

Computing Consistent Normals and Colors from Photometric Data

H. Rushmeier and F. Bernardini*

IBM Thomas J. Watson Research Center
30 Saw Mill River Road
Hawthorne, NY 10532

Abstract

We present a method for computing normals and colors from multiple sets of photometric data, that are consistent with each other and an underlying lower resolution mesh. Normals are computed by locally adjusting the light source intensities using data from the underlying mesh. Colors are derived from the photometric calculations, and are adjusted by a global color registration analogous to global geometric registration.

1 Introduction

We consider the problem of scanning objects that are large relative to the smallest geometric level of detail to be represented. Three dimensional models of these objects are obtained by combining the results of hundreds of individual scans. We have developed a hybrid multiview/photometric method for scanning such objects. A multiview light striping system is used to obtain a base geometric model. A photometric system is used to obtain normals and colors at a higher spatial resolution. In this paper we present a solution for computing these colors and normals so that the results are consistent over the object, and with the underlying base geometry.

The motivation for this work is a project to acquire a three dimensional model of Michelangelo's Florentine Pietà [1]. The model is being used by art historian Jack Wasserman as part of a comprehensive study of the piece. The requirements are to obtain a model that can be manipulated to study overall problems of composition and form, and can also be examined in close detail to study tool marks and fine repairs. While there are not any substantial variations in hue on the piece, there are many variations in surface reflectance as a result of repairs, applied coatings, and wear on the sculpture. The subtle variations in the tone and color do not represent the artist's original intent, but are useful to document for studying the treatment of the piece since it was originally created in the sixteenth century.

Our project goals also included developing a relatively inexpensive scanning system in a short (i.e. months) time frame. Rather than have a custom laser scanning system fabricated, we developed a system using readily available cameras and lights. Because measurements needed to be made in many short sessions in the museum environment, our system had to rely as little as possible on high precision positioning and control of electrical power levels. Work in a similar environment has been described in [3].

Our system consisted of the Virtuoso multiview shape camera from Visual Interface [11], augmented with an in house lighting system consisting of five halogen lights (see Figure 1). The system, described in more detail in [9], allows the capture of five images with each of the five lights turned on in sequence, registered to the geometry acquired by the Virtuoso shape camera. We obtained several hundred overlapping meshes covering the sculpture. Each geometric mesh has a resolution of approximately 2mm, and an estimated sub-millimeter accuracy. A set of normals and colors are computed from the five photometric images at a resolution of between 0.5 and 1 mm. An example of the set of data acquired from a single camera pose is illustrated in Figure 1.

Variations in lighting and positioning result in slight but discernible variations in the photometric results from mesh to mesh. In the rest of this paper we discuss how we correct these variations. The normals maps are made consistent with the underlying mesh by locally adjusting the light source intensities used in the calculations, using data from the underlying mesh. By using a consistent global underlying mesh computed by registering and remeshing the underlying base meshes, the normals maps are consistent with one another, as well as with the global mesh. We also compute corrected red, green and blue maps from the photometric data. Chromatic variations from map to map are corrected with a global color registration analogous to the global geometric registration used to obtain the global base mesh.

*{holly,fausto}@watson.ibm.com




Figure 1: (a) The scanning system used for data acquisition. The Virtuoso multiview shape camera is augmented with a custom-made lighting system. Five color images, registered to the captured geometry, are taken with the five light sources turned on in sequence. (b) Triangle mesh computed by the Virtuoso software from the six striped images. (c) Four of the five photometric color pictures captured by our system.




Figure 2: Geometry of light reflectance used in Eq. 1

2 Photometric Stereo – Calculations and Problems

The photometric stereo method for computing three dimensional shapes was originally developed by Woodham [7]. The essence of the method is to take a series of images from the same camera view, each with lighting from a different, known, direction. Assuming an approximately Lambertian bidirectional reflectance distribution function (BRDF) and identical light sources, the relative reflectance and normal at each pixel location can be computed from three images. Specifically, the following set of equations are solved:

$$\frac{\rho_p}{\pi} L_o \Delta\omega \begin{bmatrix} l_{1,x} & l_{2,x} & l_{3,x} \\ l_{1,y} & l_{2,y} & l_{3,y} \\ l_{1,z} & l_{2,z} & l_{3,z} \end{bmatrix} \begin{bmatrix} n_{p,x} \\ n_{p,y} \\ n_{p,z} \end{bmatrix} = \begin{bmatrix} \alpha L_{r,1,p} \\ \alpha L_{r,2,p} \\ \alpha L_{r,3,p} \end{bmatrix} \quad (1)$$

where ρ_p/π is the BRDF at pixel p , L_o is the light source radiance, $L_{r,i,p}$ is the radiance reflected from the point visible through pixel p in image i , $\Delta\omega$ is the solid angle subtended by the light source, \mathbf{n}_p is the surface normal at p and \mathbf{l}_i is the direction vector to the (infinitely distant) i -th light source (see Figure 2). The constant α accounts for the camera scaling of the reflected radiance to a value from 0 to 255, after the images are adjusted if necessary for gamma values other than 1. The equations can be solved directly for $(\rho_p L_o \Delta\omega / \alpha \pi) \mathbf{n}_p$. Since the magnitude of \mathbf{n}_p is one, from this result we can obtain the normal vector \mathbf{n}_p and a value $\rho_{rel,p} = \rho_p L_o \Delta\omega / \alpha \pi$, that is the reflectance at pixel p relative to the other pixels in the image.

Two difficulties with this approach are the presence of a highly specular component in the BRDF of the surface being measured, and shadows. Following the approach used by other researchers (E.g. see [6]) we obtain extra images from additional light source positions. For each pixel location, pixels with very high values are not used to exclude specular reflections, and pixels with low values are not used to exclude shadows. We use relative value only rather than color to distinguish possible specular reflection [10], since there are spectrally flat regions in the statue for which such methods would fail. Pixel locations

where fewer than three images have pixel values that fall in the valid range are not used in the computations that follow.

In our system we used a system of five lights, i.e. two redundant reflectance measurements per pixel location. This number was used because it allowed a set of lights with a significant physical separation while keeping to an overall manageable size for mounting on a tripod.

In a controlled environment, the term L_o can be made the same for each source for all points on the scanned target by using controlled, identical, isotropic light sources. The term $\Delta\omega$ can be made constant for all points and all sources by placing the sources at a distance from the target that is large relative to the sizes of the light sources and the target.

In an inexpensive system used in an environment such as the museum, any number of variations in light source radiances and position can occur:

- The sources may not be identical. The radiance from individual bulbs of the same type will vary slightly, particularly for inexpensive off-the-shelf bulbs.
- The sources may not be isotropic. In our case we used halogen bulbs for their brightness and ease of use. The emitted radiance is approximately uniform in a 30 degree cone around the direction that the source is aimed. There are small variations across the scanned area, and small disturbances caused slight changes in the source direction.
- The sources may vary over time. The emittance radiance varies with the age of the bulb. It also varies with each use as a function of time since the bulb was turned on.
- The source radiance varies with the electrical power level. A regulated power supply is required to maintain a consistent source level.
- The distance from the light sources to the area being scanned varies across the piece. Requiring the distance to the light sources to be large relative to the scanned area would result in very small areas being scanned in each shot to obtain the spatial resolution required.
- As well as overall strength, the light source spectrum varies bulb to bulb, with time and with level of power input.

All of these factors can be treated by high precision characterization and control of the lighting system, and by restricting the spatial extent of data acquired for each set of photometric images. However, such precision adds greatly to the cost of the system, and to the time required to acquire the data for the full sculpture. Instead, we deal with

these problems by exploiting the existence of an underlying lower resolution mesh, and by performing a global registration of results.

3 Revised Calculation of Normals

Because of the variations from an ideal photometric setting, we need a method to solve the system in Eq. 1 expressed with variable light source radiances and solid angles. We assume that the radius d of the light sources is the same for all the sources, and that d is small (i.e. less than 10 per cent) relative to the distance $r_{i,p}$ from surface point p to each light source i so that the solid angle $\Delta\omega_i$ for each source can be approximated by $\pi d^2/r_{i,p}^2$. We denote the radiance of light source i at pixel location p with $L_{o,i,p}$. Eq. 1 can be rewritten as

$$\rho_p d^2 \begin{bmatrix} l_{1,p,x} & l_{2,p,x} & l_{3,p,x} \\ l_{1,p,y} & l_{2,p,y} & l_{3,p,y} \\ l_{1,p,z} & l_{2,p,z} & l_{3,p,z} \end{bmatrix} \begin{bmatrix} n_{p,x} \\ n_{p,y} \\ n_{p,z} \end{bmatrix} = \begin{bmatrix} \alpha L_{r,1,p} r_{1,p}^2 / L_{o,1,p} \\ \alpha L_{r,2,p} r_{2,p}^2 / L_{o,2,p} \\ \alpha L_{r,3,p} r_{3,p}^2 / L_{o,3,p} \end{bmatrix} \quad (2)$$

All of the quantities in the above equation except for α and d vary from pixel to pixel across the image.

With our hybrid system, we obtain a geometric mesh at a spatial resolution of approximately 2mm registered with each set of five photometric images obtained. The individual meshes (Figure 1) are registered to one another using a two pass method that first uses a set of laser dots projected onto the statue for an initial alignment, and then a form of Iterative Closest Point (see e.g. [4]) for a final more precise registration. The meshes are combined into one single mesh using the Ball-Pivoting Algorithm described in [5]. For each set of photometric images we have in the end a consistent underlying mesh, and a camera transformation that gives the camera parameters for the photometric image capture in terms of the global coordinate system. Given these data, we can easily compute the following for each pixel of the photometric images:

- The distances $r_{i,p}$ from the visible surface point p to each of the light sources i . While there is some error in this distance approximation, the error is well under one percent over the typical 0.75 meter distance from the surface to the camera.
- The direction $\mathbf{l}_{i,p}$ from the visible surface point to each light source.
- An approximate surface normal \mathbf{n}'_p at the visible point computed from the lower resolution surface.




Figure 3: Intermediate images used in computing normals. (a) A color image of a small section of the statue. (b) The underlying geometry computed for this section, with a reflectance of one, lit from the front. (c) A grey scale version of the section in one of the photometric images (the one with the lighting from directly in front) and (d) a grey scale image after averaging. (e) The final surface normals (as lit from the front) computed from converting all of the photometric light images into the forms shown in (c) and (d).

We can immediately use the underlying values of $\mathbf{l}_{i,p}$ and $r_{i,p}$ as good estimates for the values in Eq. 2. We need to find estimates for the relative values of $L_{o,i,p}$ that account for the temporal, directional and bulb to bulb variations of sources. To do this, we equate the relative reflected radiances from the underlying surface illuminated by ideal sources, to the relative image radiances in the neighborhood of the pixel.

Consider ideal light sources with uniform radiances L_u in the same positions as the physical light sources. If these ideal sources illuminated the underlying base surface coated with a Lambertian reflectance of 1, the reflected radiances would be:

$$\tilde{L}_{r,i,p} = (L_u d^2 / r_{i,p}^2) \mathbf{l}_{i,p} \cdot \mathbf{n}'_p \quad (3)$$

For the i -th photometric image, the reflected radiance in the neighborhood around pixel p is

$$\alpha \bar{L}_{r,i,p} = \sum_q^Q \alpha L_{r,i,q} / |Q| \quad (4)$$

where we consider a neighborhood of pixels Q that approximately represents the area on the surface of a disk with a radius equal to the base geometry resolution (i.e. for our system 2 mm). Notice that from the images, we do not have $L_{r,i,p}$ directly, but the quantity $\alpha L_{r,i,p}$ as recorded by the camera. The values $\alpha L_{r,i,q}$ are the result of recording the reflection from the true physical sources:

$$\alpha L_{r,i,q} = \rho_q \alpha L_{o,i,q} d^2 / r_{i,q}^2 \mathbf{l}_{i,q} \cdot \mathbf{n}_q \quad (5)$$

We make the assumption that although ρ_p , $L_{o,i,p}$, $r_{i,p}$, and $\mathbf{l}_{i,p}$ vary pixel by pixel over the entire image, they vary relatively little in the small neighborhood around pixel p . We also assume that the average over Q of the dot products $\mathbf{l}_{i,q} \cdot \mathbf{n}_q$ is approximately equal to the dot product of \mathbf{l}_p and the underlying surface normal \mathbf{n}'_p . Note that these

assumptions are the same basic assumptions made in all photometric stereo calculations – with the exception that we now make the assumptions in a (small) region of pixels rather than for just the surface represented by a single pixel. With these assumptions we have:

$$\alpha \bar{L}_{r,i,p} \simeq \rho_p \alpha L_{o,i,p} d^2 / r_{i,p}^2 \mathbf{l}_{i,p} \cdot \mathbf{n}'_p \quad (6)$$

We now form the ratio of the radiance reflected from the underlying surface given in Eq. 3 to the average recorded radiance in the pixel neighborhood given in Eq. 6:

$$\tilde{L}_{r,i,p} / \alpha \bar{L}_{r,i,p} \simeq L_u / \rho_p \alpha L_{o,i,p} \quad (7)$$

We can use Eq. 7 to express the unknown source radiances $L_{o,i,p}$ in terms of quantities we can compute, or which are the same for all light source directions:

$$L_{o,i,p} \simeq \frac{L_u \alpha \bar{L}_{r,i,p}}{\rho_p \alpha \tilde{L}_{r,i,p}} \quad (8)$$

Using Eq. 8 in the right hand side of Eq. 2, and simplifying terms results in:

$$\frac{L_u}{\alpha} d^2 \begin{bmatrix} l_{1,p,x} & l_{2,p,x} & l_{3,p,x} \\ l_{1,p,y} & l_{2,p,y} & l_{3,p,y} \\ l_{1,p,z} & l_{2,p,z} & l_{3,p,z} \end{bmatrix} \begin{bmatrix} n_{p,x} \\ n_{p,y} \\ n_{p,z} \end{bmatrix} = \begin{bmatrix} \alpha L_{r,1,p} r_{1,p}^2 \tilde{L}_{r,1,p} / \alpha \bar{L}_{r,1,p} \\ \alpha L_{r,2,p} r_{2,p}^2 \tilde{L}_{r,2,p} / \alpha \bar{L}_{r,2,p} \\ \alpha L_{r,3,p} r_{3,p}^2 \tilde{L}_{r,3,p} / \alpha \bar{L}_{r,3,p} \end{bmatrix} \quad (9)$$

We have left in the values of α since the quantities αL_r are the values actually available from the recorded images. Eq. 9 can be solved for a vector in the direction of the surface normal, and normalization gives the result \mathbf{n}_p . The effect of the surface reflectance ρ_p is in both the individual pixel values and the averaged pixel values. As a result the

reflectance term is canceled out of the equation, and the length of the vector in the direction of the surface normal no longer has a useful physical interpretation.

The new calculation process is illustrated in Figure 3.

The accuracy of the resulting normals depends on the validity of the assumptions made in the approximation in Eq. 6. The small size of the neighborhood used for averaging relative to the distance to the camera (2mm versus 750mm), justifies the assumptions that $L_{o,i,p}$, $r_{i,p}$ and $\mathbf{l}_{i,p}$ are uniform. The reflectance is also reasonably assumed to be uniform in a small neighborhood. Exceptions to this are abrupt changes due to a change in surface coating. However, changes in reflectance cause a consistent change in all of the photometric images and can be detected. That is, a sharp increase in reflectance results in an increase in reflected radiance for all light source directions. The main assumption is the equivalence of the average of the dot product of the light direction and the normals on the high resolution surface and the dot product of the light direction with the normal of the underlying lower resolution surface.

To evaluate the assumption of equating the average and single dot products we consider the simple two dimensional system shown in Figure 4. We consider a horizontal surface, of width 2, that is approximated by the low resolution surface connecting the points $(0, 0)$ and $(2, 0)$. The surface is illuminated by two lights: \mathbf{l}_0 in direction $(0, 1)$, and \mathbf{l}_1 in direction $(-\sqrt{2}/2, \sqrt{2}/2)$. We apply our calculation method to a higher resolution surface that consists of two facets formed by a displacement h in the center of the surface. We examine the accuracy of the result we obtain for the normal of the left facet as a function of h .

Specifically, we compute the following quantities:

- The underlying surface normal $\mathbf{n}' = (0, 1)$ for all h .
- The true facets normals $\mathbf{n}_0 = (-h, 1)/\sqrt{1+h^2}$ and $\mathbf{n}_1 = (h, 1)/\sqrt{1+h^2}$
- The reflected light for each facet for each light, $L_{r,i,j} = \mathbf{l}_i \cdot \mathbf{n}_j$
- The approximate surface normal from Eq. 9, using $\bar{L}_{r,i} = L_{r,i,0} + L_{r,i,1}$ and $\tilde{L}_{r,i} = \mathbf{l}_i \cdot \mathbf{n}'$ and setting the constant factor to one.

The results of these calculations are shown in Figure 4. The plot compares the angle between the estimated and true normal from using the underlying surface normal, and from using the approximate normal from the photometric method. The non-linearity of the dot products with normal vectors accounts for the errors in the photometric method, and for the asymmetry between negative and positive values for h . Although the error in the photometric method





Figure 4: (a) The two dimensional system used to evaluate the validity of assumptions made in computing the normals. (b) Results of applying our method for normals computation.

increases with the absolute value of h , the photometric estimate of the normal is in all the cases a significant improvement over using the underlying normal.

The above calculations do not account for interreflections that would occur when h is negative. In our application we are measuring a surface with low reflectance and so have relatively little indirect illumination from interreflections. For surfaces with high reflectance and steep slopes in the high resolution surface there would be additional error in the computed normals.

4 Color Calculations

Given an estimate of the surface normal at each pixel, we can compute versions of the photometric input images with the shading effects due to the directional lighting removed. By taking a weighted average of these corrected images, we obtain an estimate of reflectances in the red, green and blue color channels. Unlike scans with polychromatic laser [2], we can not obtain the absolute value of the reflectance in each channel, or a precise spectral definition of each channel. However, using some additional spot measurements of spectral reflectance we can adjust the corrected images to approximate absolute reflectance in defined spectral ranges.

Denoting the wavelength of light as λ and using the subscript C to denote the color channel that is either red, green or blue, the recorded color value $\alpha L_{r,i,p,C}$ of each pixel p in each photometric image i is given by:

$$\alpha L_{r,i,p,C} = \int_0^\infty \rho(\lambda) L_{o,i,p}(\lambda) (d^2/r_{i,p}^2) \mathbf{l}_{i,p} \cdot \mathbf{n}_p S_C(\lambda) d\lambda \quad (10)$$

where $S_C(\lambda)$ is the spectral sensitivity of the camera in channel C .

Using the normals from the previous section, we compute an approximate relative reflectance ρ_{rel} for each channel from:

$$\rho_{rel,i,p,C} = \alpha L_{r,i,p,C} T_{i,p}^2 / \mathbf{l}_{i,p} \cdot \mathbf{n}_p \quad (11)$$

This reflectance is relative to other pixels in the same channel in the same image. As in the case of the calculation of normals, we do not use high or low pixel values to avoid using data in areas of specular reflection and shadow. Unlike the normals calculations, we do not need a pixel estimate from three images to estimate the relative reflectance. Because we can afford to use less of the image, we avoid including the effects of directionally varying light source intensity by only using pixels that fall in a small cone around the direction the light source is aimed.

The five relative reflectance images computed with Eq. 11 are adjusted relative to one another to compensate for variations in light source strength. The pixels for which

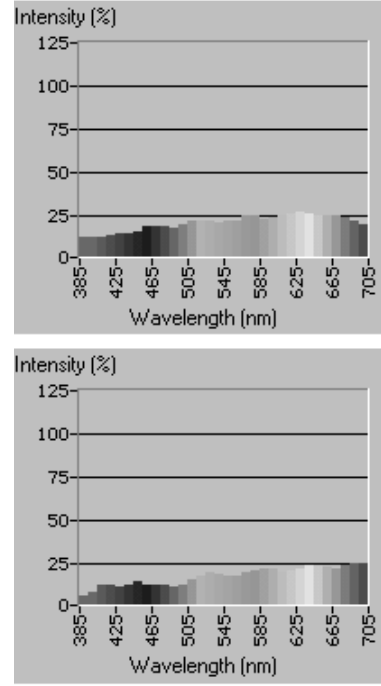


Figure 5: Typical reflectance spectra measured on the statue.

all five of the images have non-zero values are found, and the averages for these overlapping pixels are computed. The overall level of the images is adjusted by making all of the average values equal.

Once the images levels are adjusted, the five images are combined with a weighted average to produce a single corrected RGB image for the camera position. The weight assigned to each pixel in each image increases with distance to the nearest black edge in the image. This weighting is used to avoid small but sudden changes at edges where each photometric image stops contributing to the combined average.

Clearly the relative reflectances computed from Eq. 11 are not scaled to the actual percentage of visible light that would be reflected. The images also still contain the effects of the light source spectra and the camera spectral sensitivity. Furthermore, the red, green and blue channels are not defined in terms of wavelengths. To deal with these problems, we made separate spot measurements of spectral reflectance using the Colortron color ruler [8]. Two typical spectra measured on the sculpture are shown in Figure 5. The spectral reflectances are simple functions in this case, and we can readily define spectral bounds for red, green, and blue for which the spectral reflectances vary little from the average for the range. That is we can define $\lambda_{C,min}$

and $\lambda_{C,max}$ and compute

$$\rho_{ave,C} = \int_{\lambda_{C,min}}^{\lambda_{C,max}} \rho(\lambda) / (\lambda_{C,max} - \lambda_{C,min}) d\lambda \quad (12)$$

By taking the ratio of $\rho_{ave,C}$ and $\rho_{rel,C}$ at the location in the image where the spot measurement was taken, we can estimate the scaling for light source and camera spectral effects in the image. Using the scaling for each channel in the image, we can obtain at least approximate values for the absolute values of reflectance in clearly defined spectral ranges. We are still performing tests to better quantify the quality of the estimates we can obtain.

Once we have a few images that have been adjusted to approximate reflectances in the three color channels, we can adjust all of the images for which $\rho_{rel,C}$ has been computed by requiring the colors in matching points in overlapping images to be the same. As in the case of geometric registration, it is inadequate to simply perform these adjustments pairwise, since small errors accumulate. Instead we do a simultaneous, global registration.

Forming equations for the color registration is facilitated by the availability of the points used in the initial geometric alignment. One result of the initial alignment is a list of points in overlapping images that represent the same location on the base geometry. We compute a color for each of these points by taking the average of a small neighborhood of pixels in the corrected image. We use a neighborhood rather than a point sample since very small alignment errors (e.g. of the size of one pixel projected onto the surface) can produce a very large color error when there are abrupt changes in colors between pixels. Let $\beta_{C,k}$ be the level adjustment for the k -th corrected photometric image in the C channel. Each matching pair of points gives an equation of the following form for images m and n :

$$\rho_{rel,C,m} \beta_{C,m} - \rho_{rel,C,n} \beta_{C,n} = 0 \quad (13)$$

There are far many more matching points than corrected images, so a least squares solution is used to compute the values of β . In Eq. 13 no level has been set for the system. We could set the value of β for the image (or images) that have been adjusted with the Colortron measurements to one. However, we want the errors in the results to be distributed evenly over the model, rather than being skewed by distance from a particular image. To accomplish this, the following equation is added to the set:

$$\sum_k^N \beta_{C,k} = N \quad (14)$$

where N is the total number of images. Using N on the right hand side of Eq. 14 produces values of β around 1 (i.e.

from about 0.6 to 1.4 in our tests). After all of the images have been adjusted to one another, global correction values for red, green and blue can be made to match the Colortron measurements.

5 Results and Conclusions

The results of using our new method for computing colors and normals are shown Figures 6 and 7. Figure 6 shows a small 20cm-wide section of the statue. The top row of images shows the low resolution geometry lit from below, directly in front and from the right. The lower row of images shows the same section with the normals computed using the photometric method lit from the same directions. Figure 6 demonstrates that the only difference between the original base surface and the photometric results are high spatial frequency details. The appearance of the details varies a great deal with lighting direction, and so could not be represented with a simple texture map.

Figure 7 demonstrates that the normals and colors computed from various camera positions are consistent with one another. Pictures (a) and (b) in Figure 7 show 72 meshes with their default texture maps, used to reconstruct a 2.25m tall section of the statue. The patchwork effect is due to the varying lighting conditions as the camera is moved around the statue. Image (c) shows the model with the computed photometric normals. The seams between the individual meshes are not visible. The next image shows the color reflectances computed from the 72 camera viewpoints. This is the result of the global color registration only, the final adjustments to true spectral quantities have not been made.¹ Image (e) shows a rendering of the model with photometric normals and colors.

We have presented a method for successfully computing consistent normals and colors from photometric data from many different camera viewpoints. We are continuing to refine our color calculations to better represent the actual spectral characteristics of the measured surface. Our method produces many overlapping normals and color maps for each area on the surface. We are exploring methods for selecting and weighting the best maps to use in each area.

References

- [1] J. Abouaf. The Florentine Pietà: Can visualization solve the 450-year-old mystery? *IEEE Computer Graphics and Applications*, 19(1):6–10, February 1999.
- [2] R. Baribeau, M. Rioux, and G. Godin. Color reflectance modeling using a polychromatic laser range sensor. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 1992.

¹Note that after true spectral values are computed, color accurate images will have to be computed taking into account the color display characteristics under which the images are being viewed to accurately represent the statue.




Figure 6: A comparison of the low resolution model (top row) and the same model augmented by the normals computed with our method. The three columns from left to right represent the same section of the statue lit from below, directly in front and from the right.




Figure 7: Results of our method applied to the Pietà data. From left to right: (a) A montage of sample photometric images of the statue, showing the wide range of variation in illumination conditions. (b) Inconsistencies are evident in the patchwork effect resulting from stitching together the meshes and textures. (c) Photometric normals applied to the base mesh. (d) Color reflectance computed with our method. (e) A rendering of the combined normals and color maps.

- [3] J.-A. Beraldin, F. Blais, L. Cournoyer, M. Rioux, F. Bernier, and N. Harrison. Portable digital 3-D imaging system for remote sites. In *Proceedings of the 1998 IEEE International Symposium on Circuits and Systems*, pages 488–493, 1998.
- [4] R. Bergevin, M. Soucy, H. Gagnon, and D. Laurendeau. Towards a general multi-view registration technique. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 18(5):540–547, May 1996.
- [5] F. Bernardini, J. Mittleman, H. Rushmeier, C. Silva, and G. Taubin. The ball-pivoting algorithm for surface reconstruction. Submitted for publication, 1999.
- [6] R. Epstein, A. Yuille, and P. Belhumeur. Learning object representations from lighting variations. In *ECCV 96 International Workshop*, pages 179–199, April 1996.
- [7] B. K. P. Horn and M. J. Brooks. *Shape from Shading*. MIT Press, 1989.
- [8] Light Source, Inc. Colortron. <http://www.ls.com/colortron.html>.
- [9] H. Rushmeier, F. Bernardini, J. Mittleman, and G. Taubin. Acquiring input for rendering at appropriate levels of detail: Digitizing a Pietà. In *Proceedings of the Ninth Eurographics Rendering Workshop*, pages 81–92, June 1998.
- [10] Y. Sato, M. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *Computer Graphics (SIGGRAPH '97 Proceedings)*, pages 379–388, August 1997.
- [11] Visual Interface, Inc. Virtuoso. <http://www.visint.com/>.