

## CONSERVATION MONITORING OF A HERITAGE CEILING BY PHOTOMETRIC STEREO

Lindsay MacDonald<sup>a</sup>, Ian Gibb<sup>b</sup> and Stuart Robson<sup>a</sup>

<sup>a</sup> Photogrammetry, 3D Imaging and Metrology Research Centre, University College London WC1E 6BT

<sup>b</sup> Conservation Collection Care, Hampton Court Palace, Historic Royal Palaces, Surrey KT8 9AU  
ucfslwm@ucl.ac.uk

Commission V, WG V/2

**KEY WORDS:** Architecture, conservation, monitoring, camera, imaging, reconstruction.

### ABSTRACT

Photographic survey techniques were employed to monitor the condition of the paintwork on the ceiling of the Queen's Staircase at Hampton Court Palace. Illumination was provided by a flash, mounted on an 8 metre telescopic mast, raised to approximately 60 cm below the ceiling, with images were taken from a fixed camera position on the floor of the landing at a range of 6.5m. The photometric stereo method was applied to images from ten lamp positions, to calculate surface normals and a depth map. Cross-sections at the estimated surface resolution of 7.7 pixels/mm, achieved a depth (Z axis) resolution of approximately 100 microns.

### 1. CONSERVATION REQUIREMENTS

The Queen's Staircase at Hampton Court Palace, London, is part of a grand scheme of rebuilding commissioned by George II. The room is rectangular of dimensions 12m by 21m and has a grand staircase rising to a landing of 12m by 5.5m. The ceiling height is 12.5m above the ground floor at the centre and 8m above the principal floor landing. Daylight is admitted through three large windows and one smaller high window facing south.

The painted ceiling, constructed in lath and plaster, is part of the decorative scheme completed by William Kent in 1734. The painting is executed in oil bound pigment and depicts a marble dome supported by a cornice and figures, featuring a central garter star. Extensive use is made of *trompe l'oeil* techniques with coffered panels, mouldings and elevated height to enhance the grandeur. The design is centred on a Garter Star and supported by grisaille figures and elaborate ornament (Fig. 1).

In 2001, flaking was noted on the painted ceiling. The flaking was unlike that observed on other wallpaintings at the palace as areas of instability were spread over the entire surface of the ceiling and resulted in the detachment of large brittle flakes of diameter up to 20 mm. Treatment with a conventional material for wallpainting conservation failed to solve the problem. Since 2004 the approach of Historic Royal Palaces (calling on the expertise of a multidisciplinary team of specialists) has been one of cautious intervention, coupled with continued research and recording to understand the causes of deterioration and to devise a suitable long-term stabilisation strategy, including stabilisation of the air flow and temperature in the room.

Paint sampling established the sequence of preparation layers over the previous plaster finish as follows: a red pigmented organic medium sealant, followed by three ground layers (off-white, cream, then grey). Crucially, all samples feature a loosely bound or unbound chalk layer beneath the red sealant and over the original plaster (Davies, 2006). This appears to be the interface where the delamination of the paint takes place; even during sampling the sample flake would frequently come apart from the plaster substrate at this layer.

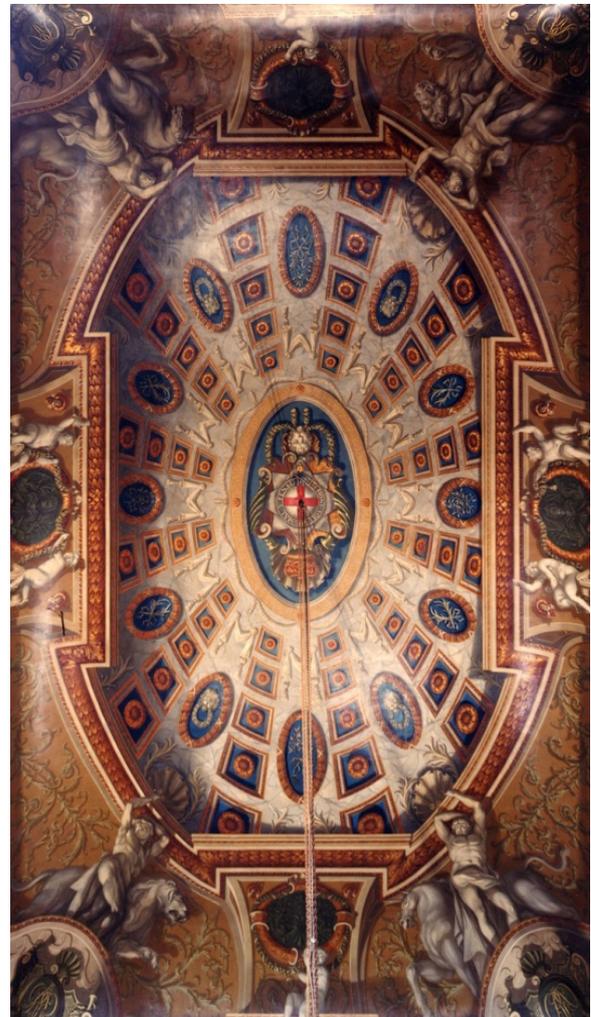


Figure 1. Rectified view of the painted ceiling over the Queen's Staircase at Hampton Court Palace.

Whilst environmental improvements can be assessed through the existing sensor network, monitoring of paint condition has always been more difficult because close examination of the ceiling causes major disruption to visitors. Previous surveys and treatment trials were carried out using either a cherry-picker (unstable working position and disruptive) or a fixed scaffold platform under the whole ceiling (expensive and very disruptive). A detailed survey from a scaffold platform in January 2010 produced rectified photography of the whole ceiling, and mapped the regions of greatest paint loss. The worst affected areas were found to be in the centre and above the landing. Possible agents of the flaking process include:

- Flooding from the apartment above in 1970s;
- Thermal convection caused by the central chandelier and a ventilated heater;
- Turbulent airflow from the double doors on the ground floor and at either end of landing;
- Humidity produced by visitors and changing atmospheric conditions;
- Vibration in structure (considered less likely).

The requirement of the Conservation Team at HCP is to monitor selected areas of the ceiling over a period of time to look for evidence of surface deformation, particularly bubbling or the formation of fine cracks that might precede the detachment of flakes of paint. The requested resolution of fine detail is 0.1mm (i.e. 100 micron), which would require 20 points/mm image resolution. This is beyond the capabilities of any laser scanner operating from the landing level – typical resolution at this range would be 3-5 mm.

The aim of the investigation described in this paper is therefore to develop a method of imaging the paint surface which can detect the flaking process as it progresses and also be deployed at long range, preferably from floor level with minimal effect on public access and palace operations. By repeating the imaging procedure at regular intervals comparison of images will provide information on rates of change, such as crack growth, surface deformation and flake formation.

## 2. PHOTOGRAPHIC TRIALS

Photographic trials were undertaken for the feasibility of capturing suitable images by digital photography from the landing level. A previous study by Liang *et al* (2008) with a camera mounted on a telescope was inconclusive. In this case a Nikon D200 camera with a 55-300 mm zoom lens was set up on a tripod approximately 2m above the landing floor level, at a distance of 6.5m below the ceiling (Fig. 2). Images were captured directly from the camera through a USB interface into a laptop PC. The ambient illumination of the ceiling was a combination of the bluish daylight coming through the windows and the yellowish light from the tungsten bulbs in the chandelier. Additional illumination was provided by a hand-held flash, synchronised by a radio link. Exposure times ranged from 1 sec to 4 sec at ISO 400 speed, changing the relative contributions of the ambient and flash sources. At maximum magnification of the lens (zoom set to 300 mm focal length) the field of view is 5°20'. The half-angle of 2°40' at a distance of 6m would represent a distance on ceiling of 28 cm, so the full image of 3:2 aspect ratio would cover an area of 56x38 cm. For the image size of 3872x2592 pixels this represents a surface resolution of 6.9 pixels/mm (approx. 150 microns per pixel).

Further investigation was made into the effect of raking light on the surface of the ceiling, and whether quantitative information could be derived from images captured by this means. The camera was again mounted on a tripod, at a height of 2m above



Figure 2. Upward view of flash mounted on top of 8m mast.

the landing floor. The flash light was mounted on an 8m telescopic mast, enabling it to be raised to less than 1 metre below the ceiling (Fig. 2). The mast consists of six concentric 18 gauge aluminium tubes, ranging from 2" dia (outermost) to 1" dia (innermost), each of length 1.5 metres, supported by a tripod base. When fully raised, the top of the mast is 7.75 metres above floor level. The flash unit and radio receiver were both mounted on a 1/4" dia screw, bolted onto the top section. Their combined weight, as payload, was 620 grams. The height of the mounted flash above the centre of the pole was 10 cm, hence 7.85 m (= 25'9") above the floor.

The previous rectified photographic survey of the ceiling in 2010 identified many places subject to damage. An area in Tile 67 was chosen, depicting two branches of laurel tied by a ribbon within an oval gilded laurel garland (Fig. 3). The target area was located on the longitudinal axis of the ceiling above the centre of the landing rail. The camera was placed directly below, pointing vertically upwards. The flash light when fully raised was at approximately the same level as the top of the cornice. Images were taken from the fixed camera position with the zoom lens set at 300mm, while the flash light was moved successively to eleven locations around the floor of the landing, with five locations in each of two concentric semicircles plus one at the centre (close to the camera axis).



Figure 3. (left) Target area on ceiling; (right) locations of damage marked in red from previous survey in 2010.

### 3. REGISTRATION OF IMAGES

The sequence of 11 images of the oval panel on the ceiling took more than an hour to capture. Although the tripod holding the camera remained in the same position, it took an average of 7 minutes to set up each successive position of the telescopic mast holding the flash light. The high magnification of the image through the 300mm lens made the process very sensitive to small perturbations of the camera and its tripod, and resulted in misalignment of the images (Fig. 4).

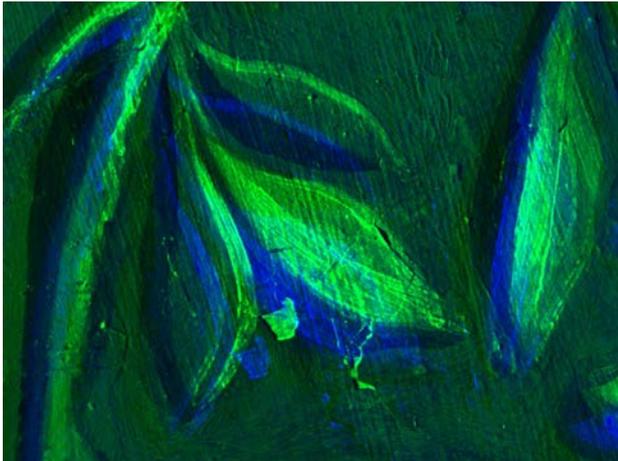


Figure 4. The effect of the misalignment of images is clear in this false-coloured composite of a detail of image 1 (green) and image 5 (blue). The offset is 45 pixels in X and 73 pixels in Y.

To use the images subsequently for determining surface normals by photometric stereo, it is essential that all images are in accurate pixel registration. In the general case of the camera moving to different locations within the room, it would require a full projective transform to register each image to a common geometric reference frame. In this case, however, the movement of the camera was small enough that the effect on the image appears as a simple two-dimensional translation.

The method for registration used 2D correlation, provided by the Matlab function `corr2`, taking a small square of size 101x101 pixels of the reference image, and (in effect) sliding it around an enclosing region of 201x201 pixels from the test image to find the best fit (Fig. 5). At each successive pixel coordinate position the correlation coefficient  $r$  was calculated:

$$r = \frac{\sum_m \sum_n (A_{mn} - \bar{A})(B_{mn} - \bar{B})}{\sqrt{\left( \sum_m \sum_n (A_{mn} - \bar{A})^2 \right) \left( \sum_m \sum_n (B_{mn} - \bar{B})^2 \right)}} \quad (1)$$

This formula has the advantage of being independent of the brightness and contrast of both images  $A$  and  $B$ , because it takes differences from the respective means and normalises the intensity values. The images were first converted from the Nikon raw (nef) format to 16-bit linear (tiff) format files, using the utility DCRAW and the green channel extracted. Thus all registration calculations were performed on 16-bit linear monochrome data. To speed up the computation the correlation procedure was done in two stages. First a 'coarse' fit was determined by calculating the correlation at intervals of four pixels in both X and Y over a 201x201 pixel area, using filtered versions of both images (a 5x5 pixel box mean) to reduce the influence of noise. In the second stage the fit was determined

for every pixel of the unfiltered images within a local 11x11 pixel area around the maximum position found in the first stage.

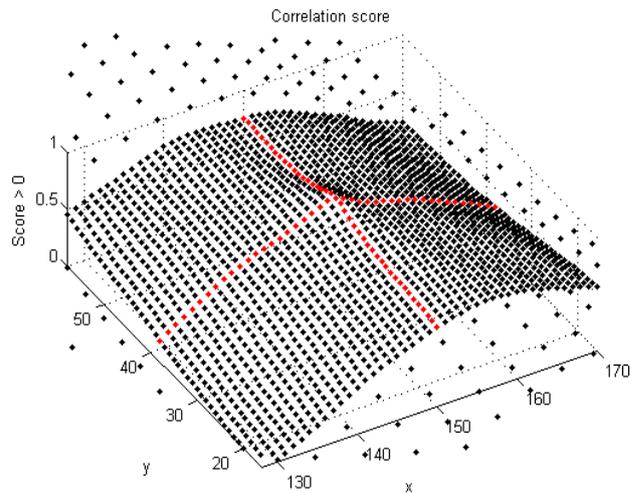


Figure 5. Surface of correlation scores for offsets of  $\pm 21$  pixels from maximum. Red lines indicate maximum row and column.

To make the procedure more robust, the calculation was repeated at 150 locations across the image pair, on a 15x10 grid with centres separated by 220 pixels in each dimension. The corresponding translation vectors between images 1 and 5 are plotted in Fig. 6. It is evident that the great majority are parallel and of equal length. This indicates a simple translation – if there were any rotation the vector lengths would vary systematically across the image. Some vectors are clearly different, however, because the correlation process may go wrong if there is insufficient image detail or excessive noise (particularly in the dark regions) or if the shadow formation from the different illumination directions is too dominant.



Figure 6. Monochrome of image 1 with translation vectors for image 5, calculated at 150 positions in a 15x10 grid.

A histogram of the 150 values of X and Y offsets for image pair 1 and 5 (Fig. 7) shows a strong central tendency, with approximately 80% of all values within a range of 2 pixel units. The median of the distribution therefore gives a very good estimate of the offset to the nearest pixel.

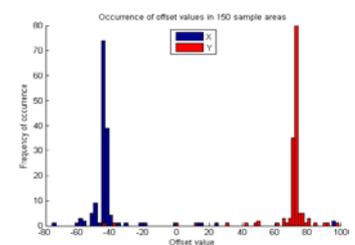


Figure 7. Histogram of X and Y offsets at 150 locations.

The offsets were calculated for all images with respect to image 7 as reference, with results in Table 1. The mean and standard deviation were calculated from the 120 central values between the 10<sup>th</sup> and 90<sup>th</sup> percentiles, which effectively removed the outliers. The only large spread remaining is for image 11, which was illuminated along the imaging axis, and is dominated by specular highlights. Yet even in this case the correlation procedure with median found the correct offset value.

Image	X (med)	Y (med)	X (mean)	Y (mean)	X (stdev)	Y (stdev)
1	48	-64	47.81	-63.83	0.76	0.68
2	48	-63	47.89	-63.08	0.67	0.43
3	48	-63	47.58	-62.71	0.50	0.47
4	48	-62	47.92	-61.67	0.54	0.47
5	3	9	3.00	9.00	0.00	0.00
6	1	0	0.79	0.00	0.41	0.00
7	0	0	0.00	0.00	0.00	0.00
8	28	45	28.18	44.50	0.39	0.50
9	20	36	20.02	36.27	0.65	0.44
10	19	36	19.13	36.29	0.68	0.54
11	21	37	19.25	28.03	22.22	25.55

Table 1. Median, mean and standard deviation of each image with respect to image 7.

The translation offsets were validated by processing all images and comparing them in sequence, displayed full screen with an image viewer. There was no visible displacement between successive images. The accuracy is also clear from a false-colour composite (Fig. 8).

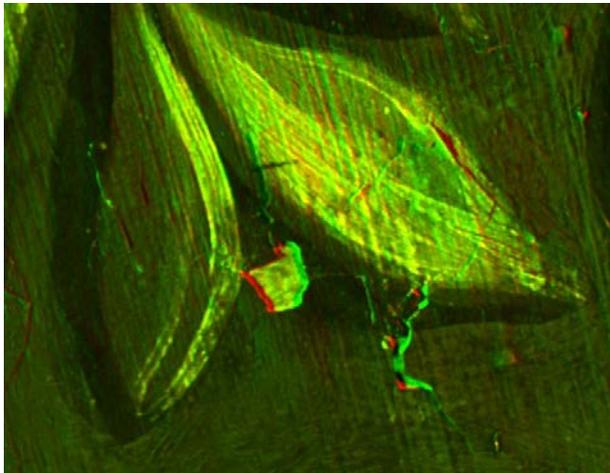


Figure 8. False-colour composite of same pair of images as in Fig. 4, after alignment. The intensity differences caused by illumination from the opposite directions are apparent.

#### 4. SURFACE FROM PHOTOMETRIC STEREO

For a Lambertian surface, from which the incident light is scattered equally in all directions, the luminance of the reflected light is given by the vector dot product:

$$L_r = \rho \mathbf{L}_i \cdot \mathbf{n} = \rho |L_i| \cos \alpha \quad (2)$$

where  $L_r$  is the luminance of the diffusely reflected light (with no angular dependence),  $\rho$  is the maximum surface reflectance (or albedo),  $\mathbf{L}_i$  is the incident light vector,  $\mathbf{n}$  is the unit normal of the surface, and  $\alpha$  is the angle between  $\mathbf{L}_i$  and  $\mathbf{n}$ .

Because the normal vector has three components, at least three equations are needed to solve the system. This can be achieved by illuminating the surface for successive images from three

different lighting directions with incident light vectors  $\mathbf{L}^1, \mathbf{L}^2, \mathbf{L}^3$ . This system can be written as:

$$I^k = \rho \mathbf{L}^k \cdot \mathbf{n} \quad (3)$$

where  $k=1,2,3$  are the three lighting directions. The three observed intensity values  $I^k$  of the reflected light can be stacked to form the 3x1 intensity vector  $\mathbf{I} = (I^1, I^2, I^3)$ , and the incident light vectors can be stacked row-wise giving the 3x3 lighting matrix  $\mathbf{L} = (\mathbf{L}^1, \mathbf{L}^2, \mathbf{L}^3)$ . Then Eq. (3) can be rewritten as:

$$\mathbf{I} = \rho \mathbf{L} \cdot \mathbf{n} \quad (4)$$

Provided that the lighting directions  $\mathbf{L}^k$  are not coplanar, the matrix  $\mathbf{L}$  can be inverted, giving:

$$\rho \mathbf{n} = \mathbf{L}^{-1} \mathbf{I} \quad (5)$$

Since the normal vector  $\mathbf{n}$  is unitary, both its direction and albedo  $\rho$  (modulus) can be recovered.

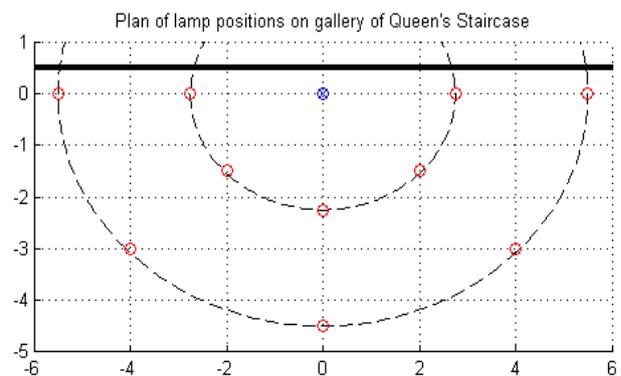


Figure 9. Ten coordinate positions of the flash lamp (red) relative to the camera (blue). All units are in metres.

The successive positions of the flash light on its mast over the floor area of the landing of the Queen's Staircase were plotted on Cartesian coordinates, with the camera position at the origin (Fig. 9). They fall approximately onto two concentric ellipses. The lamp was at a constant distance of 62 cm below the ceiling. From these dimensions the cosines with respect to the three axes were calculated (Table 2). The angle of incidence of the light onto the surface was in the range 6°–8° for lamp positions in the outer ellipse and 13°–15° in the inner.

Pos	X	Y	Z	R	theta	phi
1	-5.5	0.0	0.62	5.53	180.0	6.4
2	-4.0	-3.0	0.62	5.04	216.9	7.1
3	0.0	-4.5	0.62	4.54	270.0	7.8
4	4.0	-3.0	0.62	5.04	323.1	7.1
5	5.5	0.0	0.62	5.53	0.0	6.4
6	2.8	0.0	0.62	2.82	0.0	12.7
7	2.0	-1.5	0.62	2.58	323.1	13.9
8	0.0	-2.3	0.62	2.33	270.0	15.4
9	-2.0	-1.5	0.62	2.58	216.9	13.9
10	-2.8	0.0	0.62	2.82	180.0	12.7

Table 2. Coordinates and angles for ten positions of flash light.

The green channel was extracted from each of the registered images. All monochrome images (16-bit linear) were then scaled by finding the 99.5<sup>th</sup> percentile value from the histogram and setting it to 90% of white (i.e. 0.9\*65535 = 58982). This produced a more homogeneous image set, compensating for the different illuminance levels (because of differing lamp distances), different camera settings ( $f$ /stop and ISO and exposure time) and non-Lambertian (semi-gloss) surface.



Figure 10. (left) Area from lower centre of image with red square showing the detail of 300x300 pixels (right).

Images from ten lamp positions were used in the calculation. The eleventh image, with the flash adjacent to the camera axis, was not used because the high specular reflectance makes it very different from the other ten. Although in theory there are  ${}^{10}C_3 = (10*9*8)/(3*2*1) = 120$  possible combinations of three lamps out of ten, in practice some of these cannot be used because two or three of the lamps lie in the same plane, causing the matrix  $L$  in Eq. (4) to be singular. Applying the criterion that the condition number  $\text{cond}(L) < 10$ , reduced the number of valid three-lamp combinations to 37.

Normals for X,Y,Z were calculated using Eq. (5) for each of the 37 selected three-lamp combinations for every pixel position. This produced normal values with a considerable spread (Fig. 11), because the differentiation process is very subject to noise. The median of the distribution was taken independently for each of X,Y,Z following the method of an earlier study (MacDonald and Robson, 2010) in which it was found that the median gave a good estimate of the true value of the normal. The method has been applied successfully to analysis of ancient Egyptian clay tablets, with both positive relief (raised moulded characters) and negative relief (incised characters).

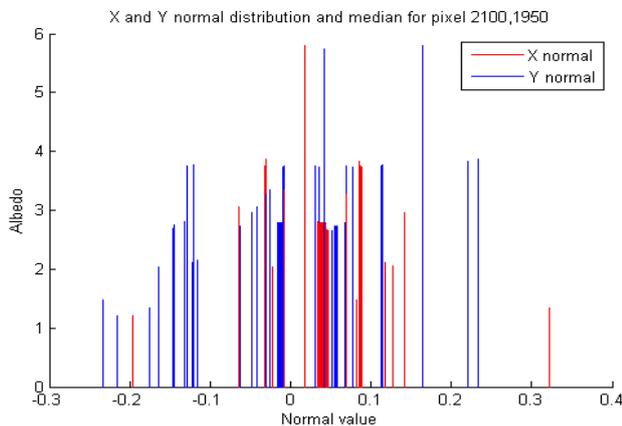
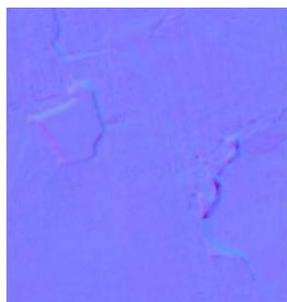


Figure 11. Distribution of values for the X normal (red) and Y normal (blue) for a single pixel in the image.

A false-colour image of the normals of the detail in Fig. 10 shows that the depth information is rather subtle, but the outlines of the missing paint flake on the left and the crack on the right are clearly visible.

Figure 12. Representation of normals, where R,G,B channels depict the X, Y and Z normals.



Plotting the normals in horizontal and vertical cross-sections shows how small is the modulation of the normal direction. It also reveals that the camera axis was not quite perpendicular to the plane of the ceiling: the X normal rises steadily across the horizontal section, while the Y normal rises down the vertical section. Zooming in on the normals within the detail of Fig. 10 shows that they do detect the edges of the features (Fig. 13), particularly the left and right edges of the paint flake (a positive normal indicates a negative slope).

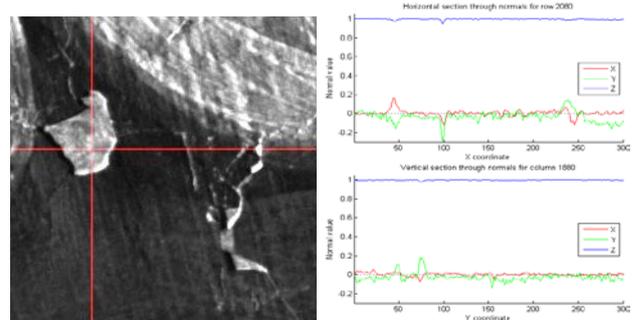


Figure 13. (left) Monochrome image of 300x300 pixel detail (same as Fig. 10), with locations of horizontal and vertical cross-sections shown by red lines; (right) normals along corresponding sections.

A depth map was derived from the surface normal, solving the Poisson equation by two-dimensional integration, using the implementation by Kovesi (2005) of the Frankot and Chellappa (1988) algorithm. The surface gradients are projected onto a set of integrable basis functions, in this case the Fourier basis functions. This algorithm has the advantages of a simple expression in Matlab and high robustness to noise.

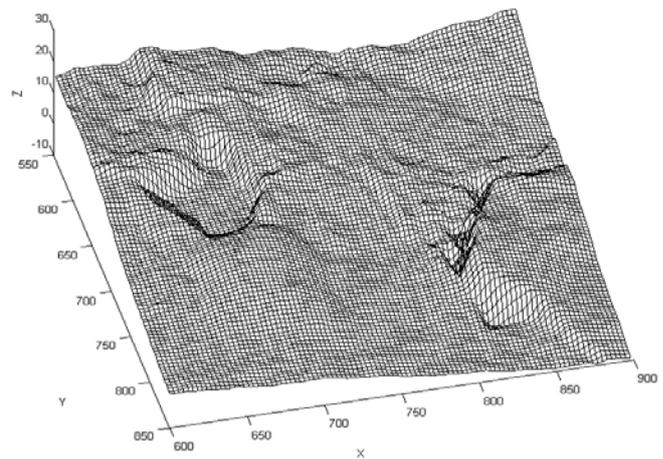


Figure 14. Reconstructed surface of 300x300 pixel area.

Figure 14 shows an oblique view of the reconstructed surface of the image area in Fig. 13, in which the region corresponding to the missing flake appears clearly as a pit and the raised flake as a ridge. Taking cross-sections of the reconstructed surface along the same lines as in Fig. 13 yields two elevations of the surface (Fig. 15). The Z axis has the same pixel dimensions as the X and Y axes. The downward overall slope of both sections is the consequence of the non-perpendicularity of the camera to the surface (assumed to be planar), resulting in an offset from zero of the mean of the normals. The rising slope of the normals, for example in the X normals across the image width, indicates a change from negative to positive slope. This could be a genuine curvature in the ceiling but is more likely to be an artefact of

either the image formation process (e.g. barrel distortion in the lens) or the geometry of the normal calculation (angle of incidence of light and also angle of reflected ray to camera are not parallel but change continuously across the image field).

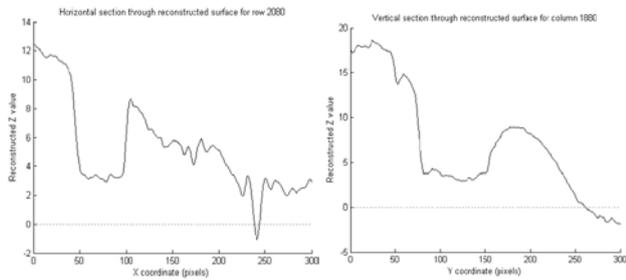


Figure 15. (left) Horizontal and (right) vertical cross-sections through the reconstructed surface.

The normals were corrected by finding a line of best fit (Matlab function `polyfit`) through the normal X data for every row and taking the median of the slope and offset values of all lines within the image region. The value predicted by the regression line across the row was then subtracted from the X normal value, giving the residual for each coordinate position. A similar correction was applied to Y normals down each column. This procedure effectively removes the macro geometry (orientation and curvature of the ceiling relative to the camera) and retains the micro geometry (local cracks, bulges and flakes).

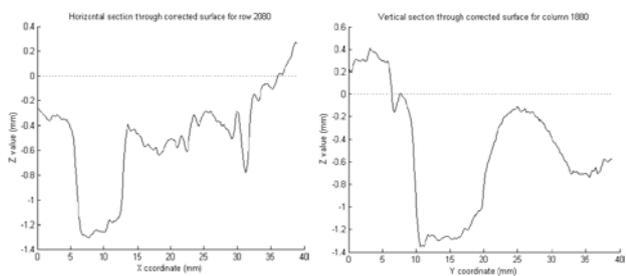


Figure 16. (left) Horizontal and (right) vertical cross-sections after correction of the normals.

After surface reconstruction of the corrected normal values, the cross-sections of Fig. 16 were obtained. The axes in these figures have been converted to physical units of mm, using the estimated surface resolution of 7.7 pixels/mm. Note that the depth (Z axis) resolution achieved by the photometric stereo technique is finer than 0.2 mm (i.e. 200 microns). With a better lens (400mm and higher MTF), combined with more careful placement of the flash light in successive positions, a depth resolution of 100 microns could readily be achieved. This is an order of magnitude better than most 3D laser scanners, and compares very favourably with the Z-axis resolution achieved by the best (and most expensive) 3D scanners, operating at a range of 6 metres (Table 3).

Scanner model	Accuracy at 6m	Noise at 6m
FARO Focus 3D Laser Scanner	±2 mm	500 micron
Leica HDS 7000 Laser Scanner	100 micron	400 micron
Nikon Laser Radar MV330	63 micron	25 micron

Table 3. Accuracy and noise of three commercial 3D laser scanners (data from manufacturers' product sheets).

## 5. CONCLUSIONS

This study has demonstrated the feasibility of using the photometric stereo technique with a set of images captured by a fixed camera and movable flash light on a mast. Local features on the ceiling could be identified, and a good estimate of their size and depth obtained, including cross-sections through the 3D surface topography. By repeating the procedure at regular time intervals (e.g. monthly) it would certainly be possible to visualise and perform quantitative analysis of changes in the condition of the surface. The advantages of the technique are that it uses relatively cheap equipment and achieves a remarkably high depth resolution. The disadvantage is that it is rather laborious to move the lighting mast to a number of successive lighting positions for each location being inspected, and if there were many such locations the time and effort required might be prohibitive.

A modified lighting apparatus and technique may be required for locations on the ceiling above the main floor area of the Queen's Staircase (i.e. not above the landing), where the height is greater than 12m. The ideal system would have a motorised platform enabling a flash light and/or camera and/or 3D scanner to be raised and orientated close to the ceiling under remote control by the operator at ground level, with a live view through the camera to facilitate location of the areas to be inspected. Such a system was developed in the EC-FP5-IST project VITRA, with a telescopic mast raised by pneumatic air pressure to heights of up to 15m above floor level (MacDonald, 2006). It was tested for conservation inspection of corroded metal panels on the ceiling at the monument Walhalla in Bavaria. This apparatus would certainly have provided superior capability for 2D imaging of the ceiling at Hampton Court, but because of its size and weight would have been more difficult to bring onto the site and to deploy in a public area. Future research will need to be directed toward new ways of achieving high resolution of a 3D surface with light-weight and low-cost components.

## REFERENCES

- Davies J (2006) Microscopy and archival research: interpreting results within the context of historical records and traditional practice, *InFocus*, Issue 3, pp. 32-45, Royal Microscopical Society, Oxford.
- Frankot RT and Chellappa R (1988) A Method for Enforcing Integrability in Shape from Shading, *IEEE Patt. Anal. Mach. Intell.*, **10**(4) 439-451.
- Kovesi P (2005) Shapelets Correlated with Surface Normals Produce Surfaces, *Proc. Tenth IEEE Intl Conf on Computer Vision (ICCV'05)*, Vol. 2, 994-1001. Software library from: <http://www.csse.uwa.edu.au/~pk/Research/MatlabFns/>
- Liang H, Keita K, Vajzovic T, Julien-Lees S, Spooner J and Zhang Q (2008) PRISMS: remote high resolution in situ multispectral imaging of wall paintings, *Intl. Council of Museums, Committee for Conservation (ICOM-CC) Triennial Conf.*, Vol. I, pp. 353-358.
- MacDonald LW (2006) A Robotic System for Digital Photography, *Proc. SPIE-IS&T Electronic Imaging Conf. on Digital Photography II*, SPIE Vol. 6069.
- MacDonald LW and Robson S (2010) Polynomial texture mapping and 3D representation, *Proc. ISPRS Commission V Symp. 'Close Range Image Measurement Techniques'*, Newcastle.