

Spread-Sheet Execution of Shape-from-Shading for Human Back Surface Measurement

Harvey MITCHELL*
University of Newcastle, Newcastle, Australia
<http://dx.doi.org/10.15221/13.203>

Abstract

Shape-from-shading is a shape measurement technique which features simple data collection. However, it is restricted in the types of objects it can measure. These characteristics do make it suitable for some types of body surface measurement, but it does not yet seem to have found many applications. This paper looks at whether the simple data collection can be augmented by equally straightforward data processing to make it feasible for scoliosis monitoring. Experiments have confirmed the basic concept of back surface measurement by SFS, but that further practical refinement is needed.

Keywords: shape-from-shading, human back measurement.

1. Introduction

Shape-from-shading uses just a single image from a camera to determine the three-dimensional shape of an object. The reflectance levels in the image are used to infer the gradients or slopes across the visible area of the object, and these slopes are integrated to reconstruct the shape of that object. The method's peculiar advantage is that the data collection is simple: it requires nothing more than a single image. In principle at least, the camera can be quite ordinary and cheap, although that may mean that high accuracy can be harder to achieve; better cameras (in terms of both radiometric and geometric characteristics) do facilitate higher accuracy. The substantial disadvantage of SFS for surface shape measurement in general is that the method is restricted to certain suitable surfaces. Firstly, the surfaces must have even light (and preferably matt) colouring, and secondly objects must, in practice, have smoothly continuous physical texture.

SFS does not seem to have found many measurement applications, but it is arguable that SFS should be valuable when simple and instant imaging is demanded, and when that requirement for suitable evenly coloured and matt textured surfaces without discontinuities is not restrictive. These circumstances seem to make it suitable for applications in body surface shape measurement. In particular, it is seen to be appealing for some body measurement tasks for which expensive processing is not justified. Body shape measurement for apparel fitting may be appropriate, but in this work, emphasis has been given to medical measurements when simple measurement is appropriate, notably when only occasional measurement is needed, when expensive (and space-consuming) equipment does not have to be purchased or stored, nor its use relearnt at each occasional use. This is suited to determining the topography of backs in particular, by either general practitioners or specialists (or their technicians) who wish to monitor or screen for scoliosis. It is appropriate when measurement is not crucial to saving life. Allied medical applications may be found in the measurement of casts, orthotics, prostheses, and some motion monitoring.

There are other features that deserve to be recognised. The method's requirement for just a single image means that the object does not have to be held perfectly still during the measurement, which can be crucial for most human measurement. Given that people move and change stance, and even change body shape over time, high accuracy in these cases can be pointless. The method does not allow measurement through clothing, but medical applications have the advantage that patients may be prepared to disrobe. Imagery could be processed by technical assistants or by a bureau service. A flash can be inconvenient, but is adopted in this work to avoid the complication of users setting-up any special lighting, or even changing room lighting.

This project is primarily aimed at developing SFS for back surface measurement for which a major design constraint is imaging which is quick and simple, using a cheap and familiar camera with a flash, to achieve three-dimensional measurement to an acceptable accuracy in the three directions. It is assumed that simple (and cheap) imaging is such a valuable feature of SFS, that SFS must be associated with processing which is equally simple to execute. Accordingly, a principal goal of this work is to determine whether it is possible to execute image processing simply - using a spread-sheet routine.

* harvey.mitchell@newcastle.edu.au; +61- 249216055; www.newcastle.edu.au

2. Theory

SFS theory assumes Lambertian reflection, without any mirror-like specular reflection. According to the Lambertian theory, the strength of the light which is returned to an observer from an object point which is lit by a beam of light is a function of the cosine of the total of two angles: the incident angle (between the surface normal and the ray from the light source to the surface point of interest) and the reflected angle (between the surface normal and the ray to the sensor relative to the surface point), [1]. There is no assumption that the reflected angle is related to the incident angle. In this case, when a flash source of light is used, the light beam is assumed to emanate from the lens centre, and to return on the same path, so the incident angle, i , and the reflected angle are in fact the same. Then, the reflectance strength R at the small surface element which is imaged by any pixel, is a function of $2i$, as in equation 1,

$$R = R_0 \cos(2i) \quad (1)$$

where R_0 is most easily described as the maximum reflection obtained at the point where i is zero. If the angle i for any pixel is determined from R by equation 1, (provided that R_0 is known) allow the direction of the normal to the surface element to be related to the direction of the image plane, by ordinary three-dimensional image geometry. (This use of a single image distinguishes SFS from "photometric stereo" and photogrammetry, both of which require multiple images.)

3. Practicalities

The practical viability of SFS is affected by a number of matters. In particular, it is theoretically impossible to deduce a three-dimensional shape from a two dimensional image, unless some assumptions are made about the object: each pixel of the image supplies just one piece of information, whereas the determination of three-dimensional shape requires two pieces of information about the direction of the normals at point on the object (a slope in two directions for example, or the slope of the surface and the direction of that slope). There are means of overcoming this problem, but they can be upset by discontinuities in the object. As well, it is crucial to recognise that the above equation models the variation in reflectance from a single point if the sensor moves; it does not describe the variation in reflectance for different points across the object as recorded at the different image pixels. As well, surface shape reconstruction by integration of surface normal directions can be problematical as errors will accumulate. Discontinuous surfaces are virtually impossible to measure. Scale is not automatically provided. These problems and other practical problems have revealed themselves in the course of the experiments carried out for this work, and many are reported below, following an outline of some of the experiments.

4. Experiments

Experience with SFS theory and practice has been gathered by imaging various objects, at scales from a few millimetres to a metre, using three different cameras: a high quality Flexiscope Piccolo intraoral dental camera, which has two in-built light sources close to the lens; a Canon EOS30D single-lens-reflex digital camera with a 28 mm Canon lens, with in-built flash; and a Panasonic Lumix DMC-ZR1 also with in-built flash.

Experimental results are broadly similar for all cameras, but this report refers to three enlightening experiments based on the latter two cameras. The Lumix camera has a 640 x 480 pixel sensor, but images were resized to 144 x 108 for easier handling at the experimental stage; the Canon camera has a 3504 x 2336 pixel sensor, but images were resized to 144 x 96.

6.1. Case 1: Concept verification with a plane object

A plane object was used to test the theory and its execution within the spread-sheet programs, while appreciating that this does not verify that measurement of complex surfaces will be successful. The image of the plane painted wall is shown in Figure 1. Details of the experiment are shown in Table 1.

The slope of the plane (given as the angle between camera sensor plane and plane of object) was measured to be 30° , and it is seen here to be given by the inverse tan of the height difference on the reconstructed object shown in figure 3 (55 pixels) divided by the horizontal length of the object (115 pixels). That is, slope γ is given by

$$\begin{aligned} \gamma &= \tan^{-1}(55/115) \\ &= 26^\circ \end{aligned}$$

Figures 2 and 3 show the image with reflectance levels quantified, and the final reconstructed surface.

Table 1: Details of the experiment involving a plane object.

Camera:	Panasonic Lumix DMC-ZR1
Object:	Vertical wall, painted off-white, at a distance of 1300 mm from the camera.
Measured angle between camera sensor plane and plane of object:	$30^\circ \pm 1^\circ$
Image size after resizing for simpler processing:	144 pixels x 108 pixels
Focal length setting, from calibration (assuming image size as above):	130 pixels
R_0 :	131.50
Minimum reflectance threshold:	45
Surface height allocated when reflectance level is less than the threshold:	-25



Fig. 1: Original of image of the plane object, a painted wall.

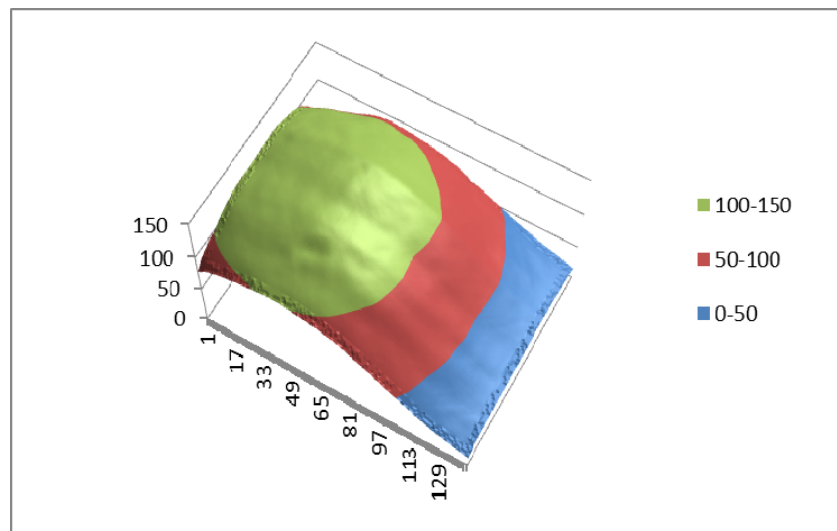


Fig. 2: Image of plane object, depicted with reflectance levels shown in three-dimensions. Reflectances range from 0 to almost 132. (Image from Microsoft Excel spread-sheet.)

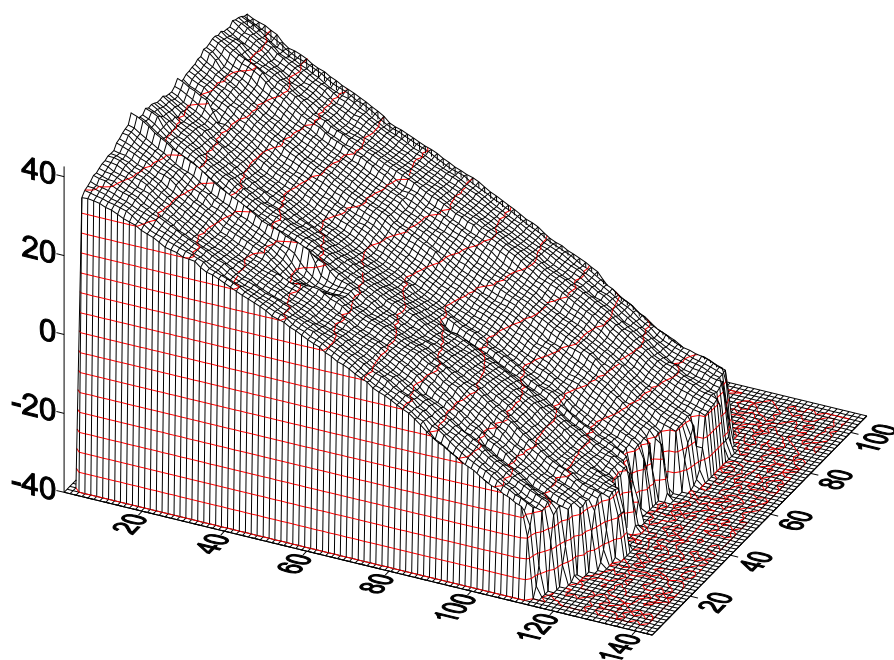


Fig. 3: Reconstructed object shape, deduced by Microsoft Excel spread-sheet using data depicted in Table 1 and Figure 2. The two elements in the flash lighting show as furrows along the object. Image from Surfer by Golden Software version 7.

It is noticeable that in Figure 1 is not easy to see grey level gradations, but figure 2 shows the smooth and substantial gradations. Figure 3 shows that a plane object has been obtained. In detail, figure 3 also reveals that i) part of the object has been excluded by the reflectance threshold of 45 grey-levels; ii) the two elements in the flash lighting show as furrows along the object iii) there is a sensor blemish whose effect is apparent in figure 2 and in the surface results.

6.2. Case 2: Concept verification with a cylinder

Details of the experiment are shown in Table 2; the image of the cylinder coated with white paper is shown in Figure 4. The cylinder has a radius of 265 mm, and the circumference covered by the imaging is about 425 mm; so we can deduce that the angle at the centre is $425/(2.\pi.265) \times 360^\circ$, i.e. close to 90° : see Figure 6. It is again noticeable that the two elements in the flash lighting show as furrows along the object in figure 6.

Table 1: Details of the experiment involving a cylindrical object.

Camera	Panasonic Lumix DMC-ZR1
Object:	Cylinder: 90° sector, covered with white paper.
Image size after resizing for simpler processing:	144 pixels x 108 pixels
Focal length setting, from calibration (assuming image size as above):	130 pixels
R_0 :	191.50
Minimum reflectance threshold:	0
Surface height allocated when reflectance level is less than the threshold:	0



Fig. 4: Original image of cylinder, of radius 265 mm.




Fig. 5: Three-dimensional representation of the original image of the cylindrical object showing the variation of reflectance levels from 0 to almost 150. Image from Surfer by Golden Software version 7.




Fig. 6: Deduced cylindrical object, from computation by Microsoft Excel spread-sheet, with first 20 columns edited out. Image from Surfer by Golden Software version 7. A bias in the surface is attributed to the flash being on the right-hand side of the lens. the two elements in the flash lighting show as furrows along the object.

6.3. Case 3: Concept verification on a human back

The first figure shows a human back (Figure 7), on which grey levels are not clear to the naked eye, but Figure 8 shows the same figure in five ranges of grey levels (depicted by a Microsoft Excel spread-sheet diagram), while Figure 9 shows the resulting surface shape, as depicted by the Microsoft Excel spread-sheet which is used for processing, in a few minutes of manual work.

Table 3: Details of the experiment involving a human subject.

Camera:	Canon EOS30D single-lens-reflex digital camera with a 28 mm Canon lens
Object:	Human subject.
Image size after resizing for simpler processing:	144 pixels x 96 pixels
Focal length setting, from calibration (assuming image size as above):	135 pixels
R_0 :	226
Minimum reflectance threshold:	100
Surface height allocated when reflectance level is less than the threshold:	-30



Figure 7: Flash image of a human back.




Fig. 8: Flash image of a human back in Figure 1 but in terms of five levels of reflectances, (as depicted by Microsoft Excel spread-sheet).




Fig. 9: The shape deduced from the image in Figure 1 by SFS, (by Microsoft Excel spread-sheet). The presented shape is not simply a reproduction of the image's grey-levels but the three-dimensional shape, reconstructed by an integration of the surface normal directions derived from the single image two-dimensional grey-level image.

5. Discussion

These experiments and others that have been undertaken confirm that the underlying theory of SFS using a Lambertian reflection model is appropriate for the objects which have been studied. It is worth appreciating that both images represented in Figure 1 and Figure 4 look similar, but one represents a plane and the other a cylinder, and indeed the results shown in Figures 3 and Figures 6 and 7 distinguish the plane from the cylinder, even though they exhibit low accuracy and some systematic flaws.

Although in practice the method essentially works, various practical difficulties remain if accuracy and reliability are to be assured. Firstly, a number of problem areas were foreseen, but in reality, these do not seem to have been problematic:

- The solution can be expected to be weak because variations of surface normal direction when the surface is facing the camera axis produce only a small change in the cosine of angle i .
- R_0 is assumed to be constant for each image only, i.e. a value must be determined for each new image.
- Lambertian rather than specular reflection may not always be valid for the subject surfaces. Moreover, the optical texture of the object must be even. Otherwise, small blemishes in an object surface can destroy the logic of the solution, while, similarly, "bright spots", due to specular reflection, can destroy the solution completely at those points and disrupt the surface shape solution significantly. Moreover, bright spots due to specular reflection are most likely to occur at the point where R_0 is to be determined. Fortunately, human skin does seem to be a good surface for reflection.
- Concave surfaces can reflect back into themselves. Again, it is fortunate that human body shapes tend to be convex.

Secondly, equipment flaws were seen to contribute to inaccuracy. Primarily, cameras are not perfect.

- Significant levels of image noise (e.g., about 2 grey-levels over a typical range of 150) must be assumed; (noise in the vicinity of the maximum reflectance can create cases when $R > R_0$, causing cosines to be greater than 1.0).
- The solution involves the lens's focussed distance, (i.e. the lens's fixed focal length corrected for any lens movement involved in focussing on the object), so camera calibration is required. This has been undertaken using the theory above with a simple object of known shape, and the deduced distance has not been the value given by the camera's specifications.

- Consistent blemishes can occur in the sensor, revealing themselves as systematic surface indentations or lumps; see Figure 3.
- Recorded grey-levels, R , are known to not necessarily to be linearly related to the true reflectance levels, but this has not been attended to in this work so far.
- Camera lens distortion is not necessarily insignificant, but this has so far not been attended to in this work so far.
- Flash lighting is not necessarily even. This has been very apparent in this work; see striping in the surface deduced in Figure 3 and 6 because of the flash's lighting elements.
- Errors can be caused by assuming incorrectly that the flash is a point light source coincident with the lens centre; see Figure 6.

It is assumed that most of these problems can be reduced by using better quality equipment including, for example, a ring flash surrounding the lens and/or by calibration. However, in this work, the cheaper cameras were seen as essential.

It has been found in practice that the flash camera in a semi-dark environment is satisfactory, and circumvents the inconvenience of working in the dark or having a special black background. Backgrounds are eliminated satisfactorily by setting a grey-level threshold; see Table 1, 2 and 3.

One of the frustrating aspects of SFS is that surface results from SFS can look like simple reproduction of the 2D grey levels as a 3D object, whereas in fact the algorithm does deduce three-dimensional surface normal information from the single two-dimensional grey-level image. The surface shape is reconstructed by an integration of the surface normal directions.

Most development time has been spent on expediting the algorithm and the quick result; accuracy needs to be refined and tested, having been only partly confirmed so far, using a limited range of geometrically-shaped objects.

6. Conclusions

SFS has been seen as ideal for some body surface measurement tasks, and in this case, it is proposed specifically for observing scoliosis. The aspect of the work described in this paper has been the search for simple calculation of SFS results, on the assumption that the applications of interest demand overall simplicity. Results are not perfect, but not discouraging. The SFS theory has been verified, but it is seen that accuracy depends on the selection of objects, and moreover on the selection and/or calibration of the camera being used. The imaging is quick, but more importantly here, solutions have been obtained straightforwardly via a spread-sheet.

References

1. Wikipedia, 2013. "Lambert's cosine law" (accessed 2013) <http://en.wikipedia.org/>