositions 8 and 9, integrability of the underlying normal field can be written as:

$$\begin{aligned}
& \hat{p}_v = \hat{q}_u \text{ over } \Omega, \\
& \iff \left(\frac{p}{f - up - vq} \right)_v = \left(\frac{q}{f - up - vq} \right)_u \text{ over } \Omega, \\
& \iff fp_v - vqp_v + vpq_v - fq_u + upq_u - uqp_u = 0 \text{ over } \Omega, \\
& \iff \begin{pmatrix} p_v \\ q_v \\ 0 \end{pmatrix}^\top \begin{pmatrix} 0 \\ -f \\ v \end{pmatrix} \times \begin{pmatrix} p \\ q \\ -1 \end{pmatrix} + \begin{pmatrix} p_u \\ q_u \\ 0 \end{pmatrix}^\top \begin{pmatrix} -f \\ 0 \\ u \end{pmatrix} \times \begin{pmatrix} p \\ q \\ -1 \end{pmatrix} = 0 \text{ over } \Omega, \tag{86}
\end{aligned}$$

where \times denotes the cross-product.

Let $\mathbf{d} = [c_2, c_3, c_4]^\top$. We have: $\frac{c_2}{c_4} = -p$ and $\frac{c_3}{c_4} = -q$ according to the form (16)

of \mathbf{m} , thus $-\frac{\mathbf{d}}{c_4} = \begin{pmatrix} p \\ q \\ -1 \end{pmatrix}$.

Let $\mathbf{w}_1 = [0, -f, v]^\top$ and $\mathbf{w}_2 = [-f, 0, u]^\top$. According to the integrability constraint (86):

$$\begin{aligned}
& \left(\frac{-\mathbf{d}}{c_4} \right)_v^\top \mathbf{w}_1 \times \left(\frac{-\mathbf{d}}{c_4} \right) + \left(\frac{-\mathbf{d}}{c_4} \right)_u^\top \mathbf{w}_2 \times \left(\frac{-\mathbf{d}}{c_4} \right) = 0 \text{ over } \Omega, \\
& \iff \left(-\frac{c_4 \mathbf{d}_v - c_{4v} \mathbf{d}}{c_4^2} \right)^\top \mathbf{w}_1 \times \left(\frac{-\mathbf{d}}{c_4} \right) + \left(-\frac{c_4 \mathbf{d}_u - c_{4u} \mathbf{d}}{c_4^2} \right)^\top \mathbf{w}_2 \times \left(\frac{-\mathbf{d}}{c_4} \right) = 0 \text{ over } \Omega. \tag{87}
\end{aligned}$$

Multiplying Eq. (87) by c_4^3 :

$$(c_4 \mathbf{d}_v - c_{4v} \mathbf{d})^\top \mathbf{w}_1 \times \mathbf{d} + (c_4 \mathbf{d}_u - c_{4u} \mathbf{d})^\top \mathbf{w}_2 \times \mathbf{d} = 0 \text{ over } \Omega. \tag{88}$$

In addition, $(\mathbf{w}_1 \times \mathbf{d}) \perp \mathbf{d}$ and $(\mathbf{w}_2 \times \mathbf{d}) \perp \mathbf{d}$, thus the following relationship holds over Ω :

$$\begin{aligned}
& c_4 \mathbf{d}_v^\top (\mathbf{w}_1 \times \mathbf{d}) + c_4 \mathbf{d}_u^\top (\mathbf{w}_2 \times \mathbf{d}) = 0, \\
& \iff \mathbf{d}_v^\top (\mathbf{w}_1 \times \mathbf{d}) + \mathbf{d}_u^\top (\mathbf{w}_2 \times \mathbf{d}) = 0, \\
& \iff u(c_{2u}c_3 - c_2c_{3u}) + v(c_{2v}c_3 - c_2c_{3v}) \\
& \quad + f(c_{2v}c_4 - c_2c_{4v}) - f(c_{3u}c_4 - c_3c_{4u}) = 0. \tag{89}
\end{aligned}$$

which concludes the proof. \square