

Representation of Cultural Objects by Image Sets with Directional Illumination

Lindsay W. MacDonald^(✉)

3D Impact Research Group, Faculty of Engineering, University College London,
London, England

`lindsay.macdonald@ucl.ac.uk`

Abstract. In a dome illumination system, many different images can be captured in pixel register from the same viewpoint, each illuminated from a different direction. This is a much richer representation than a single image, and has many applications in cultural heritage for the digitising and display of objects that are flattish with surface relief, such as coins, medals, fossils, rock art, incised tablets, bas reliefs, engravings, canvas paintings, etc. The image sets can be used in three ways: (1) visualisation by interactive movement of a virtual light source over the enclosing hemisphere; (2) 3D reconstruction of the object surface; (3) modelling of the specular highlights from the surface and hence realistic rendering.

Keywords: Illumination · Surface normals · 3D reconstruction · Specular model

1 Dome Photography

It has long been recognised that illumination incident at low angles can help to visualise the relief on surfaces. In the study of canvas paintings it has been used for examining the artist's technique [1], identification of retouching [2] and detection of forgery [3]. In astronomy the oblique direction of the sun's rays has shown the structure of craters on the moon [4]. In archaeology raking light has been used to reveal inscriptions on marble and stone that were otherwise invisible [5]. In palaeographic studies of incised wooden and lead curse tablets from the Roman empire, directional illumination has been used to enhance the marks left by the stylus [6]. Generally it has been employed in an empirical way, with the observer or photographer moving the object relative to the light source (or vice versa) until the desired effect was achieved.

Directional lighting has also been key to the measurement of angular reflectance distributions from the surfaces of materials. In a gonireflectometer, by moving the source of illumination relative to the sample and/or detector, the bidirectional reflectance distribution function (BRDF) can be obtained. Ward developed an automated system for BRDF measurement with a movable light source and rotating sample, under a hemispherical mirror [7]. Malzbender showed how directional illumination could be used in a systematic way for digital photography [8]. He built an illumination dome at HP Labs from an acrylic hemisphere of diameter 18 inches (45 cm) with 24 fixed flash lights. The camera was mounted at the 'north pole' and the object placed on a horizontal surface beneath. This enabled sets of 24 images to be taken in pixel register.

The UCL Dome is an acrylic hemisphere of nominal diameter 1030 mm, fitted with 64 flash lamps, each mounted on a separate circuit board (Fig. 1). The lamps are distributed around the hemisphere, arranged in three tiers of 16, one tier of 12, and one tier of 4 lights at approximately equal intervals. The lowest tier produces raking light across the equatorial plane ($<10^\circ$), whereas the highest tier is nearly polar ($>80^\circ$). The Nikon D200 digital camera is mounted on a rigid steel frame above the dome.




Fig. 1. (left) Hemispherical dome with the camera mounted above the north pole, with 64 flash lights on circuit boards, connected by ‘daisy chain’ ribbon cables; (right) flash lamp firing

Because both the camera mounting point and the lamp positions are fixed, the dome geometry can be characterised precisely. Although the original concept design called for the flash lights to be placed at regular intervals over the surface of the hemisphere, the positions of the lamps in the actual dome, as constructed, differ from the ideal. Three techniques were employed for the geometric calibration of flash light positions in the dome: (1) the shadow cast by a vertical pin onto graph paper; (2) multi-image photogrammetry with retro-reflective targets; and (3) multi-image photogrammetry using the flash lights themselves as targets. It was found that although photogrammetric methods could locate individual target coordinates to an accuracy of 20 microns, the uncertainty of locating the centroids of the flash lights was approximately 1.5 mm [9]. This result is considered satisfactory for photometric imaging purposes.




Fig. 2. Coordinates of 64 flash lamp centroids plotted on hemispherical dome, with representation of lens and sensor (top)

2 Visualisation of Surfaces

The set of images from the dome can be used to visualize the effect of moving a virtual light source over the object, illuminating its surface from any angle in the hemisphere. The question is how to interpolate the 64 angles of the lights in the dome to achieve a continuous movement? One approach would be to make an azimuthal equidistant projection of the lamp coordinates onto the equatorial plane and then to triangulate the network (Fig. 3). The image intensity could then be estimated as a weighted linear combination of the three nearest neighbours of the projected virtual light source. Better results could be obtained by fitting surface patches to the intensity distribution.

An alternative approach is to fit a continuous function to all intensity values over the hemisphere. Malzbender showed that the intensity distribution over all angles of the hemisphere could be approximated by a biquadratic function with six parameters, in a method he called polynomial texture mapping (PTM). Singular value decomposition (SVD) is applied to determine the projection of each of the lamp vectors onto the biquadratic components, and then regression with least-squares minimisation to obtain the six coefficients for each pixel [9]. PTM assumes separability of the reconstruction function, with a constant ‘base colour’ per pixel modulated by an angle-dependent luminance factor:

$$I = L(\Theta_i, \Phi_i, u, v)R(u, v) \quad (1)$$

for $R(u, v)$ and similarly for $G(u, v)$ and $B(u, v)$. The dependence of the luminance on light direction is modelled by the biquadratic function:

$$L(u, v; l_u, l_v) = a_0 l_u^2 + a_1 l_v^2 + a_2 l_u l_v + a_3 l_u + a_4 l_v + a_5 \quad (2)$$

where (l_u, l_v) are projections of the normalised light vector into the local texture coordinate system (u, v) and L is the resultant luminance. A separate set of six coefficients $(a_0 - a_5)$ is fitted to the image data for each pixel and stored in the PTM file at the same spatial resolution as each of the original images. For reconstruction in the viewer software, the position of the virtual light source is expressed in coordinates (u, v) and the intensity of every pixel calculated by Eq. (2). The PTM has the same spatial resolution as each of the original images, but has a low resolution in the angular space of incident illumination, because the n directions of the image set are approximated by only 6 coefficients at each pixel.

PTM has found favour with the museum and cultural heritage community because it provides a convenient and attractive way to visualise objects in collections. The interactive control of lighting direction in the viewer software facilitates perception of the surface structure compared to static photographs, thereby enhancing the legibility of surface relief and inscriptions [10]. The illusion of a 3D surface lit by a movable light source is compelling, even though there is no underlying 3D representation.




Fig. 3. Delaunay triangulation of X,Y coordinates of azimuthal equidistant projection of dome lamps




Fig. 4. Angular distributions of intensity, plotted in the azimuthal equidistant projection, for four approximations to the measured intensity distribution. The black dots represent actual reflected intensity of the 64 lamps. All intensities are normalised to a maximum of 100.

An improved method of fitting the directional distributions was introduced by Gautron *et al* [11] by limiting the domain of the orthogonal basis functions of spherical harmonics to a hemisphere instead of the full sphere. These hemispherical harmonic (HSH) functions provide a more compact and accurate way of representing hemispherical distributions than the biquadratic function used in PTM. They have since been widely adopted for a variety of computer graphic applications where only half of the spherical distribution needs to be modelled, such as the representation of BRDFs, environment map rendering of non-diffuse surfaces and global-illumination computation.

HSH components are expressed as functions of angles for azimuth θ and colatitude φ over the hemisphere [12]. Good results are obtained with sixteen components, which include four first-order, five second-order and seven third-order terms. Fitting the coefficients follows the same SVD-based procedure as for PTM, and can be applied to the image luminance (weighted sum of R,G,B channels) to provide the angular modulation at each pixel of a constant R,G,B colour value. These components can be conveniently visualised by projecting the hemisphere onto a plane through an azimuthal equidistant projection. Comparison of the HSH and PTM renderings (Fig. 4) shows that HSH (2nd-order with nine coefficients) gives a better representation of the directionality of the surface, with higher contrast for local gradients.

How many images are needed to give an accurate rendering of the angular reflectance distribution in PTM or HSH? Ideally one should capture images with illumination from all necessary angles but no more. The challenge is to reduce the number of photographic image samples that need to be acquired while preserving the power of

the digital model to represent the object realistically. A full BRDF analysis requires a systematic sampling of a four-dimensional space, with both illumination and view angles able to range over the full hemisphere. In the PTM and RTI scenarios, the view direction is always fixed (usually at the zenith, perpendicular to the centre of the object surface) and only the illumination direction is variable.

Gunawardane *et al* analysed the sampling of both view and lighting directions and whether methods for interpolation could be improved if both view and lighting information were available [13]. They conducted a data-driven study in which a test object was illuminated by a hemispherical dome with 64 tungsten lights. The object sat on a turntable and a full set of 64 images was captured for 360 rotational angles of the object, at 1° intervals for a total of 23,040 images. The complete image set was then sub-sampled for intervals of

both lighting and view angles and the errors calculated between HSH fittings of both the full and subsampled image sets. The results (Fig. 5) indicated that the minimum number of lamps is approximately 10, 20, 36 and 56 for the 1st to 4th harmonic orders respectively. The pitfall in interpolation of images from different lighting directions was found to be that errors in flow vectors caused pixels to move to incorrect positions, producing visible tearing artifacts and structural discontinuities.

Drew *et al* observed that for non-Lambertian phenomena matrix factorisation methods can produce inaccurate surface normals and lighting directions [14]. Because the basic PTM method relies on a matte surface and linear regression, it fails to model phenomena such as inter-reflections, specularities and shadows. Increasing the degree of the PTM polynomial model, for example by the use of HSH basis functions, may help to model these effects but at the expense of degrading the interpolated results at non-sampled light directions due to over-fitting. For finding the matte part of the photometric model, they used the Least Median of Squares (LMS) method, which provides automatic identification of outliers, both specular highlights and shadows. Knowledge of the inlier pixel values means that the recovered surface albedo, chromaticity and surface normals are robust, in the sense of ignoring outlier contributions and thus more accurately mapping surface reflectance, colour and shape. They also altered the polynomial used in PTM so as to generate a subset of three regression coefficients that is exactly correct in the case when the inliers are Lambertian.

Brady *et al* developed an alternative method of visualizing the relief of incised tablets, called shadow stereo, after observing how a professional palaeographer examined a stylus tablet, holding it horizontally on his upturned palm to be illuminated at a grazing angle, and slowly rotating it to change the angle of elevation [15].




Fig. 5. Error vs number of lights for four orders of hemispherical harmonics (Gunawardane, 2009)

3 3D Reconstruction

The image sets captured in the dome contain information about the geometry of the object surface. The photometric stereo technique enables the normal at each point to be determined for a single viewpoint, using the principle that the intensity of reflected light depends on the angle of incidence of the light onto the surface and the reflectance factor. With a perfectly Lambertian surface and in the absence of noise, only three intensity values $[I_1, I_2, I_3]$ from non-coplanar light sources with unit direction vectors $[\mathbf{L}_1, \mathbf{L}_2, \mathbf{L}_3]$ would be sufficient to solve for both the normal direction \mathbf{N} and the surface albedo ρ :

$$I_i = \rho \mathbf{L}_i \cdot \mathbf{N} = \rho |L_i| \cos \alpha_i \quad (3)$$

where α_i is the angle between the normal and lamp vector i . In practice, normals calculated in this way from three light directions exhibit an unacceptable level of noise and vary widely according to the particular combination of lamps selected. Better results can be obtained for noisy image data by calculating normals for many triplets of light sources. By selecting suitable combinations of three lamps, candidates for the normal can be calculated for every pixel. For a non-Lambertian surface, however, the above method gives incorrect results, because the effect of surface gloss is to exaggerate the apparent gradient of the surface.

A new method for estimating normals has been developed, which is robust and adapts to the presence of both shadows and surface gloss [12]. First all of the intensity values at a pixel are extracted from the image set and treated as a vector. The intensity values are then sorted into ascending order and the cumulative sum calculated. The subset of lamps is selected for which the normalised cumulative values lie between two thresholds, nominally 0.10 and 0.25. These thresholds are chosen to select a region of the sorted distribution that follows the slope of the cumulative sorted cosine, i.e. related to the diffuse component of the reflection. Fig. 6 shows the results of applying the technique to a 19th-century terracotta roundel of Chopin, and the derived gradients $P = (\partial I / \partial x)$ and $Q = (\partial I / \partial y)$, encoded in false colour.



Fig. 6. (left) Photometric normal vectors in false colour (Nx red, Ny green, Nz blue); (right) False colour composite of gradients (P red, Q green)

A novel method for integrating the gradients to reconstruct height was introduced by Frankot & Chellappa [16], using the Fourier transform to regularise (i.e. to enforce integrability of) the gradients in the frequency domain. This is neatly implemented in Matlab by a few lines of code, taken from the library developed by Kovesi [17].




Fig. 7. (left) Oblique view and (right) elevation of Chopin surface reconstructed by the basic Frankot & Chellappa integration of gradients in the frequency domain

Applying this technique to the Chopin gradients (Fig. 6 right) yields a 3D surface that is continuous and is recognisably Chopin, but is distorted over the whole area with the height greatly amplified. Fig. 7 (left) shows an oblique view, which looks very plausible, but when the same structure is viewed in elevation (Fig. 7 right) it is seen that the height range is from -27.6 to $+79.9$, an overall maximum height of 107.5 mm, compared with the true maximum height above the baseplane of 22.5 mm. Also there is a false undulation of the base with a period of approximately one cycle over the whole width. The problem is that although the photometric gradients give a good representation of the spatial frequencies in the surface, right up to the Nyquist frequency, they are not accurate for very low frequencies of a few cycles over the full object diameter. Such frequencies are represented in the Fourier plane by only a few sample points close to the (shifted) origin. Errors in these frequencies can result in ‘curl’ or ‘heave’ in the baseplane, even though the superimposed higher spatial frequencies may be accurate.




Fig. 8. (left) Using a height measuring gauge; (right) Heights of selected points in mm, superimposed on an image of the Chopin terracotta taken under all lights in Tier 3

The solution is to replace the inaccurate low frequencies of the photometric normals by the more accurate low frequencies of a surface constructed from a few known heights [18]. This can be conveniently achieved from the values measured by a digital height gauge (Fig. 8) by first interpolating them to produce a smooth ‘hump’ and then transforming into the frequency domain by an FFT (Fig. 9).




Fig. 9. (left) Smooth surface of hump produced by interpolation of measured points; (centre and right) Log(power) distribution of spatial frequencies of hump and photometric gradients

The low spatial frequencies of the gradients from the Frankot-Chellappa integration are replaced by the corresponding frequencies from the hump. Rather than an abrupt change at a given threshold frequency, they are blended over a radial distance in the range 1.5 to 4.0 pixels by a linear interpolation function (Fig. 10 left). The power of the high frequencies from the gradients is scaled by the ratio of the low/high power in the region affected, in order to maintain the correct overall power distribution.




Fig. 10. (left) Cross-over of low and high frequency components; (right) Reconstructed surface

Differences between the 16 values measured by the height gauge and the reconstructed values at the same positions are well distributed (Fig. 11), ranging from -1.31 to +1.92 mm with mean zero, mean-of-absolutes 0.61 and stdev 0.83. Even better results can be obtained by using the point cloud from a 3D scanner to provide the geometric basis of the low frequency components in the reconstruction.




Fig. 11. Measured and reconstructed heights and their differences (mm)

4 Specular Modelling

The image sets captured in the dome, illuminated from 64 known directions over the hemisphere, contain information about the directionality of reflection from an object surface. The aim is to model the luminance variation at each point on the object surface as a function of the angle of illumination, in such a way that the reconstructed images are indistinguishable from the original photographs. This would also enable views of the object to be ‘relit’ for a continuous range of illumination angles in between those of the fixed lamps in the dome.

The decorative test object used in this study was a polished brass dish, 125 mm in diameter, embellished in the Damascene fashion with inlaid copper and silver arabesques (Fig. 12 left). There is a significant trade-off in choosing the exposure setting when photographing the object in the dome: if too low then most of the non-specular pixels are of very low intensity (as in this case) and hence greatly affected by sensor noise; if too high then most of the specular pixels are over-exposed, producing the maximum output value and causing blooming in neighbouring pixels by spill-over of photoelectrons in the sensor.



Fig. 12. Damascene dish: (left) image illuminated by four lamps in tier 5 of dome; (centre) normal vectors in false colour; (right) albedo

The normal vector and albedo were computed for every pixel (Fig. 12 centre and right). The normals are quite subtle because most of the surface is horizontal and the relief of the decoration is shallow. The albedo is surprisingly dark and chromatic, representing the diffuse ‘base colour’ of the metal without any specular component.

The ‘specular quotient’ is calculated as the ratio between the actual intensity for each lamp direction and the intensity that would be produced by a perfect diffuser in the same direction. The more shiny the surface, the greater the quotient value (Fig. 13 left). The specular direction vector is calculated as a weighted sum of the lamp vectors exceeding a threshold, multiplied by the corresponding specular quotient values. The same weighted sum gives the colour of the specular reflection (Fig. 13 right). For most materials this would be the same colour as the illumination, i.e. white, but for metals the specular component carries the colour of the metal. Here for the Damascene dish the colours of the brass, copper and silver are clearly defined.



Fig. 13. (left) Specular quotient = ratio of specular/diffuse intensities; (centre) specular colour; (right) Specular direction vectors in false colour

The resulting specular direction vectors (Fig. 13 right) have the same general appearance as the normals (Fig. 12 centre) but are more chromatic because the specular gradients are greater with respect to the view vector. In conventional practice in computational photography it is almost universally assumed that the specular angle should be exactly double that of the normal, and for a perfect mirror this would of course be true. But the surfaces of real objects have a meso-structure with fine texture and granularity. One pixel as sampled by the camera may span a number of micro-facets at different angles, which reflect light differently from the incident illumination.

The approach taken here is to use the ideal specular (at double the angle of the normal) as a guide to where the specular angle should be. A weighted sum is taken of all lamp vectors within a cone of 45° around this direction, weighted by their quotient values. It is clear from scatter-plotting the specular vs normal angles for a random selection of 10,000 pixels (Fig. 14) that there is a considerable amount of variation around the line of slope 2 (i.e. specular angle = $2 \times$ normal angle), which is a genuine indication of the roughness of the surface. Pixels with low values of specular quotient (blue in the figure) generally have a greater scatter. Some clustering onto the five tier angles of the dome is evident in the pixels of high quotient values (red in the figure). The figure suggests that the maximum normal angle that can be quantified by the photometric stereo technique is $c.35^\circ$, with corresponding maximum specular angle of $c.70^\circ$.

In the general case the bidirectional reflectance distribution function (BRDF) has four degrees of freedom, giving the reflectance of the surface at any viewpoint when illuminated from any direction. In the case of dome imaging, however, the viewpoint is fixed with the camera always at the 'north pole' of the hemisphere and the object lying in the equatorial plane. So the problem is simplified to finding a two-dimensional function of the reflectance factor toward the camera, given the normal




Fig. 14. Specular vs normal angles, classified by quotient value

and lamp vectors. A further simplification is to assume that the function of reflectance is isotropic and therefore rotationally symmetric, i.e. dependent only on the radial angle ω from the peak but not on the phase angle around the peak. The required function needs to be positive, continuous and monotonic, with a peak at $\omega = 0$ and asymptotic to zero as $\omega \rightarrow 90^\circ$ (excluding the Fresnel component at grazing angles).

The model adopted to fit the specular peak is based on the Lorentzian function, because it naturally conforms to the observed shape and is mathematically convenient [12]. In particular the broad flanks enable the scattered light at perispecular angles to be modelled more effectively than the Gaussian function, which falls too quickly to zero. The comparison can be seen by fitting both functions empirically to the distribution (Fig. 15). An offset in the X axis has been made to accommodate the horizontal scatter, and the scale factors (divisors of X value) are different. But it is clear that the Gaussian approaches zero too rapidly and therefore underestimates the reflected intensity in the critical intermediate angles between peak and flank. The Lorentzian can be written as a function of three variables:

$$f(x) = \frac{a}{1 + \left(\frac{x}{s}\right)^2} + c \quad (4)$$

where a is the amplitude of the peak, s is the scale factor (horizontal spread), and c is a constant (uplift).

5 Specular Classification

A detail of the specular colour image of the Damascene dish shows clearly that the specular highlights of the three metallic components carry the colour of the metal (Fig. 16 left). Scatter-plotting 10,000 points chosen at random by their colours and locations in RGB space shows (Fig. 16 right) that they lie in an oblate region around the long diagonal of the colour cube, i.e. the neutral axis. There is a surprising amount of colour variation for what appears to be a surface composed of only three materials, and the tonal variation is continuous from black to the lightest points at about 0.7 of full range.

The pixel colours are converted from RGB via XYZ to CIELAB, assuming the sRGB colour space and the CIE standard 2° observer with D65 white point. Plotting the same points on the a^*-b^* chromatic plane shows that the colours of the three metals, brass, copper and silver, have distinctly different hue angles (Fig. 17). This provides the opportunity to segment the image pixels into four categories, corresponding to the three metals plus black.




Fig. 15. Comparison of Gaussian and Lorentzian functions against a real distribution of reflected intensities




Fig. 16. (left) Detail of specular colour image of Damascus dish, size 500x500 pixels; (right) scatter plot of 10,000 pixels in R,G,B colour space

The simple way to classify is by hue angle around the centroid of the three category centres, where reddish colours in the range $[-45^\circ, +70^\circ]$ correspond to copper; the yellow-greenish colours in $[70^\circ, 180^\circ]$ to brass; and the slightly bluish colours in $[180^\circ, -45^\circ]$ to silver. A more effective method is to categorise each pixel by its nearest distance to one of the focal colours for the three metals. The resulting ‘posterised’ image is equivalent to a K-means classification with four cluster centres (including black).



Fig. 17. (left) 10,000 pixels plotted on the CIELAB a^*-b^* chromatic plane, with centroids for the three metals; (centre) classification by hue angle; (right) classification by nearest colour

The ability to classify different regions of a heterogeneous surface according to their gloss enables each region to be modelled and rendered in a different way. This is an important capability for objects that are made of multiple materials, such as inlays, and also for objects that were once homogeneous but have weathered variably across the surface. It is interesting to consider whether metals could be classified in the same way.

In an attempt to differentiate the specular curves of the three metals in the Damascus dish, the map generated by image classification (Fig. 17 right) was used to select 500 random samples of each of the three metals. Curves were fitted by the Lorentzian model and plotted in superimposition (Fig. 18 left). All three sets of curves show a




Fig. 18. (left) Specular curves fitted by the Lorentzian model to 500 pixels of silver on the Damascene dish; (right) Characteristic curves for the three metals for peak amplitudes >15

similar behaviour, with a few cases having peak values of very high amplitude, in excess of 500, but the majority much lower.

Taking the median of each parameter in the sets for each metal for those cases where the amplitude exceeds 15, and using the median parameter values in the Lorentzian model gives the indicative curves of Fig. 18 right. The ordering of amplitude is: silver highest, copper second and brass lowest, but the relative differences are small and the variance of the curves for individual pixels is so great that these curves could not be used as a reliable diagnostic to determine the type of metal. The colour in the albedo and specular highlights is a much more reliable guide. All that one can say in this case is that the freshly polished silver is likely to be slightly brighter in the specular highlights than the other two metals.

6 Conclusion

A set of images in pixel register under controlled directional lighting provides a much richer representation of an object than a single image, because it contains information about both the topography and specularity of the surface. With appropriate metadata, including the directions of the incident light sources, camera position and lens distortion, such a dataset can be considered as a valid archival representation of the object, with many applications for education, conservation and interpretation.

Through an interactive visualisation technique, such as PTM, the sense of materiality of the object can be conveyed much more strongly than through a static image display. This is an example of how, as Witcomb says, “multimedia installations in museums can enhance ... the ‘affective’ possibilities of objects” [19]. They can “act as releasers of memory in much the same way as objects can make unconscious memories conscious. This they achieve through their power to affect us by ‘touching’ us or ‘moving’ us.”

Acknowledgements. Thanks to colleagues in the 3DImpact Research Group, especially Stuart Robson, Mona Hess and Ali Ahmadabadian, for assistance and encouragement in this research.

References

1. Rosen, D., Marceau, H.: A Study in the Use of Photographs in the Identification of Paintings. *Technical Studies in the Field of the Fine Arts* **6**(2), 75–105 (1937)
2. Ruhemann, H.: Criteria for distinguishing additions from original paint. *Studies in Conservation* **3**(4), 145–161 (1958)
3. Stolow, N.: The Canadian art fraud case. *Canadian art* 94 (1964)
4. Phillips, J.: Notices of Some Parts of the Surface of the Moon. *Phil. Trans. Royal Society of London* **158**, 333–345 (1868)
5. Engelbach, R.: A foundation scene of the second dynasty. *The Journal of Egyptian Archaeology*, 183–184 (1934)
6. Bowman, A.K., Brady, J.M., Tomlin, R.S.O.: Imaging incised documents. *Literary and Linguistic Computing* **12**(3), 169–176 (1997)
7. Ward, G.J.: Measuring and modeling anisotropic reflection. In: *Proc. 19th Conf. on Computer Graphics (Siggraph)*, pp. 265–272. ACM (1992)
8. Malzbender, T., Gelb, D., Wolters, H.: Polynomial texture maps. In: *Proc. 28th Conf. on Computer Graphics (Siggraph)*, pp. 519–528. ACM (2001)
9. MacDonald, L.W., Ahmadabadian, A.H., Robson, S.: Determining the coordinates of lamps in an illumination dome. In: *Proc. SPIE Conf. on Videometrics, Munich* (2015)
10. Earl, G., Martinez, K., Malzbender, T.: Archaeological applications of polynomial texture mapping: analysis, conservation and representation. *J. Archaeological Science* **37**(8), 2040–2050 (2010)
11. Gautron, P., Krivanek, J., Pattanaik, S., Bouatouch, K.: A novel hemispherical basis for accurate and efficient rendering. In: *Eurographics Symposium on Rendering Techniques*, pp. 321–330 (2004)
12. MacDonald, L.W.: Colour and directionality in surface reflectance. In: *Proc. Conf. on Artificial Intelligence and Simulation of Behaviour (AISB)*. Goldsmiths College (April 2014)
13. Gunawardane, P., Wang, O., Scher, S., Rickards, I., Davis, J., Malzbender, T.: Optimized image sampling for view and light interpolation. In: *Proc. 10th Intl. Conf. on Virtual Reality, Archaeology and Cultural Heritage (VAST)*, Eurographics, vol. 10, pp. 93–100 (2009)
14. Drew, M., Hel-Or, Y., Malzbender, T., Hajari, N.: Robust estimation of surface properties and interpolation of shadow/specularity components. *Image and Vision Computing* **30**(4), 317–331 (2012)
15. Brady, M., et al.: Shadow stereo, image filtering and constraint propagation. In: Bowman, A.K., Brady, M. (eds) *Images and Artefacts of the Ancient World*, pp. 15–30 (2005)
16. Frankot, R.T., Chellappa, R.: A method for enforcing integrability in shape from shading algorithms. *IEEE Trans. on Pattern Analysis & Machine Intell.* **10**(4), 439–451 (1988)
17. Kovesi, P.: Shapelets correlated with surface normals produce surfaces. In: *Proc. Tenth IEEE Intl. Conf. on Computer Vision*, vol. 2, pp. 994–1001 (2005)
18. Tomimaga, R., Ujike, H., Horiuchi, T.: Surface reconstruction of oil paintings for digital archiving. In: *Proc. IEEE Southwest Symp. on Image Analysis and Interpretn.*, pp. 173–176 (2010)
19. Witcomb, A.: The materiality of virtual technologies: A new approach to thinking about the impact of multimedia in museums. In: Cameron, F., Kenderdine, S. (eds) *Theorizing Digital Cultural Heritage: A Critical Discourse*, pp. 35–48. The MIT Press (2007)