Photometric stereo without multiple images

Mark S. Drew *

School of Computing Science, Simon Fraser University
Vancouver, B.C. Canada V5A 1S6

ABSTRACT

Photometric Stereo (PMS) recovers orientation vectors from a set of graylevel images. Under orthography, when the lights are *unknown*, and for a single uniform Lambertian surface, one can recover surface normals up to an unknown overall orthogonal transformation. The same situation obtains if, instead of three graylevel images, one uses a single RGB image taken with at least three point or extended *colored* lights impinging on the surface at once. Then using a robust technique and the constraints among the resulting three effective lighting vectors one can recover effective lights as well as normals, with no unknown rotation.

However, in the case of a non-Lambertian object, PMS reduces to the idea of using a lookup table (LUT) based on a calibration sphere. Here, we show that a LUT can also be used in the many-colored-lights paradigm, eliminating the need for three separate images as in standard PMS. As well, we show how to transform a calibration sphere made of a particular material into a theoretical sphere for a cognate material similar in its specular properties but of a different color. In particular, we postulate that a LUT developed from one human's skin can be used for any other person; problems arising from shadows, hair, eyes, etc. are automatically eliminated using robust statistics.

Results are shown using both synthetic and real images.

Keywords: Physics-Based Vision; Shape; Shape Representation; Color; Dichromatic Model; Neutral Interface; Lambertian; Reflectance

1 INTRODUCTION

Photometric Stereo (PMS) is a straightforward method for recovering surface normals when a surface is Lambertian, and lights are known.^{1,2} When lights are *unknown*, normals can still be recovered up to an unknown overall orthogonal transformation by fitting an ellipsoid to the scatterplot of 3-vector points formed from the intensities of three graylevel images of the surface illuminated by lights from three different directions (or, equivalently, for each narrow-band channel in a spectrally separated RGB image illuminated at one time by narrowband R, G, and B lights from different directions).³ However, the same recovery can also be carried out by utilizing a (small or large, point or extended) set of unknown *colored* lights with no restriction on spectral breadth of lights or filters. For because a Lambertian surface effectively replaces an extended light source with an equivalent point source (Ref.[4], p.237), and in this case does so in each RGB channel, one can plot RGB values in color space and fit an ellipsoid to such a scatterplot.^{5,6} If specularities or shadowed regions are present, a robust regression can pinpoint the image locations of inliers that fit the ellipsoid well enough, and in fact one can further disambiguate specularities from other outlier pixels.⁷

To find the unknown orthogonal matrix, one can use either an argument based on integrability^{6,8} or, more directly, find three effective color-strength-direction vectors $\mathbf{f}_k$, $k = 1..3$ by making use of constraints amongst

*Email: mark@cs.sfu.ca; WWW: <http://fas.sfu.ca/cs/people/Faculty/Drew/index.html>; Telephone: (604) 291-4682; Fax: (604) 291-3045

them: once the color space ellipsoid is determined, one knows the three 3-vectors $\mathbf{f}_k$ up to a rotation. Normalizing them, we know the three angles between each of three pairs of these vectors. But these angles can be determined if we additionally carry out a Pentland-source-direction type of statistical analysis⁹ for each of the RGB channels. The additional information provides enough constraint to solve for the vectors $\mathbf{f}_k$ themselves.¹⁰ A similar analysis can also be carried out with only two light sources present.¹¹

For a *non-Lambertian* surface, one can show that for a general m -lobed reflectance map it is possible to uniquely invert a set of PMS equations to recover surface orientation from a set of images.¹² Below we show that this proof does *not* carry over to the color image regime. For graylevel images, for a non-Lambertian surface one can simply provide a calibration sphere and base recovery of normals on a lookup table (LUT). Here we make the simple proposal that if we image a calibration sphere in a spectrally-varying colored light environment then a similar LUT approach can be used for surface orientation recovery.

As well, once one has arrived at a LUT for a particular surface, for a particular lighting scheme, it is possible to develop a theoretical calibration sphere, including specularities, by transforming the Lambertian part of the reflectance map into a different but cognate material and adding the same specular part.

In particular, it should be possible to develop one overall calibration sphere for human faces, regardless of complexion or non-skin features. We show results below using a real face image.

In §2, we summarize the linear relation between normals and colors, for a Lambertian surface imaged under a 3-dimensional colored lighting environment. The relationship between m -lobed models and color imaging is discussed.

The LUT method for a colored-lighting situation is set forth in §3, and results are shown to be very good for both synthetic and real images for non-Lambertian reflectance maps. Ordinary lighting, with additional colored lights, is used to produce an illumination environment that varies spectrally with direction from the surface.

In §4 it is shown how to transform a calibration sphere from one surface material to another. It is shown that recovery of surface normals is fairly accurate for a typical real image.

In §5 future work is set out: we postulate that by modeling human skin in a similar fashion as other materials it may be possible to transform a calibration sphere for one face illuminated under a colored, spectrally-varying lighting environment into that for another face. Section 6 sets out conclusions and possible directions for improvement.

2 ORIENTATION-FROM-COLOR

2.1 Lambertian case

Suppose we image a single Lambertian surface with surface spectral reflectance function $S(\lambda)$ and surface normals $\mathbf{n}$. Then under conditions of distant lighting and distant viewing, the RGB camera response ρ for a camera system with RGB optical response functions $\mathbf{q}(\lambda)$ is

$$\rho = \sum_{i=1}^L (\mathbf{a}_i^T \mathbf{n}) \int E_i(\lambda) S(\lambda) \mathbf{q}(\lambda) d\lambda \quad (1)$$

where T means transpose. Here we discretize the lighting environment into a set of L lights, $i = 1..L$, with normalized direction vectors $\mathbf{a}_i$ and spectral power distributions $E_i(\lambda)$. Integration is over the visible spectrum.

We can simplify this model by writing

$$b_i \equiv \int E_i(\lambda) S(\lambda) q(\lambda) d\lambda, \quad (2)$$

so that eqn.(1) can be economically stated as⁵

$$\rho = B A n \equiv F n, \quad (3)$$

with $L \times 3$ direction matrix A , $3 \times L$ color-strength matrix B , and 3×3 color-strength-direction matrix F : $A = (a_1, a_2, \dots, a_L)^T$, $B = (b_1, b_2, \dots, b_L)$, $F \equiv B A$.

In terms of the $3 \times N$ collection $\mathcal{R}$ of RGB values for all N pixels and the $3 \times N$ collection $\mathcal{N}$ of normal vectors we have

$$\mathcal{R} = F \mathcal{N}. \quad (4)$$

Shadows, self-shadowing, and specularities will result in not all pixels obeying this relationship, however. For those pixels that do, if F is invertible with inverse G , then we can recover n from ρ . Matrix F is invertible if product matrices A and B are both full-rank, rank 3. Matrix A is a matrix of light directions and matrix B is a matrix of the colors of those lights when reflected from the surface in question. Thus F is invertible if the lighting environment varies spectrally with direction from the surface and is linearly independent.

In most cases *ab initio* we know neither n nor F . However we can find part of F by virtue of the quadratic form

$$1 \equiv n^T n = \rho^T G^T G \rho = \rho^T C \rho, \quad \text{with } C \equiv G^T G. \quad (5)$$

This states that (for Lambertian pixels that see all the lights) the scatterplot of ρ values lie on the surface of an ellipsoid in color space. If F has the singular value decomposition (SVD) $F = U A V^T$, then we also have $C = U A^{-2} U^T$ so that both U and A are known once we know C , but V is not.

However, by running a Pentland-type source direction finder in each of the three channels separately, and retaining the tilt angles found but discarding the possibly inaccurate slant angles, we can produce three more constraints and also determine V .¹⁰

2.2 Non-Lambertian case

If the surface is not necessarily Lambertian, then one can still retrieve matrix C reliably using robust statistics⁷: if at least 50% of the pixels are indeed Lambertian, and also belong to surface points that see all lights, then the Least Median of Squares (LMS) regression¹³ can find that matrix. Note that for this robust method the magnitude of pixel values plays no role — it is their preponderance in supporting a particular set of matrix values C that counts. Outliers identified by the LMS model will generally fall into one of two classes: either they are points shadowed from some of the lights or else they are specularities. Further, specularities can be associated with high norms of color vectors ρ and thus disambiguated from shadows.⁷

To find matrix C , consider a regression with six explanatory variables ($\rho_1^2, \rho_2^2, \rho_3^2, 2\rho_1\rho_2, 2\rho_1\rho_3, 2\rho_2\rho_3$) and all right-hand-sides equal to 1. How the LMS method proceeds is by randomly selecting sets of six pixels and for each selection finding an estimate of the six coefficients making up a vectorized version of (symmetric) matrix C by means of matrix inversion. The combinatorics are limited by an important theorem (see Ref.[13]) to only 3000 randomly selected cases. The regression result is taken to be that set of coefficients yielding the least median of the squared residuals for all pixels. A robust figure of merit is available to determine goodness of fit. Another error check consists of checking that all eigenvalues of the recovered matrix C are indeed positive, as they should be. Negative eigenvalues would indicate that the model (3) is inappropriate for the data considered.

Thus one determines the six elements of matrix C , leaving three undetermined degrees of freedom for matrix F , given by an arbitrary rotation of SVD matrix V .

3 PMS WITHOUT MULTIPLE IMAGES

The problem with the analysis so far is that even if specularities are detected, we have no model of how to map colors to normals within a specularity. To remedy this deficiency, suppose we adopt the straightforward approach of PMS and insert a calibration sphere into the scene.

3.1 Uniqueness

For graylevel images, it can be shown¹² that three images are sufficient for a unique inversion of the mapping from surface normals to intensity values provided that the surface's reflectance map falls in the m -lobed class. This class consists of the constant, plus $j = 1..(m-1)$ terms of the form $\Phi^j(\mathbf{p}_j^T \mathbf{n})$ with Φ^j a non-negative, monotonically increasing function. Further, the vector $\mathbf{p}_j$ must lie in the *principal plane* formed by the light direction and the vector $\mathbf{n}_S$ half-way between the illuminant direction and the viewer. E.g., the Torrance-Sparrow model¹⁴ can be written in terms of a function with $\mathbf{p} \equiv \mathbf{n}_S$. Clearly the Lambertian model also falls into the m -lobed class, with $\mathbf{p}_j \equiv \mathbf{a}$, the light source direction.

For color images, however, this uniqueness property breaks down. For consider a simple dichromatic model,¹⁵ whereby reflectance is comprised of a diffuse or body term plus a surface (specular) term. E.g., suppose ρ consists of a Lambertian plus a surface term:

$$\rho = \mathbf{F} \mathbf{n} + \sum_{i=1}^L b^i_S \Phi(\mathbf{n}^i_S{}^T \mathbf{n}) \quad (6)$$

Here b^i_S is the color of the specularity for the i^{th} light. This is a 2-lobed model.

The uniqueness proof of Ref.[12] is for a set of graylevel equations that we are comparing to eqn.(6). For such a set of $k = 1..3$ equations, that uniqueness proof rests on the idea that vectors $\mathbf{p}_{k,j}$ for the $j = 1..(m-1)$ lobes *all* lie in the principal plane for light k . For the Lambertian term in eqn.(6), we have $\mathbf{p}_{k,1} \equiv \mathbf{f}_k$, where we denote the three rows of matrix $\mathbf{F}$ by $\mathbf{f}_k^T$.

However (with lights indexed by $i = 1..L$) the surface term in eqn.(6) is written in terms of vectors $\mathbf{n}_S^i$ that do not lie in the principal plane for $\mathbf{f}_k$ — many different lights contribute to each of the three channels. Therefore the uniqueness proof of Ref.[12] does not carry over into the colored regime when the illumination environment varies spectrally with direction from the surface.

Nonetheless, a simple solution is evident: certainly if it were possible to illuminate every different sphere normal with a different color or strength of light then uniqueness would be achieved. A constructive proof, then, or at least an effective demonstration of uniqueness is available by simply running any orientation-recovery algorithm on the calibration sphere itself. If output matches input, within noise considerations, then normals are correctly recovered.

3.2 Lookup table: synthetic image

To test the lookup table idea, consider Fig.1(b), an exemplar image formed by synthetically shading the radar range image Fig.1(a).¹ In order to use real lights and surfaces, let us use the surface spectral reflectance for "light skin" as $S(\lambda)$.¹⁶ As illuminants, let us take standard illuminant D65¹⁷ times the transmission functions of, respectively, a Kodak Wratten filter #25 (Red), a neutral density filter, and a Kodak Wratten filter #47B

¹The laser range data for the bust of Mozart is due to Fridtjof Stein of the USC Institute for Robotics and Intelligent Systems.

(Blue)¹⁸ as illuminants $E_i(\lambda)$, $i = 1..3$. Let matrix $\mathbf{A}$ correspond to directions (tilt,slant) = $(0^\circ, 35^\circ)$, $(120^\circ, 35^\circ)$, and $(-120^\circ, 35^\circ)$. Let us use the camera response curves for a Sony DXC151 camera with infrared filter for functions $q(\lambda)$.

The resulting image was produced by summing Lambertian shading for each light with self-shadowing calculated using a ray-tracing algorithm. Additionally, specular reflection was added using the Phong model,¹⁹ with

$$\Phi(n_S^i T n) = (n_S^i T n)^p \quad (7)$$

for each light, with power $p = 20$. For vectors b_S^i , let us adopt the neutral interface model,²⁰ whereby the color of a specularity is simply that of the incident illuminant, filtered by the imaging system $q(\lambda)$. Additive Gaussian noise was then added with rms value $2/255$. (Color images are displayed in black and white.)

Fig.1(c) shows the result of imaging a sphere with the same reflectance map under the same illumination conditions. If the lookup table idea works for surfaces illuminated by spectrally varying illumination, then we should be able to recover the sphere itself by looking up normal vectors from its own RGB values. A useful indicator of how well we recover normals is the set of angles between actual and recovered normals. It could be the case that some RGB values are equal at several points in the sphere image, due to quantization, and in fact we find that this is the case for about 1.8% of the pixels. Using simple, trilinear interpolation we find that even though the median of error angles is zero, the mean is 0.03° . Overall, the recovery of sphere normals is excellent and we can go on to the synthetically shaded face of Fig.1(b).

Here we use RGB values from the face, Fig.1(b) to look up normals from the sphere, using Fig.1(c). Since we in fact know the actual normals for the face, we can gauge accuracy of recovery in this case as well: the median error is 0.8° and the mean error is 3.8° . Most of the error is concentrated where the derivative is steepest, around the occluding boundary. Thus recovery is seen to be very good. For future reference, let us also make note of the form of the scatterplot of RGB values for Fig.1(b): it is shown in Fig.1(d) and takes the form of an ellipsoid as in eqn.(5), with additional specular values.

In the next section, we apply the lookup table method to a real image.

3.3 Lookup table: real image

Consider Fig.2(a), an image of an orange plastic beachball. The surface of the ball is pebbled because it is made of many small cells compressed together, and as a result it also has an overall sheen of specular reflection in most environments, e.g. from mutual illumination from white walls. The reason would appear to be that the edges of the compressed pebbles of plastic reflect the illuminant specularly. The ball was imaged using a Sony DXC-930 camera and illuminated with several colored incandescent lights as well as room light and daylight from a window.

First, a test for uniqueness is performed by recovering the ball itself. Since the ball was placed on a black cloth it is straightforward to segment the image into figure and ground. Fitting an ellipse to the occluding boundary, the ball was found to have a slight eccentricity. Assuming an ellipsoidal shape and taking the dimension in the direction toward the camera to be the mean of the two dimensions perpendicular to the camera, we arrive at normal vectors for the ball surface. Then indexing into these known normals using the measured RGB values, we find that recovered normals have median and mean angular errors of 0° and 2.4° , respectively. Therefore the LUT method works correctly for a real object.

We can experimentally verify that the surface is not at base Lambertian by carrying out the robust regression for eqn.(5); the resulting value of the robust version of the coefficient of determination R^2 is too small to indicate a good fit. Thus either the base material is not Lambertian or else there are not at least 50% of the pixels for

which specular reflection is not present.

Thus one can state that the lookup table approach works very well, even for a strongly non-Lambertian surface. Fig.2(b) shows the shading field for the recovered ball, illuminated from the (0,0,1) direction.

To show how the method can be used for an *unknown* shape, consider Fig.2(c). Here a similar ball, made of the same material, is shown deflated. The normals for this image are recovered from those for the inflated ball, imaged under the same lighting, with the results shown in Fig.2(d). Here the recovered normals are shown illuminated from direction (0,0,1) The shape would appear to be well recovered.

4 TRANSFORMATION FROM ONE MATERIAL TO ANOTHER

4.1 Robust regression: RGBs to normals

Fig.3(a) shows another toy, this one a yellow very glossy ball. Fig.3(b) displays the scatterplot of RGB values for this image. The image was taken under ordinary room lighting and window daylighting, as well as several colored incandescent bulbs. This figure resembles the form of the RGB plot in Fig.1(d), which came from a theoretical reflectance map. For this surface, we can repeat the lookup table recovery described in §3.3 now for the ellipsoidal surface of the yellow ball — the results are very good: median and mean angular errors are 0° and 2.1°.

The RGB plot shows that the surface is very likely a combination of Lambertian and specular reflectance contributions, and this is verified by carrying out the robust regression for eqn.(5), yielding matrix $\mathbf{C}$. However, we can in fact do better than the regression of eqn.(5). For since we know the surface normals, we can regress for the components of $\mathbf{F}$ directly by minimizing

$$\min_{\mathbf{F}} \text{med}[\mathcal{R} - \mathbf{F} \mathcal{N}]^2 \quad (8)$$

for LMS coefficients $\mathbf{F}$.

Such a regression consists in practice of three independent multivariate regressions on three independent variables, one regression for each of the three rows in eqn.(4). The LMS method returns weights w_i for each pixel based on the value of a dispersion estimate

$$s_0 = k \left(1 + \frac{5}{N-3} \sqrt{\text{med}_i r_i^2} \right), \quad (9)$$

where r_i is the residual r for the i^{th} case and, in terms of the cumulative standard normal distribution Φ , $k = 1/\Phi^{-1}(0.75) \simeq 1.4826$. Then a point is *accepted* as corresponding to the model, and $w_i = 1$, if $|r_i/s_0| \leq 2.5$; else the point is an outlier and is rejected, $w_i = 0$.

Since here we have a multidimensional regression there result three separate values s_0 and a set of three weights for each pixel, based on whether or not the value in each channel is an inlier. The simplest approach to identifying outliers in a multidimensional regression is to AND weights at each pixel, producing an overall weight for each ρ vector (but see Ref.[21]). The resulting outlier mask is shown in Fig.3(c). We can see that both specularities and self-shadow regions are identified.

4.2 Diffuse model

With matrix $\mathbf{F}$ in hand, we can model color inside a specularity as a sum of Lambertian plus specular contributions. This will allow us to generate color as it would appear for the ‘calibration’ sphere for a *cognate* material, i.e., one with a different body color but similar specular reflectance.

First, we must identify specularities. In keeping with the proposal in Ref.[7], we identify specularities as outlier pixels that also have RGB vector norms that lie in the fourth quartile of all RGB norms. For the yellow ball, this results in identification of specularities as in Fig.3(d).

Inside specularities, we model diffuse reflection via the Lambertian contribution $\rho_L = \mathbf{F} \mathbf{n}$. Then we can assume the surface (glossy) contribution consists of the remainder $\rho_S = \rho - \rho_L$.

However, for non-specular pixels we may simply use $\rho_D = \rho$ itself for the diffuse contribution:

$$\begin{aligned} \text{In a specular region} \quad & \begin{cases} \rho_D = \rho_L = \mathbf{F} \mathbf{n} \\ \rho_S = \rho - \rho_L \end{cases} \\ \text{In a non-specular region} \quad & \begin{cases} \rho_D = \rho \\ \rho_S = 0 \end{cases} \end{aligned} \tag{10}$$

Now we can use eqn.(10) to model reflectance for some other surface, under the same colored lighting environment.

4.3 Transformation to a cognate material

Let us now examine another ball, shown in Fig.4(a). This ball is purple, but is cognate to the yellow ball in that it displays a qualitatively similar glossiness, indicating that the same base material and approximately similar dye concentrations were used.

We are again using an ellipsoidal shape, but nonetheless we mean this purple ball to be our unknown, experimental test shape. Since we can, in fact, calculate the surface normals for this new surface, we have ground truth to compare the method to. Since in general we would not know the normals for a test shape, we must fall back on eqn.(5), not eqn.(8), since the latter requires knowledge of normals. Therefore we produce an estimate of $\mathbf{C}$, not matrix $\mathbf{F}$ itself. To extract the correct root matrix from $\mathbf{C}$ we use here the constrained-slant method of Ref.[10], and run a Pentland-type source direction finder to fix the tilts of vectors $\mathbf{f}_k$.

The method for extracting matrix $\mathbf{C}$ also identifies specularities, if we make use of the high-norm rule given in Ref.[7]; specularities identified are shown in Fig.4(b) and are very similar in both location *and color* to those in Fig.3(d).

Now let us create a calibration sphere for this purple surface from that for the yellow one. For the diffuse part of such a theoretical calibration sphere we can use $\rho_L = \mathbf{F} \mathbf{n}$, where $\mathbf{F}$ is that determined for the new surface. However, we cannot know precise values near the occluding boundary since these points will certainly be in self-shadow from some of the lights that made up $\mathbf{F}$ — for this theoretical sphere we know $\mathbf{F}$, not the lights themselves. We can definitely tell when sphere normals are turned away from (beyond the terminator for) virtual lights $\mathbf{f}_k$, but not when they are turned away from any actual light. Therefore for ρ_L the best we can do is either set $\rho_{Lk} = 0$ when $\mathbf{f}_k \cdot \mathbf{n} \leq 0$, or perhaps set 3-vector $\rho_L = \mathbf{0}$ when $\mathbf{f}_k \cdot \mathbf{n} \leq 0$ in any channel. Here we use the first method.

To model surface (glossy) reflection for the purple ball, as opposed to diffuse or body reflection modeled by ρ_L , we simply add the specular contribution identified for the yellow, known ball.

How well are surface normals recovered from a LUT based on this new, synthetically created sphere? For this image the results are error angles with median 6.2° and mean 11.3° . In comparison, angular errors using the LUT method with known normals, for the purple ball reconstructing itself from actual, not theoretical ρ values, are median 0° and mean 4.2° .

Fig.4(c) shows the shading field for the recovered purple ball, illuminated from the (0,0,1) direction.

We speculate that the results arrived at here are likely worst-case, in that cognate surfaces that are not extremely dark in some of the bands will be better mapped. An example of such a surface is explored below.

5 HUMAN SKIN

Consider the face of Fig.5(a), taken under a spectrally varying illuminant environment consisting of a room light, daylight, and several colored incandescent bulbs. The RGB values for this face are shown in Fig.5(b); by examining the singular values in an SVD factorization of these RGB's one can be certain that the lighting environment is indeed rank-3. A robust regression for the ellipsoid equation (5) gives reliable results, and is shown by the fitted ellipsoid in Fig.5(b). (Other faces have similar properties.) The inlier mask that results from this regression is shown in Fig.5(c). Another indicator of goodness-of-fit is the count of how many pixels identified as diffusely reflecting give residuals close to zero; or in other words how close RGB values lie to the shell with right hand side unity given by the quadratic form (5). If most pixels lie on this shell, then one is more certain of having found the Lambertian matrix best describing the surface. Here this indicator has a high value — most inlier pixels lie close to the ellipsoidal shell, eqn.(5), in RGB space.

Specularities detected are shown in Fig.5(d). Bright values that are on-shell are not specularities; only high-norm off-shell outliers are specularities. Note that it is possible to design the lighting to attempt to avoid specularities, e.g. by placing lights far off the z -axis. However, in that case only part of the face will be illuminated by all lights, and some face patches may incorrectly be identified as specular patches when in fact they are high-norm outlier patches illuminated by only a subset of lights. In general, for near-axis lights, specularities are indeed present in faces.

Future work, then, must consist of developing a calibration sphere for skin, based on range images. A calibration sphere can be built up incrementally by adding normals for several faces. Skin color is not important, since this is captured in matrix F and can be discounted. As well, hair, jewelry etc. are discounted by the robust regression.

Clearly, using a Lambertian model one can build the diffuse part of a calibration sphere for a second person simply by utilizing the matrix F robustly determined by the second person's image. As well, as strongly suggested by the similarity of Figs.3(d) and 4(c), it is postulated that specularities can also be added from the first face to the calibration sphere for the second, viewing the second face as a cognate material to the first.

6 CONCLUSIONS

In this work we have shown that the LUT approach for carrying out PMS for non-Lambertian surfaces can be carried over to the regime in which a surface is illuminated with light that varies spectrally with orientation from the surface. Although an important uniqueness proof for PMS does not hold once colored images are considered,

we have demonstrated that surface normal recovery can still succeed very well even for highly non-Lambertian surfaces.

We have also shown how to construct a theoretical calibration sphere for one material from the known sphere for a similar, cognate material. For a surface that obeys the Lambertian and neutral interface models well, for a lighting scheme that favors (4) over many pixels, one can develop a new, theoretical calibration sphere from a known one via the scheme (10). In this work, it should be noted that we consider real (and therefore somewhat refractory) materials, and also not highly structured lighting environments — unknown colored lights added to unknown ordinary room lighting.

A particular subject of interest is how to carry out such a transformation between materials when these materials are human skin. An advantage of using robust methods for this task is that outliers include not just shadows and specularities but also hair, eyes, etc.

To improve results, a more careful interpolation scheme should be used, and more analysis should be done on how to best account for illumination effects in self-shadowed pixels. As well, the important case of rank-reduced images (cf. Ref.[11]) should be examined for the LUT approach.

References

- [1] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19:139–144, 1980.
- [2] R. J. Woodham. Gradient and curvature from the photometric-stereo method, including local confidence estimation. *J. Opt. Soc. Am. A*, 11:3050–3068, 1994.
- [3] R. J. Woodham, Y. Iwahori, and R. A. Barman. Photometric stereo: Lambertian reflectance and light sources with unknown direction and strength. Technical Report TR 91-18, University of British Columbia Department of Computing Science, 1991.
- [4] B. K. P. Horn. *Robot Vision*. MIT Press, 1986.
- [5] M.S. Drew. Shape from color. Technical Report CSS/LCCR TR 92-07, Simon Fraser University School of Computing Science, 1992. Available using ftp://fas.sfu.ca/pub/cs/techreports/1992/CSS-LCCR92-07.ps.Z.
- [6] L. L. Kontsevich, A. P. Petrov, and I. S. Vergelskaya. Reconstruction of shape from shading in color images. *J. Opt. Soc. Am. A*, 11:1047–1052, 1994.
- [7] M.S. Drew. Robust specularly detection from a single multi-illuminant color image. *CVGIP:Image Understanding*, 59:320–327, 1994.
- [8] M. S. Drew and L. L. Kontsevich. Closed-form attitude determination under spectrally varying illumination. In *Proc. IEEE Comp. Soc. Conf. on Comp. Vis. and Patt. Rec.*, pages 985–990, 1994.
- [9] A. P. Pentland. Finding the illuminant direction. *J. Opt. Soc. Am. A*, 72:448–455, 1982.
- [10] M.S. Drew. Direct solution of orientation-from-color problem using a modification of Pentland's light source direction estimator. *Computer Vision and Image Understanding*, 64:286–299, 1996.
- [11] M.S. Drew. Reduction of rank-reduced orientation-from-color problem with many unknown lights to two-image known-illuminant photometric stereo. In *International Symposium on Computer Vision, Coral Gables, FL, Nov. 21–23*, pages 419–424. IEEE, 1995.
- [12] H.D. Tagare and R.J.P. deFigueiredo. A theory of photometric stereo for a class of non-Lambertian surfaces. *IEEE Trans. Patt. Anal. and Mach. Intell.*, 13:133–152, 1991.
- [13] P. J. Rousseeuw and A. M. Leroy. *Robust Regression and Outlier Detection*. Wiley, 1987.
- [14] K.E. Torrance and E.M. Sparrow. Theory of off-specular reflection from roughened surfaces. *J. Opt. Soc. Am.*, 57:1105–1114, 1967.
- [15] S.A. Shafer. Using color to separate reflection components. *Color Res. Appl.*, 10:210–218, 1985.
- [16] C.S. McCamy, H. Marcus, and J.G. Davidson. A color-rendition chart. *J. App. Photog. Eng.*, 2:95–99, 1976.
- [17] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. Wiley, New York, 2nd edition, 1982.
- [18] Eastman Kodak Co. *Kodak Filters for Scientific and Technical Uses*, 2nd edition, 1981.
- [19] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics: Principles and Practice*. Addison-Wesley, 2nd edition, 1990.
- [20] H.-C. Lee. Method for computing the scene-illuminant chromaticity from specular highlights. *J. Opt. Soc. Am. A*, 3:1694–1699, 1986.

- [21] P. J. Rousseeuw. Private communication, 1994. Another approach is to form lengths $d_i = [\sum_{k=1}^3 (r_{ik}/s_k)^2]^{1/2}$ where r_{ik} is the residual for the i^{th} case in the k^{th} channel. Then observations with $d_i > 2.5 \text{ median}(d_i)$ could be identified as outliers. This ensures that at least half the cases remain as inliers. Also, a point which is outlying in 3-space but is not very far away in each 1-space would still be detected (unlike the approach of multiplying the three weights). Simply multiplying weights keeps fewer points for this data set and is thus more conservative; therefore the simpler approach is the one used here.

Figure 1: (a): Depth map of plaster bust. (b): Shaded with three colored lights (the figure has been segmented from the ground, which is black). (c): Sphere shaded with same three colored lights. (d): RGB points in color space: ellipsoid centered on origin with additional specular pixels.

Figure 2: (a): Orange ball under multiple illuminants. (b): Shading field for recovered orange ball, illuminated from $(0,0,1)$. (c): Deflated orange ball (brightened for display). (d): Recovered normals, illuminated from $(0,0,1)$.

Figure 3: (a): Yellow ball (histogram-equalized for display). (b): RGB values in color space. (c): Outlier mask from robust regression. (d): Specularities detected (inverse video, histogram equalized).

Figure 4: (a): Purple ball (histogram-equalized). (b): Specularities detected (inverse video, histogram equalized). (c): Shading for purple ball reconstructed from theoretical calibration sphere, illuminated from $(0,0,1)$.

(a)

(b)

(c)

(d)

Figure 5: (a): Face illuminated by spectrally varying lighting environment. (b): RGB values and fitted ellipsoid. (c): Inlier mask. (d): Specularities detected; inliers are shown in gray, outliers that are specularities are shown in white.