

Extended Photometric Sampling for Surface Shape Recovery

Felipe Hernández-Rodríguez and Mario Castelán

Centro de Investigación y de Estudios Avanzados del I.P.N.
Robotics and Advanced Manufacturing Group,
Ramos Arizpe, Coahuila, 25900, México

Abstract. Photometric sampling is a process where the surface normals of an object are estimated through the excitation of the object's surface and a rotating light source around it. The method can be regarded as a special case of photometric stereo when extensive sampling is performed in order to calculate surface normals. The classic photometric sampling approach considers only variations around the azimuth angle of the moving light source. As a consequence, additional attention has to be paid to the recovery of the light source directions and the removal of specular and shadowed regions. This paper investigates the effect of including variations around the zenith angle of the light source vector in a photometric sampling framework, developing a geometric approach to estimate the surface normal vectors. Experiments show that increasing the number of samples along the zenith variation benefits the estimation of the surface normals.

1 Introduction

The study and analysis of reflectance for recovering shape is an active research topic in computer vision and pattern recognition. Probably, the first analytical tool for understanding the relationship of lighting and surface shape was the reflectance map, proposed by Horn in [1]. The reflectance map defines the observed intensity at a surface patch (pixel) as a function of the surface gradient and the light source direction. The concept of gradient space, whose axes are defined by the first partial derivatives of the surface is introduced for associating an image intensity with local surface orientation, depending on the direction of a light source. However, reflectance behavior becomes hard to analyze when departures from Lambert's law occur, which is the case of many real world materials.

Reflectance can also be modeled by subspace analysis methods. The idea of linearly decomposing a set of images was initially motivated by the eigenfaces of Kirby and Sirovich [2], where it was observed that the first few eigenimages explained information related to variability in illumination. The method has been widely studied in the computer vision area, being the work of Kriegman et al [3], [4], [5], [6] which probably has focused the most in the subspace analysis of imagery exhibiting lighting variations. Another example can be found in [7],

where a method based on Principal Component Analysis (PCA) for the analysis of illumination variability, including shadows, was developed.

It is worth commenting on the data acquisition process commonly carried out for photometric analysis purposes. Normally, a photometric rig consisting of carefully designed static arrangements of light sources is used. The fixed nature of these settings may translate into a lack of image samples. Sophisticated lighting structures have also been developed in order to gather as much samples as possible [4], but building them is a complicated task, mainly due to space requirements. Although an extensive gathering of data does not guarantee a perfect reconstruction, the lack of sampling usually means lack of information to either invert the problem or at least recover useful statistical properties. For example, specularities (saturated pixel values) and shadows (absence of illumination due to surface relief) are hard to detect when no extra data is available [8,9,10]. In this paper, the photometric analysis is revisited from a light source mobility point of view, with the aim of minimizing uncertainties introduced by photometric inconsistency.

This paper investigates the effect of including variations around the zenith angle of the light source vector in a photometric sampling framework, developing a geometric approach to estimate the surface normal vectors. To this end, a sine-fitting based strategy is proposed two-wise. The first fitting improves the stability of the luminance response for each pixel. Later, a second sinusoidal fitting is introduced in order to cope with missing data and find the zenith angle of the surface normal. In our approach, the calibration of both direction and intensity of light sources is not considered, neither the illumination conditions of the data acquisition process are carefully controlled, which can be considered as a trade-off between calibration and sampling.

The paper is organized as follows: in Section 2, the photometric sampling approach is introduced and the relevant literature provided; in Section 3, an approach for dealing with variations in the zenith angle of the light source direction is described; experiments are presented in Section 4 and finally conclusions are described in Section 5.

2 Photometric Analysis

The Lambertian reflectance law states that, at image coordinates (u, v) , the observed irradiance intensity $i(u, v)$ is proportional to the cosine between the surface normal vector $\mathbf{n}(u, v) \in \mathcal{R}^3$ and the light source direction vector $\mathbf{l}(u, v) \in \mathcal{R}^3$. The surface normal is the vector orthogonal to the tangent plane over the surface height $h(u, v)$, with first partial derivatives $p(u, v) = \frac{\partial h(u, v)}{\partial u}$ and $q(u, v) = \frac{\partial h(u, v)}{\partial v}$. For the sake of clarity, it will be assumed that events occur independently over each pixel location and omit the (u, v) notation. The image irradiance equation can now be simply formulated as

$$i = \langle \mathbf{n}, \mathbf{l} \rangle, \quad (1)$$

with “ $\langle \cdot \rangle$ ” denoting inner product. Note that the surface normal is $\mathbf{n} = [p, q, -1]^T$, while its norm $\|\mathbf{n}\|$ roughly represents the degree of lighting reflectance of the surface material at each surface patch. This scalar is commonly referred to as albedo. When the albedo is equal to one, surface reflectance behaves as a perfect Lambertian diffuser.

Other common assumptions in photometric based techniques are that the light source use the parallel source model in which the light source direction is at infinity (it illuminates evenly all the observed points in the object) and its norm is also equal to one. It is important to mention that Eq. 1 denotes the simplest BRDF, i.e., a function attempting to model the behavior of the outgoing luminance before the incoming light.

2.1 Using Additional Light Sources

Estimating the surface normal from a single image pixel value is extremely difficult, even with a good approximation of the light source direction and under perfectly Lambertian reflectance. This problem is known as Shape-from-Shading (SFS) [1] and it has been extensively studied for decades as it tries to deal with the simplest 3D shape recovery problem: using a single image view. The SFS formulation was the first to treat the image formation process as a photometric phenomenon.

SFS is an ill-posed problem, i.e., the knowledge of the intensity pixel value i and the light source vector \mathbf{l} does not suffice to uniquely determine the surface normal \mathbf{n} in Eq. 1. Knowledge about at least two additional non collinear light sources constrains the problem to a unique solution. From the system of equations $i_1 = \langle \mathbf{n}, \mathbf{l}_1 \rangle$, $i_2 = \langle \mathbf{n}, \mathbf{l}_2 \rangle$ and $i_3 = \langle \mathbf{n}, \mathbf{l}_3 \rangle$, the problem can be expressed as a matrix inversion, i.e.,

$$\mathbf{n} = \mathbf{L}^{-1}\mathbf{i}, \quad (2)$$

with $\mathbf{L} = [\mathbf{l}_1^T, \mathbf{l}_2^T, \mathbf{l}_3^T]$ and $\mathbf{i} = [i_1, i_2, i_3]^T$. Note how for more than three light source vectors, a generalized pseudoinverse of \mathbf{L} also provides a unique solution.

The photometric stereo method (PSM) is resumed in Eq. 2. The method states that, given at least three images of a fixed pose object illuminated by three known non-collinear light source vectors, in the absence of shadows and under Lambertian reflectance, it is possible to uniquely estimate the surface normal for each surface patch. The estimation of the light source direction, however, is a challenging problem in this scheme.

A different manner to estimate surface normals from imagery is using a statistical approach, i.e., through the orthogonal diagonalization or the singular value decomposition of a matrix of long image vectors. As shown in [4], the first three left singular vectors are related to the main changes in illumination, which may describe the orthogonal relationship of the surface normal vectors with the surface patches. Additional restrictions on the integration method should be included as, normally, the estimated surface normals depend on the long matrix vector and therefore on the sample images to compose this matrix.

2.2 Photometric Sampling

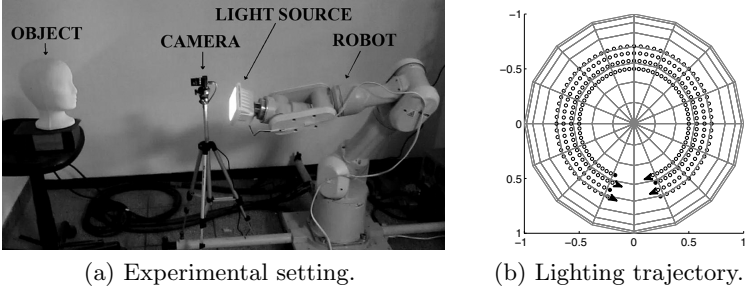
The term *photometric sampling* was introduced by Nayar et al. in [11] as an effort to recover shape by relating the sampling of a photometric function with surface orientation, reflectance and light sources. To this end, the authors proposed to measure (sample) the photometric response of a still object while a light source, placed on a turntable, moved around a single circular path. The main idea of the method was to obtain the response, at the pixel level, of the surface normal interacting with a moving light source vector. Given the circular nature of the path, the expected response of this interaction is, ideally, a sinusoidal function (see Figure 1 (b)). Unfortunately, the reflectance of real world objects includes artifacts such as specular and shadowed regions, departing from the expected sinusoidal behavior and therefore breaking the Lambertian assumption.

For this reason, efforts have been oriented to dealing with departures from Lambert's law. For example, Saito et al. [12] and more recently Liu et al. [13] borrowed the main idea of the photometric sampling to estimate the surface normals of a specular object. Assuming knowledge of the light source direction, the former described a least-squares fitting procedure applied on each RGB channel in order to avoid specular regions and calculate surface normals [12], while a similar procedure using the Hough transform for sine fitting was proposed by the latter [13]. Later, Miyasaki and Ikeuchi studied the effect of identifying shadows and specular reflections as outliers [14]. Aided by the Singular Value Decomposition (SVD) and the graph-cut algorithm, the method excluded outliers from the database before performing PSM, i.e., the pixels in the image that contained the outliers were removed to later apply SVD in a sub-matrix. A robust PCA approach for removing specularities and shadows in photometric stereo databases was also proposed in [15].

A common feature in the methods described above [12,13,14] is that they focus on the variation of the azimuth angle of the light source, not considering the reflectance response regarding variations on the zenith angle, i.e., these approaches would only use one of the circular paths shown in Figure 1 (b). In this paper, we explore the effect of including variations around the zenith angle of the light source vector in a photometric sampling framework. To this end, instead of performing fitting on a single luminance vector, the method introduced in the present paper includes the sampling of several circles around the object to get more reliable information, while a sine function is fitted on the singular vectors of the luminance matrix to improve the quality of the database. This extension allows us to develop a geometric approach for surface normal calculation, as it will be described in the next section.

3 Extended Photometric Sampling

In accordance with the gathered images during the data acquisition process, the intensities of a pixel will draw a sinusoidal function depending on the illumination changes generated around the object. This sinusoidal behavior depends on the interaction of the surface normal with the the light source direction variation.



(a) Experimental setting.

(b) Lighting trajectory.

Fig. 1. Data acquisition process scheme. In (a), the elements of the database acquisition scheme are shown, while (b) presents a projection of the light source positions, for the four zenith and the sixty-five azimuth angles. The separation between samples was 5° .

The sine curve can be decomposed in the three parameters amplitude \mathbf{A} , phase \mathbf{B} and shift \mathbf{C} as follows:

$$I(\theta) = \mathbf{A} \sin(\theta + \mathbf{B}) + \mathbf{C}, \quad (3)$$

where $I(\theta)$ is the pixel intensity at each θ variation in azimuth.

In our experiments, the data acquisition process was performed by the Mitsubishi PA-10 robotic arm in order to generate the illumination changes in concentric patterns around the optical axis of the camera. Figure 1 (b) shows the illumination patterns applied to the data acquisition process and the general scheme for data acquisition used by the robot (a).

A database of $k = k_1 \times k_2$ images of the observed object under k different light source directions is constructed, with k_1 azimuth angles and k_2 zenith angles of the light source direction vector. It is important to remark that, due to the concentric circle patterns, the set of sampled images should reflect the two main variations (i.e., changes related to azimuth and zenith) of the light source direction vector at each trajectory point.

The photometric analysis commences by generating, for each pixel, a matrix $\mathbf{M}_{k_1 \times k_2}$ storing the k pixel intensity values recorded at each (azimuth, zenith) configuration pair. This matrix, which will be called the *luminance matrix*, contains the pixel reflectance history along the two main variations of the light source trajectory. For every pixel, the observed reflectance may be decomposed by the principal axis of the luminance matrix. The study of these axis allows identifying regions which best fit a sinusoidal behavior, i.e., close to a Lambertian behavior. The signal is finally corrected once the sine curve parameters have been calculated and the signal replaced with a sine function. We use SVD to decompose the luminance matrix,

$$\mathbf{M} = \mathbf{U}_{k_1 \times r} \mathbf{\Sigma}_{r \times r} \mathbf{V}_{r \times k_2}^T, \quad (4)$$

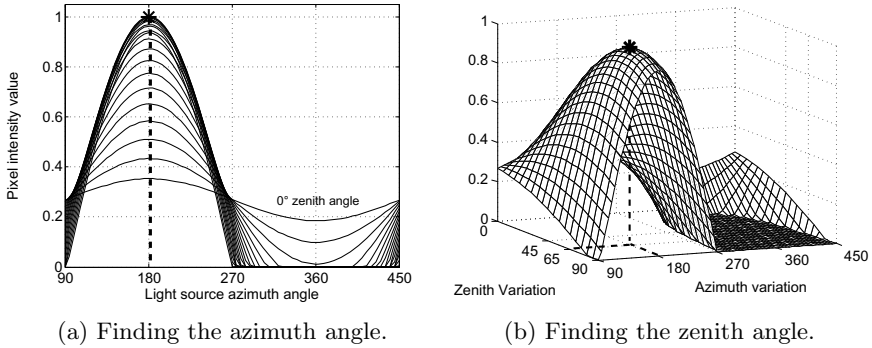



Fig. 2. Lighting variation (luminance matrix) for a pixel of a Lambertian sphere. In (a), a projection of the luminance matrix illustrates the chosen azimuth angle of the surface normal related to the pixel. A 3D plot of the luminance matrix shows how the zenith angle of the surface normal is determined in (b). For both cases, the selected angular values are shown with a dashed line. The black star is the maximum intensity value of the pixel, then the surface normal can be computed for the $(180^\circ, 65^\circ)$ angular pair.

where $r = \text{rank}(\mathbf{M})$. The column and row spaces of \mathbf{M} are decomposed into the orthogonal basis \mathbf{U} (left singular vectors) and \mathbf{V} (right singular vectors), respectively. The singular values, contained in the diagonal elements of $\mathbf{\Sigma}$ explain the degree of retained variability in both right and left singular vectors. We studied the effect of performing sine fitting in the singular vectors of the luminance matrix of all image pixels along the database. The singular vectors of the luminance matrix are preferred against the raw column vectors of the matrix since artifacts or outliers will tend to be filtered out due to the weighting of singular values.


Since real images may include noisy variations in reflectance, using least squares for calculating the three sine parameters (amplitude, phase and shift) leads to poor estimations. In order to overcome this problem, we developed a RANSAC-based fitting technique [16], which randomly adjusts sine functions to the data, choosing the sine parameters with the major consensus. Once the sine parameters are estimated over each of the singular vectors, a new luminance matrix $\mathbf{M}' = \mathbf{U}'\mathbf{\Sigma}'\mathbf{V}'^T$ is generated to improve photometric consistency on the database. The new fitted columns of \mathbf{U}' and \mathbf{V}' contain the sine-fitted singular vectors from the original matrices \mathbf{U} and \mathbf{V} in Eq. 4.

3.1 Surface Normal Estimation

We propose a geometric strategy to determine the azimuth-zenith angle pair $(\theta - \phi)$ from the new sine-adjusted matrix \mathbf{M}' . Recall that the Cartesian surface normal can be transformed by $[n_x, n_y, n_z] = [\cos(\theta)\sin(\phi); \sin(\theta)\sin(\phi); \cos(\theta)]$. In Figure 2, plots of the luminance matrix depict how the angular values of the



(a) Finding the azimuth angle.



(b) Finding the zenith angle.

Fig. 3. Lighting variation (luminance matrix) for a pixel of a Lambertian sphere with missing data. The azimuth angle can be selected from the maximum values of the illumination response (a). These values are then used to fit a sine function from which the maximum value corresponds to the sought zenith angle (b). As in Figure 2, the selected angular values are shown with a dashed line. The black star is the maximum intensity value of the pixel, then the surface normal can be computed for the $(180^\circ, 65^\circ)$ angular pair.

surface normal are selected. To illustrate the case, we use a synthetic luminance matrix containing the complete range of angular values for both the azimuth (360°) and zenith (90°). In (a), it can be noticed that the azimuth angle corresponds to the maximum response of the sinusoidal function found in the columns of the luminance matrix. This maximum value is related to collinearity with the light source direction vector and the surface normal vector. The calculation of the zenith angle is depicted in (b) as the global maximum of the surface function related to the luminance matrix.

For the databases used in our experiments, the luminance response for the complete azimuth-zenith range was not available (due to articular limitations of the robotic arm), therefore, a strategy to cope with missing data was developed. The main idea here is to predict the maximum pixel intensity from the sine-adjusted luminance matrix.

The algorithm follows three steps (see Figure 3 for a visual explanation):

1. Find the maximum value for each intensity sine curve.
2. Fit a sine function onto the set of maximum values selected on the step 1.
3. Set the maximum value of the fitted sine function as the zenith value of the surface normal.

In the next section, experimental evaluation considering two different reflectance responses will be provided, applying the extended photometric sampling approach in order to recover a height map.



Fig. 4. Some sample images from the acquired databases. Top row: matte-painted styrofoam mannequin (for a Lambertian-like reflectance). Bottom row: styrofoam mannequin (for a specular-like reflectance). For both rows, variations along the x -axis are shown.

4 Experiments

The proposed method has been applied on two databases. The first database (matte-painted styrofoam mannequin) is related to Lambertian reflectance. Some of the images in this database are shown in the top row of Figure 4. The second database resembles the specular case, i.e., the same styrofoam mannequin was used with no matte painting, as styrofoam highly reflects lighting variations. The bottom row of Figure 4 presents an image subset from this database. The experimental setting is shown in Figure 1 (a). The illumination angles for generating the database ranged from 110° to 430° in azimuth and from 15° to 30° in zenith, each sample being separated by 10° and 5° , respectively. Each database contained 65 image samples per circle in 4 concentric circles, for a total of 260 image samples per database. Every pixel in the image was analyzed in order to estimate the surface normal from the lighting variations.

For the purposes of experimental comparison, a surface normal field was obtained in three different ways: using our proposed extended photometric sampling, using statistical PSM and using classic PSM. The normal fields were later integrated using the Frankot and Chellappa integration method [17] with Discrete Cosine Transform (DCT) basis functions. All surface normals were set unit length before integration. The ground truth surface for the mannequin was obtained using a 3D laser range scanner Vivid 910 Konika Minolta. Figure 5 presents different integrated height maps and the ground truth. The first row of the image shows results related to the matte-painted mannequin while the second row depicts results under the specular mannequin. For all the cases, a frontal illumination image of the database is warped onto the height maps in order to emphasize the visual comparison of the recovered surfaces and its accordance with the ground truth surface.

It can be noticed from Figure 5 that the extended photometric sampling results shown in (b) deliver a better shape agreement with the ground truth than the PSM results shown in (c) and (d). Recall that the data acquisition process was carried out in a non-structured ambient, i.e., the reflectance of the

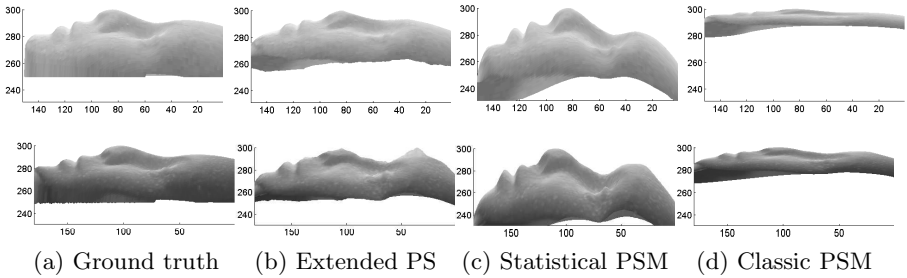


Fig. 5. Recovered height maps from surface normal fields. The first row of the image depicts the experiment with the matte-painted styrofoam mannequin, while the second row presents results on the unpainted (specular-like) styrofoam mannequin.

material and the ambient effects distort the lighting appearance of the object. Under these conditions, unless a careful calibration is performed in both positioning and intensity of the light sources, the poor performance of classic and statistical PSM is expected. This is one of the main motivations of our approach: increasing the number of samples against performing calibration tasks. The extended photometric sampling is then aimed at relaxing the lighting calibration of the data acquisition stage while adding variations in the zenith angle of the light source. Nonetheless, Figure 5 also reveals that the extended photometric sampling struggles to find a correct surface normal approximation for the specular case, as it can be observed along the forehead and nose areas of the mannequin (second row). This may be due to the fact that such areas maintained a high degree of specularity during the whole sampling process, from which the sinusoidal fitting was not enough to approximate the Lambertian reflectance of the material. This suggests that additional constraints should be included into the fitting scheme, in order to improve robustness of fitting.

5 Conclusions

We have presented an extended version of the photometric sampling method for estimating surface normal maps. The extension consists of the inclusion of variations towards the zenith angle of the light source direction. A strategy based on fitting sinusoidal functions has been described, on one hand, maximizing stability in the reflectance response of the data, on the other hand, calculating the azimuth and zenith angles of the surface normals by finding the maximum values of the luminance matrix. The method has proved to deliver reasonable approximations of the object’s surface, compared to results provided by a 3D laser range scanner, while avoiding calibration of light sources and controlled conditions in the acquisition environment. As future work we are planning to improve robustness of sine fitting by adding constraints to deal with specular and shaded regions.

References

1. Horn, B.K.P., Sjöberg, R.W.: Calculating the reflectance map. In: A.I. Memo, pp. 215–244. MIT (1979)
2. Sirovich, L., Kirby, M.: Low-dimensional procedure for the characterization of human faces. *Journal of the Optical Society of America* 4, 519–524 (1987)
3. Lee, K.-C., Ho, J., Kriegman, D.: Ninepoints of light: Acquiring subspaces for face recognition under variable lighting. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pp. 519–526 (2001)
4. Georgiades, A., Belhumeur, D., Kriegman, D.: From few to many: Illumination cone models for face recognition under variable lighting and pose. *IEEE Trans. Pattern Analysis and Machine Intelligence* 23(6), 634–660 (2001)
5. Ho, J., Yang, M., Lim, J., Lee, K., Kriegman, D.: Clustering appearances of objects under varying illumination conditions. In: Proc. IEEE Conf. on Computer Vision and Pattern Recognition, pp. 11–18 (2003)
6. Lee, K.-C., Ho, J., Kriegman, D.: Acquiring linear subspaces for face recognition under variable lighting. *IEEE Trans. Pattern Analysis and Machine Intelligence* 27(5), 684–698 (2005)
7. Ramamoorthi, R.: Analytic pca construction for theoretical analysis of lighting variability, including attached shadows, in a single image of a convex lambertian object. *IEEE Trans. Pattern Analysis and Machine Intelligence* 24, 1322–1333 (2002)
8. Barsky, S., Petrou, M.: The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Trans. Pattern Analysis and Machine Intelligence* 25(10), 1239–1252 (2003)
9. Argyriou, V., Petrou, M.: Recursive photometric stereo when multiple shadows and highlights are present. In: Proc. CVPR, pp. 1–6 (2008)
10. Hernández, C., Vogiatzis, G., Cipolla, R.: Shadows in Three-Source Photometric Stereo. In: Forsyth, D., Torr, P., Zisserman, A. (eds.) ECCV 2008, Part I. LNCS, vol. 5302, pp. 290–303. Springer, Heidelberg (2008)
11. Nayar, S.K., Ikeuchi, K., Kanade, T.: Shape and reflectance from an image sequence generated using extended source. In: Proceedings of IEEE ICRA, vol. 1, pp. 28–35 (1989)
12. Saito, H., Omata, K., Ozawa, S.: Recovery of shape and surface reflectance of specular object from relative rotation of light source. *Image and Vision Computing* 21, 777–787 (2003)
13. Liu, R., Han, J.: Recovering surface normal of specular object by hough transform method. *IET Computer Vision* 4(2), 129–137 (2010)
14. Miyasaki, D., Ikeuchi, K.: Photometric stereo under unknown light sources using robust svd with missing data. In: Proc. ICIP, pp. 4057–4060 (2010)
15. Candès, E.J., Li, X., Ma, Y., Wright, J.: Robust Principal Component Analysis?, ArXiv e-prints (2009)
16. Fischler, M.A., Bolles, R.C.: Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Commun. ACM* 24, 381–395 (1981)
17. Frankot, R., Chellappa, R.: A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. Pattern Analysis and Machine Intelligence* 10, 438–451 (1988)