

Shape in a Box

Graham D. Finlayson and Christopher Powell

School of Computing Sciences, University of East Anglia, UK

Abstract. Many techniques have been developed in computer vision to recover three-dimensional shape from two-dimensional images. These techniques impose various combinations of assumptions/restrictions of conditions to produce a representation of shape (e.g. a depth/height map). Although great progress has been made it is a problem which remains far from solved, with most methods requiring a non-passive imaging environment. In this paper we develop on a variant of photometric stereo called “Shape from color” (SFC). We remove the restriction of known, direct light sources by exploiting mutual illumination; we simply take pictures of objects within a colourful box, hence “Shape in a Box”. We discuss the engineering process used to develop our set-up and demonstrate experimentally that our passive imaging environment recovers shape to the same accuracy as SFC. A second contribution of this paper is to benchmark our approach using real objects with known ground truth, including some 3D printed objects.

Keywords: photometric stereo, mutual illumination, shape recovery

1 Introduction

Recovery of three-dimensional shape from two-dimensional images has been an active area of research since the inception of computer vision. Common shape recovery techniques include shape-from-X techniques, such as shape-from-stereo (structure-from-stereo) [8, 6, 31] and shape-from-shading [33, 18, 12]. Other methods include recovery through the use of intrinsic image properties [2, 1] and the focus of this paper, photometric stereo [32, 3].

Classic photometric stereo [32] recovers per-pixel surface normals of a static, convex, Lambertian object by capturing three images (from the same camera position) of the object illuminated by three independent lights (Fig. 1a). The intensity of pixels in the three images forms a linear relationship with the surface normals of the object (Fig. 1b) and thus the normals can be recovered.

We are interested in shape recovery techniques which work in a passive environment and can be performed using a single image. Our research begins with an existing single image variant of photometric stereo, “Shape from color” [9–11] (also known as “spectrally multiplexed photometric stereo” [14] or “photometric stereo with colour lights” [16] - we shall hereafter refer to it as “SFC”). This method can recover the surface normals of a Lambertian object from a single colour image assuming the object is illuminated by three spectrally distinct lights

simultaneously (Fig. 1d). If the colour and directions of each light source are sufficiently independent, there exists a linear relationship between the colour values of a pixel on the surface and its normal (Fig. 1e). Essentially, the three colour channels of a single image provide the same information as the three images used in classic photometric stereo. For both classic photometric stereo and SFC, recovered normals are converted to x and y derivatives and this gradient field is reintegrated to recover a height map (Fig. 1c,f).

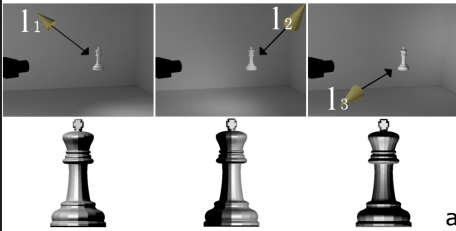


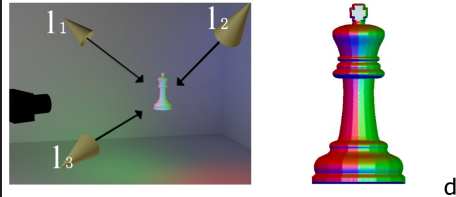

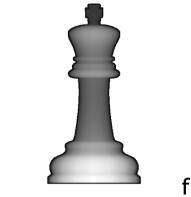
	Input Image/s			Recovered Normals	Recovered Height
Photometric Stereo					
Shape from Colour					

Fig. 1. In parts **a** and **d**, light sources are labelled as l_1 , l_2 , l_3 in the scene diagrams. In **a**, below each diagram is the image of the object. In **d**, the image is to the right of the scene diagram. Surface normals in **b** and **e** are displayed using the convention of normal maps in graphics (red = x -axis, green = y -axis, blue = z -axis). Height maps, **c** and **f** are in grayscale (white = maximum height, black = minimum height). The same colour coding is used throughout this paper.

The key contribution of this paper is to demonstrate that SFC can work without the requirement of three or more colourful lights. We make the simple observation that the light impinging on any point in a scene is often complex in nature: it is a combination of direct light sources and mutual illumination (light reflected from other surfaces). Thus we propose that the SFC method can be extended to work in any environment which provides appropriate, spectrally varying illumination. Our “Shape in a box” (SiaB) method captures the lighting environment of a specially engineered box by imaging a chrome sphere and uses this information to recover shape. The motivation for our research is to move shape recovery out of the laboratory and into a passive environment. SiaB is presented here as a first step towards achieving this goal.

Another novel contribution of this paper is that we measure shape recovery against “absolute” known ground truth 3D objects. First we use papercraft ob-

jects (available to print from many websites¹). Second, we print objects using a 3D printer. Experiments demonstrate that SiaB achieves the same level of accuracy as SFC.

Section two contains a review of literature relevant to our work. Section three details the engineering underlying our SiaB method. Section four contains experimental results which compare the accuracy produced by SiaB with SFC by benchmarking against real 3D ground truth. Section five contains a brief conclusion and proposals for further development of our method.

2 Related work

2.1 Classic photometric stereo

In classic photometric stereo [32] Woodham proposes that the surface normals of a convex, Lambertian object can be recovered from three images. It was shown that if the object is illuminated by three, distant, point light sources; then there is a linear relationship between the three sets of pixel values and the object surface normals.

To understand how shape can be recovered, let us denote the direction of each light as a vector \mathbf{e} , we have three (by assumption linearly independent) vectors \mathbf{e}_1 , \mathbf{e}_2 and \mathbf{e}_3 . With respect to Lambert’s law, a point on a surface with normal $\mathbf{n} = [n_x \ n_y \ n_z]^t$ illuminated by light source \mathbf{e} results in a pixel value p ,

$$p = \alpha(\mathbf{e} \cdot \mathbf{n}), \quad (1)$$

where α accounts for surface albedo (assumed to be constant, so we shall hereafter absorb it into the \mathbf{e} term). Let us use the notation \mathbf{p}_i to denote the i th triple of pixel responses (one for each light) and \mathbf{n}_i denote the corresponding i th scene surface normal. We group the image responses, the lighting directions and the scene surface normals into matrices P , E and N respectively,

$$P = [\mathbf{p}_1 \ \mathbf{p}_2 \ \dots \ \mathbf{p}_n], \ E^t = [\mathbf{e}_1 \ \mathbf{e}_2 \ \mathbf{e}_3], \ N = [\mathbf{n}_1 \ \mathbf{n}_2 \ \dots \ \mathbf{n}_n]. \quad (2)$$

Under the assumption that the surface in question has uniform, Lambertian reflectance, there exists a linear relationship between the light reflected at each point on the surface (captured pixel values) and the orientation of the surface at each point,

$$P = E N. \quad (3)$$

Since we know E and P , Woodham observed, we can recover N :

$$N = E^{-1}P. \quad (4)$$

Even when all the underlying assumptions hold, fine-tuning of the experimental design [29, 4] is essential for the best recovery. Of course for real data

¹ e.g. "Paper Models of Polyhedra" - <http://www.korthalsaltes.com/>

the underlying assumptions may not hold, for example the presence of shadows and specular highlights are problematic. Methods exist in the literature which in essence, extend the classic Woodham approach so that shape can be recovered for more general conditions [3, 17]. However, we do not comment on them further here, except to remark that the same extensions are applicable to the method developed in this paper. The focus of our work is only to present a simple algorithm for shape recovery in a passive environment.

2.2 Shape from color (SFC)

Again let $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ denote light direction vectors. Additionally let us denote the colour of reflected light as $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3$. Values of \mathbf{b}_i are given by

$$\mathbf{b}_i = \int I_i(\lambda) S(\lambda) \mathbf{q}(\lambda) d\lambda, \quad (5)$$

where $I_i(\lambda)$ is the spectral power distribution of the i th light, $S(\lambda)$ is the spectral reflectance function of the (Lambertian) surface and $\mathbf{q}(\lambda)$ represents the camera sensitivities. In effect \mathbf{b}_i is the RGB of a flat, frontally presented calibration surface with the same albedo as the object to be measured under the i th light.

Taking the values of $\mathbf{e}_i, \mathbf{b}_i$ and the surface normal at a pixel \mathbf{n} , it follows that the RGB camera response \mathbf{c} at that pixel is given by

$$\mathbf{c} = (\mathbf{e}_1 \cdot \mathbf{n}) \mathbf{b}_1 + (\mathbf{e}_2 \cdot \mathbf{n}) \mathbf{b}_2 + (\mathbf{e}_3 \cdot \mathbf{n}) \mathbf{b}_3. \quad (6)$$

Grouping vectors \mathbf{e} and \mathbf{b} into matrices E and B

$$E = [\mathbf{e}_1 \ \mathbf{e}_2 \ \mathbf{e}_3]^t, \quad B = [\mathbf{b}_1 \ \mathbf{b}_2 \ \mathbf{b}_3]. \quad (7)$$

Then equation (6) can be rewritten as

$$\mathbf{c} = F \mathbf{n} \equiv B E \mathbf{n}. \quad (8)$$

As F and \mathbf{c} are known, the surface normals can be recovered:

$$\mathbf{n} = F^{-1} \mathbf{c}. \quad (9)$$

Using SFC, only one image is required to recover shape. Another advantage of this method is that B and E do not have to be determined separately. Rather, a perfect Lambertian reflector of known shape can be placed in the scene and the linear relationship between recorded camera RGBs and surface normals can be recovered directly. Thus F can be found via a calibration step.

SFC can estimate shape for smooth objects and can be performed in real-time without the need for processing of temporal information. Accordingly it has found applications in face capture [30] and shape recovery in video sequences of non-rigid surfaces [16, 5]. Though we note that as with photometric stereo, finely-tuned experimental design is again necessary for optimal recovery [22].

3 Shape in a Box (SiaB)

A key idea in this paper is to substitute the multiple lights used in SFC (and the related lab-based restrictions on shape measurement) using the scene illumination environment directly in a modified SFC algorithm. Crucially, we need to ensure the environment is sufficiently varying to support shape recovery. We achieve this by using the mutual illumination from the coloured walls of a simple triangular box. In section 3.3 we carry out a graphics-based simulation to determine the required geometry of the box). Our long-term ambition is to recover shape in a room with suitable mutual illumination (e.g. a room with colourful walls - see figure 3b).

3.1 Calibration

In theory calibration is simple. Like Johnson and Adelson [20] we could simply place a perfect, spherically shaped Lambertian reflector in a scene (a “spherical probe”) and take an image. We could then solve for the linear transform relating the image RGBs to the surface normals of the sphere (section 2.2). However, this approach is not used directly here. No surface is perfectly Lambertian and unlike SFC we cannot choose lighting directions to - for example - minimise the appearance of specular highlights in the calibration image.

Instead we propose taking a picture of a chrome sphere to measure the light from all angles. Then we use spherical harmonic basis functions [26, 27, 15] to simulate the graphical model of a perfectly spherical, perfectly Lambertian reflector with the illumination environment from the chrome sphere (Fig. 2). Given the known spherical shape and synthesised image the linear transform relating image colour to surface normals is easily found [9].

It is true that the illumination environment will vary slightly according to the object placed inside the box (due to the interaction of light between the box and the object). At present our method does not account for any ‘bouncing’ of light rays, though the quality of our results show that in practice, the effects of this are small.

3.2 HDR image capture

In figure 2 we show the top of our box with the chrome sphere inside. In order to capture the full range of environment lighting we capture the image at five exposure levels. The reader can see that the light from the ambient environment (white in the middle of the sphere) is very bright. The exposures are blended into a HDR image [7, 21] prior to calculating the perfect Lambertian sphere.

3.3 The Box

Light entering through the top of the box will inevitably be of higher intensity than the light reflected from the colourful interior of the box. Objects placed




Fig. 2. Here we illustrate the scene calibration process. Left: The chrome sphere is placed inside the colourful box. Centre: Examples of bracketed exposures of the chrome sphere which are used to construct a HDR image. Right: A Lambertian sphere is synthesised under the same lighting conditions.

at the bottom need to receive sufficient mutual illumination from the sides and not have that illumination be lost in the strength of the light from the exterior environment. Equally we do not wish the box to be too deep as in this case the object placed at its bottom would be too dimly lit. So, we need a box that creates a lighting environment which meets these colour sufficiency and light intensity requirements.

To find a suitable box geometry we rendered a series of synthetic images of a chrome sphere (with 15cm diameter) placed in a box where the length and height of the triangular box walls varied. To measure the “quality” of each lighting environment we took an image of the chrome sphere and generated its Lambertian counterpart as described in section 3.1. From equation (8) we can calculate the matrix F which allows us to transform RGB values to surface normals and vice versa (i.e. $\mathbf{c} = F \mathbf{n}$). It is important that in recovering $\mathbf{n} = F^{-1} \mathbf{c}$, the inverse F^{-1} is well conditioned. That is to say, if we perturb our measured RGB values, \mathbf{c} , by a small amount ϵ , we would like

$$\hat{\mathbf{n}} = F^{-1}(\mathbf{c} + \epsilon) \simeq \mathbf{n}. \quad (10)$$

The condition number of a matrix F , $k(F) \in [1, \infty]$, is a measure of how good this approximation is (i.e. if $k(F) = 1$, the inverse is maximally stable). In the worst case when $k(F) \geq 10$, then $\hat{\mathbf{n}}$ can be about 10% different from \mathbf{n} and this is the criteria we chose to build our box. We sought an illumination environment that supports a condition number no more than 10. The $k(F)$ generated from our synthesised boxes are shown in figure 3. As a compromise between practical considerations and condition number we settled on a box with side length 80cm and height 60cm.

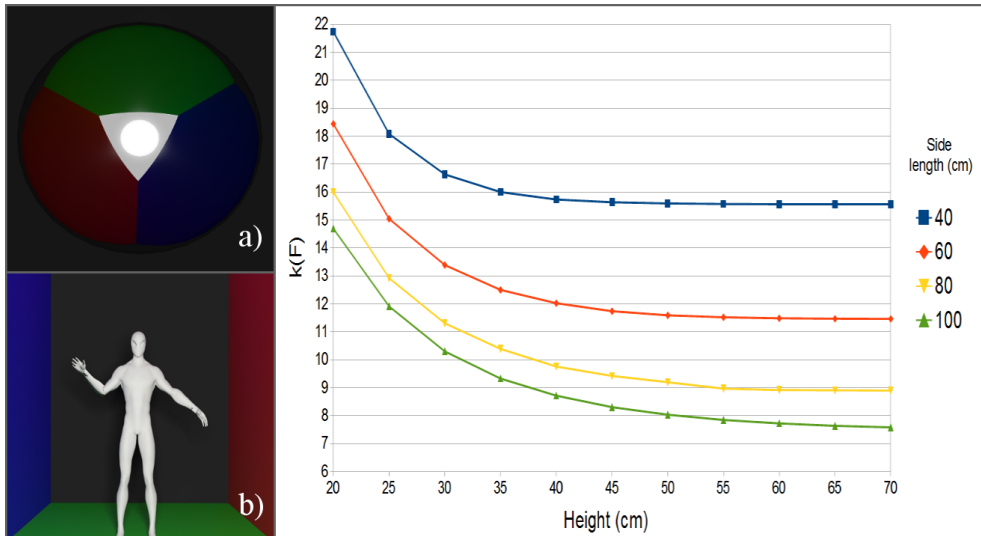


Fig. 3. a) An example of a synthetic chrome sphere from our box engineering experiments. The graph shows the results of the synthetic experiments on triangular boxes of various dimensions. Four different side lengths were each tested with the same varying heights. b) A synthetic human inside a coloured room.

3.4 Recovering height

Once surface normals have been recovered they are converted into height maps. From equation (9) we have obtained a vector field \mathbf{n} consisting of recovered surface normals. Each point in the vector field has three components, denoting the direction of the surface normal in the x, y and z axis, that is to say $\mathbf{n} = [n_x \ n_y \ n_z]$. These surface normals can be converted into a gradient field which corresponds to some surface Z by taking the ratios of the x and y components of the surface normals with their z component [19]

$$\begin{aligned}
 p &= \frac{\partial Z}{\partial x} = -\frac{n_x}{n_z}, \\
 q &= \frac{\partial Z}{\partial y} = -\frac{n_y}{n_z}.
 \end{aligned}
 \tag{11}$$

For both SFC and SiaB, it is almost certain that the underlying assumptions (e.g. perfectly Lambertian reflectance) do not hold. Accordingly the gradients calculated in (11) are usually not integrable. There may not exist a height map $Z(x, y)$ which corresponds exactly to the gradients. Thus we seek the integrable surface function $\hat{Z}(x, y)$ which is the closest approximation to $Z(x, y)$, in the sense that the derivatives of $\hat{Z}(x, y)$ are as close to those in equation (11) as possible. We can find \hat{Z} by solving Poisson's equation [28]

$$\nabla^2 \hat{z} = \frac{\partial p}{\partial x} + \frac{\partial q}{\partial y} \quad (12)$$

There have been many methods developed for the reconstruction of height from gradient fields. In this paper we present results achieved by Frankot & Chelappa’s Fourier-based method [13] and Kovessi’s shapelet correlation approach [23]. We also present the results of a direct “Jacobi-type” reintegrator. Specifically, since we assume that an object can be segmented from the background, we know the occluding contour of the shape. Thus, in effect our reintegration problem has Dirichlet boundary conditions of complex shape. Our “Jacobi-type” method is similar to that presented in [25]).

We do not go into further detail on the methods here but remark that gradient field reconstruction is an active area of research and direct the reader to [24] for a review of the topic.

4 Experiments

In our experiments we wished to measure the accuracy of recovered shape against ground truth. Often this is achieved through comparison against an existing shape recovery method deemed to be accurate or through the use of synthetic data. Novelly, here we benchmark on recovery of objects whose actual 3D shape is a priori known to good accuracy.

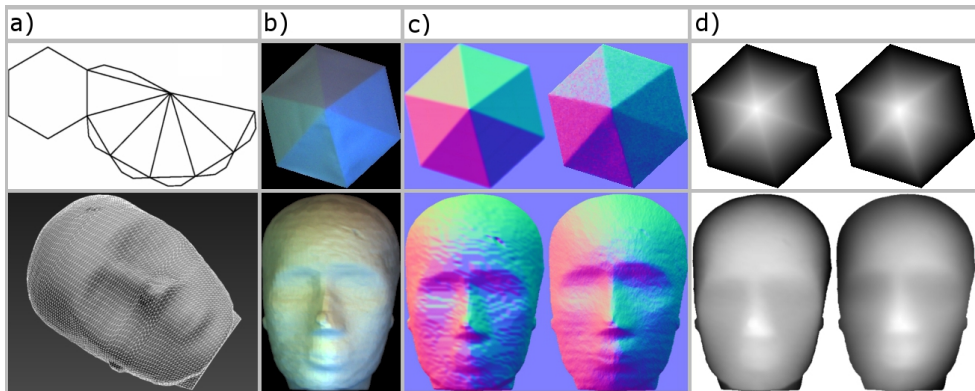


Fig. 4. Top row: a simple papercraft object. Bottom: a 3D-printed object. a) Source files, a papercraft template and a 3D model file. b) Captured images. c) left = true normal map, right = recovered normal map. d) left = true height map, right = recovered height map (using Jacobi iteration method).

4.1 Ground truth

Our first object dataset is built from “papercraft” templates. These consist of templates which can be printed and fabricated. While the true 3D shape of

the papercraft object is known, there may be some small fabrication errors. We benchmark against the perfect 3D model.

Our second ground truth set comprises of 3D printouts of custom meshes. Our ZCorp 450 3D printer prints objects by binding together successive layers of a proprietary powder (resulting in approximately Lambertian reflectance). As with the papercraft objects, there can be small discrepancies between the printed object and the source file. We benchmark against the actual 3D model source files. Examples of both types of object can be seen in figures 4 and 5.

4.2 Results

A chrome sphere is placed in the box and the linear relationship between image RGBs and sphere normals is found (section 3.1). Then an object is placed in the same location and its captured RGB values mapped to surface normals (object images can be seen in figure 5). Normals are converted to a gradient field which is reintegrated to give a height map (section 3.4).




Fig. 5. Images of our experiment subjects (3D-printed objects in *italics*, papercraft in normal type). Top row, left to right: *Bust*, *Cone*, *Face*, *Pyr4*, *Pyr5*. Bottom row, left to right: *Pyr6*, *Sph*, *Star*, *Tetra*, *Trach*.

Object	SiaB Accuracy (%)			SFC Accuracy (%)		
	Shapelet	FC	Jacobi	Shapelet	FC	Jacobi
Bust	89.14	88.13	92.94	87.59	89.04	82.25
Cone	84.01	86.00	95.90	92.71	92.01	97.36
Face	92.05	89.79	85.75	88.09	88.17	80.95
Pyr4	81.47	90.85	97.18	94.45	93.19	97.44
Pyr5	85.19	88.72	97.39	93.92	92.80	97.50
Pyr6	88.41	85.25	97.09	94.72	91.06	95.78
Sph	87.51	91.06	89.42	88.41	91.14	84.03
Star	85.76	77.56	94.62	94.74	89.94	94.06
Tetra	85.26	84.23	95.38	96.08	95.03	96.54
Trach	79.63	76.71	81.45	76.84	77.29	60.32
Average	85.84	85.83	92.71	90.76	89.97	88.61

Table 1. Height map accuracies as percentage values. “Shapelet” is Kovesei’s method [23], “FC” is Frankot & Chellappa’s algorithm [13] and “Jacobi” is our Jacobi-type method [28].

In table 1 ground truth and recovered height maps were both scaled to unit height and root mean squared error was calculated. The recovery percentage accuracy is calculated as shown in equation 13, where i and j refer respectively to the rows and columns of the true height map Z and the recovered height map \hat{Z} ; m and n are the row and column lengths.

$$accuracy(\hat{Z}) = 100 - 100 \sqrt{\frac{\sum_{i=1}^n \sum_{j=1}^m (Z_{(i,j)} - \hat{Z}_{(i,j)})^2}{nm}} \quad (13)$$

The average over all three reintegration methods yields a figure of 88.13% accurate height recovery from SiaB and 89.78% using the SFC experimental set-up. In the majority of cases, the iterative Jacobi reintegration method achieves the most accurate result of the three. An example of a recovered height map can be seen in figure 6.

5 Conclusion

Inferring 3D shape from images remains a much studied problem in computer vision. In this paper we have extended a photometric stereo technique, “Shape from color” [9] to achieve accurate shape recovery of Lambertian objects without the restriction of spectrally-varied, direct light sources. We have instead designed a passive imaging environment (the titular box) which generates a sufficiently spectrally-varied lighting environment through the effects of mutual illumination. We calibrate our set-up by imaging a chrome sphere inside the box and then recover the surface normals of unknown objects in the same environment. With our “Shape in a box” method we recover shape to the same accuracy as shape from color. It is expected that by enlarging the box (e.g. to the size of a small

room) it will be possible to recover the shape of an entire person; indeed initial simulations show this to be a promising course for future research (Fig. 3b).




Fig. 6. Renders of recovered height maps for the *Bust* object, using the Jacobi iteration reintegration method. Left is the ground truth, centre is Shape in a Box and right is using Shape from Color.

References

1. Barron, J.T., Malik, J.: Shape, albedo, and illumination from a single image of an unknown object. 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 334–341 (2012)
2. Barron, J.T., Malik, J.: Shape, illumination, and reflectance from shading. Tech. rep., UC Berkeley (2013)
3. Barsky, S., Petrou, M.: The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25(10), 1239–1252 (2003)
4. Barsky, S., Petrou, M.: Design issues for a colour photometric stereo system. *Journal of Mathematical Imaging and Vision* 24(1), 143–162 (2006)
5. Brostow, G.J., Hernández, C., Vogiatzis, G., Stenger, B., Cipolla, R.: Video normals from colored lights. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 33(10), 2104–2114 (2011)
6. Brown, M.Z., Burschka, D., Hager, G.D.: Advances in computational stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 25(8), 993–1008 (2003)
7. Debevec, P.E., Malik, J.: Recovering high dynamic range radiance maps from photographs. *Proceedings of SIGGRAPH 2007* pp. 369–378 (2007)
8. Dhond, U.R., Aggarwal, J.K.: Structure from stereo—a review. *IEEE Transactions on Systems Man and Cybernetics* 19(6), 1489–1510 (1989)
9. Drew, M.S.: Shape from color. Tech. rep., Simon Fraser University (1992)
10. Drew, M.S.: Photometric stereo without multiple images. *Electronic Imaging '97* pp. 369–380 (1997)
11. Drew, M.S., Brill, M.H.: Color from shape from color: A simple formalism with known light sources. *Journal of the Optical Society of America (JOSA)* 17(8), 1371–1381 (2000)
12. Durou, J.D., Falcone, M., Sagona, M.: Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding* 109(1), 22–43 (2008)

13. Frankot, R.T., Chellappa, R.: A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10(4), 439–451 (1988)
14. Fyffe, G., Yu, X., Debevec, P.: Single-shot photometric stereo by spectral multiplexing. 2011 IEEE International Conference on Computational Photography (ICCP) pp. 1–6 (2011)
15. Green, R.: Spherical harmonic lighting: The gritty details. *Archives of the Game Developers Conference* (2003)
16. Hernández, C., Vogiatzis, G., Brostow, G.J., Stenger, B., Cipolla, R.: Non-rigid photometric stereo with colored lights. *IEEE 11th International Conference on Computer Vision (ICCV)* pp. 1–8 (2007)
17. Hernández, C., Vogiatzis, G., Cipolla, R.: Shadows in three-source photometric stereo. *European Conference on Computer Vision (ECCV)* pp. 290–303 (2008)
18. Horn, B.K.: Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. *Tech. rep.*, Massachusetts Institute of Technology (1970)
19. Horn, B.K.: Understanding image intensities. *Artificial Intelligence* 8(2), 201–231 (1977)
20. Johnson, M.K., Adelson, E.H.: Shape estimation in natural illumination. 2011 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) pp. 2553–2560 (2011)
21. Kirk, K., Andersen, H.J.: Noise characterization of weighting schemes for combination of multiple exposures. *British Machine Vision Conference (BMVC)* pp. 1129–1138 (2006)
22. Klaudiny, M., Hilton, A.: Error analysis of photometric stereo with colour lights. *Pattern Recognition Letters* (2014), <http://www.sciencedirect.com/science/article/pii/S0167865513005114>
23. Kovessi, P.: Shapelets correlated with surface normals produce surfaces. *Tenth IEEE International Conference on Computer Vision (ICCV)* 2, 994–1001 (2005)
24. Patel, V.M., Chellappa, R.: Approximation methods for the recovery of shapes and images from gradients. *Applied and Numerical Harmonic Analysis* pp. 377–398 (2013)
25. Pérez, P., Gangnet, M., Blake, A.: Poisson image editing. *ACM Transactions on Graphics (TOG)* 22(3), 313–318 (2003)
26. Ramamoorthi, R., Hanrahan, P.: An efficient representation for irradiance environment maps. *Proceedings of SIGGRAPH 2001* pp. 497–500 (2001)
27. Schönfeld, V.: Spherical harmonics. *Tech. rep.*, RWTH Aachen University (2005)
28. Simchony, T., Chellappa, R., Shao, M.: Direct analytical methods for solving poisson equations in computer vision problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 12(5), 435–446 (1990)
29. Sun, J., Smith, M., Smith, L., Farooq, A.: Examining the uncertainty of the recovered surface normal in three light photometric stereo. *Image and Vision Computing* 25(7), 1073–1079 (2007)
30. Vogiatzis, G., Hernández, C.: Self-calibrated, multi-spectral photometric stereo for 3d face capture. *International Journal of Computer Vision* 97(1), 91–103 (2012)
31. Weng, J., Huang, T.S., Ahuja, N.: Motion and structure from image sequences. Springer Publishing Company (2012)
32. Woodham, R.J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* 19(1), 139–144 (1980)
33. Zhang, R., Tsai, P.S., Cryer, J.E., Shah, M.: Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 21(8), 690–706 (1999)