Photometric Stereo for Archeological Inscriptions (sketches_0338)

Per Einarsson, Tim Hawkins, Paul Debevec*

University of Southern California Institute for Creative Technologies

We describe a low-cost system for acquiring high-resolution geometry and reflectance properties using photometric stereo. This technique was used to model ancient inscriptions on the Parthenon in Athens, Greece and proved to be fast, simple and robust.

The principle of photometric stereo [Woodham 1980] is to take three or more images of an object from the same camera view, each with lighting from a different, known direction. Based on the brightness values in the different images, one can estimate a surface normal and a reflectance value by solving the following equation at each pixel:

$$L_i = E_i \frac{\rho_d}{\pi} \mathbf{n} \cdot \mathbf{l}_i \tag{1}$$

where L_i is the pixel value in image i, E_i is the radiant intensity, ρ_d/π is the BRDF, $\mathbf{n}$ is the surface normal, and $\mathbf{l}_i$ is the vector towards the light source.

A challenge with photometric stereo is the estimation of the light source position and of the radiant intensity at each pixel, which varies from pixel to pixel due to source anisotropy and changing distance from the source to the surface point. Hence, the technique has mostly been used in controlled laboratory settings. Recently [Malzbender et al. 2001; Rushmeier and Bernardini 1999] have showed new applications for photometric data. We have constructed a low-cost device suitable for use in the field, that allows photometric data to be captured quickly using only a digital camera with a remotely mounted flash. This enables us to easily reconstruct geometry and reflectance properties.

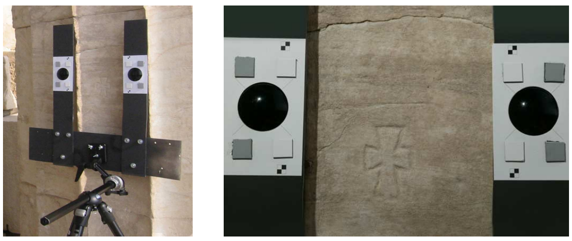

(a) (b)

Figure 1: (a) The adjustable light-capturing frame includes fiducial markers, reflectance standards and two glossy black spheres. (b) One of the input image for an inscription on a column of the Parthenon.

The light-capturing frame (Fig. 1) consists of fiducials and MacBeth color checker chart samples, from which the camera's position and the incident radiant intensity can be estimated, as well as two glossy black spheres used to indicate the position of the light source. To reconstruct each object, we re-adjusted the size of the frame and placed it around the geometry. We took approximately ten images, each lit with a remotely mounted camera flash at a different position pointed towards the center of the frame. For each set of images, we also took an additional image without the flash, which was subtracted from the other images to remove any ambient lighting. The full process including setup took approximately 20 minutes for each inscription.

To compute a surface normal map, we first determine the light vector $\boldsymbol{l}$ by shooting rays from the camera's center toward the observed reflections of the light source in the two spheres, reflecting these rays off the spheres, and intersecting the two resultant rays. The radiant intensity E is estimated at every pixel by interpolating the pixel values of the reflectance standards and dividing by the cosine of the angle to the light source position. We then form an overdetermined system of equations (1) to solve for the surface normal $\mathbf{n}$ using SVD. Assuming a Lambertian BRDF, we treat ρ_d as a constant and discard it since the result is normalized. Once the normal map has been estimated, we compute a 3D mesh using coarse-to-fine gradient descent optimization to compute depth values for all pixels that are consistent with the recovered normals. To remove any low-frequency warping introduced in the geometric reconstruction, we apply a high-pass filter to the resulting geometry.

In addition to geometry, we also compute reflectance properties for each object. Using the recovered normals, we estimate a diffuse albedo value, ρ_d, at every pixel by solving the following equation:

$$\rho_d = \sum_{i=1}^{n} \frac{\alpha_i L_i}{E_i \mathbf{n} \cdot \mathbf{l}_i} \Big/ \sum_{i=1}^{n} \alpha_i \tag{2}$$

where n is the number of input images and α_i is a weighting factor for image i. Because we expect the diffuse albedo to be most robustly estimated from brighter pixel values, we let $\alpha_i = L_i$.

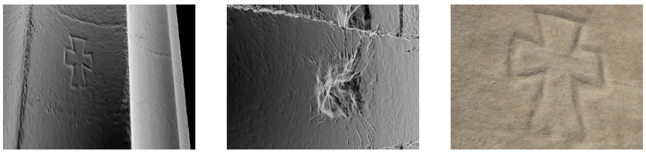

(a) (b) (c)

Figure 2: (a) and (b) Reconstructed local geometry embossed on the global geometry. (c) Rendering using estimated albedo map.

We applied this technique on four inscriptions and a cannonball impact on the Parthenon and used the results for some of the shots in the SIGGRAPH 2004 Electronic Theater film "The Parthenon" (Fig. 2). While the global geometry was known, (columns and flat walls), we found that our technique proved to be successful in reconstructing detailed local geometry, producing sub-millimeter resolution geometry with no observable noise. Moreover, it is low-cost, fast and suitable for use in the field.

References

MALZBENDER, T., GELB, D., AND WOLTERS, H. 2001. Polynomial texture maps. *Proceedings of SIGGRAPH 2001* (August), 519–528.

RUSHMEIER, H., AND BERNARDINI, F. 1999. Computing consistent normals and colors from photometric data. In *Proc. of the Second Intl. Conf. on 3D Digital Imaging and Modeling - 3DIM'99*, 99–108.

WOODHAM, R. J. 1980. Photometric method for determining surface orientation from multiple images. *Optical Engineering 19*, 1, 139–144.

*e-mail:einarsson, timh, debevec@ict.usc.edu