

Mutual illumination photometric stereo without calibration

Christopher Powell & Graham D. Finlayson; University of East Anglia; Norwich, UK.

Abstract

Recovering three-dimensional shape from two-dimensional images is a long-standing problem in computer vision. First proposed in the 1980s, photometric stereo has matured to the point that accurate recovery of complex shapes and surfaces has become achievable. However, such methods typically demand multiple image captures, highly controlled scene conditions and elaborate experiment designs which require calibration. Building on previous work, we propose using a variant of photometric stereo which needs only a single image of an object in a colourful environment and we now remove the requirement of a calibration step. Instead we build an entirely synthetic graphics model of our capture environment (a colourful box) and carry out a synthetic calibration. The validity of this approach is demonstrated by benchmarking real world experiments against ground truth data and comparison with previous work.

Introduction

Since the inception of computer vision, recovering three-dimensional shape from two-dimensional images has been an active area of research. Beyond the direct benefit of reproducing virtual representations of real-world objects (e.g. 3D printing [15, 25]), shape recovery research is also beneficial for other areas of computer vision such as object recognition [23] or motion analysis [27].

It is an intrinsically difficult problem, requiring compensation for many factors such as lighting, motion, capture geometry, shadows, highlights or interreflections caused by complex shape. As such many different methods have been developed such as shape-from-X techniques (e.g. shape-from-stereo [7, 9, 36] and shape-from-shading [13, 19, 39]), exploitation of intrinsic image properties [2, 3], structured light [30, 40] and our selected approach - photometric stereo [37].

Typically photometric stereo methods require multiple images, however we are interested in recovering shape from a single image capture. As such we focus on a variant called 'Shape from Color' [10–12] (also known as 'spectrally multiplexed photometric stereo' [17] or 'photometric stereo with colour lights' [18], hereafter we refer to it as 'SFC') which requires only a single colour image of an object illuminated by three spectrally distinct light sources simultaneously.

In previous work we removed the SFC restriction of specific light sources and demonstrated that the method functions with equivalent accuracy by instead exploiting the mutual illumination of a colourful environment - which we entitled 'Shape in a Box' (SiaB) [14]. As illustrated in Fig. 1 we measured the lighting environment by imaging a mirrored sphere and used this information to recover the shape of unknown objects in the same environment.

Ultimately we are interested in moving photometric stereo out of the laboratory and into the real world - hence the portable box. However in uncontrolled conditions the lighting is liable to




Figure 1. Our experimental set-up. Light reflects off the box walls and onto the object inside, providing sufficient spectral variance across the surface for shape recovery.

constantly change. Thus methods which require a separate calibration step in such environments are subject to additional error introduced by changes in the scene over time. Another limitation of our previous calibration-based method is that the lighting environment was measured with a chrome sphere (Fig. 1). In such a closed environment the sphere itself contributes to the lighting. If we remove the sphere and proceed to image a different object inside the box, then some error has been immediately introduced into the system. To address these issues, here we remove the restriction of requiring an additional calibration step.

Typically in shape recovery it is assumed for the sake of simplicity that images are captured with orthographic projection, however we exploit the fact that in reality captured images are usually subject to weakly perspective projection. We build a lighting model from the environment visible in the image (i.e. the box walls which can be seen in figure 1) and use this information to estimate the 3D surface of the Lambertian object contained within the same image. We capture images of objects, validating their accuracy using ground truth data and comparison against the results of our previous method which required a separate calibration step.

In section two we give a brief summary of relevant work in photometric stereo, focusing particularly on SFC-based approaches, techniques which require no calibration and attempts at photometric stereo in natural illumination. Section three contains details of our experiment set-up and implementation; results

and discussion appears in section four before the paper concludes in section five.

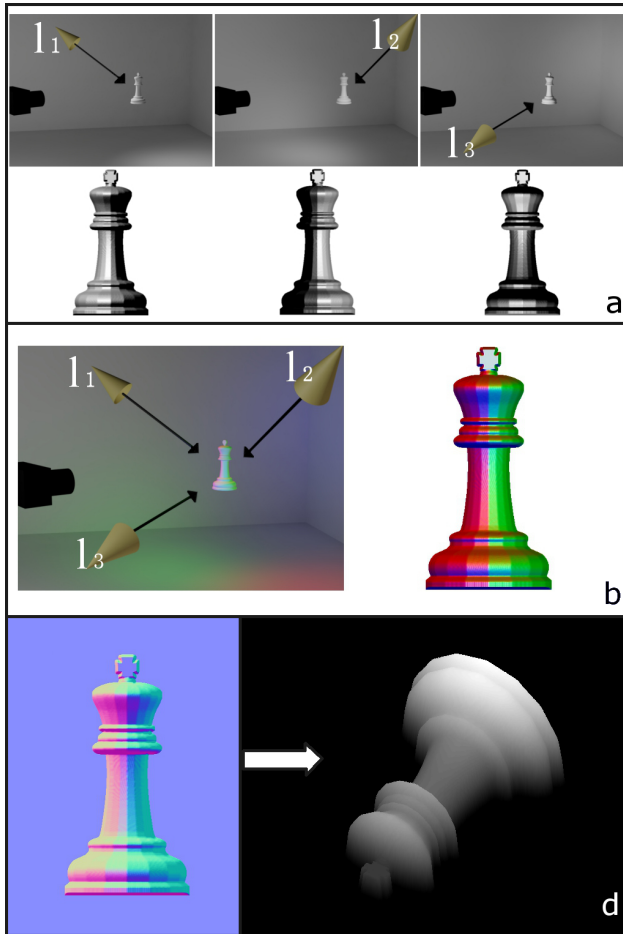


Figure 2. In parts **a** (classic photometric stereo) and **b** (SFC), light sources are labelled as l_1 , l_2 , l_3 in the scene diagrams. In **a**, below each diagram is a captured image of the object. In **b**, the image is to the right of the scene diagram. Recovered surface normals in **c** are displayed using the convention of normal maps in graphics with red, green and blue corresponding approximately to positive values in the x , y and z axis respectively. The integrated height map in **d** is textured with white corresponding to maximum height (1) and black being minimum height (0).

Related work

First proposed by Woodham, classic photometric stereo [37] recovers per-pixel surface normals of a static, convex, Lambertian object by capturing three images (from the same camera position) of the object as illuminated by three independent lights (Fig. 2a). The intensity of pixels in the three images forms a linear relationship with the surface normals of the objects (Fig. 2c) and thus they can be recovered (more detail is given in the ‘Method’ section).

Since Woodham, photometric stereo has progressed a great deal. Advanced systems are now capable of accurately recovering the shape of dynamic surfaces which exhibit multiple reflectance properties [34]. However, such methods place strict constraints on the imaging environment and require multiple image captures. We are interested in performing shape recovery from a single image

in unknown, passive environments.

In our work we extend a method called ‘Shape from color’ (SFC) [10–12]. SFC requires only a single colour image of an object illuminated by three spectrally distinct light sources simultaneously (Fig. 2b) to recover the surface normals of a Lambertian object (Fig. 2c). For both the classic method and SFC, these recovered normals are converted to x and y derivatives and the resulting gradient field is reintegrated to produce a height map (Fig. 2d). The mathematics of this process (reintegrating gradient fields) is beyond the scope of this paper. But [16, 24, 33] present commonly used methods and a review of the topic can be found in [28].

The ability to retrieve shape from individual image captures has led to the multispectral SFC technique being used primarily in the recovery of dynamic shapes. Examples include capture of human faces [35] and recovery from video sequences of non-rigid surfaces [1, 6, 18]. These approaches produce good results however they all require calibration and human subjects must be subjected to multiple lights directed towards them - something which we wish to avoid with our use of a passive imaging environment. The method presented in [1] is also capable of recovering the shape of surfaces with multiple chromaticities, however this requires an additional camera.

An example of colour photometric stereo without an explicit calibration step can be seen in [31] wherein a computer screen acts as the light source. A coloured pattern is displayed and shape is recovered from a video stream in real-time. Though again, we are interested in allowing for more natural scene conditions whereas this technique requires a dark room for good quality recovery (as well as the additional requirement of a computer screen). Other calibration-free methods include [4, 32], however these require many images.

An interesting application of SFC is the capture of micro-geometry shown in [20, 22]. Colour photometric stereo methods mostly assume Lambertian surface reflectance, however [20, 22] use a novel approach to ignore surface reflectance properties by pressing objects into a layer of clear elastomer which effectively ‘paints’ all objects to possess the same reflectance. This highly engineered solution recovers incredibly detailed surface information. However the elastomer layer means that only objects which are roughly planar and smaller than the elastomer (i.e. a few inches in diameter) can be recovered. Although we make the common assumption of Lambertian reflectance, our method can easily be extended to objects of any shape or size and does not require specialist equipment (see Fig. 6 for an example involving little in the way of resources and engineering).

Photometric stereo outside of laboratory conditions is still a relatively unexplored subject. Johnson and Adelson [21] recover shape in natural illumination conditions from a single image, however they require calibration via imaging of a sphere comprised of the same material as the objects to be recovered (obviously this is limiting). Promising results are shown for objects with singular chromaticity in [26] and multiple chromaticity surfaces are recovered under natural illumination in [38]. However both of these methods require a calibration step using a mirrored sphere (as we did in our previous work) and [38] requires multiple images. In comparison to these, our method has an additional constraint in that we require the presence of colourful mutual illumination in the capture environment, however we do not need a

calibration step.

Method

Classic photometric stereo

In classic photometric stereo [37] Woodham proposed 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 a light source as a vector \mathbf{e} . With respect to Lambert's law, a point on a surface with normal $\mathbf{n} = [n_x \ n_y \ n_z]^t$ illuminated by \mathbf{e} results in a pixel value p ,

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

where α is a scalar representing surface albedo (for the purpose of this example we assume $\alpha = 1$ i.e. a white surface). Let us use the notation \mathbf{p}_i to denote the i th triple of pixel responses, three light sources $\mathbf{e}_1, \mathbf{e}_2$ and \mathbf{e}_3) and let \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,

$$\begin{aligned} P &= \begin{bmatrix} \mathbf{p}_1 & \mathbf{p}_2 & \dots & \mathbf{p}_n \end{bmatrix}, \\ E^t &= \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_2 & \mathbf{e}_3 \end{bmatrix}, \\ N &= \begin{bmatrix} \mathbf{n}_1 & \mathbf{n}_2 & \dots & \mathbf{n}_n \end{bmatrix}. \end{aligned} \quad (2)$$

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,

$$P = EN. \quad (3)$$

Since P is known and calibration can give us E , we can recover N :

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

Normals are then transcribed into x and y derivatives and the resulting gradient field is reintegrated, we refer the reader to [28] for further detail.

Shape from color

Illuminating an object with three spectrally distinct light sources simultaneously, we can essentially use the RGB colour channels of a single image in place of the three individual greyscale images of the classic approach. If we denote the measured spectral intensities of the lights as $\mathbf{b}_1, \mathbf{b}_2$ and \mathbf{b}_3 and group them in the same fashion as equation 2,

$$B = \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 & \mathbf{b}_3 \end{bmatrix}. \quad (5)$$

Then given an RGB pixel value \mathbf{c} , we can recover the corresponding surface normal \mathbf{n} at the pixel:

$$\mathbf{n} = F^{-1} \mathbf{c}, \quad (6)$$

where $F = BE$. Simply put, a transform between RGB pixel values and surface normals is established. Typically, this transform is calculated in a calibration step by imaging an object for which surface normals and reflectance are known.

Shape in a box

The ideal experimental conditions for SFC are in a calibrated, laboratory setting. In our previous work we replaced the three spectrally distinct lights of SFC with a colourful environment (Fig. 1). The box was designed by simulating SFC performance for a graphical model of the proposed capture environment. Using synthetic data we calculated that mutual illumination from the box walls should cast sufficient spectrally varied light over the surface of an object to provide a well-conditioned matrix transform F if the box had sufficient depth. We then showed this to be true for a set of real world objects of known shape.

Like SFC, in SiaB we needed to calculate the 3x3 transformation matrix F which relates the RGB values to surface normals. To calculate the transform we performed a calibration step. We measured the lighting environment by capturing an image of a mirrored sphere inside our box (Fig. 1). Using spherical harmonics [29] we then synthesised the appearance of a perfectly Lambertian equivalent. Thus we were able to use Eqn. 6 to acquire a set of RGB pixel values which corresponded to a set of known surface normals and could calculate F .

However, the requirement of a separate calibration step is (as well as being inconvenient) inadequate for environments with constantly changing light - i.e. the majority of scenes outside controlled laboratory settings where we may want to place our box. Further, the chrome sphere itself impacts on the lighting environment causing the process to be inherently inaccurate.

Eliminating calibration

Typically in shape recovery, images are assumed to be captured with orthographic projection. However in practice they are usually subject to a weakly perspective projection. As can be seen in Fig. 1, this was the case for our experimental set-up. If the images truly were orthographic, we would not be able to see the interior walls of the box. It is this fact that we exploit here for the purpose of skipping the calibration step.




Figure 3. Illustrating the 3D space in which the box exists e.g. In a captured image, point A would have $[x, y, z]$ coordinates roughly equal to $[0.5, 1.0, 1.0]$.



Figure 4. Part a) shows an example of a captured image. In b) the walls have been segmented. A synthetic Lambertian sphere is shown in c) - rendered using the surface normal and average RGB value from each wall as light sources.

Firstly let us consider the captured image to have unit height and width. If we consider any points along the top of the box to have a z-coordinate of one (and accordingly, the bottom to have $z = 0$), then we have a three-dimensional space in which the box resides (Fig. 3).

Segmenting the box walls (Fig. 4 b) gives us three sets of points which exist on planes in our 3D space. Given three points on a plane, it is a straightforward matter to calculate the surface normal of each wall. Thus we obtain the direction vectors of three “light sources”.

Next we simply take the average RGB pixel value for each wall and treat these values as the spectral properties of our three “light sources”. Additionally we represent external light entering the box by a fourth white light source aligned with the z-axis. Thus we have four light source direction vectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ and $[0, 0, 1]$ and four corresponding light source colour values $\mathbf{b}_1, \mathbf{b}_2, \mathbf{b}_3$ and $[1, 1, 1]$. As in equation 2, stacking these vectors into matrices gives us,

$$\begin{aligned} B &= \begin{bmatrix} \mathbf{b}_1 & \mathbf{b}_2 & \mathbf{b}_3 & [1, 1, 1] \end{bmatrix}, \\ E^t &= \begin{bmatrix} \mathbf{e}_1 & \mathbf{e}_2 & \mathbf{e}_3 & [0, 0, 1] \end{bmatrix}. \end{aligned} \quad (7)$$

We can then calculate a matrix transform F as shown in equation 6. The effect of this lighting model can be observed in Fig. 4c. Although we are making some naive assumptions, results show that this method holds up well in practice.

Results

It was shown in previous work [14] that our SiaB method produces results which are roughly equivalent in accuracy to those achieved by SFC. Here we reproduce results captured in the same conditions but without the requirement for a calibration step.

In our experiments we use objects which have known shape. These include constructed ‘papercraft’ geometric primitives and a more complex, 3D-printed face (Fig. 5). Ground truth for the primitives can be obtained via a few simple measurements and the original model file provides the true shape of the face. As such, in Tables 1 and 2 we are able to calculate the accuracy of recovered surface normals and height maps. The experiment objects and a height map comparison for the more complex object can be seen in Fig. 5. All images were captured using a Canon Powershot G11.

Object	With Calibration	Without Calibration
Cone	2.34	5.16
Face	10.41	11.54
Pyr4	3.4	3.96
Pyr5	3.47	4.24
Pyr6	6.85	2.81
Sphere	4.89	6.95
Tetra	3.45	4.66
Average	4.97	5.61

Table 1: Angular error of recovered surface normals.

Object	With Calibration	Without Calibration
Cone	95.9	94.04
Face	85.75	83.19
Pyr4	97.18	96.04
Pyr5	97.39	94.37
Pyr6	97.09	97.9
Sphere	89.42	87.46
Tetra	95.38	95.2
Average	94.02	92.6

Table 2: Height map accuracies as percentage values.

In Table 2 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 8, 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.

$$accuracy(\hat{Z}) = 100 - 100 \sqrt{\frac{\sum_{i=1}^n \sum_{j=1}^m (Z_{(i,j)} - \hat{Z}_{(i,j)})^2}{nm}} \quad (8)$$

As can be seen in Tables 1 and 2, without a calibration step there is a slight loss of accuracy for both surface normals and height maps. However it is a small increase in error which is to be expected given the severely decreased amount of lighting information used (i.e. four estimated light sources in place of the 100,000+ taken by imaging the mirrored sphere). There are several obvious additional problems with this simplistic model which

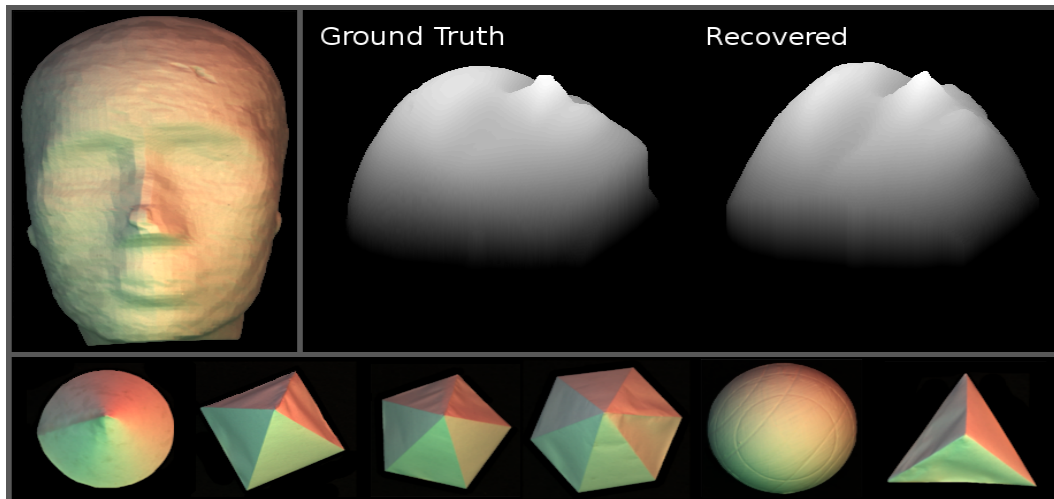


Figure 5. Captured images for all objects and height maps for the 'Face' object (Left to right: cone, pyr4, pyr5, pyr6, sphere, tetra).



Figure 6. Preliminary results demonstrating the resilience of our method to less ideal environments. A lampshade illuminated by indirect illumination from three t-shirts. Shape was recovered from a single rendered image (i.e. jpg) rather than linear raw data with good accuracy.

could be addressed in future work, though it still functions to an acceptable degree even outside of our engineered box solution (Fig. 6).

Firstly, here we have naively used each wall of the box to calculate a distant point light source. However, the box walls are clearly not distant point light sources. As such a lot of potentially useful lighting information is being discarded. Additionally, there is no mechanism in place to account for the interreflection of light within the environment (we are effectively assuming a maximum of one bounce before contact with the object). Both of these issues could be handled by a more complex system - e.g. a radiosity based method [8] wherein the environment would be divided into patches and transfer of light energy between them could be calculated.

Another problem likely to be impacting our results is the assumption of a single source of light external to the box (which we model as being a white light source aligned with the z-axis). For

any real environment in which the box is placed, the truth will invariably differ from this. It would be more accurate to allow for multiple external lights, the direction and spectral properties of which could be calculated either through modification of the box (for example, placing strips with known reflectance around the top); or alternatively the gradient of light cast across the walls could be used to provide cues to external conditions [5].

However, despite these caveats, by building a simple box and without the need for a separate calibration step, we can recover reasonably accurate models of even quite complex shapes.

Conclusion

We have proposed a method of shape recovery which requires only a single image, a colourful environment and has no requirement for a separate calibration step. The initial results shown here demonstrate a slight decrease in accuracy in comparison to our previous calibrated experiments. However given the simple nature of the solution presented here, these results are still impressive. It is the firm belief of the authors that further work will yield greatly improved results.

References

- [1] Robert Anderson, Björn Stenger, and Roberto Cipolla. Color photometric stereo for multicolored surfaces. In *IEEE International Conference on Computer Vision*, pages 2182–2189, 2011.
- [2] Jonathan T Barron and Jitendra Malik. Shape, albedo, and illumination from a single image of an unknown object. *IEEE Conference on Computer Vision and Pattern Recognition*, pages 334–341, 2012.
- [3] Jonathan T Barron and Jitendra Malik. Shape, illumination, and reflectance from shading. Technical report, UC Berkeley, 2013.
- [4] Ronen Basri, David Jacobs, and Ira Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of Computer Vision*, 72(3):239–257, 2007.
- [5] Bas Boom, Sergio Orts-Escolano, Xi Ning, Steven McDonagh, Peter Sandilands, and Robert B. Fisher. Point light source estimation based on scenes recorded by a RGB-D camera, 2013.
- [6] Gabriel J Brostow, Carlos Hernández, George Vogiatzis, Björn Stenger, and Roberto Cipolla. Video normals from colored lights.

- IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(10):2104–2114, 2011.
- [7] Myron Z Brown, Darius Burschka, and Gregory D Hager. Advances in computational stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(8):993–1008, 2003.
- [8] Michael F Cohen and Donald P Greenberg. The hemi-cube: A radiosity solution for complex environments. In *ACM SIGGRAPH Computer Graphics*, volume 19, pages 31–40, 1985.
- [9] Umesh R Dhond and Jake K Aggarwal. Structure from stereo—a review. *IEEE Transactions on Systems Man and Cybernetics*, 19(6):1489–1510, 1989.
- [10] Mark S Drew. Shape from color. Technical report, Simon Fraser University, 1992.
- [11] Mark S Drew. Photometric stereo without multiple images. *Electronic Imaging*, pages 369–380, 1997.
- [12] Mark S Drew and Michael H Brill. Color from shape from color: A simple formalism with known light sources. *Journal of the Optical Society of America A*, 17(8):1371–1381, 2000.
- [13] Jean-Denis Durou, Maurizio Falcone, and Manuela Sagona. Numerical methods for shape-from-shading: A new survey with benchmarks. *Computer Vision and Image Understanding*, 109(1):22–43, 2008.
- [14] Graham D. Finlayson and Christopher Powell. Shape in a box. *4th Color and Photometry in Computer Vision Workshop*, 2014.
- [15] Mark Fisher, Christopher Applegate, Mohammad Ryalat, Stephen Laycock, Mark Hulse, Daniel Emmens, and Duncan Bell. Evaluation of 3-d printed immobilisation shells for head and neck imrt. *Open Journal of Radiology*, 4(04):322, 2014.
- [16] Robert T. Frankot and Rama Chellappa. A method for enforcing integrability in shape from shading algorithms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 10(4):439–451, 1988.
- [17] Graham Fyffe, Xueming Yu, and Paul Debevec. Single-shot photometric stereo by spectral multiplexing. In *IEEE International Conference on Computational Photography*, pages 1–6, 2011.
- [18] Carlos Hernández, George Vogiatzis, Gabriel J Brostow, Björn Stenger, and Roberto Cipolla. Non-rigid photometric stereo with colored lights. *IEEE 11th International Conference on Computer Vision (ICCV)*, pages 1–8, 2007.
- [19] Berthold KP Horn. Shape from shading: A method for obtaining the shape of a smooth opaque object from one view. Technical report, Massachusetts Institute of Technology, 1970.
- [20] Micah K Johnson and Edward H Adelson. Retrographic sensing for the measurement of surface texture and shape. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1070–1077, 2009.
- [21] Micah K Johnson and Edward H Adelson. Shape estimation in natural illumination. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2553–2560, 2011.
- [22] Micah K Johnson, Forrester Cole, Alvin Raj, and Edward H Adelson. Microgeometry capture using an elastomeric sensor. In *ACM Transactions on Graphics*, volume 30, page 46, 2011.
- [23] Kevin Karsch, Zicheng Liao, Jason Rock, Jonathan T Barron, and Derek Hoiem. Boundary cues for 3d object shape recovery. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2163–2170, 2013.
- [24] Peter Kovési. Shapelets correlated with surface normals produce surfaces. In *IEEE International Conference on Computer Vision*, volume 2, pages 994–1001, 2005.
- [25] SD Laycock, GD Bell, N Corps, DB Mortimore, G Cox, S May, and I Finkel. Using a combination of micro-computed tomography, cad and 3d printing techniques to reconstruct incomplete 19th-century cantonese chess pieces. *Journal on Computing and Cultural Heritage*, 7(4):25, 2015.
- [26] Geoffrey Oxholm and Ko Nishino. Shape and reflectance from natural illumination. In *European Conference on Computer Vision*, pages 528–541, 2012.
- [27] Hailang Pan, Hongwen Huo, Guoqin Cui, and Shengyong Chen. Modeling for deformable body and motion analysis: A review. *Mathematical Problems in Engineering*, 2013, 2013.
- [28] Vishal M Patel and Rama Chellappa. Approximation methods for the recovery of shapes and images from gradients. In *Excursions in Harmonic Analysis, Volume 1*, pages 377–398, 2013.
- [29] Ravi Ramamoorthi and Pat Hanrahan. An efficient representation for irradiance environment maps. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 497–500, 2001.
- [30] Daniel Scharstein and Richard Szeliski. High-accuracy stereo depth maps using structured light. In *IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, volume 1, pages I–195, 2003.
- [31] Grant Schindler. Photometric stereo via computer screen lighting for real-time surface reconstruction. In *International Symposium on 3D Data Processing, Visualization and Transmission*, 2008.
- [32] Boxin Shi, Yasuyuki Matsushita, Yichen Wei, Chao Xu, and Ping Tan. Self-calibrating photometric stereo. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1118–1125, 2010.
- [33] Tal Simchony, Rama Chellappa, and M Shao. Direct analytical methods for solving poisson equations in computer vision problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(5):435–446, 1990.
- [34] Daniel Vlastic, Pieter Peers, Ilya Baran, Paul Debevec, Jovan Popović, Szymon Rusinkiewicz, and Wojciech Matusik. Dynamic shape capture using multi-view photometric stereo. In *ACM Transactions on Graphics*, volume 28, page 174, 2009.
- [35] George Vogiatzis and Carlos Hernández. Self-calibrated, multi-spectral photometric stereo for 3d face capture. *International Journal of Computer Vision*, 97(1):91–103, 2012.
- [36] Juyang Weng, Thomas S Huang, and Narendra Ahuja. *Motion and structure from image sequences*. 2012.
- [37] Robert J Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [38] Lap-Fai Yu, Sai-Kit Yeung, Yu-Wing Tai, Demetri Terzopoulos, and Tony F Chan. Outdoor photometric stereo. In *IEEE International Conference on Computational Photography*, pages 1–8, 2013.
- [39] Ruo Zhang, P-S Tsai, James Edwin Cryer, and Mubarak Shah. Shape-from-shading: a survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21(8):690–706, 1999.
- [40] Yueyi Zhang, Zhiwei Xiong, Zhe Yang, and Feng Wu. Real-time scalable depth sensing with hybrid structured light illumination. *IEEE Transactions on Image Processing*, 23(1):97–109, 2014.

Author Biography

Christopher Powell received his BSc in Computing Science from the University of East Anglia (2010) and is currently a PhD student at the same institute. His research interests include physics-based computer vision, computer graphics and digital heritage.

Graham Finlayson is a Professor of Computing Sciences at the University of East Anglia. His research interests span colour, physics-based computer vision, image processing and the engineering required to embed technology in devices. Unusually, for an academic, Prof Finlayson's research is implemented in numerous commercial products ranging from standalone image enhancement software to the processing pipelines found in professional SLR and mobile phone cameras. Professor Finlayson has published over 250 refereed papers and he has also co-authored 30 patents. Professor Finlayson is the recipient of the Philip Leverhulme prize (2002) and a Royal Society-Wolfson Merit award (2008). The Royal Photographic Society awarded him the Davies medal for his contributions to the photographic industry in 2009. He is a fellow of the Royal Photographic Society, the Society for Imaging Science and Technology and the Institution for Engineering Technology. He is a member of the IEEE.