Shape From Shading via the Fusion of Specular and Lambertian Image Components

James J. Clark Alan L. Yuille

Division of Applied Sciences
Harvard University
Cambridge, MA

Abstract

We present a closed form analytical solution to the problem of obtaining surface normal information from the specular and Lambertian components of the image of a surface. We show that this algebraic approach to the fusion of data is extremely sensitive to noise and we therefore provide two alternative approaches based on the minimization of energy functionals. The first approach weights the specular and Lambertian information uniformly over the image with respect to a smoothness constraint. The second approach weights the specular and Lambertian components adaptively according to a measure of the sensitivity of the algebraically derived surface normals to noise in the measured image components. This results in a greater dependence on the smoothness constraint in the parts of the image where the surface reconstruction process is most sensitive to noise and provides a more accurate reconstruction of the surface than the uniform weighting technique.

1 An Algebraic Approach to Fusing Specular and Lambertian Reflectance Data

The vision process that we examine in this paper is that of obtaining object shape from information about the specular and Lambertian components of an image.

We will assume that we have available two spatially and temporally registered images, corresponding to the specular and Lambertian components of the surface reflectance. It is possible to extract the specular and Lambertian reflectance components from a color image. If we assume a dichromatic model of surface reflectance [5, 6], it is possible to use color information to distinguish between these two reflectance components [3, 7]. Another approach for extracting the specular and Lambertian reflectance components involves the use of polarizing filters. This approach was described in [8].

However the specular and Lambertian surface reflectance components are obtained, they can directly provide surface normal values, assuming the following simple reflectance models. A simple model for the reflectance function of specularities is given by

$$E_s(x,y) = (\hat{k} \cdot \hat{h}(x,y))^m = R_s(\hat{n}) \tag{1}$$

where $E_s(x,y)$ is the specular image intensity (normalized to lie between 0 and 1), $\hat{k}$ is a unit vector in the viewer direction, $\hat{s}$ is a unit vector in the light source direction, $\hat{n}$ is the surface unit normal vector and $\hat{h}$ is a unit vector that depends on $\hat{s}$ and $\hat{n}$

$$\hat{h}(x,y) = 2(\hat{n}(x,y) \cdot \hat{s})\hat{n}(x,y) - \hat{s} \tag{2}$$

This research was supported in part by the Brown/Harvard/MIT Center for Intelligent Control Systems, under A.R.O. grant number DAA103-86-K-0170.

The parameter m is a number corresponding to the sharpness of the specularity. We also have the Lambertian shading component $E_l(\hat{n})$ given by the Lambertian model

$$E_l(x,y) = (\hat{n}(x,y) \cdot \hat{s}) = R_l(\hat{n}) \tag{3}$$

The functions $R_s(\hat{n})$ and $R_l(\hat{n})$ are the *reflectance maps* for specular and Lambertian surfaces respectively. Each of the image reflectance components is, by itself, insufficient for the task of obtaining a unique solution to the shape extraction problem, as in each case we have only one equation for the two unknowns (the x and y components of the surface unit normal vectors).

Using either the color based or polarization based methods of extracting specular and Lambertian reflectance we can obtain both E_s and E_l. We therefore now have two equations for the surface normal, one from each reflectance component, and so theoretically there is enough information to determine the two components of the surface normal uniquely.

Combining the equations for the surface normal vectors as a function of the specular and Lambertian components we can see that (after some algebra):

$$\hat{n}(\hat{x}) = \alpha(\hat{x})\hat{s} + \beta(\hat{x})\hat{k} + \gamma(\hat{x})\hat{s} \times \hat{k} \tag{4}$$

where

$$\alpha = \frac{1}{s^2\Omega}\left(E_l(\vec{x}) - \frac{E_s^{\frac{1}{m}}(\hat{x})c\Omega}{2E_l(\hat{x})} - \frac{c^2\Omega}{2E_l(\hat{x})}\right) \tag{5}$$

$$\beta = \frac{1}{s^2\Omega}\left(\frac{E_s^{\frac{1}{m}}(\vec{x})}{2E_l(\hat{x})} + \frac{c\Omega}{2E_l(\hat{x})} - c\Omega E_l(\vec{x})\right) \tag{6}$$

$$\gamma = \pm\frac{\left(4s^2\Omega E_l^2 - 4E_l^4 - (E_s^{\frac{1}{m}} + c\Omega)^2 + 4E_l^2(E_s^{\frac{1}{m}} + c\Omega)c\Omega\right)^{1/2}}{2E_l s\Omega} \tag{7}$$

with $c\Omega = \cos\Omega = \hat{k} \cdot \hat{s}$ being the cosine of the angle between the viewer direction and the light source, and $s\Omega = \sin\Omega$. Note that there is an ambiguity in the sign of γ, which will manifest in an ambiguity in the sign of the component of the unit normal vector in the direction of $\hat{s} \times \hat{k}$. This ambiguity *cannot* be resolved using only the specular and Lambertian image components, since both $R_s(\hat{n})$ and $R_l(\hat{n})$ are symmetrical about $\hat{n} = \hat{s} \times \hat{k}$. Thus, in order to obtain a unique surface normal, we must provide more information than just E_s and E_l. In our implementation this information was put in by hand (i.e. the sign of γ was predetermined from knowledge of the shape of the synthetic surface used).

To investigate the uncertainty of the computed shape we see how small changes in the input data affect the output. Suppose we have perturbations δE_l and δE_s in the input data. We then get changes in the output of the fusional module given by

$$\delta\hat{n} = (\delta\alpha)\hat{s} + (\delta\beta)\hat{k} + (\delta\gamma)\hat{s} \times \hat{k} \tag{8}$$

where

$$\delta\alpha = \frac{1}{s^2\Omega}\left(\left[1 + \frac{E_s^{1/m}c\Omega}{2E_l^2} + \frac{c^2\Omega}{2E_l^2}\right]\delta E_l - \left[\frac{c\Omega}{2mE_lE_s^{(m-1)/m}}\right]\delta E_s\right) \quad (9)$$

$$\delta\beta = \frac{1}{s^2\Omega}\left(-\left[c\Omega + \frac{c\Omega}{2E_l^2} + \frac{E_s^{1/m}}{2E_l^2}\right]\delta E_l + \left[\frac{1}{2mE_lE_s^{(m-1)/m}}\right]\delta E_s\right) \quad (10)$$

$$\delta\gamma = \frac{-1}{\tau s^2\Omega}\left(\left[\frac{(E_s^{1/m}+c\Omega)^2 - 4E_l^4}{2E_l^2}\right]\delta E_l + \left[\frac{2c\Omega E_l^2 - E_s^{1/m} - c\Omega}{2mE_lE_s^{(m-1)/m}}\right]\delta E_s\right) \quad (11)$$

where

$$\tau = \pm\left(4s^2\Omega E_l^2 - 4E_l^4 - (E_s^{1/m}+c\Omega)^2 + 4E_l^2(E_s^{1/m}+c\Omega)c\Omega\right)^{1/2} \quad (12)$$

It is clear from these equations that there will be instabilities as $\Omega \mapsto 0$, $E_l \mapsto 0$, $E_s \mapsto 0$ and as $\tau \mapsto 0$. Since the instability as $\tau \mapsto 0$ occurs when $\hat{s}$, $\hat{n}$ and $\hat{k}$ are coplanar and $\gamma = 0$, it is non-generic in the sense that any small perturbation of $\hat{n}$ will move one away from the locus of instability. The instability as $\Omega \mapsto 0$, however, cannot be avoided as this case occurs when $\hat{k}$ and $\hat{s}$ are parallel and occurs independently of the value of $\hat{n}$. In this situation information about only one component of $\hat{n}$ is available. Similarly the instabilities as $E_l \mapsto 0$ and $E_s \mapsto 0$ cannot be avoided.

A simple approach to the generation of uncertainty measures is to look at the product of the uncertainty in the inputs with the sensitivity of the computed surface normal to changes in the inputs. That is, we linearize the variation of $\delta\hat{n}$ with respect to δE_s and δE_l to give:

$$\delta\hat{n} = \left(\frac{\partial\hat{n}(E_s, E_l, \hat{s}, \hat{k})}{\partial E_s}\right)\delta E_s + \left(\frac{\partial\hat{n}(E_s, E_l, \hat{s}, \hat{k})}{\partial E_l}\right)\delta E_l \quad (13)$$

where $\frac{\partial\hat{n}}{\partial E_s}$ and $\frac{\partial\hat{n}}{\partial E_l}$ can be determined from equation (5) and the equations for $\delta\alpha, \delta\beta$, and $\delta\gamma$ to give:

$$\frac{\partial\hat{n}}{\partial E_s} = \frac{\left[\hat{k} - \hat{s}c\Omega + \frac{1}{\tau}(\hat{s}\times\hat{k})(2c\Omega E_l^2 - E_s^{\frac{1}{m}} - c\Omega)\right]}{2mE_lE_s^{\frac{m-1}{m}}s\Omega} \quad (14)$$

$$\frac{\partial\hat{n}}{\partial E_l} = \frac{\hat{s}(2E_l^2 + E_s^{\frac{1}{m}}c\Omega + c^2\Omega) - \hat{k}(2E_l^2c\Omega + c\Omega + E_s^{\frac{1}{m}})}{2E_l^2s\Omega} - \frac{(\hat{s}\times\hat{k})((E_s^{\frac{1}{m}}+c\Omega)^2 - 4E_l^4)}{2\tau E_l^2 s\Omega} \quad (15)$$

Our algorithm was run on a synthetic image of a sphere created using the specular and Lambertian reflectance models of equations (1) and (3), where it was assumed that the extraction of the specular and Lambertian components has already been performed, and that we have a module for computing the sign of γ. We took, for the experiment, $m = 15$, $\hat{s} = (-1, 0, 1)/\sqrt{2}$, $\hat{k} = (0, 0, 1)$. It was assumed that the noise in the measurements of the specular and Lambertian image components was white zero mean gaussian with variances $\sigma_s^2 = .05$ and $\sigma_l^2 = .025$. Note that the Lambertian measurement is less noisy than the specular measurement. The algorithm, however, does not use this information in any way, as both measurements are absolutely required and we cannot trade one off for the other. The original shaded image is shown in figure 1. The resulting surface obtained from the algorithm is shown in figure 2. To generate figure 2, we have moved the position of the light source to $\hat{s} = (-1, 1, 1)/\sqrt{3}$ to highlight the errors in the surface normals. Note that the surface has been reconstructed only over a fraction of the surface visible in the image. This is due to the fact that the specular component E_s becomes shadowed before the Lambertian component does. Since both image components are required for the surface reconstruction process, the surface normals are not computed in the specular shadow area, even though the surface is visible due to the Lambertian component. The reconstruction is essentially random noise. The apparent step in the intensity of the reconstructed surface is due to the imposition of the correct sign of γ. This creates a bias in the mean value of the value of the y component of the reconstructed surface normal map, leading to the difference in the mean brightness for $y < 0$ and $y > 0$. The maximum error in the magnitude of the surface normal was 1.22, and the errors were more or less uniformly distributed.

One problem that was encountered in the application of the algebraic algorithm to shape recovery from noisy images was a lack of stability. This lack of stability is due to the fact that the range of the mapping R from surface normals $\hat{n}$ to image components E_s, E_l does not cover the area $(0,1)\times(0,1)$. Thus, there are pairs of E_s, E_l values in the area $(0,1)\times(0,1)$ for which there is no corresponding surface normal. For such (E_s, E_l) pairs the value of γ will be imaginary. In a noise free situation the pathological image component pairs that do not have corresponding surface normals will not be observed. If the values of E_s and E_l are noisy, however, then these values can be observed. It is clear that these pathological values are most likely to arise when the surface normals map to E_s, E_l values near the boundary of the region containing allowable values. The form of the allowable brightness component pairs is described in [2]. In our implementation a heuristic, described in detail in [2], was used to associate a unique surface normal with a pathological image pair. This heuristic has the drawback of allowing the algorithm to behave discontinuously to some smooth variations in the input (i.e. when we have image pairs that are near the boundary of the allowable region).

The shape from shading algorithm we have just described illustrates a major drawback of all algebraic fusional algorithms. In such algorithms *all* sources of information are absolutely required. There is no way to reduce the dependence of the solution on a unreliable source of data. Thus, in situations where one, or more, of the sources of data is very unreliable, the output of the module will likewise be very unreliable. In order to be able to minimize the effect of an unreliable source of data in such a case we must replace the unreliable source of data with some other, more reliable, source of information. Since, in the application that we are examining, the specular and Lambertian image components are the only scene dependent sources of information we have available, any additional sources of information must be in the form of *a priori* (scene independent) constraints.

2 Fusing Specular and Lambertian Information Through Energy Functional Minimization

As we have noted, there is no way, in general, to embed more constraints in an algebraic formulation of a fusional problem than are needed to obtain a unique solution. A Bayesian or energy function formulation of data fusion, however, has the advantage over the algebraic method that we can add additional constraints on the solution. This approach will result in an algorithm which is more robust away from specularities than the algebraic algorithm that was described in the preceding section.

To show how we can implement such an algorithm let us begin by casting the algebraic solution of the previous section as an energy function minimization problem. Consider:

$$\hat{n} \leftarrow \min_{\hat{n}} \int \left[(E_L - \hat{n}\cdot\hat{s})^2 + (E_s^{\frac{1}{m}} - \hat{k}\cdot\hat{h})^2 + \mu(||\hat{n}|| - 1)^2\right] dA \quad (16)$$

This equation is similar to the one described by Horn and Brooks [4]. The third term in the integrand is required to maintain the length of the surface normals near unity.

It is evident, in the noise free case, that the global minimum of this energy functional occurs when $\hat{n}(\hat{x})$ is given by equation (5). Hence the solution obtained through the energy minimization process is the same

as for the algebraic process. The energy function approach, however, has the advantage that it is possible to embed additional constraints on the problem, which can be used to reduce dependence on unreliable data. For example we could include a smoothness term, of the form $||\nabla\hat{n}||^2 = (|\partial\hat{n}/\partial x|^2 + |\partial\hat{n}/\partial y|^2)$ into the energy functional given above.

The smoothness term will reduce the effect of noisy inputs on the resulting surface normal values. In addition, since the smoothness constraint allows the solution of the problem even if only one of the sources of data are present (i.e. either E_s only or E_l only), we can *weight* the relative contributions of the input sources of data. Thus, if we know that one source of data is much more reliable than the other, then we can weight that source more highly than the unreliable source. For example, we want to find $\hat{n}(x,y)$ which minimizes:

$$\int [w_l(E_l - \hat{n}\cdot\hat{s})^2 + w_s(E_s - (\hat{k}\cdot\hat{h})^m)^2 + \mu(||\hat{n}||-1)^2 + \lambda(||\nabla\hat{n}||^2)]dA \quad (17)$$

where w_l and w_s are weights related to the reliability associated with the input data sources E_l and E_s. The weight λ determines the amount of smoothing done on the resulting surface normal field.

The energy based fusion algorithm described above was run on the same synthetic sphere data as was the algebraic algorithm in the previous section. The iterative approach described in [1] is used to minimize the energy functional as follows. First, the Euler equations for the surface normal components associated with the energy functional are found to be:

$$w_s(E_s - R_s(\hat{n}))\frac{\partial R_s}{\partial\hat{n}} + w_l(E_l - R_l(\hat{n}))\frac{\partial R_l}{\partial\hat{n}} + \lambda\nabla^2\hat{n} - \mu\hat{n} = 0 \quad (18)$$

where

$$\frac{\partial R_l}{\partial\hat{n}} = \hat{s} \quad (19)$$

and

$$\frac{\partial R_s}{\partial\hat{n}} = [2\hat{k}(\hat{n}\cdot\hat{s}) + 2\hat{s}(\hat{n}\cdot\hat{k})]m[2(\hat{n}\cdot\hat{k})(\hat{n}\cdot\hat{s}) - (\hat{s}\cdot\hat{k})]^{m-1} \quad (20)$$

We then discretize the Euler equations and rearrange according to the procedure given in [1] to yield the iteration equations:

$$\vec{m}_{ij}^{(k+1)} = \bar{n}_{ij}^{(k)} + \frac{\epsilon^2}{4\lambda}[w_s(E_{s_{ij}} - R_s(\hat{n}_{ij}^{(k)}))\frac{\partial R_s(\hat{n}_{ij}^{(k)})}{\partial\hat{n}} + w_l(E_{l_{ij}} - R_l(\hat{n}_{ij}^{(k)}))\frac{\partial R_l}{\partial\hat{n}}] \quad (21)$$

$$\bar{n}_{ij}^{(k)} = \frac{1}{4}(\hat{n}_{i,j+1}^{(k)} + \hat{n}_{i,j-1}^{(k)} + \hat{n}_{i+1,j}^{(k)} + \hat{n}_{i-1,j}^{(k)}) \quad (22)$$

$$\hat{n}_{ij}^{(k+1)} = \frac{\vec{m}_{ij}^{(k+1)}}{||\vec{m}_{ij}^{(k+1)}||} \quad (23)$$

The purpose of the third equation above is to ensure that the surface normal vector obtained has unit magnitude. Thus, the term involving μ does not appear in the iterative formulation. The indices (ij) index into the spatial lattice of the input arrays E_s and E_l. The index k is an iteration counter.

The above iterative method requires that boundary conditions on the solution be specified. These are typically obtained from the occluding contour and the self-shadow line. If the reflectance function at the self-shadow line is Lambertian only, then we have that $\hat{n}\cdot\hat{s} = 0$. This constraint, along with the constraint imposed by the perpendicular to the self-shadow line in the image plane is enough to specify the surface normal along the self-shadow line (this assumes that the image plane is not parallel to the illumination direction, and that the self-shadow boundary lies in a plane that is perpendicular to the illumination direction [5]). We will ignore, for the purposes of this paper, the problem of distinguishing self-shadow lines from occluding contours.

An important aspect to consider in the implementation of the iterative solution to the Euler equation is the determination of the weights, w_s and w_l. Ideally, we would like w_s and w_l to reflect our confidence in the specular and Lambertian data sources. In practice, however, the value of these weights relative to the λ term affect both the amount of smoothness applied to the reconstructed surace normal field and the stability of the iterative process. As the value of λ decreases the increment added to $\bar{n}_{ij}^{(k)}$ in the iteration equation becomes larger. If λ is too small the iteration will become unstable and convergence will not be attained. If, on the other hand, the value of λ is too large, the increment added to $\bar{n}_{ij}^{(k)}$ will be negligble and the resulting solution will depend only on the boundary conditions. Thus, we would like to adjust λ so that it is small enough to prevent excessive smoothing of the surface, but still large enough to ensure that the iteration process converges. If we had only one data source, the determination of a suitable value for λ would be fairly straighforward; a trial and error process could be used, at the very least, to determine, for a given weight on the data term, the minimum value of λ for which the iteration converges. With two or more sources of data, however, the problem of determining λ is complicated by the presence of multiple terms being added to $\bar{n}_{ij}^{(k)}$. We must ensure that each of these terms always be less than a certain level, else instability will arise. The simplest way to do this is to set the weights on each of the terms (i.e. w_s and w_l) to be equal to the inverse of the maximum possible value of the data dependent update term, and then determine λ correspondingly (i.e. by trial and error given these w_s and w_l values). That is, in our case we would have:

$$w_s = \frac{1}{\max|(E_s - R_s)\frac{\partial R_s}{\partial\hat{n}}|}, \quad w_l = \frac{1}{\max|(E_l - R_l)\frac{\partial R_l}{\partial\hat{n}}|} \quad (24)$$

The max operator in the above equations is taken over all possible normal vectors (i.e all unit vectors) and over the three components of the argument of the max operator (i.e. we use the component with largest absolute value). We can determine these weights by assuming that $\max(E_s - R_s) = \max(E_l - R_l) = 1$ and differentiating $\partial R_s/\partial\hat{n}$ and $\partial R_l/\partial\hat{n}$ with respect to $\hat{n}$ and set the result equal to zero to obtain the maximum value attained by $\partial R_s/\partial\hat{n}$ and $\partial R_l/\partial\hat{n}$. Doing so for the illumination parameters of the example we have been using of the sphere yields:

$$w_l = \sqrt{2} = 1.414, \quad w_s = \frac{[\sqrt{2}]^{14}}{15\sqrt{2}} = 6.034 \quad (25)$$

This approach, however, does not allow us the liberty of trading off the Lambertian data over the specular data if it is known that one data source is more reliable than the other. We can, however, use the weights given in the above equations as upper bounds, and use the information regarding the relative reliabilities of the data sources to reduce one of them. For example, if the uncertainty in E_s was twice that of E_l we would want the relative weighting of E_s to be half that of E_l. In order to keep the amount of smoothing as low as possible we would set w_l to be the upper bound give above in equation (24) and set w_s to be half of the upper bound. Hence for our example in which the variance in the noise added to E_s is 0.05, and the variance of the noise added to E_l is 0.025 we would set $w_s = 3.017$ and keep w_l at its upper bound, 1.414. With these weights it was found that λ needed to be at least 35 to ensure stability of the iterative process.

A shaded view of the resulting reconstructed surface is shown in figure 3. Compare this surface to that produced by the algebraic algorithm. Note that we have surface normal values over the entire illuminated portion of the sphere, as the smoothness constraint along with the Lambertian component is sufficient to solve for the surface normals in the specular shadow region. The maximum surface normal error magnitude, after 1100 iterations, is about 0.73, corresponding to an angular error of about 36 degrees. The average error is about 0.085, corresponding to an average angular error of about 5 degrees. Observe that the largest part of the error in the resulting surface surface normals occurs near the specular shadow boundary, where the specular data is mostly noise. Note the distortion of the specularity (which for the second light source direction happens to lie near the specular shadow line for the first light source position). The best reconstruction of the surface occurs near the maximum of the specular image.

Whatever the drawbacks of the regularization type approach to the solution of the shape extraction problem, it is clearly better than the algebraic approach. The improvements over the algebraic method are mainly due to the effect of smoothing on the surface normal field, the elimination of the stability problems inherent in the algebraic method and in a very small amount to the relative weighting of the Lambertian data over the specular data. In addition, the surface area over which the surface normal is obtained is greater than in the algebraic method since we can obtain solutions even where there is no specular information, which was not possible in the algebraic method.

In order to improve our surface reconstruction even further we need to use a data fusion algorithm that will control the weighting applied to the input data as a function of its local reliability, rather than only globally applying weighting values.

3 An Adaptive Weighting Approach To Shape From Shading

The data fusion algorithm described in the previous section uses only the information about the reliability of the input data E_s and E_l in determining weights for the two sources of data. It does not use any information regarding the *sensitivity* of the shape from shading algorithm to errors in the inputs. It therefore can have problems in areas of the image where the sensitivity of the shape from shading algorithm to errors in the input data is high (such as near the specular shadow boundary). Ideally we would like to reduce the weighting of the data consistency terms when the sensitivity to errors in the data is high, as well as when the errors in the data are large. We could consider having a separate module whose task was to take in the values of E_s and E_l and compute the shape from shading algorithm sensitivities. These sensitivities would then be fed into the fusional module. The sensitivities would be used by the fusional module to *adaptively* determine weights to be applied to the data consistency terms for E_s and E_l.

The energy function formulation of the adaptive fusion approach to shape from shading is identical to that of the previous section (equation (17)), however, now the weights w_s and w_l are no longer be constant over the image but are now some function of the sensitivities to noise that were derived earlier (i.e. equations 14 and 15). We must still be careful to ensure that the iterative solution to the Euler equation is stable, by proper scaling of the weight values. As in the previous section we do this by setting the maximum value of w_s and w_l to the values given in equation (24) and reduce these weights by an amount related to the relative uncertainty in the data. It is not clear exactly what form should be used to express the weight values in terms of the sensitivities, save for a general rule that the weights should decrease as the noise sensitivities increase. We use the following law:

$$w_s(x,y) = w_{s_{max}} \left[\frac{\min(\sigma_l^2, \sigma_s^2)}{\sigma_s^2}\right] \frac{1}{1+\log\left(1+\left|\frac{\partial \hat{n}(x,y)}{\partial E_s}\right|\right)} \tag{26}$$

$$w_l(x,y) = w_{l_{max}} \left[\frac{\min(\sigma_l^2, \sigma_s^2)}{\sigma_l^2}\right] \frac{1}{1+\log\left(1+\left|\frac{\partial \hat{n}(x,y)}{\partial E_l}\right|\right)} \tag{27}$$

where $\partial\hat{n}/\partial E_s$ and $\partial\hat{n}/\partial E_l$ are given by equations (14) and (15), and where $w_{s_{max}}$ and $W_{l_{max}}$ are the upper bounds given in equation (24). The log function was used to compress the sensitivity functions in order to retain some data dependent terms in the energy function near the boundaries, otherwise the surface would be excessively smoothed near the boundary of the sphere and near the shadow line.

To compare the performance of the adaptive fusion algorithm with the two weakly coupled algorithms described earlier, we ran the adaptive algorithm on the sphere data that was used in the previous sections. The exact iterative equations are used as in the implementation of the energy functional algorithm described in the previous section. In the present case, however, the weights w_s and w_l were functions of position as indicated by equations (26) and (27). A shaded view of the resulting surface is shown in figure 4. Compare this to the results of the previous two methods. The reconstruction is very good. There is significant improvement towards the boundary of the specular zone due to the reduction of w_s in this region. The maximum surface error normal error magnitude for this algorithm was 0.17 after 1500 iterations, corresponding to a maximum angular error of 9.6 degrees. The average surface normal error magnitude was 0.057, corresponding to an average angular error of 3 degrees. The effect of the specular data is felt most strongly near the specularity. To see this, we compared the error map obtained with the adaptive fusion algorithm with the error map (not shown) produced by running the adaptive algorithm with w_s fixed to be zero everywhere (i.e. no influence by the specular component, only the Lambertian component). The error is slightly lower near the specularity when we take the specular data into account than when we do not. The maximum surface normal error magnitude in the Lambertian only case, after 1500 iterations was again 0.17, while the average surface normal error magnitude was 0.062.

4 Discussion

We have presented a data fusion approach to the determination of object surface normals given the specular and Lambertian components of the illumination reflecting from the object surface. We derived a closed form expression for the components of the surface normal vectors as functions of a known light source direction, known surface albedos, and known specular and Lambertian reflectance images.

The sensitivities of the components of the algebraically reconstructed surface normal vectors to perturbations in the values of the specular and Lambertian reflectance images were derived. It was shown that the algebraic surface reconstruction process is ill-posed in the respect that a small change in the input data can result in a large change in the derived surface normal vector. This results from the fact that there can exist, in the presence of noise, pairs of values of the specular and Lambertian images that do not have any corresponding surface normal vector. In order to assign such a surface normal vector it is neccessary to employ a heuristic, such as finding the surface normal vector which has the corresponding theoretical specular/Lambertian reflectance value pair which is nearest to the actual pair of values. The application of this heuristic can cause the aforementioned ill-posedness.

The algebraic fusional method was observed to perform very poorly on noisy synthetic data, and is expected to work poorly in any application in which there is noise present. A fusional method involving the minimization of an energy functional was introduced. This method had the advantage over the algebraic method that the recovered surface covers the whole of the visible portion of the sphere, as reconstruction was possible in the regions containing only a Lambertian image component. In addition, the stability of the energy functional based algorithm was much better than the algebraic method due to the regularizing properties of the smoothness term.

The first energy functional minimization fusion method that was presented weighted the Lambertian and specular data consistency terms, with respect to a smoothing term, uniformly over the image area. This weighting procedure, however, does not take into account that the sensitivity of the surface normal reconstruction process varies over the image. It is preferable to perform more smoothing in regions of the image wherein this sensitivity is high than where it is low. We presented an *adaptively weighted* energy functional minimization based fusion algorithm wherein the weights on the specular and Lambertian data consistency terms were made to be functions of the sensitivity of the algebraic surface normal reconstruction to noise in the specular and Lambertian images. This algorithm reduces the error in the reconstructed surface normal in the regions of low specular and Lambertian reflectance than in the uniform weighting algorithm.

References

[1] Brooks, M.J., and Horn, B.K.P., "Shape and source from shading", MIT AI Lab Memo 820, January 1985

[2] Clark, J.J., and Yuille, A.L., **Data Fusion for Sensory Information Processing Systems,** Kluwer Academic Publishers, Boston, MA, 1990

[3] Gershon, R., **The Use of Color in Computation Vision,** PhD Thesis, Department of Computer Science, University of Toronto, 1987

[4] Horn, B.K.P., and Brooks, M.J., "The variational approach to shape from shading.", MIT AI Lab memo No. 813, March 1985

[5] Klinker, G.J., Shafer, S.A., and Kanade, T., "Color image analysis with an intrinsic reflection model", Proceedings of the 2nd International Conference on Computer Vision, Tampa, Florida, Dec. 1988, pp 292-296

[6] Lee, H-C., "Method for computing the scene-illuminant chromaticity from specular highlights", Journal of the Optical Society of America A, Vol. 3, No. 10, pp 1694-1699, 1986

[7] Shafer, S., "Using colour to separate reflection components", Technical Report TR-136 Computer Science Dept, University of Rochester.

[8] Wolff, L.B., "Using polarization to separate reflectance components", Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, San Diego, 1989, pp 361-369

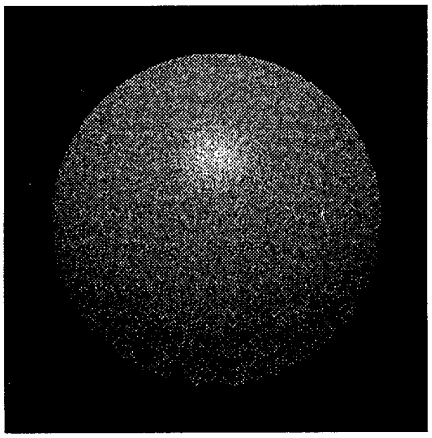

Figure 1: The original image (both specular and Lambertian components).

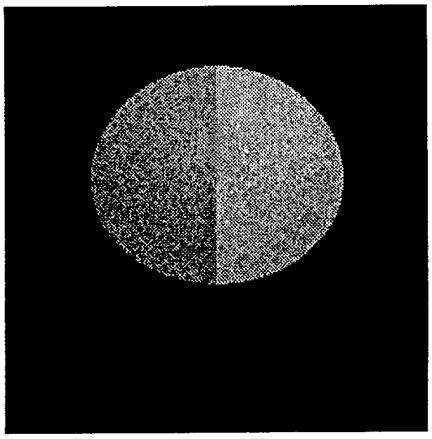

Figure 2: The algebraic reconstruction.

Figure 3: The surface obtained via energy function minimization.

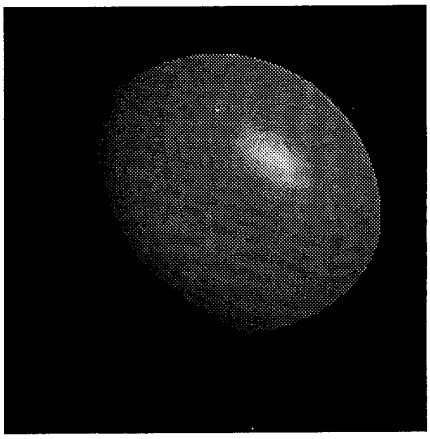

Figure 4: The surface obtained with the adaptive energy minimization scheme.