

Colour and Directionality in Surface Reflectance

Lindsay W. MacDonald¹

Abstract. The reflectance of light from the surface of an object is modelled by a new technique, using components extracted from a set of 64 images taken in an illumination dome. Techniques of photometric stereo are used to determine albedo and surface normals at each pixel. The ratio of actual to diffuse reflectance is modelled to find the specular angle and an angular distribution of intensity based on a Lorentzian function. Displaying images in a custom software viewer demonstrates the close appearance of the model to each original photograph when illuminated from the same angle.

1 INTRODUCTION

There is nothing absolute about the colour of an object surface: it changes continually with illumination and orientation. Given many images of the same object, even under the same light from different angles, how can one say what is the ‘true’ surface colour? Colorimetry specifies the colour of an object as the product of the illuminating power by the reflectance factor of the surface by the sensitivity of the observer, integrated over all wavelengths of the visible spectrum. This is the basis of the ubiquitous CIE system, but it relies on the assumption that the surface is perfectly matte so that every point reflects the incident light equally in all directions, i.e. that it is perfectly Lambertian. But in fact almost every real material exhibits some angular dependence in the way it reflects light, and this must be taken into account when modelling the appearance of the object, by adding a gloss component to the underlying diffuse colour. The added light may appear as a sheen over the surface or as localised specular highlights, but its effect is to modulate the lightness and thereby to change the colour stimulus.



Figure 1. Test object used in this study, the ‘Roman medallion’.

The dome imaging system at UCL enables sets of images of an object to be taken with illumination from different directions. A hemisphere of 1 metre diameter is fitted with 64 flash lights, calibrated so that the geometric centroid of every light source is known to within 3mm. A Nikon D200 digital camera fixed at the ‘north pole’ captures a series of 64 colour images, each illuminated by flash light from a different direction, all in pixel register. This provides a much richer dataset than a single image, and opens the way for characterising the texture, gloss and geometry of the object surface and for visualising its appearance in any lighting.

For this study a ‘Roman medallion’ was chosen as the test object (Fig. 1). This is modern, made in Rome, cast from brass and painted to look like antique bronze. It is 10 cm in diameter, c.15 mm thick, and weighs 674 grams. The inscription DN HONORI VSPFAVG is an abbreviation of D(ominus) N(oster) HONORI-VS P(ius) F(elix) AVG(ustus). It depicts the Emperor Honorius, who ruled from 395 to 423 AD in the declining years of the Roman Empire. The design is probably based on a gold solidus coin, minted in Ravenna.

2 PTM VISUALISATION

Polynomial texture mapping (PTM) was developed in 2001 by Malzbender [1], who showed that the intensity distribution over all angles of the hemisphere above a fixed object could be approximated by a biquadratic function with six parameters. 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. Malzbender also developed software utilities for fitting and viewing the PTM representations, and these have been made publicly available by HP [2]. The method has found favour with the museum and cultural heritage community because it provides a convenient and attractive way to visualise objects in collections. To the end-user the viewer gives a compelling illusion of a 3D surface lit by a movable light source, even though there is no underlying 3D representation.

PTM has the ability to represent arbitrary geometric shadowing and diffuse shading effects across a surface. It is assumed that the chromaticity is independent of angle, and that only the luminance of each pixel varies with light source direction. This enables separability of the reconstruction function, with a constant 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 luminance on light direction is modelled by a biquadratic function:

¹ Dept. of Civil, Environmental and Geomatic Engineering, University College London, email: ucfslw@ucl.ac.uk

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

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. Thus a separate set of 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 (Fig. 2).



Figure 2. Six coefficients of PTM luminance, shown as image planes. Mid-grey represents zero, with positive values lighter and negative values darker.

3 HEMISPHERICAL HARMONICS (HSH)

An improved method of fitting the directional luminance distribution was introduced in 2004 by Gautron *et al* [3]. It is a transformation of the Associated Legendre Polynomials which form the basis of spherical harmonics used in geodesic and planetary modelling. By limiting the domain of the orthogonal basis functions to a hemisphere, HSH functions provide a more compact and accurate way of representing hemispherical distributions. They have been widely adopted for computer graphic applications such as 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 co-latitude φ over the hemisphere:

$$\theta = \text{acos}\left(\sqrt{1 - l_u^2 - l_v^2}\right) \text{ and } \varphi = \text{atan2}(l_v, l_u) \quad (3)$$

Good results are obtained with sixteen components, which include four first-order, five second-order and seven third-order terms:

$$\begin{aligned} h_0 &= \sqrt{1/2\pi} & (4) \\ h_1 &= \sqrt{6/\pi} \cos(\varphi) \sqrt{\cos(\theta) - \cos^2(\theta)} \\ h_2 &= \sqrt{3/2\pi} (2\cos(\theta) - 1) \\ h_3 &= \sqrt{6/\pi} \sin(\varphi) \sqrt{\cos(\theta) - \cos^2(\theta)} \\ h_4 &= \sqrt{30/\pi} \cos(2\varphi) (\cos^2(\theta) - \cos(\theta)) \\ h_5 &= \sqrt{30/\pi} \cos(\varphi) (-1 + 2\cos(\theta) \sqrt{\cos(\theta) - \cos^2(\theta)}) \\ h_6 &= \sqrt{5/2\pi} (1 - 6(\cos(\theta) - \cos^2(\theta))) \\ h_7 &= \sqrt{30/\pi} \sin(\varphi) (-1 + 2\cos(\theta) \sqrt{\cos(\theta) - \cos^2(\theta)}) \\ h_8 &= \sqrt{30/\pi} \sin(2\varphi) (\cos^2(\theta) - \cos(\theta)) \\ h_9 &= \sqrt{140/\pi} \cos(3\varphi) (\sqrt{\cos(\theta) - \cos^2(\theta)})^3 \\ h_{10} &= \sqrt{210/\pi} \cos(2\varphi) (1 - 2\cos(\theta)) (\cos(\theta) - \cos^2(\theta)) \\ h_{11} &= \sqrt{84/\pi} \cos(\varphi) \sqrt{\cos(\theta) - \cos^2(\theta)} (1 - 5(\cos(\theta) - \cos^2(\theta))) \\ h_{12} &= \sqrt{7/2\pi} (12\cos(\theta) - 1 - 30\cos^2(\theta) + 20\cos^3(\theta)) \\ h_{13} &= \sqrt{84/\pi} \sin(\varphi) \sqrt{\cos(\theta) - \cos^2(\theta)} (1 - 5(\cos(\theta) - \cos^2(\theta))) \\ h_{14} &= \sqrt{210/\pi} \sin(2\varphi) (1 - 2\cos(\theta)) (\cos(\theta) - \cos^2(\theta)) \\ h_{15} &= \sqrt{140/\pi} \sin(3\varphi) (\sqrt{\cos(\theta) - \cos^2(\theta)})^3 \end{aligned}$$

These components can be conveniently visualised by projecting the hemisphere onto a plane through an azimuthal equidistant projection (Fig. 3). Apart from the first constant term, all functions are equally balanced between positive and negative excursions.




Figure 3. The first sixteen modes of hemispherical harmonics, shown in azimuthal equidistant projection onto the plane. Positive values are indicated by green and negative by blue.

Fitting of the coefficients follows the same procedure as for PTM, and can be applied to the luminance channel to provide the angular modulation at each pixel of a constant R,G,B colour value (Fig. 4).



Figure 4. Sixteen coefficients of HSH luminance, shown as image planes.

PTM and HSH can be regarded as two members of a more general family of Reflectance Transform Imaging (RTI) methods, in which the distribution of light reflected from an object surface can be modelled as functions of space, angle, spectrum and time. An open-source fitter and viewer that supports both PTM and HSH image formats has been developed and is used and supported by a worldwide community of practitioners, particularly in cultural heritage [4]. Key to the broader adoption of RTI has been the development of Highlight-RTI, in which a glossy sphere is placed in the scene so that the direction of incident illumination can be inferred from the coordinates of the highlight in each image. This obviates the need for a dome system, and enables the photographs to be taken *in situ*. A good review of applications in archaeology is given by Earl *et al* [5].




Figure 5. An 8x8 mosaic of a detail of images from 64 lamps in dome.

4 PHOTOMETRIC NORMALS

The photometric stereo technique enables the normal at each point of the object surface to be determined for a single viewpoint, using the principle that the intensity of the reflected light depends on the angle of incidence of the light onto the surface. With a perfectly Lambertian surface and in the absence of noise, only three intensity values from non-coplanar light sources L_i would be sufficient to solve for both normal direction \mathbf{N} and surface albedo ρ :

$$L_r = \rho \mathbf{L}_i \cdot \mathbf{N} = \rho |\mathbf{L}_i| \cos \alpha \quad (5)$$

In practice, better results can be obtained for noisy image data by taking the median of results for many triplets of light sources [6], but this does not give the correct solution for a non-Lambertian surface. Consider the geometry of Fig. 6 at point P on a curved shiny surface. The view vector \mathbf{V} , passing through the perspective centre of the camera lens, subtends an angle α with the normal \mathbf{N} . Any lamp near the specular vector \mathbf{S} , at an angle of α from \mathbf{N} and 2α from \mathbf{V} , will produce an intensity in the image greater than would be produced by a matte surface. In the solution of the photometric stereo equations (5) this is interpreted as if the surface normal were orientated closer to \mathbf{S} than it actually is, producing the distorted normal \mathbf{N}' . Thus the effect of surface gloss is to exaggerate the apparent gradient of the surface (dotted line).




Figure 6. Geometry of specular reflection, showing distortion of normal.

Real surfaces are non-Lambertian and non-planar, and so exhibit both gloss and self-shadowing from oblique lighting. Fig. 5 shows a mosaic of a detail of 200x200 pixels from all 64 images of the medallion, taken in the UCL dome under 64 different lighting directions. Although the mean brightness increases with illumination angle, as one would expect, the shadowing is severe at low incident light angles, especially for the lowest tier of lamps with an elevation less than 10° . Also the shiny surface produces very bright areas at different positions in the images, requiring (ideally) a high dynamic range (HDR) imaging strategy. The problem is therefore how to select the best subset of all the intensity values at each pixel to avoid both shadows and specular highlights.

Various ways of compensating for shadows and highlights in photometric stereo have been proposed. Julia *et al* [7] had an alternation technique to decompose the measurement matrix, containing the intensity images, into separate surface and light source matrices. The shadows and saturated regions were considered as missing data, and so did not influence the results. Chandraker *et al* [8] used a shadow labelling algorithm based on fast graph cuts as input to a photometric stereo algorithm using multiple light sources to enhance surface coverage and reconstruction accuracy. Drew *et al* [9] used the inliers found from luminance regression, and were able to identify both specular and shadow pixels, model their contribution, and interpolate data for out-of-sample lighting directions.

In this paper a new method for estimating normals is presented, which is robust and adapts to the presence of both shadows and surface gloss. First all the intensity values at a pixel are extracted from the image set and treated as a vector P . As an example a pixel is selected from the detail of Fig. 5 on the cheekbone just below the eye. Plotting the 64 values against lamp number (Fig. 7) shows the wide variation in intensity for different lamp angles, with peaks for lamp numbers 13, 30, 46 and 59. Also plotted for comparison (in magenta) is the intensity that would be expected from a perfect Lambertian surface of the same albedo illuminated by the same lamps.

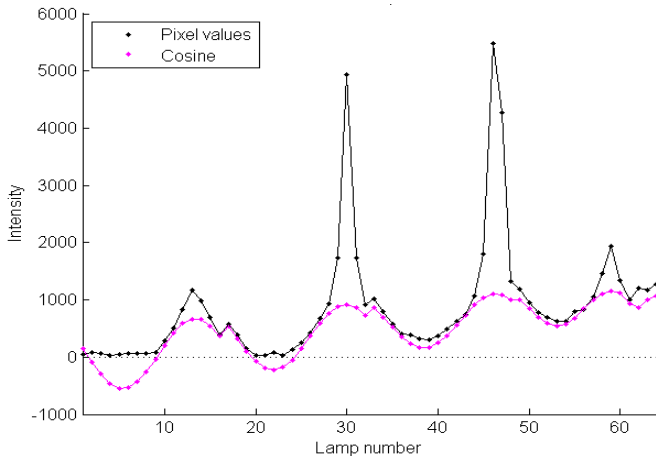


Figure 7. Actual intensity values at a single pixel for 64 illumination directions, compared with a similarly-illuminated perfect diffuser.

The intensity values are then sorted into ascending order and the cumulative sum calculated (Fig. 8). The subset of lamps is selected for which the normalised cumulative values lie between two thresholds, set to 0.1 and 0.2. In the example for this pixel, eight lamps are selected. Regression is then applied over the subset of lamp vectors to solve Eq. (5) to estimate the normal and albedo.

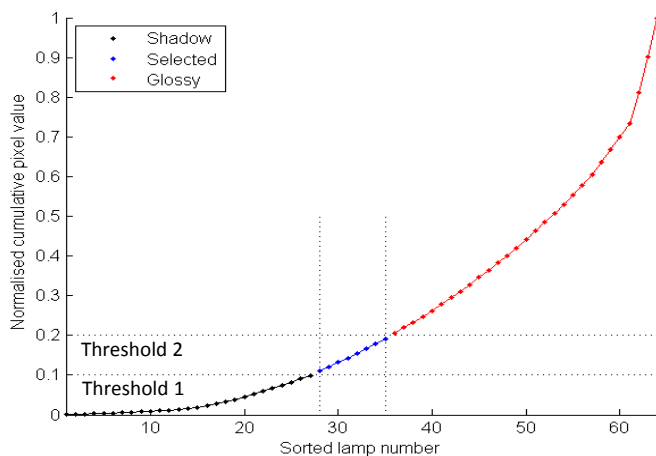


Figure 8. Cumulative distribution of sorted intensity values

The significance of the procedure can be seen by plotting sorted intensity values (Fig. 9). The sorted cosine values (plotted in magenta) show the baseline of the diffuse component. The selected subset lies above all cases where the pixel is in shadow, but below cases where there is an added component of gloss. The colour albedo arising from this procedure is a good approximation to the intrinsic diffuse component of surface reflectance.

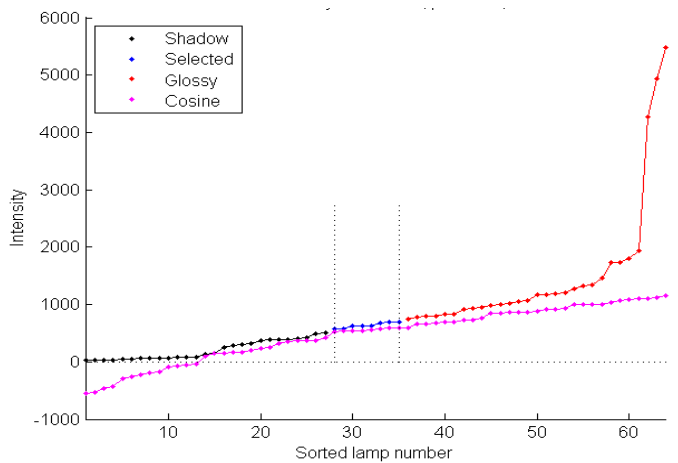


Figure 9. Sorted intensity distribution, showing the selected subset of lamps.

The subset of lamps selected by the procedure can also be visualized by plotting 3D vectors of length proportional to the pixel intensity corresponding to each lamp (Fig. 10). The heavy black line represents the view vector V , the green line the normal N , and the red line the ideal specular S . The black, blue and red lines correspond to the points in Figs. 8 and 9 for the 64 lamps. It is apparent that the selected lamp vectors (shown in blue) lie in an umbrella formation, well away from the normal where the influence of gloss is low.

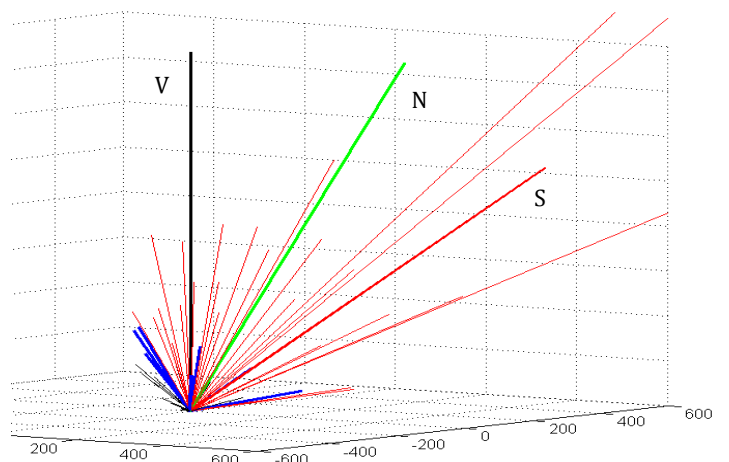


Figure 10. Visualisation of intensities as 3D vectors.

The resulting albedo (Fig. 11) shows the 'body colour' of the object without any specularly. The normals image, shown here in false colour, gives a good estimate of the surface normal at every pixel.



Figure 11. Computed albedo (left) and normals (right).

5 DETERMINING SPECULAR ANGLE

The general direction of the specular vector is clear from Fig. 10, where the intensities from the three lamps with the nearest directions are much greater than any others. The figure also indicates the coarseness inherent in sampling angular space with only 64 points over the hemisphere, where every adjacent pair of lamps is separated by an interval of 20–25 degrees.

The strategy adopted in this study is to calculate the ‘specular quotient’, i.e. the ratio between the actual intensity for each lamp direction and the intensity that would be produced by a perfect diffuser for the same direction. (This would be the ratio of black divided by magenta values for the pixel illustrated in Fig. 7.) The more glossy the surface, the greater the quotient value (Fig. 12 left). The resulting specular vectors (Fig. 12 right) have the same general appearance as the normals (Fig. 11) but are more chromatic because the specular gradients are greater with respect to the view vector.



Figure 12. (left) Ratio of specular/diffuse intensities; (right) Specular vectors.

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 35° around this direction, weighted by their quotient values.




Figure 13. Specular vs normal angles, classified by quotient value.

It is clear from plotting the specular vs normal angles for a random selection of 10,000 pixels (Fig. 13) that although the majority of pixels lie close to the line of slope 2 (i.e. specular angle = 2x normal angle), there is a considerable amount of scatter 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). Plotting the specular and normal angles along a horizontal section of the image (Fig. 14) illustrates their relationship, and also shows that the maximum normal angle that can be quantified by the photometric stereo technique is c.35°, with corresponding maximum specular angle of c.70°.

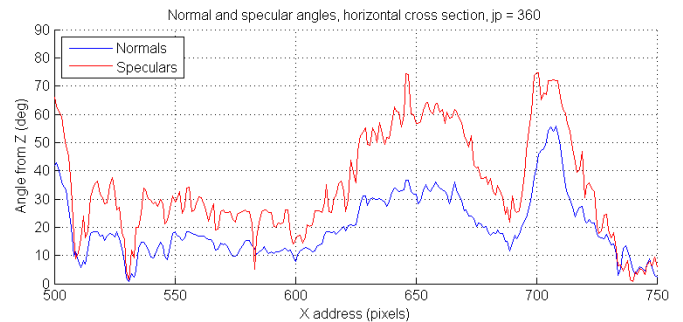


Figure 14. Specular and normal angles along a horizontal image section.

The specular colour is also calculated for each pixel using the same weighting coefficients as multipliers on the R,G,B values of the pixel in those images corresponding to the selected lamp vectors. This follows the dichromatic model of reflectance developed by Shafer [10]. The specular colour for glossy surfaces on the majority of materials is the colour of the illumination, but for metals it is the colour of the metal. The effect is particularly striking for gold [11].

6 MODELLING ANGULAR DISTRIBUTION

The aim is to model the luminance variation at each point on the object surface as a function of 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 illumination angles in between those of the lamps in the dome.

A complete model would accommodate the bidirectional reflectance distribution function (BRDF), with 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 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).

Plotting the specular quotients against radial angle from the specular peak (Fig. 15) shows the typical distribution for a glossy surface, with large values at small angles falling to a ‘knee’ between 20° and 30° and then a long tail out to 90°, asymptotic to unity.

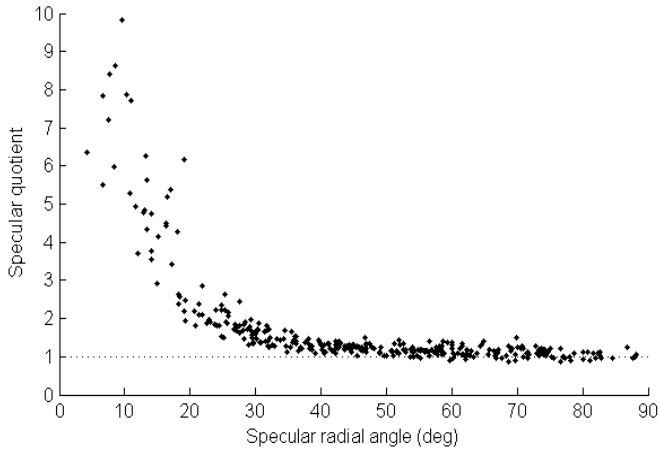


Figure 15. Distribution of specular quotients in a region of 3x3 pixels.

In the computer graphics community, motivated by the desire for realistic rendering from models, there have been many proposals for a specular function, based on a statistical distribution of the reflectance from micro-facets on the surface. The simplest function is a Gaussian $e^{-(\omega/a)^2}$ as formulated by Beckman [12] for surface scattering based on the Kirchhoff scalar diffraction theory. Blinn [13] preferred a function developed by Trowbridge and Reitz [14] based on $\cos^{-4}(\omega)$. Ward [15] chose a variant of the Gaussian based on $e^{-\tan^2(\omega)}$. Watson and Raven [16] described an empirical model called Mopaf, with a sum of Lorentzian functions to model BRDF, based on $1/(1 + (\omega/a)^2)$. More recently, Bagher *et al* [17] studied BRDF measurements of shiny materials and showed that they do not follow the Beckmann distribution. Instead the micro-facets are more likely to be aligned with the surface normal, resulting in a sharper peak. They also noted the broader flanks in real distributions and proposed a function based on $e^{-\omega/a}/\sqrt{\omega - a^2}$.

Plotting these functions together, with parameters adjusted to give the same width at half maximum amplitude (Fig. 16), shows that for small angles the Ward function is almost identical to a Gaussian and that it falls quickly to zero with little amplitude in the flanks beyond 5°. The Trowbridge & Reitz function is broader, extending to about 10°, and the Bagher function is broader still, extending to about 15°. The Lorentzian function has the broadest flanks, extending out to beyond 45°.




Figure 16. Five specular functions of angle.

The model adopted in this investigation to fit the specular peak is based on the Lorentzian function, because it naturally fits the observed shape and is mathematically convenient. During analysis of various materials it was found that the flank of the specular quotient distributions varies in slope according to the roughness of the surface, and the final version of the model therefore consists of a sum of two components, a Lorentzian peak and a linear flank, each with two parameters for amplitude and scaling:

$$f(\omega) = \frac{p_a}{1+(\omega/p_s)^2} + (f_a\omega + f_s) \quad (6)$$

where $\omega = \text{acosd}(\mathbf{L} \cdot \mathbf{S})$ is the angle in degrees between the lamp vector and specular vector. The fitted peak and flank components and the combined function are shown in Fig. 17.

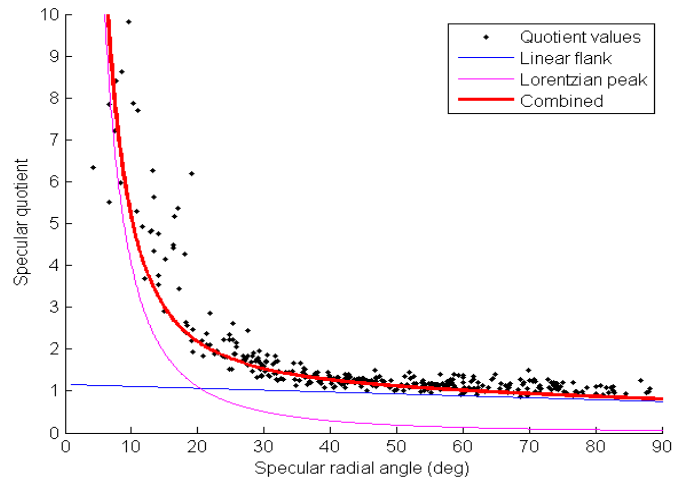


Figure 17. Lorentzian peak and linear flank functions fitted to distribution.

The complete model for a single light source takes the form of a sum of diffuse and specular terms:

$$I = \mathbf{L} \cdot \mathbf{N} A_{rgb} + (f(\omega) - 1) S_{rgb} A_m / S_m \quad (7)$$

where A_{rgb} is the albedo colour, S_{rgb} is the specular colour, A_m is the monochrome albedo (weighted sum of the R,G,B channels of A_{rgb}), and S_m is the monochrome specular intensity.



Figure 18. Image components for peak (above) and flank (below).

7 ASSESSMENT

A software viewer has been developed in Matlab, which enables direct comparison between any of the original images and the corresponding rendering methods. Fig. 19 shows six snips from the viewer screen, based on one of the set of 64 images taken in the dome, illuminated by lamp 35, in the third tier at an elevation of c.45°. The rendering with the Lorentzian function is a close simulation of the original image, and is notably better than the PTM or HSH renderings. The greater realism of the Lorentzian rendering is achieved not by changing the base colour of the diffuse component but by enhancing the local contrast to match the tonal modulation of the surface gloss.

REFERENCES

- [1] T. Malzbender, D. Gelb and H. Wolters, 'Polynomial Texture Maps', *Proc. ACM SIGGRAPH*, **28**, 519–528 (2001).
- [2] <http://www.hpl.hp.com/research/ptm/>
- [3] P. Gautron, J. Krivanek, S. Pattanaik, and K. Bouatouch, 'A Novel Hemispherical Basis for Accurate and Efficient Rendering', *Eurographics Symposium on Rendering Techniques*, 321–330 (2004).
- [4] <http://culturalheritageimaging.org/Technologies/RTI/>
- [5] G. Earl, K. Martinez, and T. Malzbender, 'Archaeological applications of polynomial texture mapping: analysis, conservation and representation', *J. Archaeological Science*, **37**(8), 2040–2050 (2010).
- [6] L. MacDonald and S. Robson, 'Polynomial texture mapping and 3D representation', *Proc. ISPRS Commission V Symp. 'Close Range Image Measurement Techniques'*, Newcastle (2010).
- [7] C. Julia, A. Sappa, F. Lumbreras, J. Serrat and A. López, 'Photometric stereo through an adapted alternation approach', *Proc. 15th IEEE International Conf. on Image Processing (ICIP)*, 1500–1503 (2008).
- [8] M. Chandraker, S. Agarwal and D. Kriegman, 'Shadow Cuts: Photometric stereo with shadows', *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 1–8 (2007).
- [9] M. Drew, Y. Hel-Or, T. Malzbender, and N. Hajari, 'Robust estimation of surface properties and interpolation of shadow/specularity components', *Image and Vision Computing*, **30**(4), 317–331 (2012).
- [10] S. Shafer, 'Using color to separate reflection components', *Color Research & Application* **10**(4) 210–218 (1985).
- [11] L. MacDonald, M.F. Guerra, R. Pillay, M. Hess, S. Quirke, S. Robson, A.H. Ahmadabadian, 'Practice-based comparison of imaging methods for visualization of toolmarks on an Egyptian Scarab', *Proc. ICISP, Cherbourg* (2014), to be published.
- [12] P. Beckmann and A. Spizzichino, *The Scattering of Electromagnetic Waves from Rough Surfaces*, Pergamon Press, NY, (1963).
- [13] J. Blinn, 'Models of light reflection for computer synthesized pictures', *ACM SIGGRAPH Computer Graphics*, **11**(2) 192–198 (1977).
- [14] T. Trowbridge and K. Reitz, 'Average irregularity representation of a roughened surface for ray reflection'. *J. Opt. Soc. Am.* **65**(5) 531–536 (1975).
- [15] G. Ward, 'Measuring and modeling anisotropic reflection', *ACM SIGGRAPH Computer Graphics* **26**(2) 265–272 (1992).
- [16] R. Watson and P. Raven, 'Comparison of measured BRDF data with parameterized reflectance models', *Proc. SPIE Conf. on Targets and Backgrounds VII: Characterization and Representation*, **4370** 159–168. (2001).
- [17] M. Bagher, C. Soler and N. Holzschuch. 'Accurate fitting of measured reflectances using a shifted gamma micro-facet distribution', *Computer Graphics Forum*, **31**(4), 1509–1518 (2012).



Figure 19. Photographic image (top left) of the Roman medallion illuminated by a single light source, and five renderings: Lorentzian (top centre); Lorentzian achromatic specular component (top right) Lambertian (lower left); PTM (lower centre); HSH (lower right).