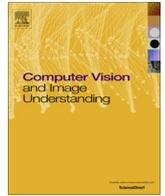




Contents lists available at ScienceDirect

Computer Vision and Image Understanding

journal homepage: www.elsevier.com/locate/cviu

Optimal illumination directions for faces and rough surfaces for single and multiple light imaging using class-specific prior knowledge [☆]

Vasileios Argyriou ^{a,*}, Stefanos Zafeiriou ^b, Maria Petrou ^b

^a Kingston University, London, United Kingdom

^b Imperial College, London, United Kingdom

ARTICLE INFO

Article history:

Received 1 November 2012

Accepted 31 January 2014

Available online xxxx

Keywords:

Photometric stereo

3D imaging

Face reconstruction

ABSTRACT

The detection of image detail variation due to changes in illumination direction is a key issue in 3D shape and texture analysis. In this paper two approaches for estimating the optimal illumination direction for maximum enhancement of image detail and maximum suppression of shadows and highlights are presented. The methods are applicable both to single image/single illumination direction imaging and to photometric stereo imaging. This paper uses class-specific prior knowledge, where the distribution of the normals of the class of surfaces is used in the optimisation. Both the Lambertian and the Phong models are considered and the theoretical development is demonstrated with experimental results for both models. For each method experiments were performed using artificial images with isotropic and anisotropic distributions of normals, followed by experiments with real faces but synthesised images. Finally, results are presented using real objects and faces with and without ground-truth.

© 2014 Elsevier Inc. All rights reserved.

1. Introduction

The variation of the intensities observed in images depends on variation in both surface reflectance and surface relief. While the reflectance properties are intrinsic to a surface, the surface relief produces a pattern of shadings that depends strongly on the direction of illumination. The appearance of a 3D surface changes drastically with illumination [8]. Different image details are enhanced for different illumination directions. The idea of photometric stereo is to use this information to recover the intrinsic surface parameters, that is, local surface orientation and albedo, independent of the illumination direction. Photometric stereo estimates relief and reflectance information using three or more images of a surface illuminated from different directions captured from a single viewpoint. The method considers the isophotes in gradient space. The intersections of the isophotes from different illumination directions identify the gradients of individual surface facets (corresponding to pixels). Due to the non-linear nature of the problem, the shape, density and orientation of the isophotes affect the accuracy with which the gradient vectors can be estimated. The characteristics of the isophotes depend on the lighting directions.

So, illumination directions play a crucial role in the quality of the produced image and they should be carefully chosen for applications in which this is possible. Sub-optimal geometric arrangements may crucially affect the reliability of the subsequently inferred information.

In this paper, two methods that estimate the optimal imaging configuration are proposed. The first one is based on maximising the level of detail in the reconstructed surfaces, revealing details and salient features of the imaged surface. In the second approach, we estimate optimal positions of the light sources by maximising the area covered by dense isophotes in gradient space for an arbitrarily shaped surface when the distribution of surface normals is roughly known. The problem has been studied before [32,15], but previous authors did not consider prior knowledge on the distribution of the normals of the facets of the surface. All previous optimisation methods effectively assumed uniform distribution of the normals. This paper uses class-specific prior knowledge, where the distribution of the normals of the class of surfaces is used in the optimisation. At this point we would like to stress that we do not expect to know the exact surface normals beforehand, as this would defeat the purpose of photometric stereo. We only assume the availability of the *statistical distribution* of the surface normals, which may be used for all objects of the same class, e.g. all human faces, although this statistical distribution might be obtained from the face of a plastic model, for example, using active sensing. The ultimate purpose is to take into consideration the statistical distribution of the normals, in order to design an optimal lighting

[☆] This paper has been recommended for acceptance by Edwin R. Hancock.

* Corresponding author. Address: Faculty of Science, Engineering and Computing, Kingston University, Penrhyn Road, Kingston upon Thames, Surrey KT1 2EE, UK. Tel.: +44 0 2084172591; fax: +44 0 2084172972.

E-mail address: Vasileios.Argyriou@kingston.ac.uk (V. Argyriou).

system for this class of object, and so achieve best reconstruction of the surfaces of *individual* objects of the same class.

This paper is organised as follows. In Section 2 previous work on photometric stereo and optimal illumination estimation is reviewed. In Section 3, we propose an overview of the proposed approach. In Section 4 the first proposed methodology is analysed both for the Lambertian and the Phong models. Also, details on how this approach can be generalised to n illumination sources are presented. In Section 5 the second approach is introduced based on the density of the isophotes in the gradient space, for Lambertian surfaces. In Section 6 initially an analysis of some preliminary results related to the estimations of some model parameters are presented for the first approach. Then in Section 7, the proposed illumination setups are compared with the one proposed in [10]. First, experiments are performed using artificial images with isotropic distribution of normals, followed by experiments with anisotropic distributions with four different dominant orientations. Subsequently, experiments with real faces but synthesised images are performed. Finally, in Section 8 results are presented using real objects with ground-truth and real faces without ground-truth but using the side views and the Hausdorff distance for evaluation. In Section 9 conclusions on the proposed methodology and the evaluation process are presented.

2. Previous work

In [41] photometric stereo was introduced. He proposed a method which was simple and efficient, but was sensitive to noise. In his method, the surface gradient can be recovered by using two photometric images, assuming that the surface albedo is already known for each point on the surface. In [9] photometric stereo was extended to four light sources, where specular reflections were discarded and estimation of surface shape could be performed by means of diffuse reflections and the use of the Lambertian model. In [27] a photometric approach which uses a linear combination of the Lambertian model was developed and an impulse specular component to obtain the shape and reflectance information for a surface. In [3–5] an algorithm for estimating the local surface gradient and albedo using four source colour photometric stereo was presented in the presence of highlights and shadows. In [22] an approach that utilises nine illumination sources was presented. It is also worth mentioning the related work presented in [31] focusing on 4-lights PS based on shape information and statistical segmentation techniques to determine which pixels are specular and which are non-specular. In [7] an algorithm for Lambertian photometric stereo in the presence of shadows is proposed based on fast graph cuts estimating per pixel light source visibility. Also, it allows images to be acquired with multiple illuminants, and there can be fewer images than light sources. In [24] a method to remove shadows from real images based on a probability shadow map is introduced and in [12] a shadow removal method is presented from a 3-band colour image finding an intrinsic reflectivity image based on assumptions of Lambertian reflectance. In the work presented in [30] a method for cast shadow removal from obliquely illuminated images of faces is suggested based on a statistical model of surface normal directions. [34] proposed a multi-light source photometric stereo system for reconstructing images of various characteristics of non-Lambertian rough surfaces with widely varying texture and specularity. An algorithm that no calibration is needed for recovery of geometry of objects with general reflectance properties from images was introduced in [17]. In [1] a technique for resolving the GBR ambiguity based on minimisation of the entropy of the recovered albedos was proposed. In [13] an uncalibrated photometric stereo technique for unknown light sources and general reflectance model was suggested. A

review of different reflectance maps proposed in the literature for modelling reflection from real-world surfaces as presented in [35] and in [28] the problem of estimating the proportions of Lambertian and specular reflection components in order to improve the quality of surface normal information recoverable using shape-from-shading was discussed. Recursive approaches for shadow and highlights estimation for advanced reconstruction using photometric stereo of any number of lights were presented in [38,39]. In [18] a technique for face recognition was presented that combines the Fisherface method with the ridgelet transform and high-speed Photometric Stereo. In [21] was shown that a plausible shape can be obtained based on two light sources and in [11] a multiple illumination technique that directly recovers the viewer-centered curvature matrix and being independent of knowledge of incident illumination orientation, local surface orientation, or diffuse surface albedo was presented. A scheme to resolve handwriting from background printing using photometric stereo to recover the surface was suggested in [25] and in [19] a photometric stereo algorithm was presented that reconstructs object shapes from multiple images, in which 3D surfaces were approximated by Legendre polynomials. In [37] an algorithm for shape from shading was introduced based on the assumption that a single input image will be matched to a second image through a uniform disparity field.

The influence of a lighting arrangement to the accuracy of surface reconstruction based on photometric stereo has been considered in the past [41,23,32,8] and suggestions for optimal illumination configurations in terms of azimuth ϕ_i and zenith θ_i angles have been reported. In [40], using reflectance maps, dense iso-intensity contours were recommend to obtain maximal accuracy.

Lee and Kuo in [23], using two reflectance maps in the case of a two image photometric stereo, deduced that it is desirable to incorporate reflectance maps that compensate each other's weaknesses (i.e. the accuracy of the reconstructed surface height is related to the slope of the reflectance map function and therefore using a second reflectance map can be combined to improve the reconstruction in certain areas), in order to determine the optimal illumination configuration.

Based on [23], it was further confirmed that the two image photometric stereo is more sensitive to the azimuth rather than the zenith angle difference and that the optimal value for azimuth difference is 90° . In [15] three light photometric stereo was considered and it was suggested that distributing the illumination azimuth angles uniformly through 360° is optimal. A theoretical analysis of Gullon's arrangement was presented in [32,33], based on the sensitivity analysis of photometric stereo and by deriving expressions of each surface normal vector with respect to image intensities.

It was discovered in [32] that the optimal azimuth and zenith angles cannot be specified without any assumptions and that the configuration that results the minimum noise is not unique. Instead, it was determined for a two lights Photometric Stereo that an orthogonal arrangement of the illumination vectors (with an angle of 90° to each other) is the only restriction to obtain the optimal configuration. In the case of the common zenith angle being constrained, the optimal values for azimuth angles for three lights Photometric Stereo were estimated and it was suggested to use 120° difference in a three-image Lambertian photometric stereo configuration.

The optimal zenith angle in the case of uniformly distributed light sources according to the azimuth was found to be around 55° , but if shadows are present the angle should be reduced [32,15]. On the contrary, if the surface is smooth and shadows are not an issue, the zenith angle can be increased. Furthermore, in [10] the above was extended for n light sources and derived

the same value for the optimal zenith angle when the sources are equally spaced in azimuth angles of $360/n$ degrees. It is also worth mentioning the similar methods presented in [14,16,29,20,42].

3. Proposed methodologies – preliminaries

It is well known that the fraction of light reflected on an object's surface in a certain direction depends upon the optical properties of the surface material. In this paper we use the Lambertian and Phong models, thus the fraction of the incident illumination reflected in a particular direction depends only on the surface normals. It has been proven in [10] that it is not possible to define the optimal imaging conditions without some prior knowledge about the imaged surface. In this work, we propose two novel approaches that separate the observed objects into classes which share common features and characteristics, such as similar distribution of the normals, and finds the optimal illumination conditions for each class of object (e.g. faces, bottles, round fruits, etc.).

If the purpose of imaging is to reveal the 3D structure of the surface, the assumption that the surface consists of a collection of facets, each with its own orientation and possibly albedo, has already been made. It is not possible then to select an optimal imaging configuration by considering either the characteristics of a single surface facet or the average characteristics of the full surface. Instead, the optimal configuration is selected by considering the *statistical distribution* of the normal vectors of the facets that make up the surface. The application of the proposed methodology is clearly for controlled environments only. For example, in an industrial inspection problem or a 3D face database creation problem, it is expected to know a priori what type of surface will be imaged. Then, from generic graphics models of surfaces of this kind, it may be assumed that the distribution of the normals of the surface is approximately known. Therefore, the proposed model is not parametric, but the required distribution of surface normals could be extracted from a generic object model of the type of object class we are interested in imaging.

At this point the relationship between the imaging resolution and the obtained surface normals should be analysed in more detail. The majority of the 3D reconstruction methods based on photometric stereo have a one-to-one relationship between the surface normals and the resolution of the acquired image. This means that each pixel corresponds to a facet with one normal. Furthermore, methods that provide subpixel reconstructions using photometric stereo are available [36] allowing lower resolution capturing devices to provide detailed estimates. It should be mentioned that this one-to-one relationship due to the lack of the correspondence problem is one of the main advantages of photometric stereo, allowing detailed reconstructions with applications to object inspection for defects, surveillance and recognition.

So, it is proposed that in case of objects (e.g. faces) that share common surface characteristics, which therefore result in similar statistical distributions of normal vectors, an optimal arrangement of the illumination directions may be obtained. The 3D surface of any face (even one of a plastic model scanned by a laser sensor) can supply us with the required distribution, as all faces are expected to have pretty similar distributions of the normals of their facets. This approach can be extended to any class of surface, not only faces, as long as they share similar facet normal distributions. With the proposed approach the optimal illumination directions are estimated once and then the obtained setup is used to capture other objects of the same class. If no prior knowledge on the distribution of the facets is available, a recursive approach may be used: a preliminary surface estimate may be obtained using a default light setting; from this estimate the optimal light setting may be

worked out and the surface may be reconstructed more accurately using images captured under the new lighting arrangement. In the sections that follow, two approaches are proposed based on the above assumptions. The first one estimates the optimal illumination directions operating in the pixel domain while the second one reasons about the isophotes in the gradient domain.

4. Method 1 – Optimal illumination setup in the pixel domain

The approach that operates in the pixel domain provides the optimal illumination setup both for a Lambertian and a Phong model. The solution is obtained by optimising an objective function consisting of three components that express three different factors that influence the performance of Photometric Stereo. The first component tries to maximise and enhance the visibility of the surface details; the second component tries to minimise the shadows; and the last one, that is applicable only to the Phong model, tries to determine the illumination setup that minimises the number of highlights (see Fig. 10). Regarding the shadow component it should be mentioned that only self-shadows were considered in the proposed model while cast shadows are not modelled in our simulations.

4.1. Lambertian model

If N_i is the normal vector of a surface facet, ρ_i its albedo and I_i the corresponding brightness value recorded for that facet, for the Lambertian model we have

$$I_i(\vec{L}) = \rho_i \vec{N}_i \cdot \vec{L} \quad (1)$$

By incorporating the albedo into the normal vector, a generalised normal $\vec{N}'_i \equiv \rho_i \vec{N}_i$ may be defined for each facet, so that

$$I_i(\vec{L}) = \vec{N}'_i \cdot \vec{L} \quad (2)$$

Treating \vec{N}' as a random variable, its distribution may be assumed known, constituting the prior knowledge we have for the problem, obtained from a training surface of similar type. In case the actual albedo of the generic model surface used is not available, it can be replaced by a uniform one without affecting the proposed methodology.

Any two neighbouring facets i and j will be recorded with brightness contrast

$$\Delta I_{ij}(\vec{L}) = (\vec{N}'_i - \vec{N}'_j) \cdot \vec{L} \equiv \Delta \vec{N}'_{ij} \cdot \vec{L} \quad (3)$$

Over an imaged surface, $\Delta \vec{N}'_{ij}$ is expected to take values from the estimated distribution of the ground truth (e.g. known class sample) as ij traces all neighbouring facets of the surface. L is the

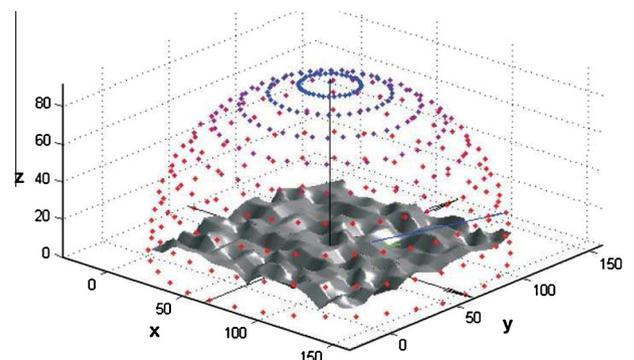


Fig. 1. Examples of possible illumination directions. Each dot on the hemisphere corresponds to a possible light l . It may be identified by its azimuth angle ϕ_l and its zenith angle θ_l .

unknown illumination vector and in our problem it is constrained on a hemisphere with radius 1 (see Fig. 1), thus its position is parameterized with two angles φ_l and θ_l .

The idea then is to select \vec{L} so that the spread of ΔI_{ij} is maximised, i.e. the detail seen in the image is maximised. If the distribution of $\Delta \vec{N}_{ij}$ is known, it is obvious that \vec{L} could be estimated using a maximisation algorithm. So, the illumination direction could be estimated by maximising the sum of the magnitudes of the right-hand side of Eq. (3) for all neighbouring facets, integrated over the surface normals, given by

$$f(\varphi_l, \theta_l) \equiv \frac{1}{S} \sum_{i, \text{ and } j \in R_i} |\Delta \vec{N}_{ij} \cdot \vec{L}| \quad (4)$$

where S is the total number of pairs (i, j) considered and R_i is the 4-neighbourhood of facet i .

If this function were to be minimised, however, we might end up with a large number of facets receiving no light (self-shadowed). This is undesirable in a photometric stereo application. To avoid self-shadows, the angle between the illumination direction \vec{L} and the surface normal \vec{N} in Eq. (1) should be in the range $(-\frac{\pi}{2}, \frac{\pi}{2})$. So, a constraint is introduced to constrain the angles in the desirable range. We define:

$$g(\varphi_l, \theta_l) \equiv \frac{1}{M} \sum_i^M W_i(\varphi_l, \theta_l) \quad (5)$$

where $W_i = 1$ if $N_i \cdot L > 0$ and $W_i = -1$ if $N_i \cdot L < 0$, and M is the number of surface normals. So, instead of maximising the number of facets with $N_i \cdot L > 0$, the difference of the number of facets with $N_i \cdot L > 0$ minus the number of facets with $N_i \cdot L < 0$ is maximised. Since the number of all facets is constant, it makes no difference either only the number of facets with positive W_i , or the difference of the two types of facet is maximised. The case of a $\pi/2$ angle is treated as neutral. Combining the two requirements (4) and (5) in a single quality function, we obtain

$$Q(\varphi_l, \theta_l) \equiv f(\varphi_l, \theta_l) + \lambda g(\varphi_l, \theta_l) \quad (6)$$

where $\lambda > 0$ is a weighting factor. Simulated annealing may be used to solve this global optimisation problem and obtain the optimal direction (φ_l, θ_l) for the light source. In more details, the obtained quality function in Eq. (6) is the input to the simulated annealing and the optimal values for the (φ_l) and (θ_l) parameters are obtained. In this way the distribution of the normals $\Delta \vec{N}$ and the self-shadow component are incorporated in the optimisation process.

4.2. Phong model

The Phong model consists of three components, namely the diffuse, specular and ambient components (Fig. 2). The diffuse component is the Lambertian model. The specular component forms a lobe of reflected light which spreads out around the specular direction and is modelled by a cosine function raised to a power. The ambient component accounts for the ambient light and the inter-reflections. This model is mathematically expressed as

$$I(\varphi_l, \theta_l) = I_x + \rho \vec{L} \cdot \vec{N} + k(\vec{R} \cdot \vec{V})^m \quad (7)$$

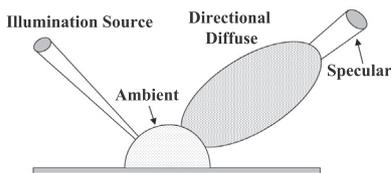


Fig. 2. Components of the Phong model.

where I_x is the ambient constant light intensity and k is the specular reflectance coefficient. \vec{V} denotes the viewing direction, \vec{R} is the perfect reflector vector, and m controls the width of the specular lobe and it was equal to 64 in our simulations. One can easily work out that direction \vec{R} is given by $\vec{R} = 2\vec{N}(\vec{N} \cdot \vec{L}) - \vec{L}$.

For our study the ambient light may be ignored, since it has the same constant value for all surface facets. Thus, the model of interest here consists of the Lambertian component and the highlight only. Since the Lambertian component is uniformly radiated, the position of the camera is not influenced by this. The Lambertian component is only influenced by the illumination direction, which may still be determined by the method discussed in Section 4.1. The position of the camera, however, should be such that the number of the highlights is minimised. Highlights are a nuisance when one wants to reconstruct a surface from photometric data, since the intensity saturates at the corresponding facets.

Since we must avoid highlights, the camera direction $\vec{V} \equiv (\varphi_v, \theta_v)$ should be selected so

$$t(\varphi_l, \theta_l) \equiv \frac{1}{M} \sum_i^M \vec{R}_i \cdot \vec{V} \quad (8)$$

is minimised. Combining (6) and (8) we obtain as quality function

$$P(\varphi_l, \theta_l) \equiv f(\varphi_l, \theta_l) + \lambda g(\varphi_l, \theta_l) - \kappa t(\varphi_l, \theta_l) \quad (9)$$

where $\kappa > 0$ is a weighting factor. Note the minus sign in front of κ , indicating that this term has to be minimised, while the other two have to be maximised. Simulated annealing may be used to solve this global optimisation problem and obtain the optimal illumination (φ_l, θ_l) and viewing (φ_v, θ_v) directions. In the Phong model case the approach is similar as in the Lambertian case but the quality function in Eq. (9) that is the input at simulated annealing contains a third component that incorporates the viewers direction in the optimisation process. Other stochastic optimisation methods may equally well be used. In our work, the directions proposed in [10] were used for initialisation.

In photometric stereo the main problem is to reconstruct the surface normals at the areas where shadows or highlights are present. Also, if the contrast is low, small irregularities, e.g. moles of a face, would not be revealed. In Eq. (9) the first term f is used to obtain the illumination direction that maximises the contrast, the second term to minimise the facets with shadows and the last one to minimise the highlights. The first and the second terms could not be combined since the contrast maximisation could be benefited by the presence of shadows and highlights.

4.3. Extension to n illumination sources

The extension to n illumination sources for use in conjunction with photometric stereo is performed using alternating optimisation. According to this approach, the optimal light source direction of the first light is estimated using (6) or (9). During this first optimisation, the remaining illumination directions are kept constant. After the first illumination direction is obtained, the second is estimated in the same way but keeping the values estimated for the first light fixed. The remaining lights are estimated with the same process. After the directions of all light sources have been computed, the whole process is repeated until convergence is reached.

In the case of the Phong model, the viewing direction is defined only once, when the direction of the last light is worked out, since there is only a single acquisition point.

The optimisation process is performed once for each type of surface under inspection, off line, and therefore there are no issues related to the complexity and the required time for this step.

5. Method 2 – Optimal illumination setup using isophotes in the gradient space

The second proposed approach to obtain the optimal illumination directions for a specific class of object (e.g. faces) that share similar distribution of normals is based on the fact that denser isophotes in the gradient space result more accurate normals during the reconstruction process. Knowing the class distribution of normals we try to estimate the illumination directions that will provide the denser isophotes to the areas with the most common and frequent normals since for most object classes a uniform distribution of normals is not the case. This idea is somehow equivalent to the idea we exploited in the previous section, since denser isophotes result from large changes in brightness between neighbouring pixels. The difference is that in that case, the term “neighbouring pixels” indicated proximity in the physical space, while here indicates proximity in the gradient space.

The reflectance characteristics of an object that are produced when a ray of light strikes its surface, (specular and diffuse) depend on the surface material, microstructure, incident wavelength, and the direction of incidence. Assuming a Lambertian surface and a viewing direction that is aligned with the z-axis of the object coordinate system, the normal vector of the surface is $(-p_n, -q_n, 1)$, where $p_n = \partial x / \partial z$ and $q_n = \partial y / \partial z$. If the corresponding illumination vector is $(-p_l, -q_l, 1)$, the image intensity is calculated by taking the normalised dot product of the two vectors and multiplying it with the surface albedo $\rho(x, y)$:

$$R(x, y) = \rho(x, y) \frac{1 + p_l p_n + q_l q_n}{\sqrt{1 + p_n^2 + q_n^2} \sqrt{1 + p_l^2 + q_l^2}} \quad (10)$$

If two images with different lighting directions are used, the orientation of each surface facet may be determined by the intersection of the iso-brightness contours (isophotes) of the two reflectance functions it lies on.

In general, the reflectance functions are non-linear, so more than one solutions are possible (contours may intersect at two points; twofold ambiguity). A third reflectance map, for a different illumination direction, is needed to resolve the ambiguity. The methodologies proposed in [10,15,23,32,33] for maximising the accuracy of isophote intersection determination, are based on the assumptions of isotropic and homogenous surfaces with uniform distribution of the (p_n, q_n) values of the surface facets in the reflectance map, resulting in the conclusion that the light sources should be symmetrically arranged in azimuth around the surface. These assumptions may be relaxed when some prior knowledge on the acquired surface is available.

Based on the fact that photometric stereo is most accurate in the regions of gradient space where the density of reflectance functions is high [41], a method that estimates the optimal illumination arrangement is proposed. This is self explanatory, since a reflectance map with higher density increases the probability for three or more isophotes to cross through a single point, providing an accurate solution for the system. Assuming a given surface material, ie fixed $\rho(x, y)$, and point light sources, the angle between the viewer and the light source is one of the main parameters that affect this accuracy. In this case, the density of reflectance map contours increases for larger zenith angles, which at the same time increases the shadowed areas. So, a compromise should be reached. A second determiner of the accuracy is the choice of the relative directions of the light sources. Assuming that a rough model of the observed object is known (e.g. an average face in case of 3D face reconstruction applications), and therefore the distribution of its facets' orientations is known, the illumination angles may be adjusted in such a way, that areas with dense

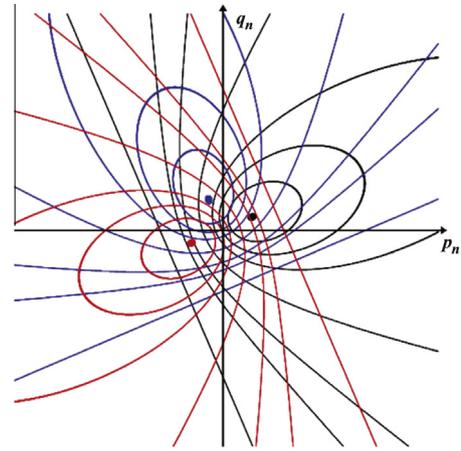


Fig. 3. Isophotes created from three illumination directions.

surface normals (p_n, q_n) will also be covered densely by isophotes (see Fig. 3).

Let us assume an illumination direction

$$\begin{aligned} \vec{L} = (x_l, y_l, z_l) &= (\sin \theta_l \cos \phi_l, \sin \theta_l \sin \phi_l, \cos \theta_l) \\ &= \frac{1}{\sqrt{p_l^2 + q_l^2 + 1}} (-p_l, -q_l, 1) \end{aligned} \quad (11)$$

where $\cos \theta_l = \frac{1}{\sqrt{p_l^2 + q_l^2 + 1}}$, $p_l = \cos \phi_l \tan \theta_l$ and $q_l = \sin \phi_l \tan \theta_l$. So, for a given surface normal $\vec{N} = \frac{1}{\sqrt{p_n^2 + q_n^2 + 1}} (-p_n, -q_n, 1)$ the intensity of this facet, given by Eq. (10), depends only on the azimuth ϕ_l and zenith θ_l angles (and the albedo).

If we assume that the (p_n, q_n) values are uniformly distributed in a square area of the reflectance map, corresponding to the case of an isotropic and uniformly rough surface, the density of the isophotes should be uniform too in the same area of the (p, q) space, in order to achieve a uniformly reliable recovery of all the surface normals. However for faces or other anisotropic surfaces, the lights should be arranged in such a way that the areas with more (p_n, q_n) points should be covered by more isophotes. So, the directions should be selected to satisfy the above requirement and they may be estimated using the following optimisation methodology.

Without any loss of generality, let us assume that our object class of interest is human faces. Let us consider that we have at our disposal the model of a 3D human face, obtained from a mannequin. This 3D surface represents the average face containing the most common facial characteristics and features.

Let us assume that the model face consists of M vertices and let an $M \times 2$ matrix \mathbf{W} contain the directions of the vertices' normals. The third component of the normal is discarded as being fully defined by the other two, due to the choice of the coordinate system we use.

So, $\mathbf{W}_i = (p_{ni}, q_{ni})$, where $p_{ni} = dx_i / dz_i$ and $q_{ni} = dy_i / dz_i$. The normals are quantised and their histogram $\mathbf{W}_h = h(\mathbf{W})$ is obtained containing m discrete entries keeping their number equal to the number of vertices (i.e. pixels on the 2D height map, with one axis for q and one for p). Some examples of \mathbf{W}_h are shown in Fig. 4.

Since the normals are discretized on the (p, q) plane, their range is known with the vectors \mathbf{P} and \mathbf{Q} containing all possible values in each axis (i.e. the range of the normals' values). For all these discrete normals, the reflectance map for a given illumination direction (p_l, q_l) may be estimated using Eq. (10) and be represented as:

$$\mathbf{R}_{ij}^{(k)} = \rho \frac{1 + p_{lk} \mathbf{P}_i + q_{lk} \mathbf{Q}_j}{\sqrt{1 + \mathbf{P}_i^2 + \mathbf{Q}_j^2} \sqrt{1 + p_{lk}^2 + q_{lk}^2}}$$

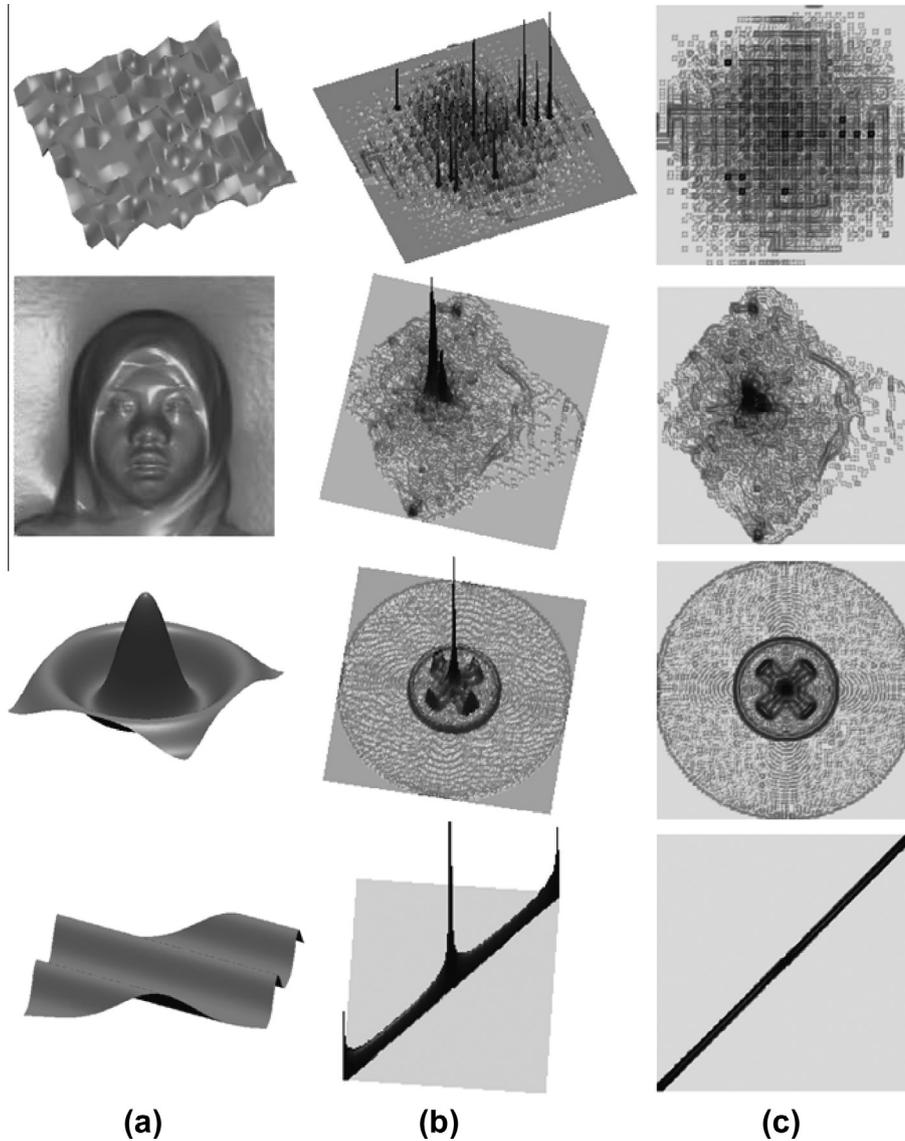


Fig. 4. An example of the histogram of the normals for (a) an isotropic surface, a face, an isotropic surface with significant altitude variation and an anisotropic surface, both in (b) 3D and (c) 2D views.

where $R_{ij}^{(k)}$ is the value of the reflectance map (isophota) at location ij for the k illumination direction.

So, the illumination directions of k lights should be selected to satisfy the above requirement. They may be estimated by maximising the sum of the gradient magnitudes of Eq. (12) for the K light sources, integrated over the whole available range of normals, given by

$$f(p_i, q_i) = \sum_{i=1}^I \sum_{j=1}^J \mathbf{w}_h(\mathbf{P}_i, \mathbf{Q}_j) \left[\sum_{k=1}^K \|\nabla R_{ij}^{(k)}\| \right] \quad (13)$$

where $\|\nabla R_{ij}^{(k)}\| = \sqrt{\left(\frac{\partial R_{ij}^{(k)}}{\partial p_i}\right)^2 + \left(\frac{\partial R_{ij}^{(k)}}{\partial q_j}\right)^2}$. An example of \mathbf{R} and $f(p_i, q_i)$ is shown in Fig. 5. We use simulated annealing to solve this global optimisation problem and obtain the optimal (p_i, q_i) values for the light sources. Once the light gradients (p_i, q_i) for each light source, that maximise our requirements have been estimated, we can obtain the directions of the lights because the zenith angle was considered constant and known.

6. Experiments and results

The idea of the experiments is to reconstruct a surface by using photometric stereo images of the surface, once captured under 'default illumination directions' (as proposed in [10]) and once captured under optimal illumination directions, as specified by the proposed methods. The aim is to show that more accurate reconstructions are obtained when the lighting directions have been selected by the proposed approaches. For all reconstructions, the standard 4-lights photometric stereo algorithm was used, without applying any techniques to exclude pixels affected by shadows or highlights. Initially, simulated surfaces are constructed from which photometric image sets are created. Also 3D models of human faces are used to perform controlled experiments. Three different textures were used to simulate albedo (see Fig. 6) on randomly generated isotropic and anisotropic surfaces. In the case of surfaces of human faces, a uniform grey colour was used, as well as their real albedo.

In more details our testing surfaces in the case of the isotropic contain a facet at each pixel location with the distribution of the

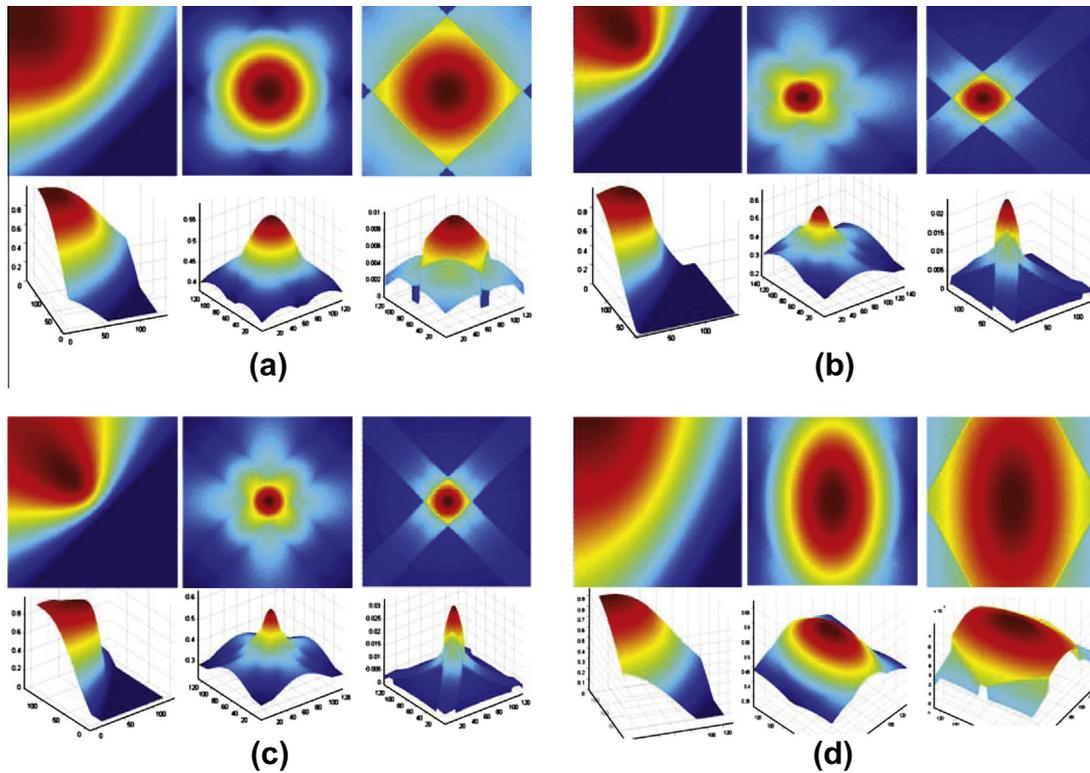


Fig. 5. An example of $\mathbf{R}_{ij}^{(k)}$, $\sum_{k=1}^K \mathbf{R}_{ij}^{(k)}$ and $\sum_{k=1}^K \|\nabla \mathbf{R}_{ij}^{(k)}\|$ for (a) an isotropic surface, (b) a face, (c) an isotropic surface with significant altitude variations and (d) an anisotropic surface both in 2D and 3D views for a given illumination direction.



Fig. 6. The three different textures used to simulate the albedo for all simulated surfaces used, except faces.

normals to be uniform. In order to obtain these surfaces we restrict the orientation of the normals to be on a hemisphere and we sample uniformly obtaining the isotropic surfaces. In the anisotropic case the surfaces contain mainly ripples of different lengths and orientations and as a result this orientation reflects on the orientation of the distribution of the normals. Also faces were used that have anisotropic distribution of normals but there is an axis of symmetry. All the simulated models were rendered in the same way using the Phong model in order to avoid unfair comparisons. Furthermore, regarding the real data faces, bottles and other objects were used with or without ground truth depending on each case. Also it should be mentioned that both the simulated isotropic and anisotropic surfaces are not perfect isotropic and anisotropic due to the random generation of facets and the cropping the obtained distributions of the normals are not completely uniform.

Before we move to the actual evaluation of the proposed methods an experiment was performed to observe the optimisation landscape of the optimisation algorithm that was utilised. Therefore, we keep three of the four lights fixed and we optimise only for the last one and then in Fig. 7 the accuracy of reconstruction as a function of the direction of the 4th light source during the

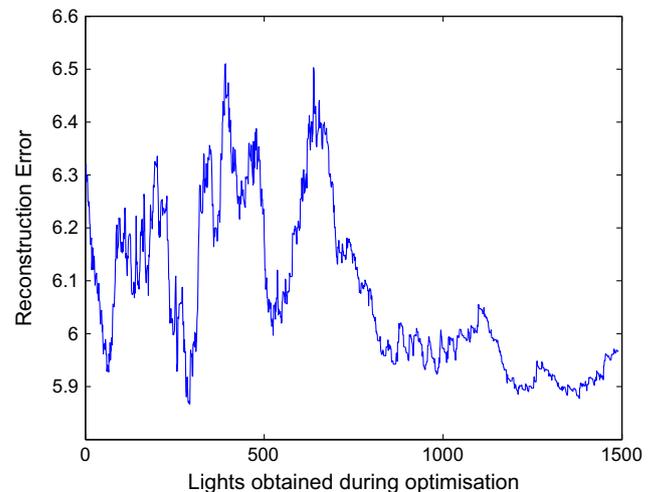


Fig. 7. The accuracy of reconstruction as a function of the direction of the 4th light source during the optimisation process.

optimisation process is plotted. From the figure it can be seen that the illumination direction stabilises after approximately 1200 iterations close to the overall minimum.

6.1. Method 1 – Parameters selection

Regarding the first method, in order to select the proper values for the weights κ and λ in Eqs. (6) and (9), some pilot experiments were performed to identify the order of magnitude of each term. The histograms of each of these terms are shown in Fig. 8. It may be inferred that all of them are of the same order of magnitude but the modes of their values are not balanced. In this work we assume that each component is equally important in the quality functions and therefore these weights are select to secure this assumption and provide a global optimal with each term contributing the same. Therefore, the weights λ and κ should be properly selected to make the modes of all terms roughly the same, so all terms play role in the optimisation. From these histograms, it is inferred that $\lambda \sim 0.5$ and $\kappa \sim 2$.

As simulated annealing works by minimising a cost function, while quality functions have been defined that have to be maximised, the cost function used in the optimisation step was $-Q(\varphi_l, \theta_l)$ or $-P(\varphi_l, \theta_l, \varphi_v, \theta_v)$.

Fig. 9 shows the plots of the minimisation of these functions for an example case. In Fig. 10 the value of each of the terms considered in the construction of quality function $P(\varphi_l, \theta_l, \varphi_v, \theta_v)$ at each pixel is shown, for specific values of the illumination and viewing directions, and for a particular example case. Furthermore, in Fig. 10 we have a visual representation of all terms f, g and t in Eqs. (6) and (9). Column (c) corresponds to term f that represents the relative orientation between neighbouring facets trying to maximise the local contrast and improve the visibility of local detail; column (d) corresponds to term g trying to minimise the self-shadows; and column (e) corresponds to term t indicating the highlights present in the image. Therefore, it is clear that Eqs. (6) and (9) depend on the contents of the input images.

7. Experiments with simulated surfaces

Experiments were performed with three sets of surface data: isotropic surfaces, anisotropic surfaces and simulated human faces with the terms isotropic and anisotropic we refer to the distribution of the surface normals. For all these surfaces the normal vector of each facet (ground truth) is known. For each set of surfaces, that share common characteristics and statistics related to the surface

normals, the directions of four light sources were calculated using the proposed methodology. A subset of surfaces for each surface class was used to obtain the optimal illumination directions for the corresponding set. The mean illumination configuration from the training subset was selected and applied to the surfaces of the test subset. The term ‘mean’ corresponds to the average four vectors of all the surfaces in the training subset of each surface class. Regarding the first approach, in order to identify the corresponding illumination directions, which had to be averaged, the k -means clustering algorithm was used in the (φ_l, θ_l) space to partition all the estimated directions into four clusters. The mean for each cluster was calculated providing the direction of one light source. Fig. 11 shows some examples of the obtained clusters for different classes of surfaces used in our experiments.

In order to compare the performance of the proposed approaches, the angular error (AE) measure suggested by Beauchemin et al. [2] was used:

$$\Psi_{AE} = \cos^{-1} \left[\frac{x_e x_c + y_e y_c + z_e z_c + 1}{\sqrt{1 + x_r^2 + y_r^2 + z_r^2} \sqrt{1 + x_e^2 + y_e^2 + z_e^2}} \right] \quad (14)$$

Here $(x_e, y_e, z_e)^T$ and $(x_r, y_r, z_r)^T$ are the estimated and the real surface normals for a particular facet, respectively.

Furthermore, since the real albedo was available, the mean absolute difference between true and worked out albedo was used to compare the performance of the proposed methodologies

$$\epsilon_{AD} = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N |A_e(i, j) - A_r(i, j)| \quad (15)$$

where A_e and A_r are the estimated and the real albedo, respectively.

Fifteen isotropic surfaces for training were used and fifteen for evaluation (see Fig. 12). The surfaces were created to contain random peaks with low altitude, in order to avoid cast-shadows and provide a uniform distribution of normals. In all cases, the simulated annealing was initialised with the ‘default’ vectors for the illumination directions and with the z -axis for the viewing direction. In Fig. 13 the worked out illumination directions are shown for both methods. It can be seen that, as expected, they are uniformly distributed about the surface. In the first row of Tables 1–3 the average angular and albedo errors for the isotropic surfaces for the Lambertian and the Phong models respectively are shown for all three textures that were used (see Fig. 6). Figs. 14 and 15 show the angular error in degrees and the albedo difference as a function of the isotropic surfaces, respectively. The mean error

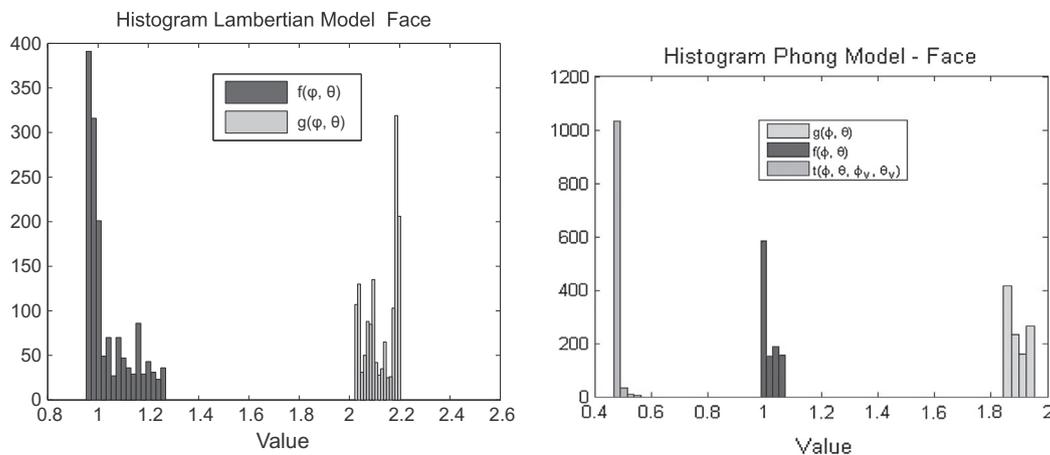


Fig. 8. Histograms of all the optimisation terms (x -axis) over all the different values of ϕ and θ for the Lambertian and Phong model of the first method, respectively, for four faces.

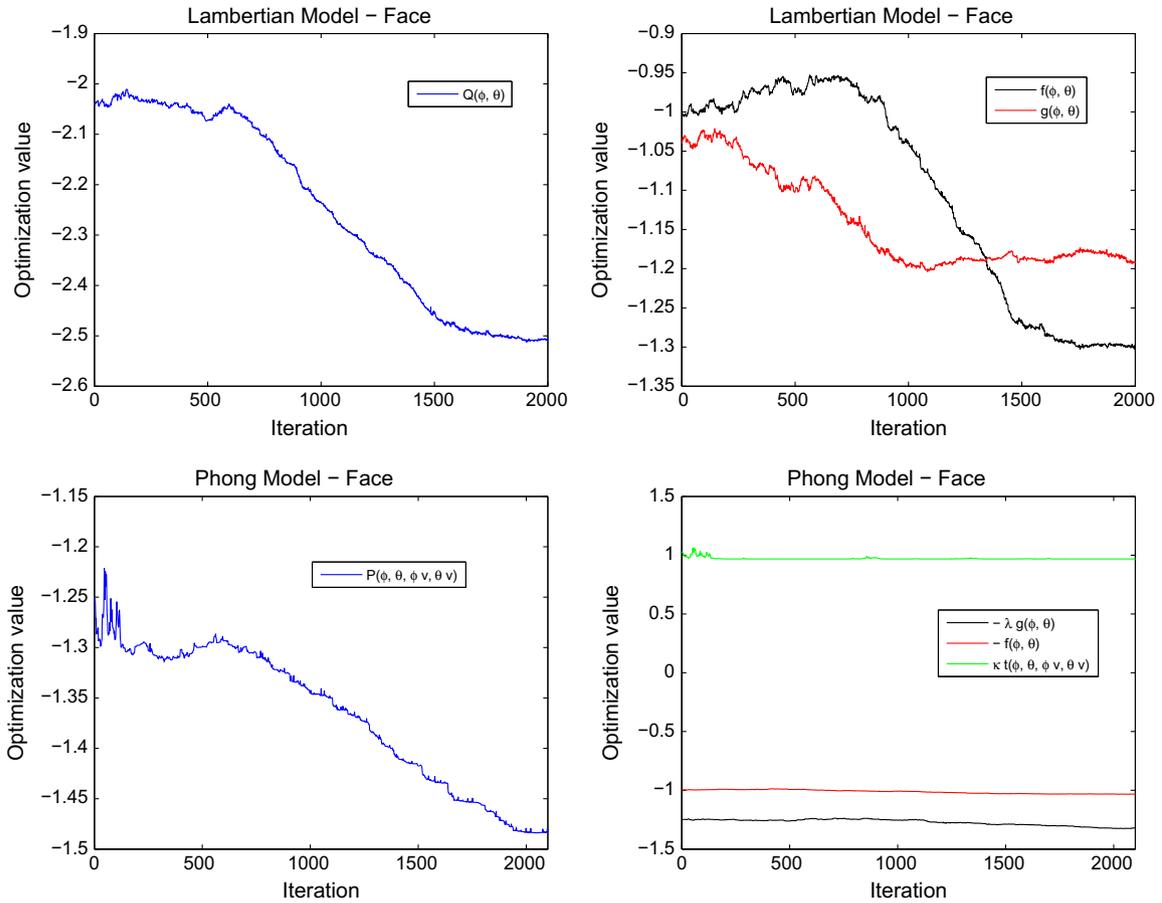


Fig. 9. Examples of cost function reduction by the optimisation process for the first method. At the top for the Lambertian model and at the bottom for the Phong model. On the left the value of each cost function as a function of the iterations, while on the right each individual term. These runs were for a human face with $\lambda = 0.75$ and $\kappa = 2$.

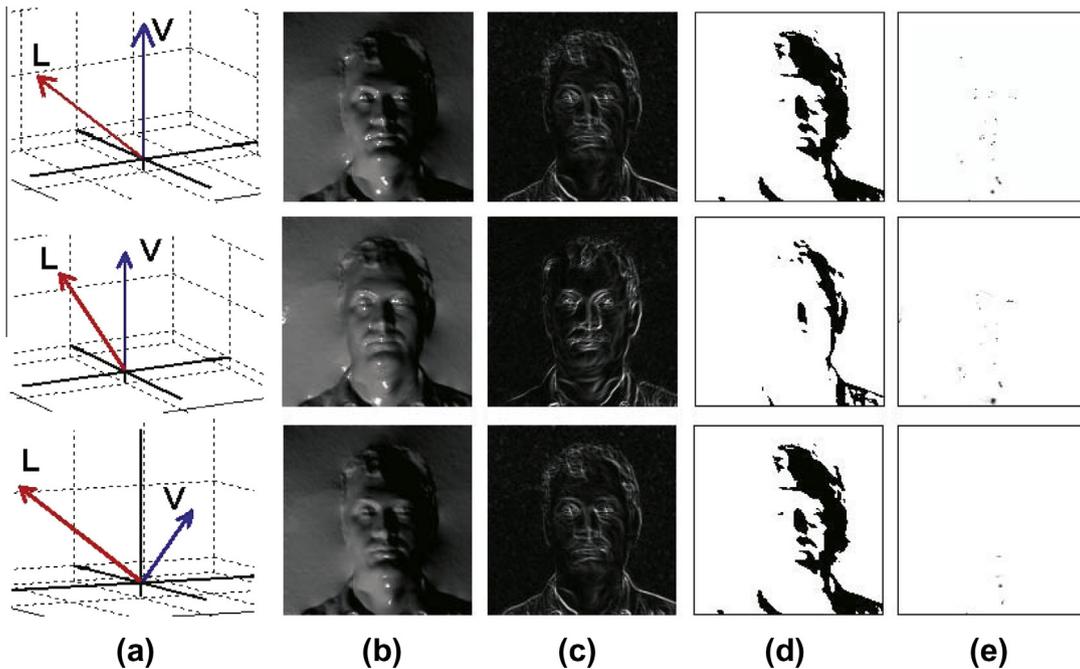


Fig. 10. (a) Selected illumination and viewing directions for the first approach. (b) 3D representation of a face. (c) Every pixel is given a grey value according to the difference in illumination it receives from the average illumination of its four nearest neighbours. (d) Black pixels are not illuminated. Their number should be minimised. (e) Black pixels are highlights. Their number should be minimised. In the 1st and 2nd rows the viewing direction is the same, while in the 1st and 3rd rows the illumination direction is the same.

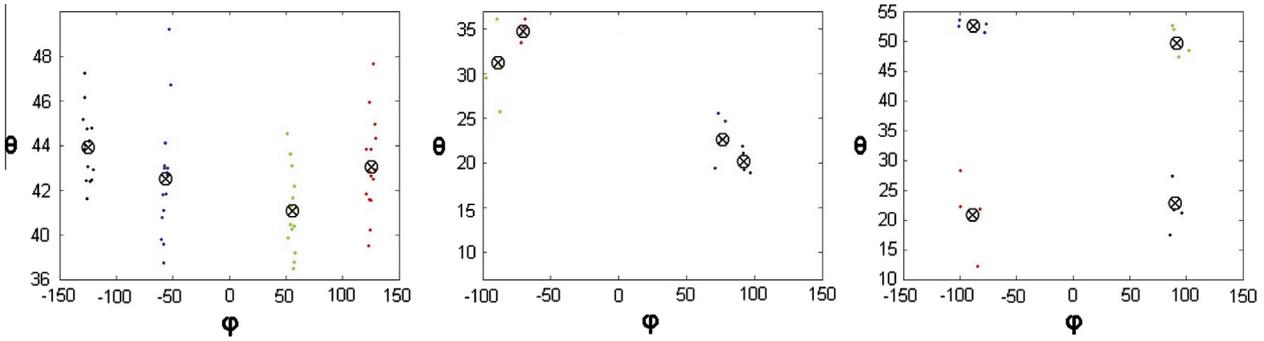


Fig. 11. Clusters of the illumination directions for different classes of surfaces. Left isotropic surfaces, middle faces and at the right anisotropic surfaces.

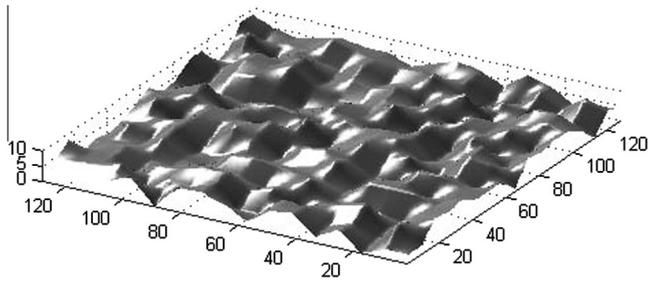


Fig. 12. An example of a simulated isotropic surface of size 128 × 128 pixels. The RMS roughness value of these surfaces was in the range between 4.5 and 5.5.

of the recovered albedo is insignificant either the default illumination directions or the worked out ones were used.

From the results about the second approach, we may observe that for the case of 4-lights photometric stereo, the lights should be located close to the corners of a square as it was expected, since the simulated data satisfy the required assumption of a uniformly rough surface. Based on Eq. (11) the azimuth angle is given by $\varphi_l = \tan^{-1}(q_l/p_l)$ indicating that the angle remains the same when the ratio of the illumination directions is constant.

Next, four sets of anisotropic surfaces (Fig. 16) were used. The selection of these examples and these orientations was mainly to allow us to demonstrate that the obtained estimates are correct and according the expectations of the proposed methods. In each set there were seven surfaces with different appearance, but all with prominent anisotropy in the same direction. From the seven surfaces, three were used for training and four for testing. The estimated illumination directions for the anisotropic surfaces are shown in Fig. 17 coinciding with the directions of the gradients and exploiting the density of the isophotes as it was expected for both approaches. The average angular error is given in Tables 1 and 2 for the Lambertian and the Phong models, respectively, in

Table 1

The mean angular error (MAE) computed over all faces and all simulated surfaces and for $\lambda = 0.75$ in the first method using the Lambertian model. In bold the best result for each case.

MAE	Default	Proposed A	Proposed B
Isotropic	1.0817	0.0000	0.0000
Anisotropic 0°	3.9902	0.8618	1.1859
Anisotropic 45°	4.2101	1.0434	1.3518
Anisotropic 135°	4.2065	1.2320	1.3627
Anisotropic 90°	4.0553	0.9650	1.0977
Faces	5.0481	0.1244	0.3315

Table 2

The mean angular error (MAE) computed over all facets and all simulated surfaces and for $\lambda = 0.75$ and $\kappa = 0.5$ using the Phong model. In bold the best result for each case.

MAE	Default	Proposed A
Isotropic	1.0817	0.0000
Anisotropic 0°	3.9902	0.9346
Anisotropic 45°	4.2101	1.0198
Anisotropic 135°	4.2065	1.0679
Anisotropic 90°	4.0553	0.9610
Faces	5.0481	0.0001

Table 3

The mean absolute difference (MAD) per pixel between the ground truth and the estimated albedo computed over all isotropic simulated surfaces for different textures and for all illumination configurations for both methodologies. In bold the best result for each case.

MAD	Texture 1	Texture 2	Texture 3
Default	0.0019	0.0021	0.0014
Proposed A	0.0000	0.0000	0.0000
Proposed B	0.0000	0.0000	0.0000

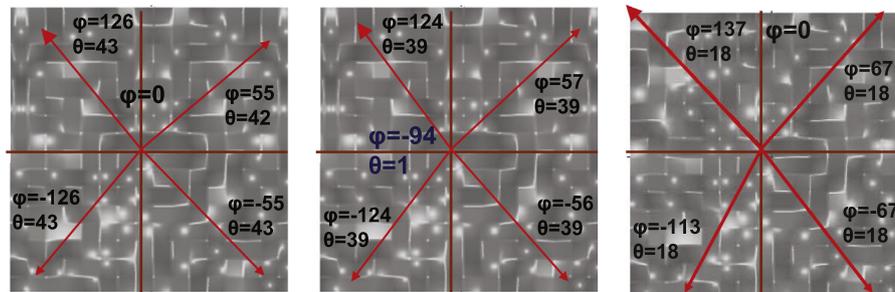


Fig. 13. Estimated illumination directions for the class of isotropic simulated surfaces using both approaches, (left) method A and Lambertian model, (middle) method A and Phong model and (right) method B and Lambertian model.

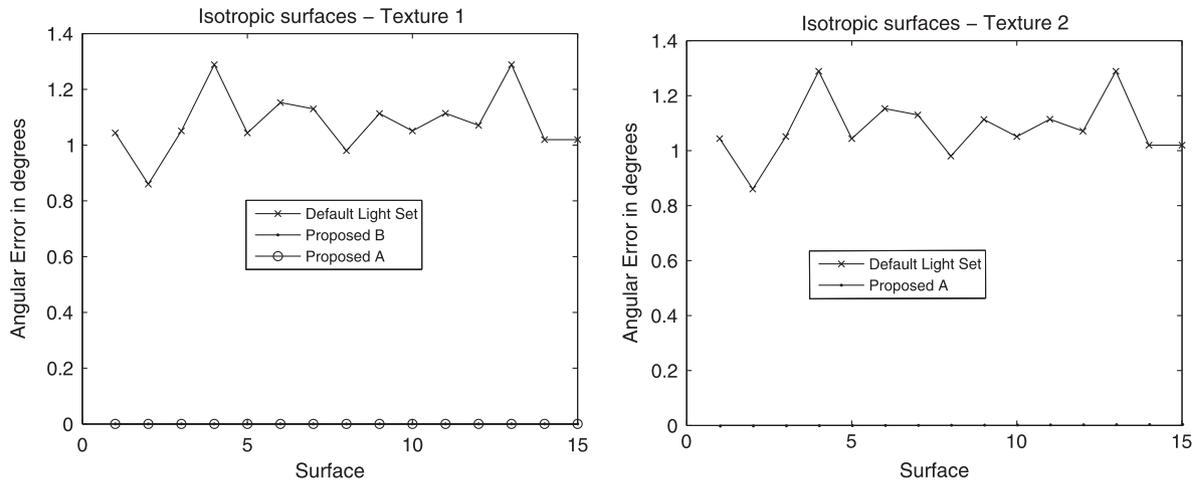


Fig. 14. Performance comparison in terms of mean angular error for each anisotropic surface using the 'default' illumination configuration and the proposed set for both approaches, (left) method A and B with Lambertian model and (right) method A and Phong model.

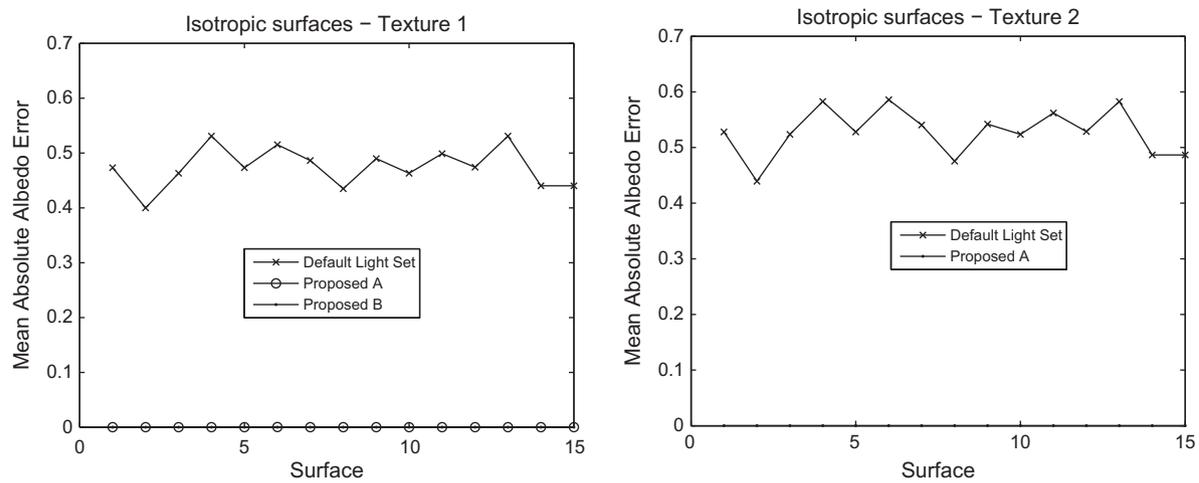


Fig. 15. Performance comparison in terms of mean absolute difference for each anisotropic surface using the 'default' illumination configuration and the proposed for both approaches, (left) method A and B with Lambertian model and (right) method A and Phong model.

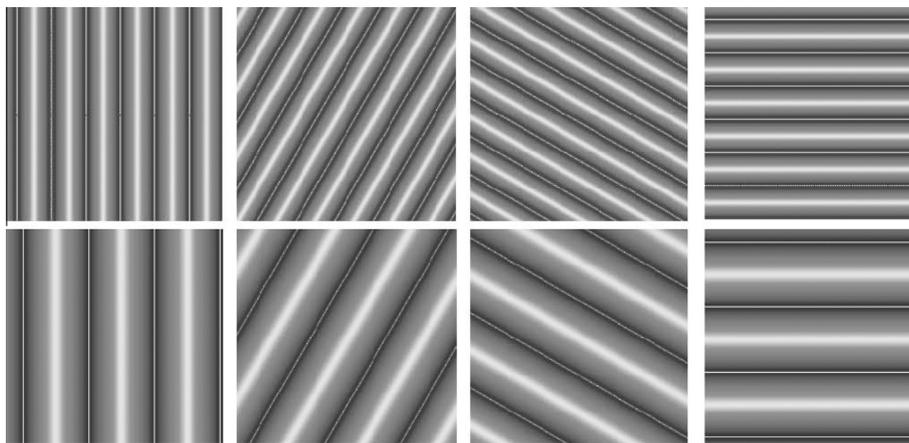


Fig. 16. Two example surfaces from each of the four sets of anisotropic surfaces used in our experiments. From left to right the main orientation of the distribution of the normals is 0° , 45° , 90° and 135° . The RMS roughness value of these surfaces was in the range between 3.8 and 14.5.

rows two to five. It should be mentioned again that the the Phong model is applicable only for the first method. The average albedo error for each of the four directions is shown in Table 4 for both

methods and the Lambertian model and in Table 5 only for the first approach and the Phong model. It may be observed that the proposed methodologies results more accurate normals with

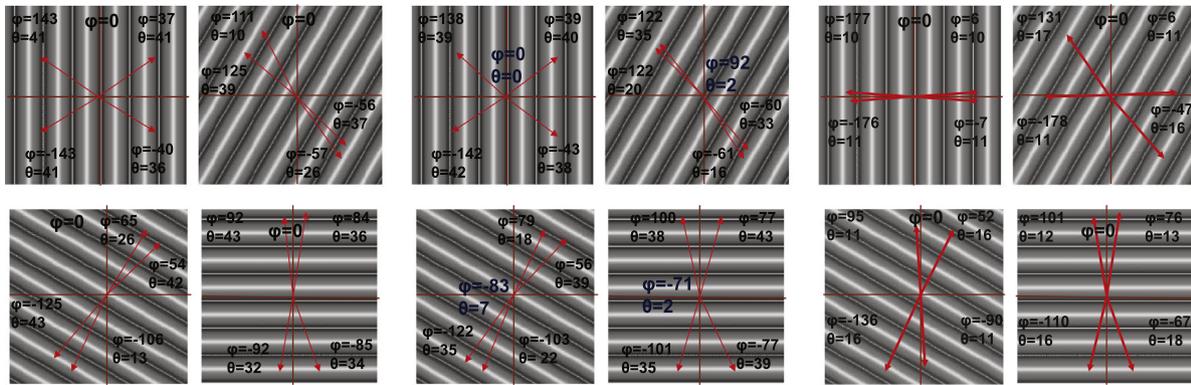


Fig. 17. Estimated illumination directions for the class of anisotropic simulated surfaces using both approaches, (left) method A and Lambertian model, (middle) method A and Phong model and (right) method B and Lambertian model.

Table 4

The mean absolute difference (MAD) per pixel in albedo computed over all anisotropic simulated surfaces for different textures and for all illumination configurations and both methods for the Lambertian model. In bold the best result for each case.

MAD	T 1	T 2	T 3	T 1	T 2	T 3
	0°			45°		
Default	0.0042	0.0046	0.0030	0.0057	0.0064	0.0042
Proposed A	0.0005	0.0006	0.0004	0.0012	0.0014	0.0009
Proposed B	0.0031	0.0035	0.0023	0.0040	0.0045	0.0029
	135°			90°		
Default	0.0057	0.0064	0.0042	0.0042	0.0047	0.0031
Proposed A	0.0051	0.0047	0.0040	0.0004	0.0004	0.0003
Proposed B	0.0040	0.0045	0.0029	0.0033	0.0036	0.0024

Table 5

The mean absolute difference (MAD) in albedo computed over all anisotropic simulated surfaces for different textures and for all illumination configurations using the Phong model. In bold the best result for each case.

MAD	T 1	T 2	T 3	T 1	T 2	T 3
	0°			45°		
Default	0.0042	0.0046	0.0030	0.0057	0.0064	0.0042
Proposed A	0.0019	0.0021	0.0014	0.0003	0.0003	0.0002
	135°			90°		
Default	0.0057	0.0064	0.0042	0.0042	0.0047	0.0031
Proposed A	0.0006	0.0006	0.0004	0.0018	0.0020	0.0013

significant error reduction. Also from the obtained results it is observed that the first proposed method outperforms in the case of textured data in general, mainly due to the first term in Eqs. (6) and (9) that takes in to account the local surface deformations while the other approaches assume it is constant providing an advantage in these cases. In figures Figs. 18 and 19 the angular and the mean absolute error for the anisotropic surfaces using the ‘default’ illumination configuration, versus the proposed one are shown. Observing the figures, we can see that the first method gives different results for the two models (Lambertian-left and Phong-right) due to the different illumination directions. Regarding the default lights, since they are the same illumination setup in both cases and the percentage of the highlights present in the images is very low compared with the total number of pixels, the average difference is almost negligible between the two models (Lambertian and Phong) for the first method. Even the average contribution of the highlights to the total error is very low, their removal is regarded essential due to the fact that highlights create local peaks that can affect significantly the visual outcome. Similar results were obtained for λ in the range [0.5,0.75] and κ in the range [1.5,2].

Further experiments were performed using faces (Fig. 20). The estimated optimal mean illumination directions of the lights are shown in Fig. 21 and all the results for each face are shown in Table 6. Observing the results it may be seen that for the faces the optimal illumination directions are from the left and the right

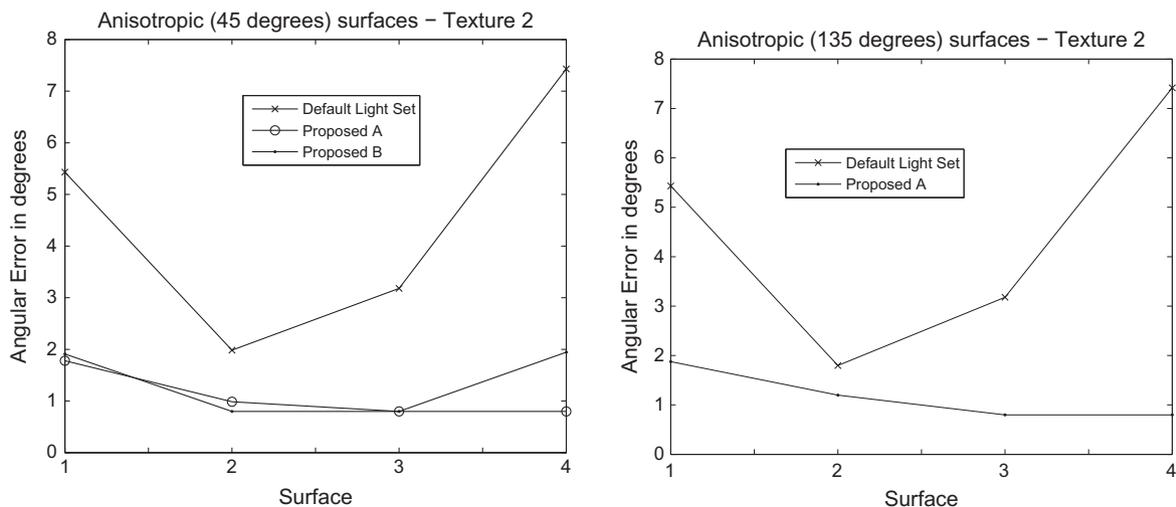


Fig. 18. Performance comparison in terms of mean angular error for each anisotropic surface using the ‘default’ illumination configuration and the proposed set for both approaches, (left) method A and B with Lambertian model and (right) method A and Phong model.

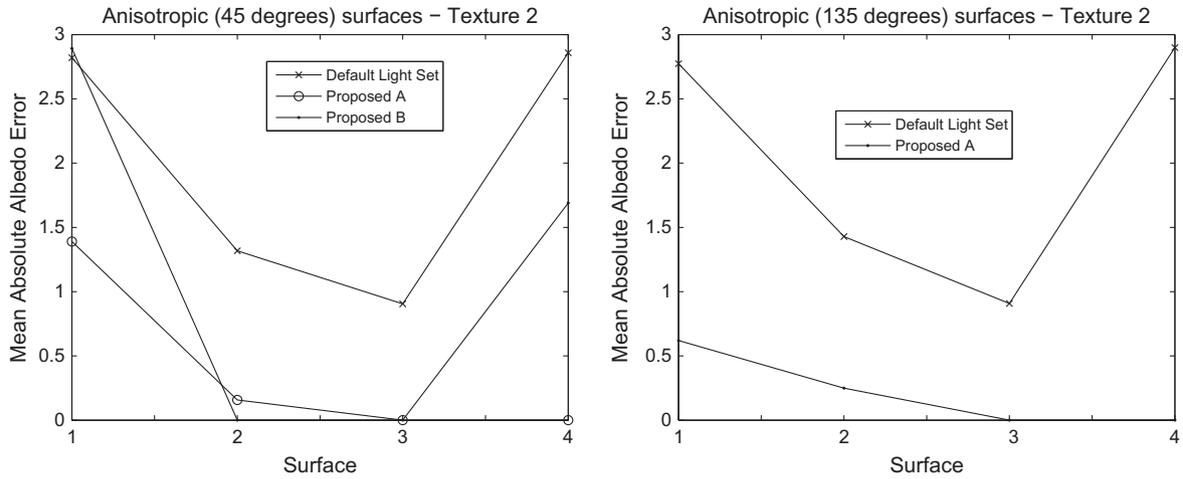


Fig. 19. Performance comparison in terms of mean absolute difference for each anisotropic surface using the 'default' illumination configuration and the proposed for both approaches, (left) method A and B with Lambertian model and (right) method A and Phong model.



Fig. 20. Examples of faces used in our experiments. The RMS roughness value of these surfaces was in the range between 28.6 and 38.1.

sides of the facial axis of symmetry, a bit lower than the chin or at the level of the eyes. The main difference is the angle θ that is smaller in the case of Phong model with a small change in the viewing direction. The differences are due to the extra component in the objective functions that tries to reduce the presence of highlights and also to specify the optimal viewing direction which is not an

Table 6

The estimated illumination directions for both approaches.

Light	φ_1, θ_1	φ_2, θ_2	φ_3, θ_3	φ_4, θ_4
Proposed A	(92°, 20°)	(78°, 22°)	(-70°, 35°)	(-89°, 31°)
Proposed B	(101°, 12°)	(76°, 13°)	(-67°, 18°)	(-110°, 16°)

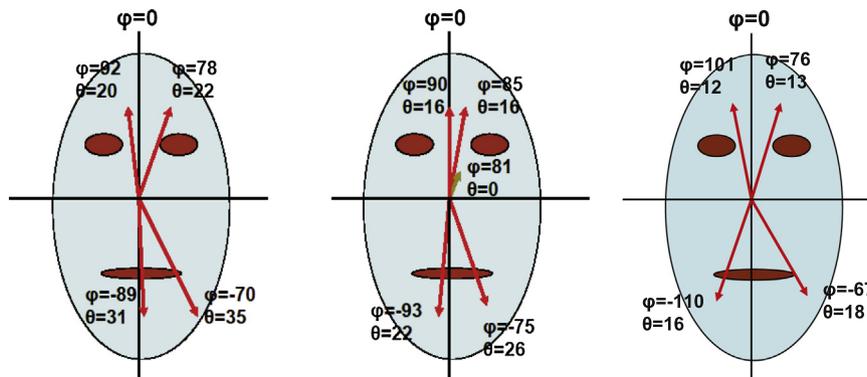


Fig. 21. Estimated illumination direction for the class face using both approaches (left) method A and Lambertian model, (middle) method A and Phong model and (right) method B and Lambertian model.

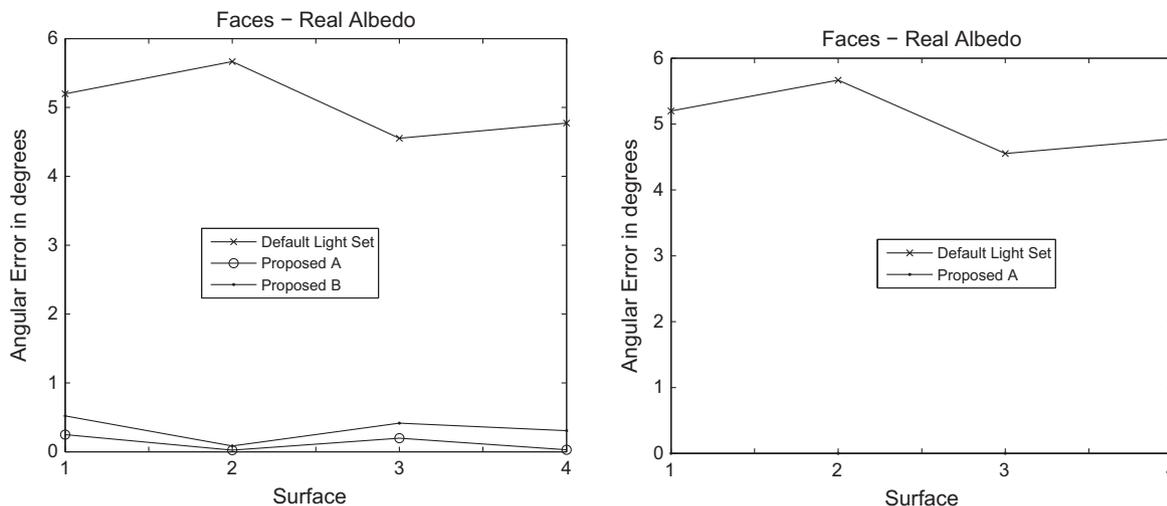


Fig. 22. Performance comparison in terms of mean angular error for each face using the ‘default’ illumination configuration and the proposed set for both approaches, (left) method A and B with Lambertian model and (right) method A and Phong model.

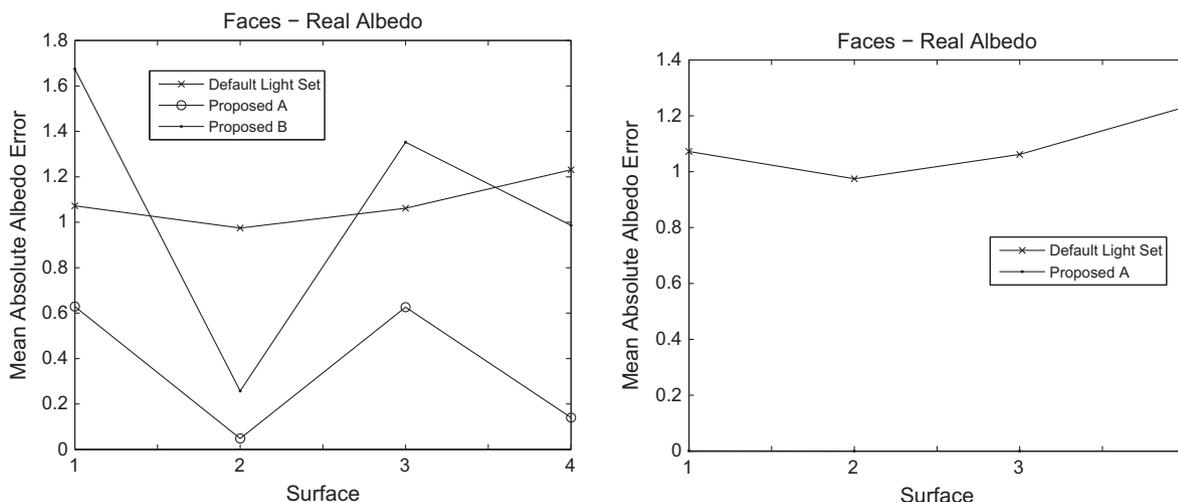


Fig. 23. Performance comparison in terms of mean absolute difference for each face using the ‘default’ illumination configuration and the proposed set for both approaches, (left) method A and B with Lambertian model and (right) method A and Phong model.

Table 7
The mean absolute difference (MAD) computed over all faces for different textures and for both methods and all the illumination configurations using the Lambertian model. Bold values indicate the method with the best performance (lower error).

MAD	Texture 1	Texture 2
Default	0.0053	0.0042
Proposed A	0.0016	0.0014
Proposed B	0.0035	0.0041

Table 8
The mean absolute difference (MAD) computed over all faces for different textures, and for all illumination configurations using the Phong model. Bold values indicate the method with the best performance (lower error).

MAD	Texture 1	Texture 2
Default	0.0053	0.0042
Proposed	0.0000	0.0000

issue for the Lambertian model. The mean angular and albedo errors for each face are shown in Figs. 22 and 23. All mean errors are shown in Tables 1, 2, 7 and 8. Similar results were obtained for λ in the range [0.5,0.75] and κ in the range [1.5,2]. This can be observed in figure Fig. 24 which shows the error versus the values of λ and κ . Observing the results of the second method, we may say that it provides illumination configurations that adjust the isophotes in order to achieve higher density at more frequented areas on the pq -plane (i.e. more frequently encountered normals).

8. Experiments with real objects and faces

Having learnt the optimal illumination directions from the simulated faces and anisotropic surfaces, now these arrangements are tested with real surfaces of similar type, for which the ground truth is available with the help of the 3dMD imaging system [26]. In particular, two mannequin faces (see Fig. 25) will be reconstructed using the optimal illumination directions reported in Tables 9 and 10 and the reconstruction will be compared with the real height maps worked by the 3dMD scanner. In addition, a bottle, which may be thought of as a ribbed anisotropic surface, will also be tested using the optimal illumination directions for

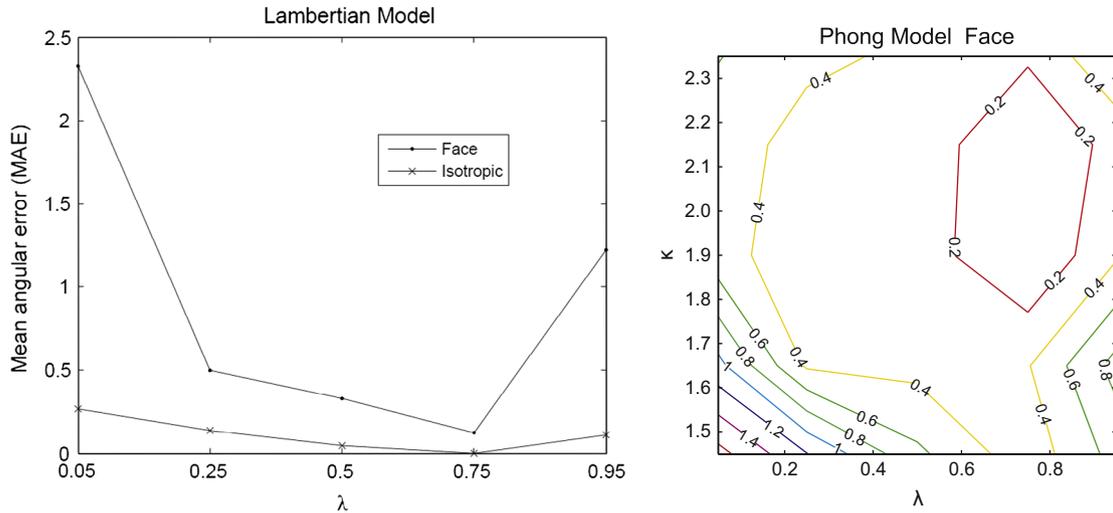


Fig. 24. The error versus the values of λ of the first method for the Lambertian model (left) and isoerror contours for the Phong model (right) for various combinations of parameter values (λ, κ).



Fig. 25. 3D surfaces obtained using the 3dMD imaging system.

such surfaces, also reported in Tables 9–12 and in Figs. 32–34. It can be observed that the optimal viewing direction is effectively at the zenith.

For these experiments, the sum of the absolute height map difference, between the estimated reconstructed surface, using Photometric Stereo, and the ground truth, obtained using the 3dMD imaging system, was used to compare the accuracy of the proposed methodology:

$$h_{EG} = \frac{1}{MN} \sum_{i=1}^M \sum_{j=1}^N |H_e(i, j) - H_r(i, j)| \quad (16)$$

where H_e and H_r are the estimated and the real height maps, respectively.

Each object was illuminated both from the default and the estimated optimal directions. The captured images are shown in Figs. 26–28. Using the images captured from the default and the estimated optimal light sources, Photometric Stereo and integration are applied in succession, in order to obtain the 3D surfaces for both illumination configurations. Using the Iterative Closest Point algorithm [6] the obtained surfaces are aligned with the

corresponding surfaces captured from the 3dMD imaging system. In order to evaluate the accuracy of the obtained surfaces, the absolute height map difference in Eq. (16) was utilised. The reconstructed 3D surfaces for both configurations are shown in Figs. 29–31. Observing the surfaces, we may say that the proposed configuration produces fewer shadows with more details revealed especially close to the ears and the chin. Also, the seams created at the borders of shadowed and non-shadowed areas are less pronounced. The average absolute height map difference is shown

Table 9 The default and the proposed illumination directions for faces, isotropic and anisotropic surfaces using the Lambertian model in Method 1.

Light	$\varphi_{l_1}, \theta_{l_1}$	$\varphi_{l_2}, \theta_{l_2}$	$\varphi_{l_3}, \theta_{l_3}$	$\varphi_{l_4}, \theta_{l_4}$
Default	(135°, 55°)	(45°, 55°)	(-45°, 55°)	(-135°, 55°)
Isotr	(126°, 43°)	(55°, 42°)	(-55°, 43°)	(-126°, 43°)
Anisot 0°	(143°, 41°)	(37°, 41°)	(-40°, 36°)	(-143°, 41°)
Anisot 135°	(125°, 39°)	(111°, 10°)	(-56°, 37°)	(-57°, 26°)
Anisot 45°	(65°, 26°)	(54°, 42°)	(-106°, 13°)	(-125°, 43°)
Anisot 90°	(92°, 43°)	(84°, 36°)	(-85°, 34°)	(-92°, 32°)
Faces	(92°, 20°)	(78°, 22°)	(-70°, 35°)	(-89°, 31°)

Table 10 The default and the proposed illumination directions for faces, isotropic and anisotropic surfaces using the Phong model in Method 1.

Light	$\varphi_{l_1}, \theta_{l_1}$	$\varphi_{l_2}, \theta_{l_2}$	$\varphi_{l_3}, \theta_{l_3}$	$\varphi_{l_4}, \theta_{l_4}$
Default	(135°, 55°)	(45°, 55°)	(-45°, 55°)	(-135°, 55°)
Isot	(124°, 39°)	(57°, 39°)	(-56°, 39°)	(-124°, 39°)
Anisot 0°	(138°, 39°)	(39°, 40°)	(-43°, 38°)	(-142°, 42°)
Anisot 135°	(122°, 35°)	(122°, 20°)	(-60°, 33°)	(-61°, 16°)
Anisot 45°	(79°, 18°)	(56°, 39°)	(-103°, 22°)	(-122°, 35°)
Anisot 90°	(100°, 38°)	(77°, 43°)	(-77°, 39°)	(-101°, 35°)
Faces	(90°, 16°)	(85°, 16°)	(-75°, 26°)	(-93°, 22°)

Table 11 The default and the proposed viewing directions for faces, isotropic and anisotropic surfaces using the Phong model in Method 1. Note that if $\theta = 0, \varphi$ is undetermined.

Camera	φ_v, θ_v
Default	(-0°)
Isot	(-94°, 1°)
Anisot 0°	(-0°)
Anisot 135°	(92°, 2°)
Anisot 45°	(-83°, 7°)
Anisot 90°	(-71°, 2°)
Faces	(-0°)

Table 12

The default and the proposed illumination directions for faces, isotropic and anisotropic surfaces based on Method 2.

Light	$\varphi_{l_1}, \theta_{l_1}$	$\varphi_{l_2}, \theta_{l_2}$	$\varphi_{l_3}, \theta_{l_3}$	$\varphi_{l_4}, \theta_{l_4}$
Default	(135°, 55°)	(45°, 55°)	(-45°, 55°)	(-135°, 55°)
Isot	(137°, 18°)	(67°, 18°)	(-67°, 18°)	(-113°, 18°)
Anisot 0°	(177°, 10°)	(6°, 10°)	(-7°, 11°)	(-176°, 11°)
Anisot 135°	(131°, 17°)	(6°, 11°)	(-47°, 16°)	(-178°, 11°)
Anisot 45°	(95°, 11°)	(52°, 16°)	(-90°, 11°)	(-136°, 16°)
Anisot 90°	(101°, 12°)	(76°, 13°)	(-67°, 18°)	(-110°, 16°)
Faces	(101°, 12°)	(76°, 13°)	(-67°, 18°)	(-110°, 16°)

in Table 13 with the proposed illumination configuration resulting in lower errors.

The proposed algorithm was further applied to the reconstruction of six real human faces and four potatoes (see Fig. 35) using photometric data captured both with the default and the proposed illumination directions. The person is assumed to be still during the acquisition stage since a high speed camera was used for the acquisition (i.e. 200 frames per second), eliminating the registration problem.

In Figs. 36 and 37 results of the reconstructed faces and potatoes obtained from the two compared illumination setups are

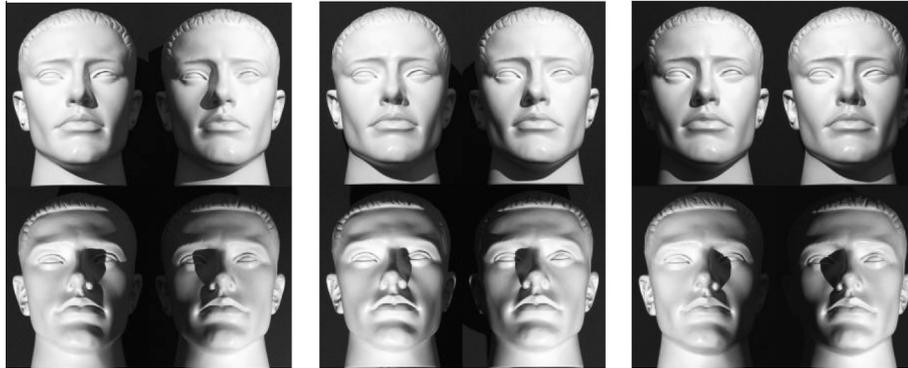


Fig. 26. Images captured using the default (left column) the estimated based on method A (middle column) and estimated based on method B (right column) illumination configurations for 'Adam'.



Fig. 27. Images captured using the default (left column) the estimated based on method A (middle column) and estimated based on method B (right column) illumination configurations for 'Eve'.



Fig. 28. Images captured using the default (left column) the estimated based on method A (middle column) and estimated based on method B (right column) illumination configurations for 'coke'.

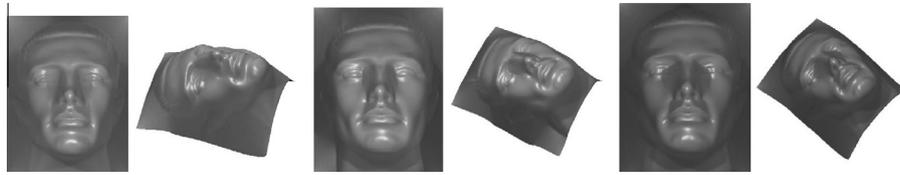


Fig. 29. 3D surfaces obtained using the default (left column) the proposed based on method A (middle column) and proposed based on method B (right column) illumination configuration for 'Adam'.

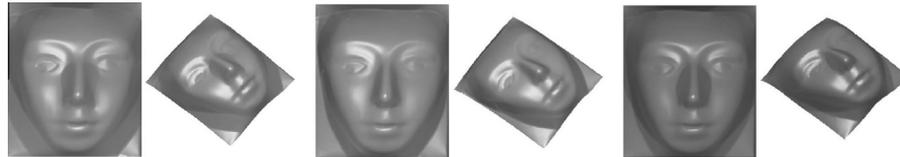


Fig. 30. 3D surfaces obtained using the default (left column) the proposed based on method A (middle column) and proposed based on method B (right column) illumination configuration for 'Eve'.

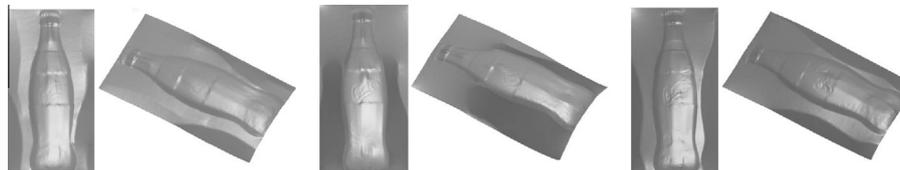


Fig. 31. 3D surfaces obtained using the default (left column) the proposed based on method A (middle column) and proposed based on method B (right column) illumination configuration for 'coke'.

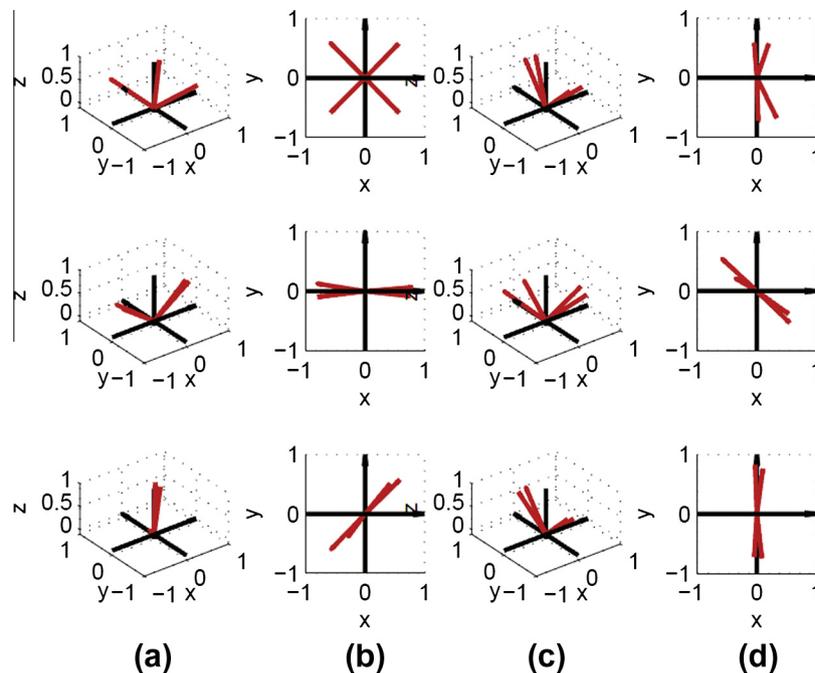


Fig. 32. An example of the illumination and viewing directions using the first method in 3D representation (columns (a) and (c)) and in 2D top view (columns (b) and (d)) using the Lambertian model for isotropic surfaces (top left), faces (top right), an isotropic 0° (mid left), an isotropic 135° (mid right), an isotropic 45° (bottom left) and an isotropic 90° (bottom right).

shown. Observing the results it can be inferred that the proposed illumination directions result more accurate estimates especially at the region under the jaw indicating that the proposed illumination directions provide more accurate reconstructions for faces. Furthermore, the side view was used to reevaluate the performance of the proposed iterative approach. The background was extracted manually and the Hausdorff distance was used to compare the reconstructions with the original profiles. Tables 14 and 15 shows

Table 13

The mean absolute height map difference (MAD) computed over all three surfaces captured under real environmental conditions using the default and the proposed optimal illumination configuration. Bold values indicate the method with the best performance (lower error).

	Adam	Eve	Coke
Default	69.4904	26.7044	17.0176
Proposed A	52.0409	25.9146	11.6387
Proposed B	57.6912	23.3832	13.5134

all the results for all the faces and objects. The obtained results can provide indications to practitioners on how to arrange the capturing system selecting the appropriate number and locations of the lights. In the results presented above different configurations are proposed based on the object class trying to provide most accurate reconstructions keeping the amount of lights constant.

9. Conclusions

In this paper, two methods for estimating the optimal illumination and viewing direction were proposed assuming that prior knowledge of the statistical distribution of the surface to be reconstructed is available. The first approach estimates the

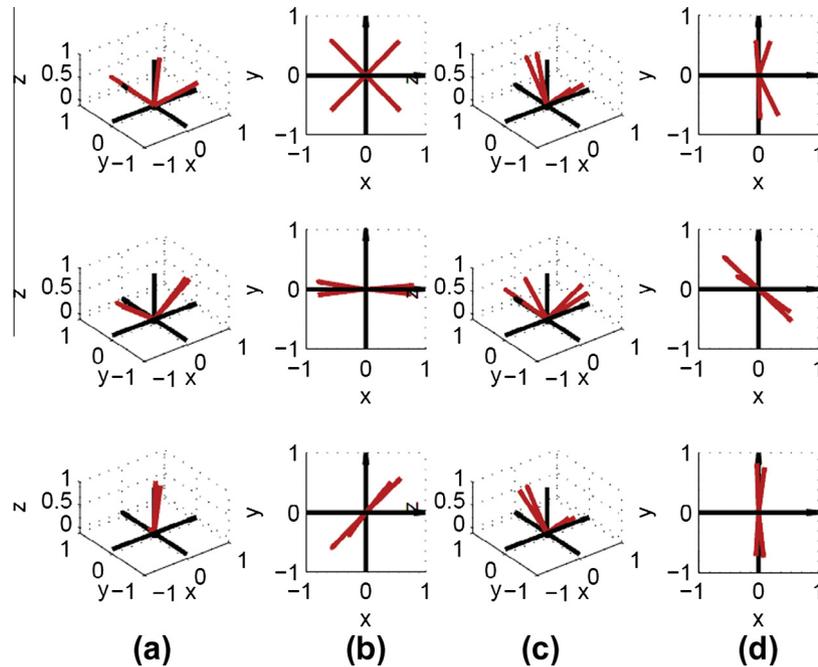


Fig. 33. An example of the illumination and viewing directions using the first method in 3D representation (columns (a) and (c)) and in 2D top view (columns (b) and (d)) using the Phong model for isotropic surfaces (top left), faces (top right), an isotropic 0° (mid left), an isotropic 135° (mid right), an isotropic 45° (bottom left) and an isotropic 90° (bottom right).

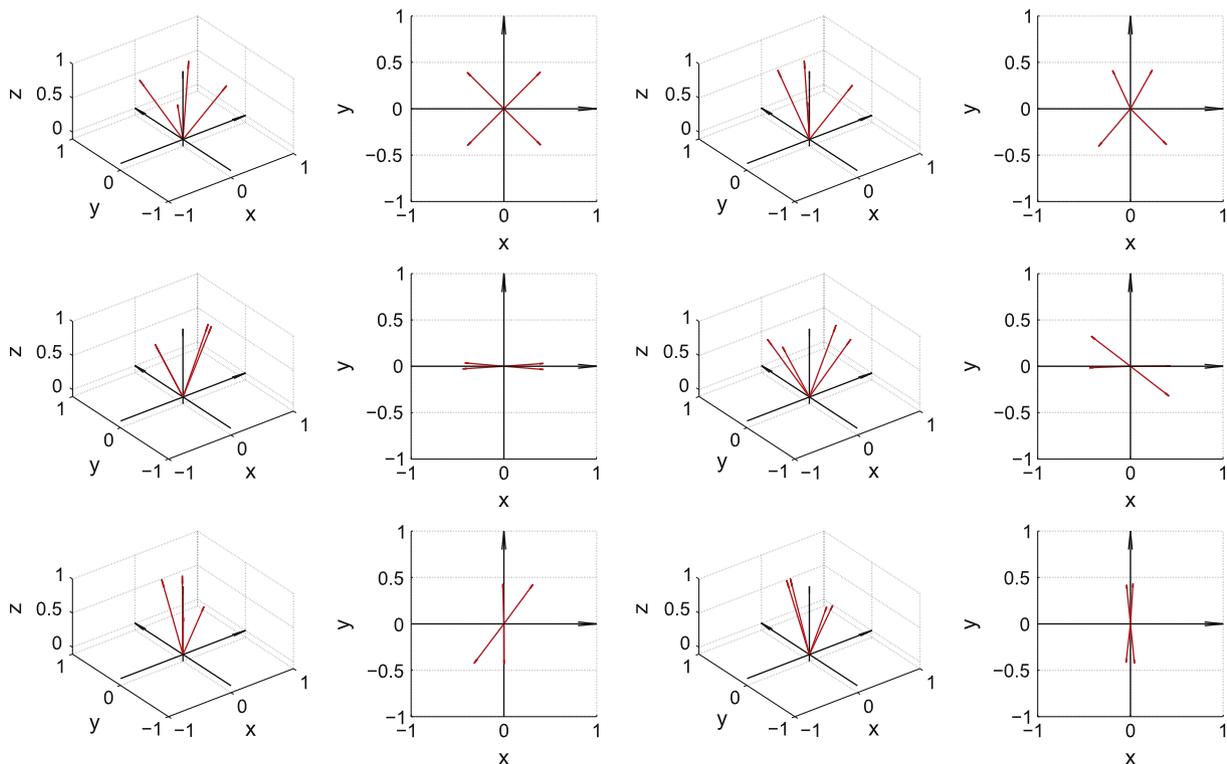


Fig. 34. An example of the illumination directions using the second method in 3D representation for isotropic surfaces (top left), faces (top right), isotropic 0° (mid left), isotropic 135° (mid right), isotropic 45° (bottom left) and isotropic 90° (bottom right) in 3D view (left column) and top view (right column).

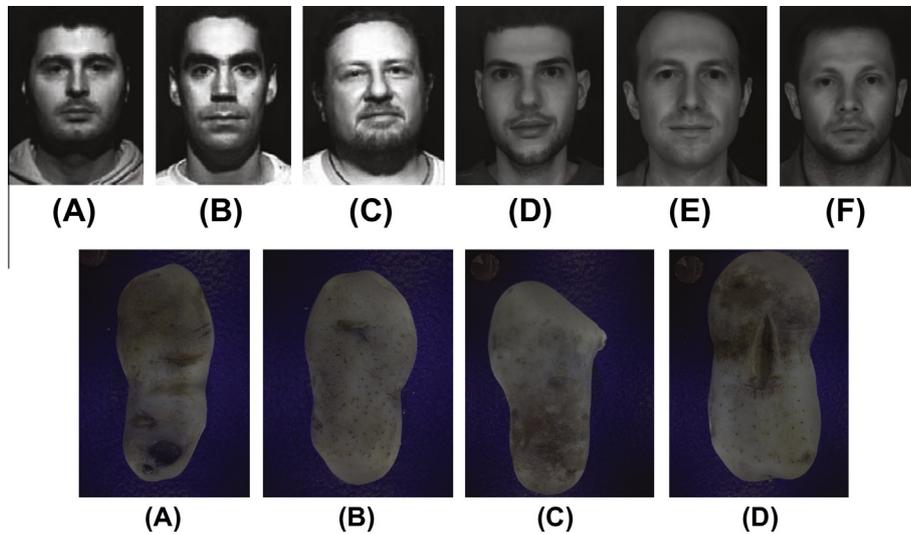


Fig. 35. Real faces and potatoes used for experiments.

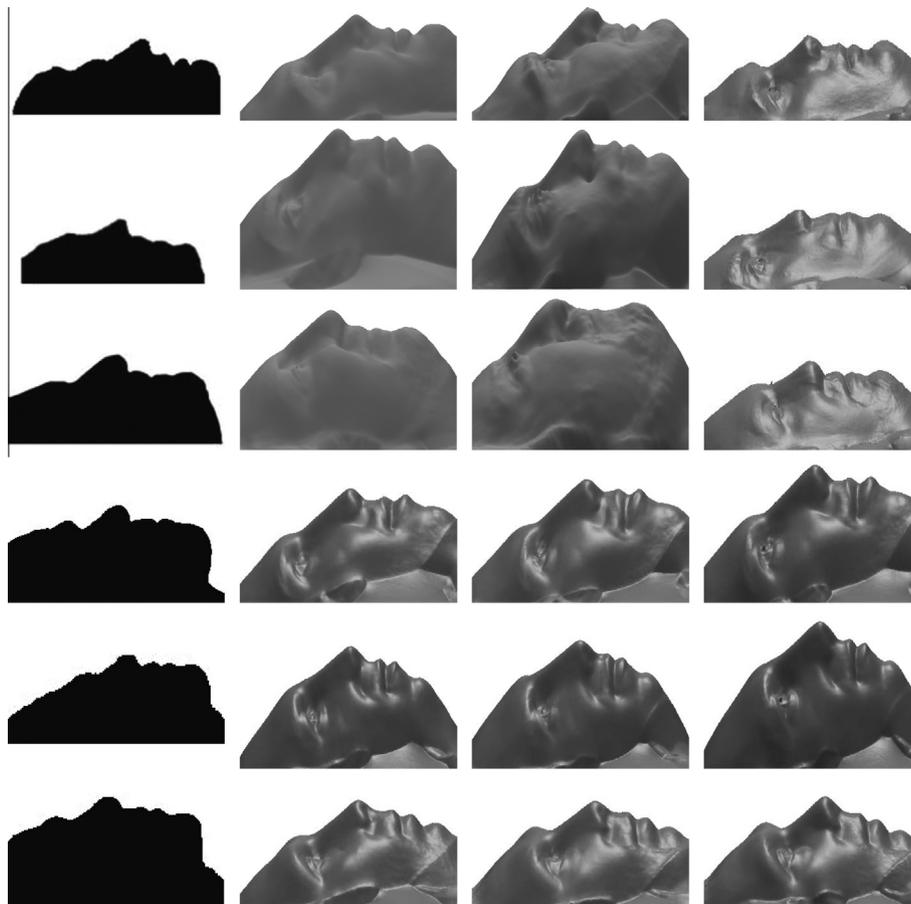


Fig. 36. From left to right, the first column corresponds to the ground truth, in the second column we have the profile view of the obtained 3D surfaces using the default illumination directions, in the third column we have the proposed method 1 and in the last one the proposed method 2.

optimal illumination conditions taking into consideration three parameters; and according to this approach the effects of shadows and highlights are minimised, while it tries to maximise the distinctiveness of surface detail. The second approach is based on the fact that denser isophotes in the gradient space result more accurate normals during the reconstruction process. This method tries to estimate the illumination directions that will

provide the denser isophotes to the areas with the most frequent normals.

Experiments with simulated and real surfaces were performed in order to evaluate the performance of the proposed schemes. From the results it could be inferred that both methods provide similar results with the first approach to outperform especially in the case of rough surfaces, mainly due to the third parameter in

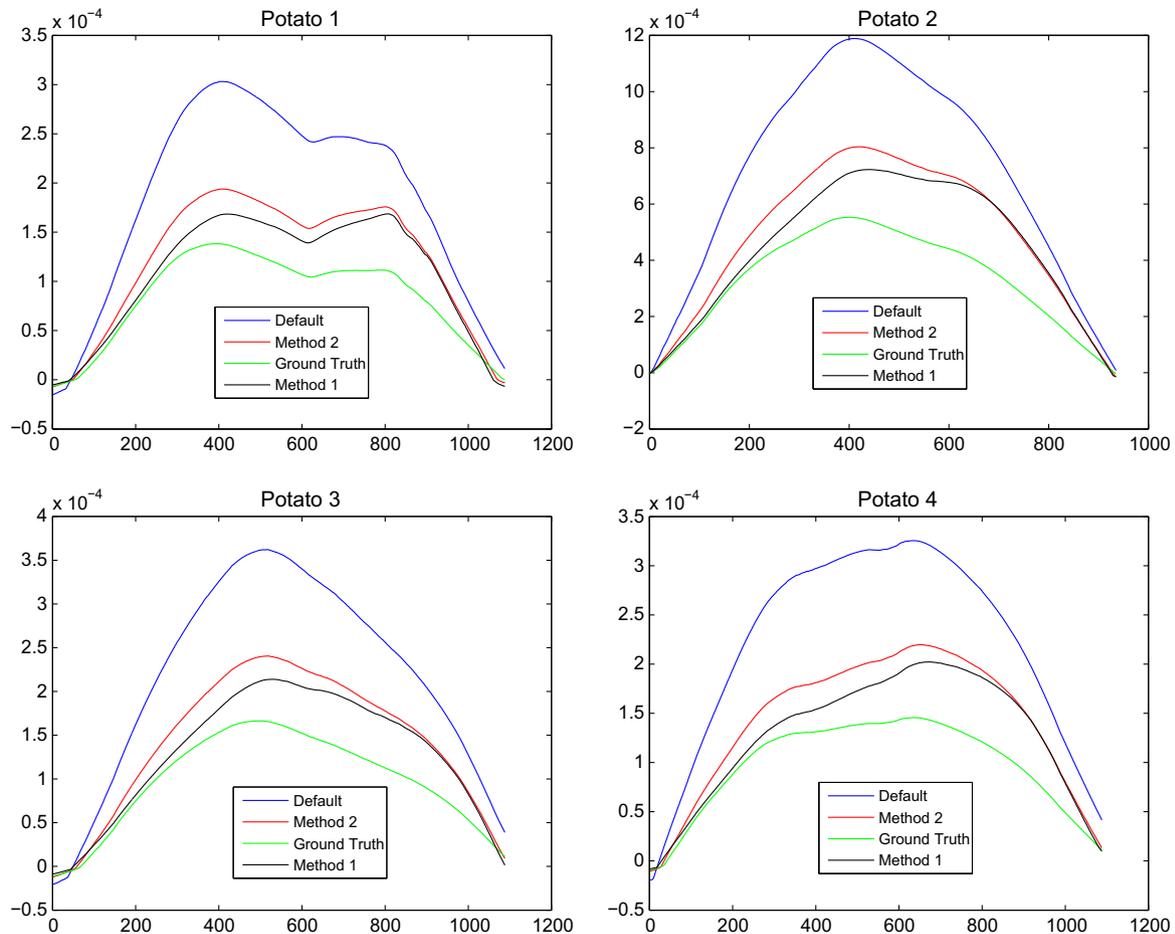


Fig. 37. The profile views of the obtained 3D surfaces using the default illumination directions, the proposed methods 1 and 2 and the ground truth for the four potatoes.

Table 14

The Hausdorff distance of the side views of the faces from the reconstructed side views for all the tested illumination directions. In bold, the best result for each face.

	Face A	Face B	Face C	Face D	Face E	Face F
Default	54.3386	54.7264	48.4010	56.70	22.45	61.92
Proposed A	41.1762	37.7257	36.2604	50.00	16.83	54.33
Proposed B	43.5891	41.2357	39.7544	50.15	18.28	54.53

Table 15

The Hausdorff distance of the side views of the potatoes from the reconstructed side views for all the tested illumination directions. In bold, the best result for each potato.

	Potato A	Potato B	Potato C	Potato D
Default	77.7950	239.2161	108.2240	111.2596
Proposed A	19.9113	74.8425	28.7071	28.5516
Proposed B	29.0997	100.3842	41.1328	41.1490

the optimisation function related to the surface roughness. Regarding the faces, the optimal directions are not uniformly distributed but the lights at the top and bottom of the face are closer compared to the sides. Similar outcome was obtained for the anisotropic surfaces but rotated according to the surface directionality. The surfaces with near isotropic facets the obtained distribution of the light sources is almost uniform, and this is further affected by the roughness of the surface. About the estimated zenith angle, it is at the same range for isotropic surfaces but in case of faces, it is significantly lower with the variations in general to depend on the morphology of the surface and the presence of cast shadows.

In order to further evaluate the obtained estimates, experiments were performed with real faces and other objects indicating the improvement in 3D reconstruction by selecting the appropriate illumination directions.

References

- [1] N. Alldrin, S. Mallick, D. Kriegman, Resolving the generalized bas-relief ambiguity by entropy minimization, in: Conf on Comp Vision and Pattern Recognition (CVPR), 2007, pp. 1–7.
- [2] J. Barron, D. Fleet, S. Beauchemin, Performance of optical flow techniques, *Int. J. Comput. Vis.* 12 (1) (1994) 43–77.
- [3] S. Barsky, M. Petrou, Shadows and highlights detection in 4-source colour photometric stereo, in: International Conference on Image Processing, vol. 3, 2001, pp. 967–970.
- [4] S. Barsky, M. Petrou, The 4-source photometric stereo technique for 3-dimensional surfaces in the presence of highlights and shadows, *IEEE Trans. Pattern Anal. Mach. Intell.* 25 (10) (2003) 1239–1252.
- [5] S. Barsky, M. Petrou, Design issues for a colour photometric stereo system, *J. Math. Imag. Vis.* 24 (1) (2006) 143–162.
- [6] P. Besl, N. McKay, A method for registration of 3-d shapes, *IEEE Trans. Pattern Anal. Mach. Intell.* 14 (2) (1992) 239–256.
- [7] M. Chandraker, S. Agarwal, D. Kriegman, Shadowcuts: photometric stereo with shadows, *CVPR* (2007).
- [8] M. Chantler, Why illuminant direction is fundamental to texture analysis, *VISP* 142 (4) (1995) 199–206.
- [9] E. Coleman, R. Jain, Obtaining 3-dimensional shape of textured and specular surfaces using 4-source photometry, *Comput. Vis. Graph. Image Process.* 18 (1982) 309–328.
- [10] O. Drbohlav, M. Chantler, On optimal light configurations in photometric stereo, in: IEEE International Conference on Computer Vision, ICCV, vol. 2, 2005, pp. 1707–1712.
- [11] J. Fan, L.B. Wolff, Surface curvature and shape reconstruction from unknown multiple illumination and integrability, *Comput. Vis. Image Understan.* 65 (2) (1997) 347–359.

- [12] G. Finlayson, M. Drew, C. Lu, Intrinsic images by entropy minimisation, in: Proc. ECCV, 2004, pp. 582–595.
- [13] A. Georghiadis, Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo, in: 9th International Conference on Computer Vision, vol. 2, 2003, pp. 816–825.
- [14] A. Georghiadis, Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo, in: Ninth IEEE International Conference on Computer Vision (ICCV'03), vol. 2, 2003, pp. 816–823.
- [15] C. Gullon, Height Recovery of Rough Surfaces from Intensity Images, Ph.D. Thesis; School of Engineering and Physical Sciences, Heriot-Watt University, 2002.
- [16] S. Gumhold, Maximum entropy light source placement, *IEEE Visual*. (2002) 275–282.
- [17] A. Hertzmann, S. Seitz, Example-based photometric stereo: shape reconstruction with general, varying brdfs, *IEEE PAMI* 27 (8) (2005) 1254–1264.
- [18] S. Kautkar, G. Atkinson, M. Smith, Face recognition in 2d and 2.5d using ridgelets and photometric stereo, *Pattern Recognit.* 45 (2012) 3317–3327.
- [19] B.H. Kim, R.H. Park, Multi-image photometric stereo using surface approximation by legendre polynomials, *Pattern Recognit.* 31 (8) (1998) 1033–1047.
- [20] C. Lee, X. Hao, A. Varshney, Light collages: lighting design for effective visualization, *IEEE Visual*. (2004) 281–288.
- [21] C.H. Lee, A. Rosenfeld, An approximation technique for photometricstereo, *Pattern Recognit. Lett.* 2 (5) (1984) 339–343.
- [22] Kuang-Chih Lee, J. Ho, D. Kriegman, Nine points of light: acquiring subspaces for face recognition under variable lighting, in: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001), vol. 1, 2001, pp. I-519–I-526.
- [23] K. Lee, C. Kuo, Shape reconstruction from photometric stereo, *Comput. Vis. Pattern Recognit.* (1992) 479–484.
- [24] M. Levine, J. Bhattacharyya, Removing shadows, *Pattern Recognit. Lett.* 26 (3) (2005) 251–265.
- [25] G. McGunnigle, M. Chantler, Resolving handwriting from background printing using photometric stereo, *Pattern Recognit.* 36 (8) (2003) 1869–1879.
- [26] N. Nadtoka, J. Edge, P. Jackson, A. Hilton, Representing dynamics of facial expressions, *CVMP* (2006).
- [27] S. Nayar, K. Ikeuchi, T. Kanade, Determining shape and reflectance of hybrid surfaces by photometric sampling, *IEEE Trans. Robot. Automat.* 6 (4) (1990) 418–431.
- [28] H. Ragheb, E. Hancock, A probabilistic framework for specular shape-from-shading, *Pattern Recognit.* 36 (2) (2003) 407–427.
- [29] R. Shackled, D. Lischinski, Automatic lighting design using a perceptual quality metric, EG 2001 Proceedings of Computer Graphics Forum, vol. 20(3), 2001, pp. 215–226.
- [30] W. Smith, E. Hancock, Estimating cast shadows using sfs and class-based surface completion, in: Proceedings of the 18th International Conference on Pattern Recognition, vol. 4, 2006, pp. 86–90.
- [31] F. Solomon, K. Ikeuchi, Extracting the shape and roughness of specular lobe objects using four light photometric stereo, *IEEE Trans. Pattern Anal. Mach. Intell.* 18 (4) (1996) 449–454.
- [32] A. Spence, M. Chantler, Optimal illumination for three-image photometric stereo acquisition of texture, in: Proceedings of the 3rd International Workshop on Texture Analysis and Synthesis, 2003, pp. 89–94.
- [33] A. Spence, M. Chantler, Optimal illumination for three-image photometric stereo using sensitivity analysis, *IEE Proc. Vis. Image Signal Process.* 153 (2) (2006) 149–159.
- [34] J. Sun, M. Smith, L. Smith, S. Midha, J. Bamber, Object surface recovery using a multi-light photometric stereo technique for non-lambertian surfaces subject to shadows and specularities, *Image Vis. Comput.* 25 (7) (2007) 1050–1057.
- [35] H. Tagare, J. deFigueiredo, A theory of photometric stereo for a class of diffuse non-lambertian surfaces, *IEEE Trans. Pattern Anal. Mach. Intell.* 13 (2) (1991) 133–152.
- [36] P. Tan, S. Lin, L. Quan, Subpixel photometric stereo, *IEEE Trans. Pattern Anal. Mach. Intell.* 30 (8) (2008) 1460–1471.
- [37] J.R. Torrao, A green's function approach to shape from shading, *Pattern Recognit.* 34 (12) (2001) 2367–2382.
- [38] V. Argyriou, S. Barsky, M. Petrou, Generalisation of photometric stereo technique to q -illuminants, *Proc. BMVC* (2008).
- [39] V. Argyriou, M. Petrou, Recursive photometric stereo when multiple shadows and highlights are present, *Proc. CVPR* (2008).
- [40] R. Woodham, Photometric stereo: a reflectance map technique for determining surface orientation from image intensities, *SPIE* 155 (1978) 136–143.
- [41] R. Woodham, Photometric method for determining surface orientation from multiple images, *Opt. Eng.* 19 (1) (1980) 139–144.
- [42] T. Zickler, P. Belhumeur, D. Kriegman, Helmholtz stereopsis: exploiting reciprocity for surface reconstruction, *Int. J. Comput. Vis.* 49 (2/3) (2002) 215–227.