O. Drbohlav and M. Chantler: On Optimal Light Configurations in Photometric Stereo. In *ICCV 2005: Proceedings of the 10th IEEE International Conference on Computer Vision*, vol. II, pp. 1707–1712, Beijing, China, 2005.

On Optimal Light Configurations in Photometric Stereo

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

Abstract

This paper develops new theory for the optimal placement of photometric stereo lighting in the presence of camera noise. We show that for three lights, any triplet of orthogonal light directions minimises the uncertainty in scaled normal computation. The assumptions are that the camera noise is additive and normally distributed, and uncertainty is defined as the expectation of squared distance of scaled normal to the ground truth. If the camera noise is of zero mean and variance σ^2, the optimal (minimum) uncertainty in the scaled normal is $3\sigma^2$. For case of $n > 3$ lights, we show that the minimum uncertainty is $9\sigma^2/n$, and identify sets of light configurations which reach this theoretical minimum.

1. Introduction

Photometric stereo [7] is a method which computes local surface orientations and reflectance at each pixel using images captured under different illumination conditions with a fixed camera. In general, the intensity i at a pixel observing a certain surface patch will depend on the orientation of the patch with respect to the viewing and illumination directions (see Fig. 1). For Lambertian surfaces [3], this intensity is given by*

$$i = \rho\kappa\cos\theta_i\,, \tag{1}$$

where ρ is the *albedo* representing the amount of light reflected back from the surface, κ is the light source intensity, and θ_i is the angle of incidence. This can be further rewritten as

$$i = \rho\kappa\mathbf{l}^\mathsf{T}\mathbf{n} = (\kappa\mathbf{l})^\mathsf{T}(\rho\mathbf{n}) = \mathbf{s}^\mathsf{T}\mathbf{b} \tag{2}$$

where $\mathbf{l}$ and $\mathbf{n}$ are the unit light and normal vectors, respectively, and vectors $\mathbf{b} = \rho\mathbf{n}$ and $\mathbf{s} = \kappa\mathbf{l}$ are the *scaled normal* and the *scaled light*, respectively. Scaled normals and scaled lights are the basic entities of Lambertian photometric stereo.

*In this article, we adopt the usual convention that 'pixel intensity' and 'radiance' are interchangeable terms.

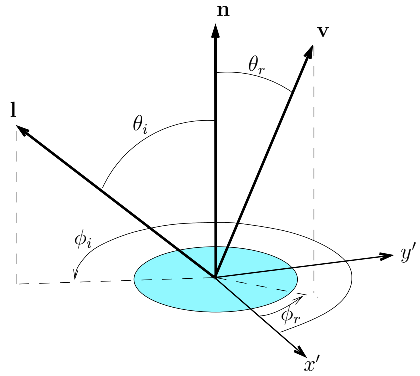

Figure 1: Reflectance geometry. In general, the reflectance depends on the orientation of the surface patch normal $\mathbf{n}$ and its reference plane (x'–y') with respect to the viewing direction $\mathbf{v}$ and illumination direction $\mathbf{l}$.

Photometric stereo is based on *inverting* the model of reflectance. For the case of Lambertian reflectance, it is obvious from (2) that having the intensities i_1, i_2 and i_3 observed for three scaled light sources $\mathbf{s}_1$, $\mathbf{s}_2$ and $\mathbf{s}_3$, there holds

$$\mathbf{i} = [i_1, i_2, i_3]^\mathsf{T} = [\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3]^\mathsf{T}\mathbf{b} = \mathbf{S}^\mathsf{T}\mathbf{b}\,, \tag{3}$$

where the vector $\mathbf{i} = [i_1, i_2, i_3]^\mathsf{T}$ stacks the intensity measurements and the matrix $\mathbf{S}$, called the *light matrix*, stacks the scaled light vectors, $\mathbf{S} = [\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3]$. Therefore, provided that $\mathbf{S}$ is invertible, it is possible to compute the scaled normal $\mathbf{b}$ as

$$\mathbf{b} = \mathbf{S}^{-\mathsf{T}}\mathbf{i}\,. \tag{4}$$

Albedo ρ can then be computed as the magnitude of this vector, while the normal orientation $\mathbf{n}$ is given by normalising this vector.

In principle, there are several sources of errors which affect the accuracy with which the scaled surface normal is determined.

In a fundamental paper on this topic, Ray, Birk and Kelley [5] identify principal sources of errors, and strategies for dealing with them. Among the error sources, there are: imprecise measurement of intensities $\mathbf{i}$ (camera sensor noise); error in calibration of light source directions; detector nonlinearity; shadows and specularities (i.e. deviation from the reflectance model assumed); spatial and spectral distribution of incident light; surface micro-structure; and optical imperfections of the surrounding environment. The authors provide equations for computing the error caused by errors in observed intensities and in the directions of lights used. For such *error analysis*, they work with gradient space representations of reflectance maps, which results in rather complicated equations. This is why Jiang and Bunke [2] later repeated the derivation using a different parametrisation to provide a simpler formulation.

While the work discussed above concentrated on error analysis, the goal of this article is to present new theory that enables the light source configurations to be determined that guarantee that the effect of the errors is minimised.

To the best of our knowledge this fundamental topic has not been addressed formally in full before.

A *qualitative* recommendation on the placement of light sources has been made by Lee and Kuo [4] who developed an algorithm for two-source photometric stereo. By examining the shape of the reflectance maps, they suggested that the tilts of light sources should be $90°$ apart.

The article closest to the work described here is that of Spence and Chantler [6]. They worked with three lights of equal slant and using *numerical optimisation* they concluded that the normals are best reconstructed when the light sources are $120°$ apart, and of slant $\sim 54.74°$. This corresponds to *orthogonal* light directions and is in full agreement with the results derived from the theory presented here. Additionally, from our formulation of the problem it follows that all orthogonal triplets of lights are optimal, relaxing the condition that their slant should be the same. This result also agrees with a brief analysis within the work of Bezděk [1] who investigated a tri-illuminant case and concluded that the three light sources should be orthogonal.

In this paper we only consider errors due to camera noise. This admittedly is just one of the possible sources of errors identified in [5]. However, we believe that the theoretical treatment developed here provides a significant step forward. We identify optimal configurations both for the case of three light sources, and for the more complex case involving four or more lights.

2. Theory

In this Section, we identify light source configurations which guarantee that the scaled normal vector $\mathbf{b}$ computed by photometric stereo is least affected by errors in intensity measurements. After describing the basic assumptions and the problem formulation in Section 2.1, we derive the optimal configurations for three lights in Section 2.2, and then generalise the theory for four or more lights in Section 2.3.

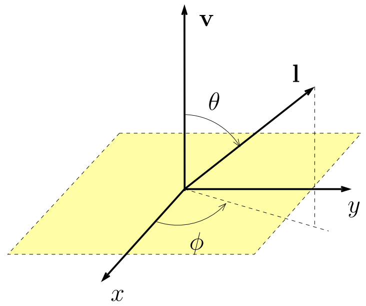

Figure 2: Slant and tilt of a light source $\mathbf{l}$: slant θ is the angle between $\mathbf{l}$ and the viewing vector $\mathbf{v}$ (coinciding with the camera axis). Tilt ϕ is the angle between the projection of $\mathbf{l}$ onto the camera plane (x–y) and the x-axis.

2.1. Assumptions and problem formulation

Slant and tilt. The direction of a light $\mathbf{l}$ is parametrised by slant and tilt defined as in Fig. 2. Slant is the angle between $\mathbf{l}$ and the vector $\mathbf{v}$ pointing towards the camera. Tilt is the angle which the projection of $\mathbf{l}$ onto the camera plane (x–y) makes with axis x. The viewing vector $\mathbf{v}$ is assumed to be vertical for notational convenience.

Camera noise. The camera sensor is noisy, with the noise being additive and normally distributed. Therefore, if the camera sensor outputs intensity i, it is the result of $i = \hat{\imath} + \Delta i$ with $\hat{\imath}$ being the true (noiseless) value and Δi being from $\mathcal{N}(0, \sigma^2)$ where $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution with the mean μ and variance σ^2.

Equal light intensities. To constrain the solution set for optimal configurations, we work under the assumption that the light sources are of equal intensities. To fix the intensity, we set $\kappa = 1$ for all light sources. To stress this assumption, we will use $\mathbf{L}$ to denote the light matrix instead of $\mathbf{S}$ (cf. (3)). Having n lights, the light matrix $\mathbf{L}$ is thus formed by columns $\mathbf{l}_k$ of unit L_2 norm:

$$\|\mathbf{l}_k\| = 1\,, \quad k = \{1, 2, \ldots, n\}\,. \tag{5}$$

Scaled normals' uncertainty. The camera noise (Δi) results in errors $\Delta\mathbf{b}$ occurring in the scaled normal estimates $\mathbf{b}$ ($\Delta\mathbf{b} = \mathbf{b} - \hat{\mathbf{b}}$ where $\hat{\mathbf{b}}$ is the ground truth of the scaled normal). The uncertainty $\epsilon(\mathbf{L})$ in scaled normal computation for a given light matrix $\mathbf{L}$ is defined as

$$\epsilon(\mathbf{L}) = E\left[\Delta\mathbf{b}^{\mathsf{T}}\Delta\mathbf{b}\right]\,, \tag{6}$$

where $E\left[\cdot\right]$ is expectation.

Problem formulation. For $n \geq 3$, find the light configuration $\mathbf{L}$ consisting of n light sources as in (5) such that the uncertainty in the scaled normal estimation computed using (6) is minimum.

2.2. Optimal configurations for $n = 3$ lights

For three light sources and with the assumptions formulated in the previous Section, Equation (4) is rewritten as

$$\mathbf{b} = \mathbf{L}^{-\mathsf{T}}\mathbf{i}\,. \tag{7}$$

Because this equation is linear, it follows that $\Delta\mathbf{b} = \mathbf{L}^{-\mathsf{T}}\Delta\mathbf{i}$, and for the uncertainty there holds (cf. (6))

$$\epsilon(\mathbf{L}) = E\left[\Delta\mathbf{i}^{\mathsf{T}}\mathbf{L}^{-1}\mathbf{L}^{-\mathsf{T}}\Delta\mathbf{i}\right]\,. \tag{8}$$

Because $E[(\Delta i_k)^2] = \sigma^2$ and as the noise in individual components of $\mathbf{i}$ is independent, we can see that this can be further rewritten to give

$$\epsilon(\mathbf{L}) = \sigma^2 \mathrm{trace}\,[\mathbf{L}^{-1}\mathbf{L}^{-\mathsf{T}}] = \sigma^2 \mathrm{trace}\,[(\mathbf{L}^{\mathsf{T}}\mathbf{L})^{-1}]\,. \tag{9}$$

Our task now will be to minimise the above, subject to the three constraints given by (5). Putting this down formally, the problem to be solved is the following:

Formulation 1 *Find the solution set* Ω^* *of all optimal light matrices* $\mathbf{L}^*$ *which satisfy*

$$\mathbf{L}^* = \underset{\mathbf{L}\in GL(3)}{\mathrm{minimise}}\ \sigma^2 \mathrm{trace}\,[(\mathbf{L}^{\mathsf{T}}\mathbf{L})^{-1}] \tag{10}$$

$$\textit{subject to}\ \mathbf{L}^{\mathsf{T}}\mathbf{L} = \begin{bmatrix} 1 & \cdot & \cdot \\ \cdot & 1 & \cdot \\ \cdot & \cdot & 1 \end{bmatrix}\,, \tag{11}$$

where $GL(3)$ is the group of 3×3 invertible matrices. In this formulation, Eq. (11) represents the same constraints as Eq. (5) in a form which better reveals the structure of the problem. Such formulation is not easily solvable though, because it seems to be non-trivial to combine constraints on elements of matrix $\mathbf{C} = \mathbf{L}^{\mathsf{T}}\mathbf{L}$ (in (11)) with optimisation in the inverse of that matrix (in (10)). We will try to bypass this problem as follows.

Formulation 2 (with weakened constraint) *Find the solution set* Ω^{**} *of all light matrices* $\mathbf{L}^{**}$ *such that*

$$\mathbf{L}^{**} = \underset{\mathbf{L}\in GL(3)}{\mathrm{minimise}}\ \sigma^2 \mathrm{trace}\,[(\mathbf{L}^{\mathsf{T}}\mathbf{L})^{-1}] \tag{12}$$

$$\textit{subject to}\ \mathrm{trace}\,[\mathbf{L}^{\mathsf{T}}\mathbf{L}] = 3\,, \tag{13}$$

and, subsequently, compute the intersection Ω *of this solution set with the set of matrices conforming to (11):*

$$\Omega = \Omega^{**} \cap \left\{ \mathbf{L} \,:\, \mathbf{L}^{\mathsf{T}}\mathbf{L} = \begin{bmatrix} 1 & \cdot & \cdot \\ \cdot & 1 & \cdot \\ \cdot & \cdot & 1 \end{bmatrix} \right\}\,. \tag{14}$$

Observation 1 (Relation of the two formulations.) The optimisation parts of the two formulation share the same objective function. They differ only in the solution candidate set. The candidate set of Formulation 2 is the superset of the candidate set of Formulation 1 because if (11) holds then (13) also holds. The optimum attained in the second case must therefore be lower or equal to the optimum of the original problem. If the intersection set Ω above is empty then it means that the problem with the modified constraint given by (13) attains a *strictly lower* minimum value compared with the original problem. In that case our effort would be wasted because solving the modified problem (Formulation 2) would not provide any useful information concerning the original problem. But if the intersection Ω is non-empty then the two problems attain the same minimum value and Ω represents *all* solutions for the original problem, i.e.

$$\Omega = \Omega^* \tag{15}$$

(for Ω see (14), for Ω^* see Formulation 1). In the following text we will show that this is indeed true. We solve the problem given by Formulation 2 and then observe not only that Ω is non-empty, but even that the solution sets of Formulations 1 and 2 are equal: $\Omega^* = \Omega^{**}$.

Before proceeding, let us present the following lemma which will help us to solve Formulation 2 easily.

Lemma 1 *Let* $P(3)$ *denote the set of all symmetric positive definite* 3×3 *matrices. The solution to the problem*

$$\mathbf{C}^* = \underset{\mathbf{C}\in P(3)}{\mathrm{minimise}}\ \sigma^2 \mathrm{trace}\,\mathbf{C}^{-1} \tag{16}$$

$$\textit{subject to}\ \mathrm{trace}\,\mathbf{C} = n\,, \tag{17}$$

is

$$\mathbf{C}^* = \mathrm{diag}\left[\frac{n}{3}, \frac{n}{3}, \frac{n}{3}\right]\,. \tag{18}$$

Proof. *The SVD decomposition of any symmetric positive definite matrix* $\mathbf{C}$ *from the solution candidate set will be*

$$\mathbf{C} = \mathbf{U}\,\mathrm{diag}[d_1, d_2, d_3]\,\mathbf{U}^{\mathsf{T}}\,, \tag{19}$$

where $\mathbf{U}$ *is an orthogonal* 3×3 *matrix and* d_k *are positive values. The SVD decomposition of the inverse* $\mathbf{C}^{-1}$ *is then*

$$\mathbf{C}^{-1} = \mathbf{U}\,\mathrm{diag}\left[\frac{1}{d_1}, \frac{1}{d_2}, \frac{1}{d_3}\right]\mathbf{U}^{\mathsf{T}}\,, \tag{20}$$

and the traces of the two matrices are

$$\mathrm{trace}\,\mathbf{C} = d_1 + d_2 + d_3\,, \tag{21}$$

$$\mathrm{trace}\,\mathbf{C}^{-1} = \frac{1}{d_1} + \frac{1}{d_2} + \frac{1}{d_3}\,. \tag{22}$$

Both latter equalities follow from the fact that the trace of a product of matrices is invariant under a cyclic permutation of the product:

$$\mathrm{trace}\,[\mathbf{AB}\ldots\mathbf{YZ}] = \mathrm{trace}\,[\mathbf{ZAB}\ldots\mathbf{Y}]\,. \quad (23)$$

Using these results, the modified optimisation function (obtained by including the constraint represented by (17) using a Lagrange multiplier λ) is

$$\sigma^2 \mathrm{trace}\,\mathbf{C}^{-1} + \lambda(\mathrm{trace}\,\mathbf{C} - n) =$$
$$= \sigma^2 \left(\frac{1}{d_1} + \frac{1}{d_2} + \frac{1}{d_3}\right) + \lambda(d_1 + d_2 + d_3 - n)\,. \quad (24)$$

The necessary condition of optimality is that all partial derivatives vanish, which leads to

$$\frac{\partial}{\partial d_k}: \qquad \sigma^2 \frac{1}{d_k^2} = \lambda, \;\; k = \{1,2,3\}, \quad (25)$$
$$\frac{\partial}{\partial \lambda}: \qquad \sum_{k=1}^{3} d_k = n\,. \quad (26)$$

From (25) it follows that all d_k^2 must be the same, and as d_k's are positive, all three, therefore, have to be equal. The second condition represented by (26) then fixes $d_k = n/3$ for all three. When this result is utilised in (19) together with the fact that $\mathbf{U}$ is orthogonal, (18) follows. □

Solution of Formulations 1 and 2. The optimisation in Formulation 2 is the same as the problem in Lemma 1 for $\mathbf{C} = \mathbf{L}^\top\mathbf{L}$ and $n = 3$. Therefore, we obtain $\mathbf{L}^\top\mathbf{L} = \mathrm{diag}[1,1,1]$ and, consequently, $\mathbf{L}$ is any triplet of orthogonal light vectors, or $\mathbf{L} \in O(3)$ where $O(3)$ is the group of orthogonal matrices. This solution indeed satisfies the constraint represented by (11), and hence this is also the solution to the original problem presented in Formulation 1. Utilising this result in (9), the minimum uncertainty attained is

$$\epsilon(\mathbf{L}) = 3\sigma^2\,. \quad (27)$$

2.3. Optimal configurations for $n \geq 3$ lights

The algebraic approach presented in Section 2.2 is easily extended to the case of more than three lights. Let us first briefly summarise how the governing equations differ from case $n = 3$.

– The basic structures are constructed as previously: the light matrix $\mathbf{L} = [\mathbf{l}_1, \mathbf{l}_2, .., \mathbf{l}_n]$ stores the lights column-wise again, and the intensity vector $\mathbf{i}$ is now an n-vector $\mathbf{i} = [i_1, i_2, .., i_n]^\top$.

– For $n > 3$ lights in a general configuration, the system of equations (3) for $\mathbf{b}$ is now obviously overdetermined and the scaled normal $\mathbf{b}$ must, therefore, be computed to minimise some predefined cost. A common choice is to require that the solution $\mathbf{b}$ must minimise the square of the relighting error

$$\epsilon_{rel} = \|\mathbf{L}^\top\mathbf{b} - \mathbf{i}\|^2\,, \quad (28)$$

which leads to the use of a pseudo-inverse (for a rectangular $k \times l$ matrix $\mathbf{A}$ of column rank l, the pseudo-inverse is $\mathbf{A}^+ = (\mathbf{A}^\top\mathbf{A})^{-1}\mathbf{A}^\top$). Here we have $\mathbf{A} = \mathbf{L}^\top$, and the normal is thus given by (cf. (7))

$$\mathbf{b} = (\mathbf{L}\mathbf{L}^\top)^{-1}\mathbf{L}\mathbf{i}. \quad (29)$$

– The definition of uncertainty represented by (6) is of course unchanged. The analogue for (8) is now

$$\begin{aligned}
\epsilon(\mathbf{L}) &= E\left[\Delta\mathbf{i}^\top\mathbf{L}^\top(\mathbf{L}\mathbf{L}^\top)^{-\top}(\mathbf{L}\mathbf{L}^\top)^{-1}\mathbf{L}\Delta\mathbf{i}\right] = \quad (30)\\
&\quad \text{(compute expectation)}\\
&= \sigma^2\mathrm{trace}\,[\mathbf{L}^\top(\mathbf{L}\mathbf{L}^\top)^{-\top}(\mathbf{L}\mathbf{L}^\top)^{-1}\mathbf{L}] = \quad (31)\\
&\quad \text{(shift one position to the left)}\\
&= \sigma^2\mathrm{trace}\,[(\mathbf{L}\mathbf{L}^\top)^{-\top}\underbrace{(\mathbf{L}\mathbf{L}^\top)^{-1}(\mathbf{L}\mathbf{L}^\top)}_{\text{identity}}] = \quad (32)\\
&= \sigma^2\mathrm{trace}\,[(\mathbf{L}\mathbf{L}^\top)^{-1}]\,. \quad (33)
\end{aligned}$$

In the above manipulation, (31) follows from (30) because of the independence of the individual components of $\Delta\mathbf{i}$ (cf. Eqs. (8) and (9)). Eq. (32) follows from (31) by the cyclic property of the trace (see (23)). Eq. (33) results straightforwardly by discarding the identity from the right-most part of the expression, and applying the fact that $\mathbf{L}\mathbf{L}^\top$ is a symmetric matrix.

We may easily check that the final expression is — for $n = 3$ — *exactly the same as in (9)*†.

– The analogue of (11) is now expressed in terms of an $n \times n$ matrix $\mathbf{L}^\top\mathbf{L}$,

$$\mathbf{L}^\top\mathbf{L} = \begin{bmatrix} 1 & \cdot & \cdots & \cdot \\ \cdot & 1 & & \cdot \\ \vdots & & \ddots & \vdots \\ \cdot & \cdot & \cdots & 1 \end{bmatrix}. \quad (34)$$

– The analogue of (13) now changes straightforwardly into

$$\mathrm{trace}\,[\mathbf{L}^\top\mathbf{L}] = n (= \mathrm{trace}\,[\mathbf{L}\mathbf{L}^\top])\,. \quad (35)$$

† A careful reader will notice that while here we have $\mathrm{trace}\,[(\mathbf{L}\mathbf{L}^\top)^{-1}]$, in (9) we use $\mathrm{trace}\,[(\mathbf{L}^\top\mathbf{L})^{-1}]$. But for $n = 3$ this *is* equal because the latter is $\mathrm{trace}\,[(\mathbf{L}^\top\mathbf{L})^{-1}] = \mathrm{trace}\,[\mathbf{L}^{-1}\mathbf{L}^{-\top}] = \mathrm{trace}\,[\mathbf{L}^{-\top}\mathbf{L}^{-1}] = \mathrm{trace}\,[(\mathbf{L}\mathbf{L}^\top)^{-1}]$.

In solving for optimal light configurations for $n > 3$, we will follow the same strategy as in the case $n = 3$. Therefore, instead of minimising (33) subject to (34) (which corresponds to Formulation 1), we will address the problem of minimising (33) subject to (35) which corresponds to Formulation 2. The solution to this problem is immediate, because it is formally analogous to the problem presented in Lemma 1 with $\mathbf{C} = \mathbf{L}\mathbf{L}^\mathsf{T}$ (and n equal to the number of lights involved)[‡]. As a result, we now have

$$\mathbf{L}\mathbf{L}^\mathsf{T} = \operatorname{diag}\left[\frac{n}{3}, \frac{n}{3}, \frac{n}{3}\right] . \tag{36}$$

What is required now is to determine which matrices satisfying (36) are also consistent with the equal intensity constraint represented by (34).

Solutions. To find a solution it suffices to find a $3 \times n$ matrix $\mathbf{L}$ whose columns have unit L_2 norms (due to equal light intensity constraint) and whose rows are mutually orthogonal and of equal L_2 norm (due to (36); note that requiring this norm to be $\sqrt{n/3}$ would be redundant). Note that all these constraints are invariant to transforming $\mathbf{L}$ by an arbitrary orthogonal transformation. This means that if some $\mathbf{L}$ is an optimal light configuration then any global rotation/reflection of the lights is also an optimal solution. As for the minimum uncertainty attained, we now have

$$\epsilon(\mathbf{L}) = \frac{9}{n}\sigma^2 , \tag{37}$$

which (of course) for $n = 3$ reduces to (27).

Observation 2 (constant slant solution.) Let the number of lights be $n \geq 3$. The following light matrices are members of the solution set:

$$\mathbf{L}_F(n) = \sqrt{\frac{2}{3}} \times$$
$$\times \begin{bmatrix} \sin\frac{2\pi}{n}1 & \sin\frac{2\pi}{n}2 & \dots & \sin\frac{2\pi}{n}(n-1) & \sin 2\pi \\ \cos\frac{2\pi}{n}1 & \cos\frac{2\pi}{n}2 & \dots & \cos\frac{2\pi}{n}(n-1) & \cos 2\pi \\ \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} & \dots & \frac{1}{\sqrt{2}} & \frac{1}{\sqrt{2}} \end{bmatrix} \tag{38}$$

which is justified as follows. The squared L_2 norm of each column is obviously 1 because the term involved is always $2/3(\sin^2 + \cos^2 + 1/2)$. Orthogonality of rows follows from the fact that the rows are actually members of an n-point Fourier basis (which is known to be orthogonal). It can also be verified that all rows have the same norm.

This observation tells us that light sources equally separated on a circle of uniform slant represent one of optimal configurations. The optimal slant is independent of n and is equal to $\operatorname{atan}\sqrt{2}$ ($\sim 54.74°$).

Observation 3 (combining optimal solutions.) Let the number of lights be n, and let $\{k_i, i = 1, 2, \dots, I\}$ be the set of integers such that $\{k_i\colon k_i \geq 3, \sum_{i=1}^{I} k_i = n\}$. Let matrix $\mathbf{L}_i(k_i)$ be any optimal light configuration for k_i lights. Then the concatenation of these matrices is an optimal solution for n lights:

$$\mathbf{L}(n) = [\mathbf{L}_1(k_1)|\mathbf{L}_2(k_2)|\dots|\mathbf{L}_I(k_I)] , \tag{39}$$

where '|' denotes concatenation. This results immediately from the fact that the rows in any $\mathbf{L}_i(k_i)$ are mutually orthogonal (hence their concatenations are orthogonal as well). They are also of the same norm, and hence the concatenated rows likewise. And, all columns in all $\mathbf{L}_i(k_i)$'s have unit length.

As an example, for $n = 6$ lights, one of the optimal configurations is given by (38). But what the current Observation tells us is that optimal configurations can also be constructed from the concatenation of *any* optimal configurations for $3 + 3$ lights. Any pair of orthogonal triplets is, therefore, also an optimal light configuration for six lights[§].

Observation 4 (adding one vertical light) Let $n \geq 4$. If we take the light matrix $\mathbf{L}_F(n-1)$ for $n-1$ lights as in (38), adjust the scale of the rows, and concatenate the result with a vertical vector $[0, 0, 1]^\mathsf{T}$, we obtain an optimal configuration for n lights. More precisely,

$$\mathbf{L} = \left(\begin{bmatrix} \sqrt{\frac{n}{n-1}} & 0 & 0 \\ 0 & \sqrt{\frac{n}{n-1}} & 0 \\ 0 & 0 & \sqrt{\frac{n-3}{n-1}} \end{bmatrix} \mathbf{L}_F(n-1) \middle| \begin{matrix} 0 \\ 0 \\ 1 \end{matrix} \right) \tag{40}$$

is a member of the solution set for n lights. This corresponds to placing $n - 1$ lights equidistantly on a constant slant circle, and putting the last light into the vertical direction. The slant for the $n - 1$ lights is dependent on n (see Fig. 3). Obviously, for $n \to \infty$ the optimal slant would attain $\sim 54.74°$ again, as in that case one vertical light would be just an 'unimportant perturbation'.

3. Summary, and Conclusion

This paper has presented important new theory that allows the optimum lighting configurations for photometric stereo

[‡]Note that although we have $n > 3$, the matrix $\mathbf{C} = \mathbf{L}\mathbf{L}^\mathsf{T}$ is again 3×3, hence in an SVD parametrisation of $\mathbf{L}\mathbf{L}^\mathsf{T}$ in (19) we still have 3 diagonal values and a 3×3 orthogonal matrix $\mathbf{U}$.

[§]Let us recall that the optimality constraints are invariant under any orthogonal transformation of a light configuration. This means that light vectors within any $\mathbf{L}_i(k_i)$ can be freely rotated/reflected by a common orthogonal transformation. For the case $n = 6$ discussed, one can even have two *equal*, arbitrarily oriented light triplets.

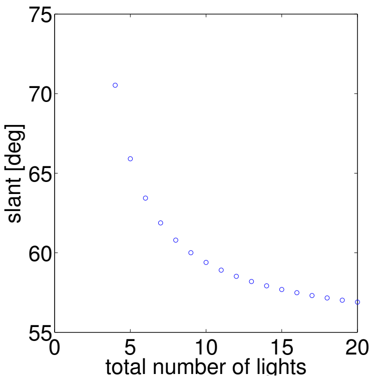

Figure 3: The dependence of optimal slant for configuration of $n \geq 4$ lights, one of which is vertical and the others residing on the constant slant circle.

to be determined in the presence of camera noise. The theory was developed initially for the three-source case and then extended to cover four or more lights. Furthermore, we have used this novel theory to establish a number of results.

We have shown that the optimal light configuration for $n = 3$ lights is an arbitrary orthogonal triplet. Specifically this includes the constant slant solution (of $\sim 54.74°$) derived in [6] using repeated numerical optimisation and empirical observation.

Interestingly enough, we showed that for $n > 3$ light sources, this slant is optimal *as well*, with light sources spaced equally in tilt by $360/n$ degrees (Observation 2). We showed that in the case of $n > 3$ lights, the optimal light configurations are numerous, including configurations with $n-1$ lights with a constant slant plus one vertical light (Observation 4). Observation 3 stressed out that optimal configurations for n lights can be constructed from optimal configurations for lesser lights by concatenation.

The question of optimal light configuration is fundamental to photometric stereo. It is therefore surprising that it has not received more attention. To our knowledge this paper is the first to provide a concise theoretical explanation on how $n \geq 3$ light sources should be placed such that the surface properties are most precisely reconstructed in the presence of camera noise.

It would be interesting to adapt this theory for other sources of error.

Acknowledgements

O. Drbohlav is supported by the Marie Curie Intra-European Fellowship No. 506053 (PhoCal) within the 6th European Community Framework Programme. M. Chantler acknowledges the support by EPSRC grant GR/S12395/01. We wish to thank Andy Spence from the Texture Lab for sharing his expertise in the optimal configuration problem with us. We are grateful to Radim Šára from the Center for Machine Perception (Czech Technical University in Prague) for pointing us to most relevant prior art. We also thank the reviewers for their recommendations regarding the future work.

References

[1] D. Bezděk. Fotometrické stereo (Photometric Stereo). Master's thesis, Charles University, Faculty of Mathematics and Physics, Prague, Czech Republic, 1998. Supervised by R. Šára. *In Czech language.*

[2] X. Y. Jiang and H. Bunke. On error analysis for surface normals determined by photometric stereo. *Signal Processing*, 23(3):221–226, 1991.

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

[4] K. M. Lee and C.-C. J. Kuo. Surface reconstruction from photometric stereo images. *Journal of Optical Society of America A*, 10(5):855–868, 1993.

[5] R. Ray, J. Birk, and R. B. Kelley. Error analysis of surface normals determined by radiometry. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 5(6):631–645, 1983.

[6] A. D. Spence and M. J. Chantler. Optimal illumination for three-image photometric stereo acquisition of texture. In *Texture 2003 : Proceedings of the 3rd international workshop on texture analysis and synthesis*, pages 89–94, 2003. Available at http://www.macs.hw.ac.uk/~mjc/texture2003/camera-ready/cam-rdy-ab024.pdf.

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