Shape from Interreflections

Shree K. Nayar, Katsushi Ikeuchi, and Takeo Kanade

The Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

All shape-from-intensity methods assume that points in a scene are only illuminated by sources of light. Most scenes consist of concave surfaces and/or concavities that result from multiple objects in the scene. In such cases, points in the scene reflect light between themselves. In the presence of these interreflections, shape-from-intensity methods produce erroneous (pseudo) estimates of shape and reflectance. The pseudo shape and reflectance estimates, however, are shown to carry information about the actual shape and reflectance of the surface. An iterative algorithm is presented that simultaneously recovers the actual shape and the actual reflectance from the pseudo estimates. The recovery algorithm works on Lambertian surfaces of arbitrary shape with possibly varying and unknown reflectance. The general behavior of the algorithm and its convergence properties are discussed. Both simulation as well as experimental results are included to demonstrate the accuracy and stability of the algorithm.

1 The Interreflection Problem

We address a challenging vision problem that has remained unsolved for the past two decades. Surface elements in a scene, when illuminated, reflect light not only in the direction of the sensor but also between themselves. This is always the case except when the scene consists of only a single convex surface. These *interreflections*, also referred to as *mutual illuminations*, can appreciably alter the appearance of the scene. *None* of the existing vision algorithms reason about, or even take into account, the effects of interreflections. Consequently, interreflections often confuse vision algorithms and cause them to produce erroneous results.

A class of vision algorithms that are particularly affected by interreflections are *shape-from-intensity* algorithms, such as, shape-from-shading [5], photometric stereo [12], and photometric sampling [9]. *All* these methods, are based on the assumption that points in the scene are illuminated only by the sources of light and not other points in the scene; interreflections are assumed *not* to exist. As a result, existing shape-from-intensity methods produce erroneous results when applied to concave surfaces and concavities that result from multiple objects in the scene. As an example, Figure 1a shows a concave Lambertian surface of constant reflectance (albedo = 0.75), and Figure 1b shows its shape extracted using photometric stereo. The inability to deal with interreflections has in the past limited the utility of shape-from-intensity methods.

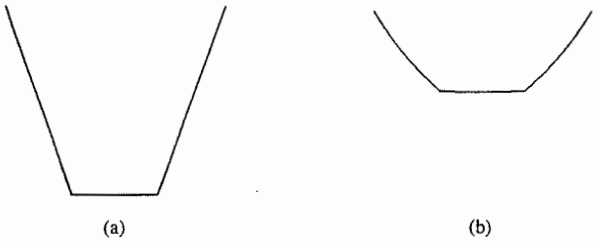

Figure 1: (a) A concave surface. (b) Its shape extracted using photometric stereo.

We identify two separate problems associated with interreflections; the *forward* (graphics) problem and the *inverse* (vision) problem. All previous work done in this area is related to the forward problem. The forward problem, involves the prediction of image brightness values given the shape and reflectance of a scene. Horn [4] discussed the changes in image intensities due to interreflections caused by polyhedral surfaces that are Lambertian in reflectance. Koenderink and van Doorn [7] formalized the interreflection process for Lambertian surfaces of arbitrary shape and varying reflectance (albedo). They proposed a solution to the forward problem in terms of the eigenfunctions of the interreflection kernel. Cohen and Greenberg [1] modeled the scene as a finite collection of Lambertian planar facets and proposed a radiosity solution to the forward problem and used it to render images for graphics. Recently, Forsyth and Zisserman [2] used a similar numerical solution to the forward problem to compare predicted image intensities with experimentally obtained image intensities for a set of Lambertian surfaces of known shape and reflectance.

Our goal here is to solve the inverse (vision) problem. Given image intensities, we wish to recover the shape and reflectance of the scene in the presence of interreflections. The inverse interreflection problem is a particularly difficult one, for in its ambiguity, it resembles the well-known "chicken and egg problem" (Figure 2). If we can model the interreflection effects, we may be able to compensate scene images for these effects and extract accurate shape information. However, it is obvious that modeling interreflections requires prior knowledge of shape and reflectance. But it is shape that we are attempting to recover! So which one comes first, shape or interreflections?

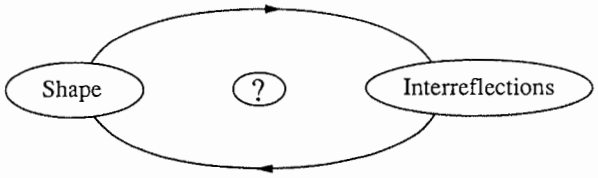

Figure 2: The inter-dependence between shape and interreflections makes shape recovery in the presence of interreflections a difficult problem.

In this paper, we present a solution to the inverse problem for Lambertian surfaces of arbitrary (but continuous) shape, with possibly varying but unknown reflectance (albedo). The shape and reflectance *recovery algorithm* works as follows. First, a *local* shape-from-intensity method is applied to the concave surface to obtain "pseudo" (erroneous) estimates of shape and reflectance. Our solution is based on the observation that the pseudo shape and reflectance, though erroneous, carry information about the actual shape and reflectance of the surface. The pseudo shape and reflectance are used to model the interreflection effects. We show that the pseudo shape is "shallower" than the actual shape and hence exhibits weaker interreflections. These interreflections are used to compensate the pseudo shape and reflectance estimates to obtain a better (more concave) shape estimate and reflectance information. This shape and reflectance is again used to model interreflections to obtain even more accurate shape and reflectance estimates from the pseudo estimates. In this manner, shape and reflectance estimates are iteratively refined to finally converge to the correct shape and reflectance. A detailed analysis of convergence is given for the simple case of two planar surface elements. Convergence for the more general case is discussed and demonstrated by numerous simulation results. Several experimental results are included to demonstrate the robustness, accuracy, and practical feasibility of the proposed algorithm.

2 Modeling Interreflections

Our solution to the inverse interreflection problem is based on the solution to the forward problem; modeling interreflections for a surface of given shape and reflectance. Hence, this section will serve as background theory for subsequent sections. The interreflection model that we describe here is primarily based on the formulation proposed by Koenderink and van Doorn [7]. All surfaces in the scene are assumed to be **Lambertian**. We will shortly see that this assumption is necessary to obtain a closed form solution to the forward interreflection problem. The Lambertian surface can have any arbitrary shape and varying reflectance, i.e. albedo value (ρ) may vary from surface point to surface point. In deriving the interreflection model, we will use radiometric concepts such as irradiance and radiance which are defined in Appendix A.1.

2.1 Analytic Forward Solution

Consider the concave surface $\mathbf{x}(u, v)$ shown in Figure 3a. We are interested in finding the radiance $L(\mathbf{x})$ of the point $\mathbf{x}$ due to the radiance $L(\mathbf{x}')$ of the point $\mathbf{x}'$. The point $\mathbf{x}$ can be illuminated by the point $\mathbf{x}'$ only if the two points can "see" each other. The visibility or *View* function is defined as:

$$View(\mathbf{x}, \mathbf{x}') = \frac{\mathbf{n}.(-\mathbf{r}) + |\,\mathbf{n}.(-\mathbf{r})\,|}{2\,|\,\mathbf{n}.(-\mathbf{r})\,|} \cdot \frac{\mathbf{n}'.\mathbf{r} + |\,\mathbf{n}'.\mathbf{r}\,|}{2\,|\,\mathbf{n}'.\mathbf{r}\,|} \qquad (1)$$

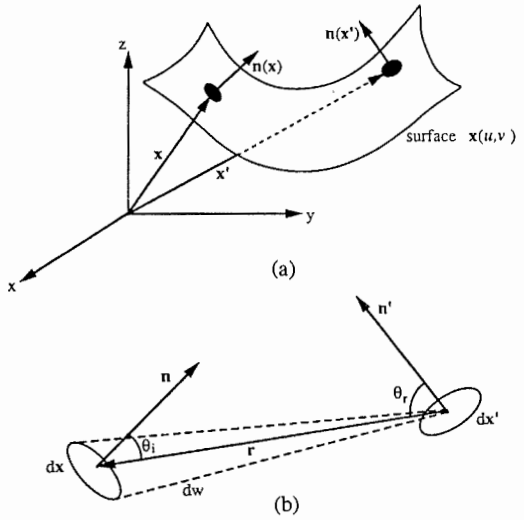

Figure 3: (a) A concave surface in three-dimensional space. (b) Two surface elements that are visible to one another.

where $\mathbf{n}$ and $\mathbf{n}'$ are unit surface normal vectors at the points $\mathbf{x}$ and $\mathbf{x}'$, respectively, and $\mathbf{r}$ is the vector from $\mathbf{x}'$ to $\mathbf{x}$. The function $View(\mathbf{x}, \mathbf{x}')$ equals unity when the two points can illuminate each other and zero otherwise. The radiance of the point $\mathbf{x}$ is related to its irradiance $E(\mathbf{x})$ as:

$$L(\mathbf{x}) = \frac{\rho(\mathbf{x})}{\pi} E(\mathbf{x}) \qquad (2)$$

where $\rho(\mathbf{x})$ is the albedo function, and the factor $\rho(\mathbf{x})/\pi$ is the bi-directional reflectance distribution function (Appendix A.1) for a Lambertian surface. The irradiance $E(\mathbf{x})$ of the surface element $d\mathbf{x}$ due to the radiance of the element $d\mathbf{x}'$ may be derived using the geometry shown in Figure 3b:

$$E(\mathbf{x}) = \left[\frac{[\mathbf{n}.(-\mathbf{r})]\ [\mathbf{n}'.\mathbf{r}]\ View(\mathbf{x}, \mathbf{x}')}{[\mathbf{r}.\mathbf{r}]^2}\right] L(\mathbf{x}')\, d\mathbf{x}' \qquad (3)$$

From equations 2 and 3, we obtain:

$$L(\mathbf{x}) = \frac{\rho(\mathbf{x})}{\pi}\left[\frac{[\mathbf{n}.(-\mathbf{r})]\ [\mathbf{n}'.\mathbf{r}]\ View(\mathbf{x}, \mathbf{x}')}{[\mathbf{r}.\mathbf{r}]^2}\right] L(\mathbf{x}')\, d\mathbf{x}' \qquad (4)$$

From a geometrical perspective, the interreflections between points $\mathbf{x}$ and $\mathbf{x}'$ are governed by the factor:

$$K(\mathbf{x}, \mathbf{x}') = \left[\frac{[\mathbf{n}.(-\mathbf{r})]\ [\mathbf{n}'.\mathbf{r}]\ View(\mathbf{x}, \mathbf{x}')}{[\mathbf{r}.\mathbf{r}]^2}\right] \qquad (5)$$

$K(\mathbf{x}, \mathbf{x}')$ is called the interreflection **kernel** and is a symmetric, positive definite function. It vanishes for surface points that do not illuminate one another either due to their orientations or due to occlusion by other points. The kernel is bounded by the geometrical constraint that no surface element can radiate in, or receive radiations from, more than a half-space around it. Hence:

$$\int\int \left\| \frac{1}{\pi} K(\mathbf{x}, \mathbf{x}') \right\|^2 d\mathbf{x}\, d\mathbf{x}' \leq 1 \qquad (6)$$

Now assume that the concave surface in Figure 3a is illuminated by a single distant point source of light in the direction $\mathbf{s}$. Then, the radiance of the surface due to the source *alone* (excluding interreflection effects) may be expressed in terms of the irradiance of the surface due to the source:

$$L_s(\mathbf{x}) = \frac{\rho(\mathbf{x})}{\pi} E_s(\mathbf{x}) \qquad (7)$$

The irradiance due to the source is the flux incident per unit area of the surface and is proportional to the cosine of the angle between the source direction and the surface normal direction, i.e. $E_s(\mathbf{x}) = k\mathbf{n}.\mathbf{s}$. The constant of proportionality k is determined by the radiant intensity of the source and its distance from the surface.

The total radiance of $\mathbf{x}$ is then expressed as a sum of the radiance due to the source and the radiance due to all other points on the surface:

$$L(\mathbf{x}) = L_s(\mathbf{x}) + \frac{\rho(\mathbf{x})}{\pi} \int K(\mathbf{x}, \mathbf{x}')\, L(\mathbf{x}')\, d\mathbf{x}' \qquad (8)$$

The above equation is referred to as the *interreflection equation*. It is similar in form to the *Fredholm's integral* [7] [4] and does not lend itself to a straightforward solution. However, if all points on the concave surface have the same reflectance ($\rho(\mathbf{x}) = \rho$), a solution to $L(\mathbf{x})$ (or the forward interreflection problem) is given by the Neumann series as:

$$L(\mathbf{x}) = L_s(\mathbf{x}) + \sum_{m=1}^{\infty} \rho^m \int K_m(\mathbf{x}, \mathbf{x}')\, L(\mathbf{x}')\, d\mathbf{x}' \qquad (9)$$

where

$$K_m(\mathbf{x}, \mathbf{x}') = \int \frac{K(\mathbf{x}, \mathbf{y})}{\pi} K_{m-1}(\mathbf{y}, \mathbf{x}')\, d\mathbf{y} \quad (m \geq 2)$$

$$\text{and} \quad K_1 = \frac{K}{\pi}$$

The following observations are made with respect to the above solution:

- It is important to note that the above solution is valid only under the Lambertian assumption. For Lambertian reflectance, the radiance of a surface point is independent of the vantage point. As a result, both $L(\mathbf{x})$ and $L(\mathbf{x}')$ are constants in equation 8 and hence a solution can be obtained.
- The solution is iterative in nature; it is an infinite sum of the kernels K_m that must each be evaluated using the previous kernel K_{m-1}.
- The solution may be interpreted as a mathematical representation of the "ray-tracing" process that is often used in the area of computer graphics. The m^{th} iteration explicitly represents the contribution of the m times interreflected rays.
- Though the Neumann series is an infinite one, the solution is guaranteed to converge to a finite value. This is because $\rho(\mathbf{x}) < 1$ for all surface points, and hence, the series diminishes to zero as m approaches infinity. This is consistent with our real-world experience; diffuse concave surfaces that exhibit interreflections never appear to be infinitely bright.

2.2 Numerical Forward Solution

Discretization of the concave surface leads to a more elegant forward solution than the Neumann series. The following solution has been previously used to render discrete images in graphics [1] and to compare experimentally obtained image intensities with predicted intensities [2]. Let us assume the surface to be comprised of m *facets* as shown in Figure 4. The radiance and albedo values of each facet i are assumed to be constant over the entire facet and equal to the radiance and albedo values at the center point $\mathbf{x}_i$ of the facet, i.e. $L_i = L(\mathbf{x}_i)$ and $\rho_i = \rho(\mathbf{x}_i)$. Then we can write equation 8 as:

$$L_i = L_{si} + \frac{\rho_i}{\pi} \sum_{j \neq i} L_j \int_{S_j} K(\mathbf{x}_i, \mathbf{x}_j)\, d\mathbf{x}_j \tag{10}$$

where S_j is the entire surface of the facet j. Let us assume that the facets are infinitesimally small, i.e. $S_j = d\mathbf{x}_j$. We define the discrete form of the interreflection kernel as:

$$K_{ij} = K(\mathbf{x}_i, \mathbf{x}_j)\, d\mathbf{x}_j \tag{11}$$

where K_{ii} is undefined and K_{ij} vanishes for facet pairs that are not visible to one another. The discrete form of the interreflection equation can therefore be written as:

$$L_i = L_{si} + \frac{\rho_i}{\pi} \sum_{j \neq i} L_j K_{ij} \tag{12}$$

We define the *facet radiance vector* as $\mathbf{L} = [\, L_1, L_2, \ldots\ldots, L_m \,]^T$ and the *source contribution vector* as $\mathbf{L_S} = [\, L_{s1}, L_{s2}, \ldots\ldots, L_{sm} \,]^T$. We also define the *albedo matrix* $\mathbf{P}$ and the *kernel matrix* $\mathbf{K}$ as:

$$\mathbf{P} = \frac{1}{\pi} \begin{bmatrix} \rho_1 & 0 & \ldots & 0 \\ 0 & \rho_2 & 0 \ldots & 0 \\ \ldots & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots \\ 0 & 0 & \ldots & \rho_m \end{bmatrix} \quad \mathbf{K} = \begin{bmatrix} 0 & K_{12} & \ldots & \ldots \\ K_{21} & 0 & \ldots & \ldots \\ \ldots & \ldots & 0 & \ldots \\ \ldots & \ldots & \ldots & 0 \ldots \\ \ldots & \ldots & \ldots & \ldots 0 \end{bmatrix} \tag{13}$$

Then, equation 12 may be written as:

$$\mathbf{L} = \mathbf{L_S} + \mathbf{PKL} \tag{14}$$

or:

$$(\mathbf{I} - \mathbf{PK})\mathbf{L} = \mathbf{L_S} \tag{15}$$

where $\mathbf{I}$ is the identity matrix. Hence, we find that discretization of the surface enables us to obtain a non-iterative, closed-form solution to the forward interreflection problem. The kernel and albedo matrices are determined by the shape and reflectance of the surface, respectively. The source direction and intensity may be used to obtain the source contribution vector $\mathbf{L_S}$. Then the radiance of the surface facets, $\mathbf{L}$, may be determined using the above equation.

Equation 14 explicitly describes the radiance of a facet as the sum of its radiance due to the source and the contributions of other facets.

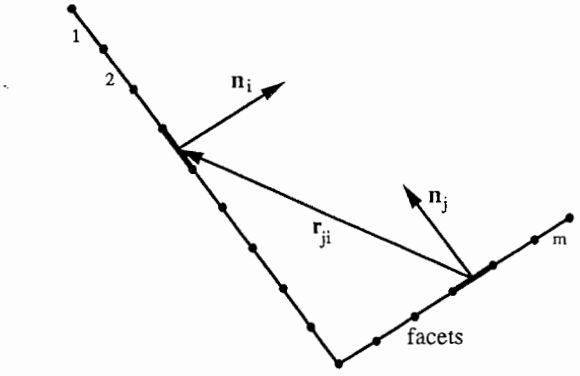

Figure 4: Modeling the surface as a collection of facets, each with its own radiance and albedo values.

Loosely speaking, this may be interpreted as a weighted averaging of radiance values in the direction of concave curvature. As a result of this effect, surface concavities tend to be visually less conspicuous.

We would like to conclude this section with a brief note on the size of individual facets. We have assumed that radiance and albedo are constant over the facet area. This assumption is valid only when the facets are planar and infinitesimally small. While solving the forward interreflection problem, we are free to select appropriate (small) facet sizes. In solving the vision problem, however, we are limited by the resolution of the sensor used to image the scene. The image brightness at a "pixel" location is assumed to be constant over the entire surface facet that the pixel represents. The area $d\mathbf{x}_j$ of the facet may be related to the area dA_j of the pixel as:

$$d\mathbf{x}_j = \frac{dA_j}{\mathbf{v}_j \cdot \mathbf{n}_j} \tag{16}$$

where $\mathbf{n}_j$ and $\mathbf{v}_j$ are the normal vector and viewing direction vector, respectively, for the facet j. For the case of orthographic image projection, the viewing vector is constant over the entire field of view, i.e. $\mathbf{v}_j = \mathbf{v}$. Using equations 5, 11, and 16, we determine the discrete kernel to be:

$$K_{ij} = \frac{[\mathbf{n}_i \cdot \mathbf{r}_{ij}]\ [\mathbf{n}_j \cdot \mathbf{r}_{ji}]}{[\mathbf{r}_{ij} \cdot \mathbf{r}_{ij}]^2} \cdot \frac{dA_j}{\mathbf{v} \cdot \mathbf{n}_j} \cdot View_{ij} \tag{17}$$

The kernel provides a good approximation only when the facets i and j are both infinitesimal and distant from each other. Note the size of a facet also depends on its tilt with respect to the viewing direction. Therefore, while using the discrete kernel we assume that facets are not viewed at angles close to the grazing angle. The above kernel represents interreflections between two facets positioned and oriented in three-dimensional space. In Appendix A.2, we have also included the kernel form for the special case of three-dimensional surfaces that have single translational symmetry.

3 The Extracted "Pseudo" Shape

The previous section showed that surface radiance values are affected by the presence of interreflections. This indicates that if a shape-from-intensity method is applied to a concave surface it is expected to produce erroneous estimates of shape. In order to generalize the inverse interreflection problem that we are attempting to solve, we assume that the reflectance of the Lambertian surface is also *unknown* and may vary from point to point. Therefore, by the term shape-from-intensity, we mean those methods that extract *both* shape (orientation) and reflectance (albedo) information. Photometric stereo [12] and photometric sampling [9] are examples of such shape-from-intensity methods. In the presence of interreflections, these shape-from-intensities methods produce erroneous shape as well as erroneous reflectance information. We refer to the extracted shape as the "pseudo shape" and the extracted reflectance as the "pseudo reflectance" of the surface. In this section, we investigate how the pseudo shape and reflectance are related to the actual shape and reflectance of the surface.

Once again, consider the surface comprised of m facets (Figure 4). Each facet i is defined as:

$$\mathbf{N}_i = \frac{\rho_i}{\pi} \mathbf{n}_i \tag{18}$$

where $\mathbf{n}_i$ and ρ_i are the unit surface normal vector and the albedo for the facet i, respectively. Therefore, the term "facet" represents both local shape and local reflectance information. The complete surface is then defined by the *facet matrix* $\mathbf{F} = [\, \mathbf{N}_1, \mathbf{N}_2, \ldots\ldots, \mathbf{N}_m \,]^T$. From the interreflection equation (equation 15) we can express the surface radiance as:

$$\mathbf{L} = (\mathbf{I} - \mathbf{PK})^{-1} \mathbf{L_S} \tag{19}$$

Since the surface is Lambertian, the source contribution vector $\mathbf{L_S}$ may be determined from the facet matrix $\mathbf{F}$ and the source direction vector $\mathbf{s} = [s_x, s_y, s_z]^T$ as:

$$\mathbf{L} = (\mathbf{I} - \mathbf{PK})^{-1} \mathbf{F} . \mathbf{s} \tag{20}$$

Now let us examine the result of applying photometric stereo to the surface. Three source directions, $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_3$, are used sequentially to illuminate the surface. We assume that all three sources are visible to

all facets on the surface. The three resulting surface radiance vectors L_1, L_2, and L_3 may be expressed as:

$$[\,L_1\,,\,L_2\,,\,L_3\,] \;=\; (\mathbf{I} - \mathbf{PK})^{-1}\,\mathbf{F}\,.\,[\,s_1\,,\,s_2\,,\,s_3\,] \qquad (21)$$

The extracted shape and reflectance information is represented by the *pseudo facet matrix* $\mathbf{F}_p$ and is computed as:

$$\mathbf{F}_p \;=\; [\,L_1\,,\,L_2\,,\,L_3\,]\,.\,[\,s_1\,,\,s_2\,,\,s_3\,]^{-1} \qquad (22)$$

From equations 21 and 22 we find that:

$$\mathbf{F}_p \;=\; (\mathbf{I} - \mathbf{PK})^{-1}\,\mathbf{F} \qquad (23)$$

Note that the albedo matrix **P**, the kernel matrix **K**, and the actual facet matrix **F** are all invariant to the direction and intensity of the illumination. Therefore, the pseudo facet matrix $\mathbf{F}_p$ is also illumination invariant, i.e. the extracted pseudo shape and pseudo reflectance are independent of the source directions used by the shape-from-intensity method to illuminate the object. We will shortly see that this property of the pseudo shape and reflectance enables us to recover actual shape and reflectance of the surface.

The i^{th} *pseudo facet*[1] in $\mathbf{F}_p$ may be written as:

$$\mathbf{N}_{p_i} \;=\; \frac{\rho_{p_i}}{\pi}\,\mathbf{n}_{p_i} \qquad (24)$$

where $\mathbf{n}_{p_i}$ and ρ_{p_i} are the pseudo surface normal and the pseudo albedo for the facet i and, in the presence of interreflections, differ from the actual surface normal and actual albedo of the facet. We make a few important observations regarding the pseudo facets:

- From equations 23 and 20, we see that the pseudo facets are also Lambertian in their reflectance!
- While the actual albedo values must satisfy the physical constraint $\rho_i < 1$, the pseudo albedo values tend to be greater than the actual values and for actual albedo values close to unity, the pseudo albedo values may even exceed unity.
- The pseudo orientations may be described as a result of the weighted averaging of actual orientations in the direction of concave curvature. Qualitatively speaking, for concave surfaces the pseudo shape may be viewed as a smoothed version of the actual shape and appears to be "shallower" than the actual shape (see Figure 1).

4 Recovering Actual Shape and Reflectance

Our objective is to *simultaneously* recover both actual shape and actual reflectance of the surface from the extracted pseudo shape and reflectance. The method is based on the observation that the pseudo shape and reflectance carry information about the actual shape and reflectance. A closed-form solution for the actual shape and reflectance, however, does not seem possible as the pseudo shape and reflectance are highly non-linear functions of the actual ones. Hence, we seek an iterative approach.

4.1 The Recovery Algorithm

Figure 5 illustrates the flow of the algorithm. At first, a local shape-from-intensity method is applied to the scene. If the scene consists of a single convex surface, the extracted pseudo shape and reflectance are simply the actual ones. However, if the scene consists of concavities resulting from concave surfaces and/or multiple objects in the scene, the pseudo shape and reflectance differ from the actual ones. As we described in the previous section, the pseudo shape is expected to be shallower (less concave but yet concave) version of the actual shape. Hence, the algorithm uses the pseudo shape and reflectance as initial guesses of the actual shape and reflectance, to obtain initial estimates of the albedo matrix **P** and the kernel matrix **K**. It is important to note that the pseudo shape and reflectance serve as conservative initial estimates, in that, they produce interreflections that are greater than zero but less than in the case of the actual shape and reflectance. The estimated **P**, **K**, and the pseudo facets $\mathbf{F}_p$ are then inserted in equation 23 to obtain

[1]These pseudo facets are different from the ones defined by Koenderink and van Doom in [7]

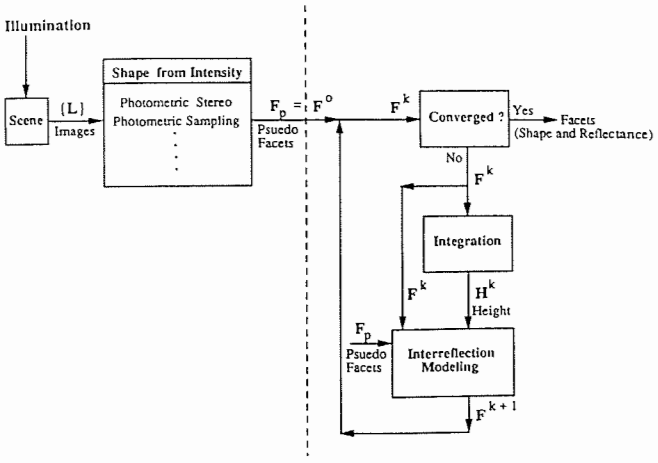

Figure 5: The shape and reflectance recovery algorithm.

the next estimate of the actual facets. This estimate of the surface is expected to be more concave than the previous estimate and is used in the next iteration to obtain an even better estimate. The algorithm may be written as:

$$\mathbf{F}^{k+1} \;=\; (\mathbf{I} - \mathbf{P}^k\mathbf{K}^k)\,\mathbf{F}_p \qquad (25)$$

$$\text{where} \quad \mathbf{F}^0 \;=\; \mathbf{F}_p$$

In the above equation, $\mathbf{P}^k = \mathbf{P}(\mathbf{F}^k)$ and $\mathbf{K}^k = \mathbf{K}(\mathbf{F}^k)$. Note that each set of estimates of the surface facets provides estimates of ***both*** shape and reflectance. With each iteration, more accurate estimates of shape and reflectance are obtained and the result finally converges at the actual shape and reflectance estimates. *The convergence properties of the algorithm will be discussed later.* We now state a few assumptions and observations related to the above algorithm.

- The surface is assumed to be continuous. Note that the interreflection kernel depends not only on the orientations of individual facets but also their relative positions. Therefore, a depth map of the scene must be reconstructed (by integration) from the orientation map computed in each iteration of the algorithm. The continuity assumption is necessary to ensure integrability of the orientation maps. It appears that discontinuities in the depth of scene points can also be handled if this information is provided by a depth measurement method, such as, stereo.
- All facets that contribute to the interreflections in the scene must be visible to the sensor. It is easy to see that if facets that are not visible to the sensor affect the radiance values of the visible facets, the kernel matrix would, in a sense, be incomplete. In such cases, the result produced by the algorithm is difficult to predict but would be close to the desired result if the facets that are not visible do not contribute substantially to the radiance of the visible facets.
- The proposed recovery algorithm may be used in conjunction with *any* local shape-from-intensity method. The shape-from-intensity method used must be capable of computing accurate estimates of both pseudo shape and pseudo reflectance[2]. The recovery algorithm is in no way related to the shape-from-intensity method used to obtain the pseudo shape and reflectance. This fact is emphasized by the dotted line shown in Figure 5.
- No extra images (measurements), in addition to the images used by the shape-from-intensity method, are needed to recover actual shape and reflectance.
- For each iteration of the above algorithm, the kernel is computed for every pair of facets in the scene. Therefore, the algorithm is of $O(M\,n^2)$ complexity, where n is the number of facets in the scene and M is the number of iterations required for shape and reflectance estimates to converge.

[2]We do not include shape-from-shading algorithms in this category as they assume that the surface has constant albedo and that this albedo value is known a-priori.

4.2 Simulation Results

Figure 7 shows simulation results for three-dimensional surfaces that are assumed to have single translation symmetry. The form of the interreflection kernel for this case is given in Appendix A.2. These simulation results are included to give the reader a feel for the behavior of the algorithm. Experimental results for the general three-dimensional case as well as the translation symmetry case will be presented in a later section. For each surface in Figure 7, the numerical forward solution (section 2.2) was used to predict the radiance of the surface from its actual shape and reflectance. Facet radiance values for two different source directions were computed and a photometric stereo algorithm was used to compute the pseudo shape and reflectance estimates. Note that the single symmetry assumption reduces the problem to a two-dimensional one and only two source directions are necessary to compute facet orientations and albedo values. The pseudo shape and reflectance are then used by the recovery algorithm (equation 25) to compute the actual shape and reflectance.

For the surfaces in Figures 7a and 7b, a constant albedo value of 0.75 over the entire surface was used to determine the pseudo shape and reflectance. For the surface shown in Figure 7c, a ramp function that varies from 0.25 to 0.95 was used to compute the pseudo shape and reflectance. For the surface shown in Figure 7d, a checker-board albedo function that varies between 0.3 and 0.7 was used to compute the pseudo shape and reflectance. In Figure 7d, some sections of the surface are occluded from other sections. While computing the interreflections, the algorithm uses geometrical reasoning to determine if two facets on the surface are occluded from each other by other facets. In all of the above cases, the recovery algorithm did not rely on prior knowledge of surface reflectance, but rather used the pseudo reflectance along with the pseudo shape to recover actual shape and actual reflectance simultaneously.

In Figures 7c and 7d the pseudo reflectances and the intermediate estimates of reflectance are also shown. For these surfaces, a convergence graph is also included that shows the *mean orientation error* $\overline{\theta_e}$, computed using all facets, for each iteration of the recovery algorithm. The orientation error at a surface point is defined as the absolute angle (in degrees) between the actual normal vector and the estimated normal vector. For all of the above surfaces, and numerous unreported simulation results, fairly accurate shape estimates are usually obtained in about 7 iterations of the algorithm.

5 Convergence

In this section we study the convergence properties of the shape and reflectance recovery algorithm. Since the interreflection equation is non-linear in the shape and reflectance of the surface, it is difficult to prove the convergence of the algorithm for the general case of arbitrary shape and reflectance. Therefore, we start by analyzing the convergence properties for the simplest case of two planar facets and later extend our analysis to the general case.

5.1 Two-Facet Case

Consider two infinitesimal planar facets of equal size that are separated by a distance r (Figure 6a). We identify an axis (dotted line) with respect to which the two facets are symmetrically oriented and positioned. The unit normal vectors $\mathbf{n}_1$ and $\mathbf{n}_2$ are coplanar and therefore are defined by just two parameters, namely, n_y and n_z. As a result of the symmetrical facet arrangement, the interreflection kernels for the two facets are equal, i.e. $K_{12} = K_{21} = \pi K$. Further, we assume that the two facets have equal albedo values, i.e. $\rho_1 = \rho_2 = \rho$. By applying a shape-from-intensity method to the two facets, their pseudo facets can be computed. From equation 23 we see that the actual facets may be expressed in terms of the pseudo facets as:

$$\begin{aligned} \mathbf{N}_1 &= \mathbf{N}_{p1} - \rho K \mathbf{N}_{p2} \\ \mathbf{N}_2 &= \mathbf{N}_{p2} - \rho K \mathbf{N}_{p1} \end{aligned} \tag{26}$$

where $\mathbf{N}_1$ and $\mathbf{N}_2$ are the actual facets and $\mathbf{N}_{p1}$ and $\mathbf{N}_{p2}$ are the pseudo facets. A graphical illustration of the above relation is shown in Figure 6b. If the recovery algorithm is applied to the pseudo facets, the result of the k^{th} iteration may be expressed as:

$$\begin{aligned} \mathbf{N}_1^{k+1} &= \mathbf{N}_{p1} - \rho^k K^k \mathbf{N}_{p2} \\ \mathbf{N}_2^{k+1} &= \mathbf{N}_{p2} - \rho^k K^k \mathbf{N}_{p1} \end{aligned} \tag{27}$$

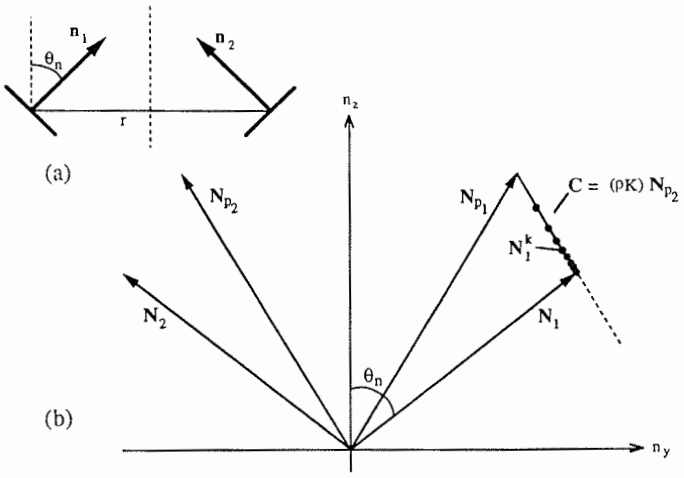

Figure 6: (a) The two-facet case. (b) The line constraint. All intermediate estimates of the facet $\mathbf{N}_1$ must lie on the line passing through the vector C.

where ρ^k and K^k are computed using the intermediate facet estimates $\mathbf{N}_1^k$ and $\mathbf{N}_2^k$. Let us focus our attention on one of the two facets, namely, $\mathbf{N}_1$. Since $\mathbf{N}_{p1}$ and $\mathbf{N}_{p2}$ are constant, new estimates of $\mathbf{N}_1$ result solely from changes in the factor $\rho^k K^k$. Since $\rho^k K^k$ is a scalar, the facet estimates $\mathbf{N}_1^k$ must lie on the line passing through the vector C (Figure 6b). This *line constraint* implies that the convergence of $\mathbf{N}_1^k$ may be studied by analyzing the convergence of $\rho^k K^k$.

We assume that the reflectance estimates ρ^k do not vary substantially from the actual reflectance ρ. This assumption is based on the observation that the pseudo reflectance results from the multiple reflections of light rays between the two facets. This process produces a pseudo reflectance, ρ^0, that maybe expressed as an infinite exponential series in the actual reflectance value ρ. Since the actual albedo must be less than unity, the higher order terms in the series may be neglected and the pseudo albedo is governed by the first few terms. Therefore, for actual albedo values that are not close to unity (say $\rho < 0.75$), the pseudo albedo may be assumed to be close to the actual albedo. Hence, we make the assumption that the pseudo albedo and all intermediate estimates of albedo in the recovery process do not vary substantially from the actual albedo value, i.e. $\rho^k \approx \rho$. Therefore, variations in the factor $\rho^k K^k$ are dominated primarily by variation in K^k.

From the geometry shown in Figure 6a, we see that the relative orientation of the two facets is determined by the tilt angle θ_n. The interreflection process, in a sense, tends to make the orientation of each facet more like that of the other facet. In other words, the pseudo facets are guaranteed to have a smaller tilt angle than the actual facets (Figure 6b). Further, the interreflection kernel K is a *monotonic* function of the tilt angle θ_n; K increases from zero to unity as θ_n varies from zero to $\pi/2$ [10]. Therefore, the first estimate of the kernel, namely, K^0, is less than the actual kernel K but yet greater than zero. Consequently, the facet estimate $\mathbf{N}_1^1$ has a greater tilt angle than the previous estimate but less than that of the actual facet. With each iteration, therefore, the kernel estimates increase in value and approach the actual kernel value, i.e. $K^0 \leq K^k \leq K^{k+1} \leq K$. Equivalently, the facet estimates, $\mathbf{N}_1^k$, start from $\mathbf{N}_{p1}$ and move along the vector C to finally converge at $\mathbf{N}_1$.

The above argument holds for all facets whose pseudo albedo values do not vary substantially over the vector C. Figure 8a shows a convergence map for the two-facet case. Each point on the map corresponds to an instance of the actual facet vector $\mathbf{N}_1$; the tilt angle of the facets vary along the concentric circles and the albedo of the facets vary radially. For all instances of $\mathbf{N}_1$, the two facets are assumed to be separated by a constant distance of $r = 1$. For each instance of $\mathbf{N}_1$, the corresponding pseudo-facet $\mathbf{N}_{p1}$ is computed using the forward interreflection solution and plotted on the convergence map for $\mathbf{N}_{p1}$ shown in Figure 8b. The recovery algorithm is independently applied to each pseudo facet and a *facet tilt error* is computed at the end of 100 iterations. The small dots in Figure 8a correspond to those actual facets for which the algorithm successfully recovers the actual facet from the corresponding pseudo facet. The large dots correspond to the facets for

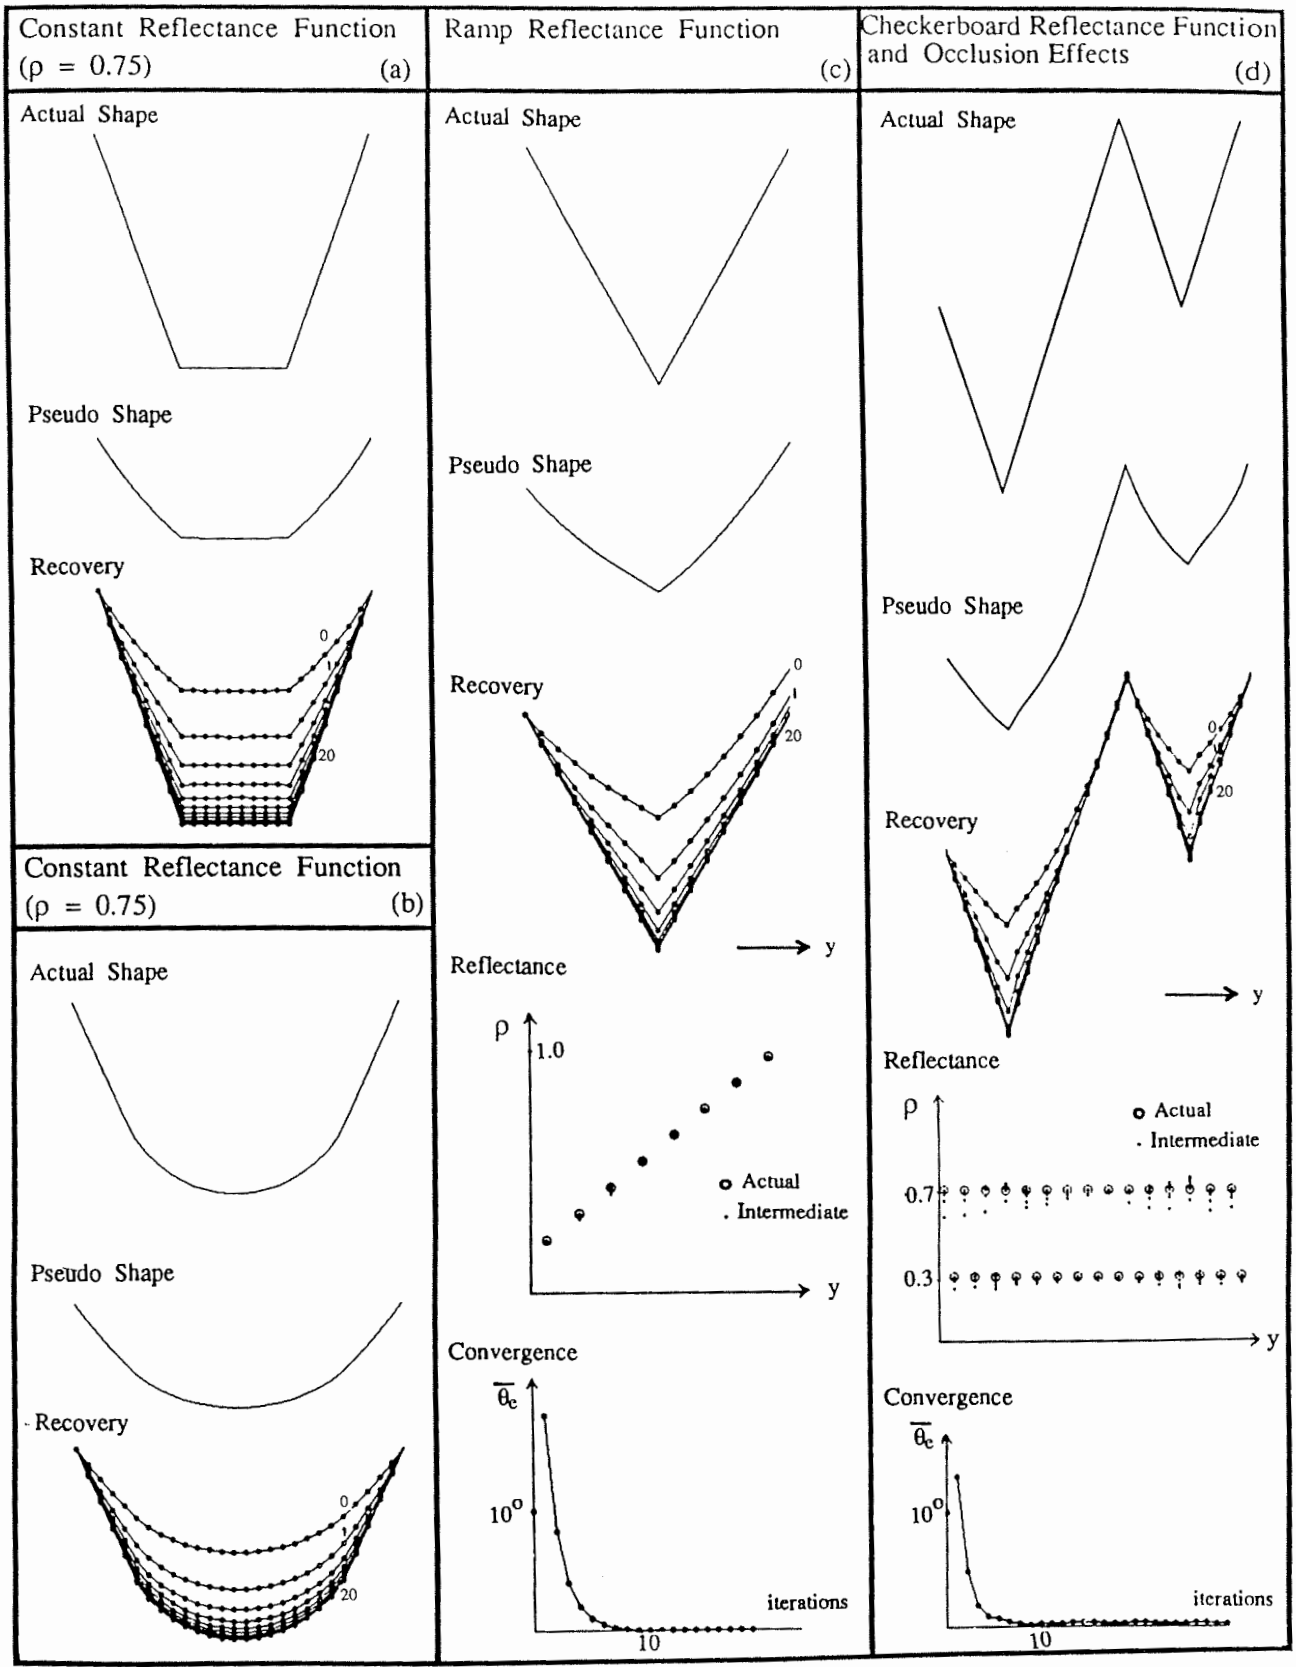

Figure 7: **Simulation Results:** For each surface, the pseudo shape and pseudo reflectance are computed from the actual shape and actual reflectance using the forward solution (section 2.2). The recovery algorithm is applied to the pseudo shape and reflectance to recover the actual shape and reflectance.

which the algorithm does not converge at the actual facet but rather at some other point. Similarly, the small dots and large dots in Figure 8b correspond to the pseudo facets that do and do not produce accurate actual facet estimates, respectively.

Note that the algorithm fails to produce accurate estimates for points that have large albedo values ($\rho > 0.75$) and large tilt angles ($\theta_n > 70$ degrees). For points in this range (large dots), the algorithm does in fact converge, but it converges to other facet vectors that lie in between the pseudo facet and the actual facet on the vector C (Figure 6b). This results from the fact that facets with two different tilt angles and two different albedo values can produce the same pseudo facets! This is also illustrated in the convergence map for the pseudo facets (Figure 8b) where constant albedo curves intersect one another, thus implying that two different actual facets can result in the same pseudo facet. Note that each facet separation distance, r, results in a different set of convergence maps.

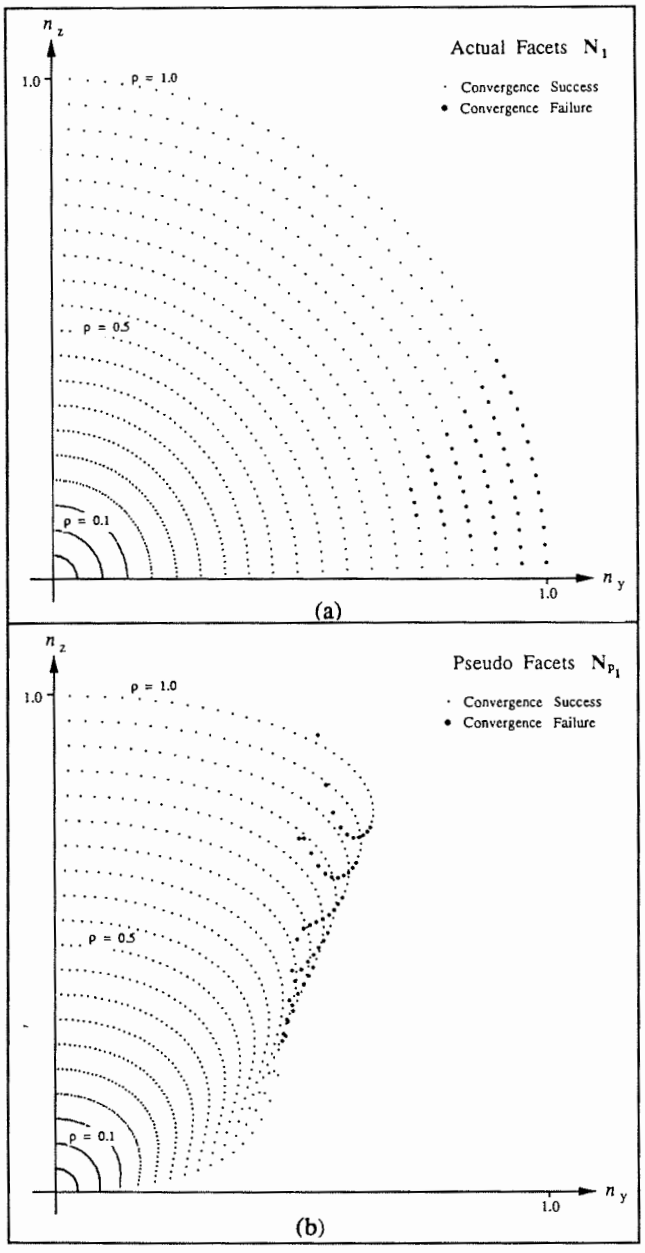

Figure 8: (a) Convergence map for actual facets N_1. $r = 1$. (b) Convergence map for pseudo facets N_{p1}. $r = 1$.

5.2 General Case

Convergence for the general case of surfaces of arbitrary shape (more than two facets) and with varying reflectance is very difficult to prove. This is because the algorithm iterates on a system of equations that are non-linear functions of the shape and reflectance of the surface. Hence, the most convincing evidence of the stability and accuracy of the algorithm lies in the the simulation results of the previous section and the experimental results to follow. However, fairly convincing qualitative statements may be made regarding the general behavior of the algorithm. The algorithm uses the pseudo shape and reflectance as an initial estimate of the actual shape and reflectance. Since the pseudo shape is less concave than, but yet resembles in shape to, the actual shape, the interreflections produced by the pseudo shape are less than in the case of the actual surface but yet greater than zero. Modeling of these non-zero interreflections is expected to produce a better estimate of shape and reflectance in the next iteration. As the shape and reflectance estimates approach the actual ones, the pseudo shape and reflectance are almost perfectly accounted for and the algorithm begins to converge. Just as in the two-facet case, two different surface shapes with two different albedo functions can produce the same pseudo shape. However, we are using both the pseudo shape as well as the pseudo reflectance together to estimate actual shape and reflectance. Hence, the probability of the algorithm converging at intermediate shape estimates is small though not zero. In general, however, lower albedo values produce lesser interreflections and therefore higher likelihood of convergence. Further, in all our simulations and experiments we confine ourselves to facets whose tilt angle are less than 70 degrees. This is because larger tilt angles produce larger facet areas that result in poorer approximations to the interreflection kernel. In some cases, such poor approximations can lead to instability; the algorithm may produce grossly erroneous intermediate estimates that finally lead to divergence, rather than convergence, of shape and reflectance.

6 Experimental Results

We have conducted experiments to demonstrate the accuracy and practical feasibility of the shape and reflectance recovery algorithm. The algorithm was applied to general three-dimensional surfaces as well as surfaces that have translational symmetry. A 512×480 CCD camera was used to view the surfaces. Three incandescent lamps were used to illuminate the surfaces from three different directions to extract surface orientation[3] and reflectance information by photometric stereo. The brightness of each source was determined by a calibration procedure that uses flat surfaces of known albedo values.

6.1 Translational Symmetry Case

Figure 9 shows results for objects with translational symmetry in a single direction. Each object is painted with dull white paint to give it a matte (Lambertian-like) reflectance. In each case, a photo of the object is shown and the horizontal line in the photo represents the surface points that were used by the recovery algorithm. The cross-sectional shape (actual shape) of the surface was determined from the known shape of the object. Due to the two-dimensional nature of the problem, only two light source directions were needed to extract pseudo shape and reflectance estimates by photometric stereo. The extracted pseudo albedo value of each facet is represented by a circle in the reflectance graph. The discrete two-dimensional kernel for the translational symmetry case (Appendix A.2) was used by the recovery algorithm to obtain the actual shape and reflectance from the pseudo ones. Intermediate shape estimates are numbered according to the iteration that produced them. For both the surfaces in Figure 9, the shape estimates converge to reasonably accurate estimates within 7 iterations of the algorithm. For each surface, the mean orientation error $\overline{\theta_e}$ (section 4.2) was computed after 25 iterations and was found to be less than 2.5 degrees. For both surfaces, the psuedo albedo values are greater than unity in the areas of maximum concavity. Note that the albedo estimates converge simultaneously with the orientation estimates.

For *convex* surfaces the pseudo shape and reflectance estimates are equal to the actual ones. Since no two facets on a convex surface are visible to one another (*View* = 0), the algorithm converges at the pseudo shape and reflectance estimates [10].

[3]Surface shape is obtained by integration of surface orientation.

6.2 General Case

Figures 10 shows a photo of an inverted pyramid. Again, the surface is painted and has a matte finish . In this case, three light source directions were used to extract pseudo shape and reflectance estimates and the general form of the discrete kernel (equation 17) was used by the recovery algorithm to extract the actual shape and reflectance. Figures 11a and 11d illustrate isometric and front views of the structure of the inverted pyramid in Figure 10. Figures 11b and 11e show the isometric and front views of the pseudo shape of the inverted pyramid extracted by photometric stereo. Figures 11c and 11f show isometric and front views of the shape produced by the recovery algorithm after 10 iterations. The convergence graph for the inverted pyramid is shown in Figure 12. The shape estimate converges in about 7 iterations with a mean orientation error $\overline{\theta_e} \approx 3$ degrees.

6.3 Discussion

The surfaces used in the experiments have high albedo values (approximately 0.9) and thus exhibit strong interreflections. Though surface albedo was not known a-priori, the algorithm was successful in extracting fairly accurate estimates of shape *and* reflectance from the pseudo estimates. Errors in the recovered shape and reflectance estimates are caused by the following factors:

- The presence of noise in the images used by the shape-from-intensity method to extract pseudo shape and reflectance estimates.
- The Lambertian assumption. Though all surfaces used in the experiments have a matte finish, they are not perfectly Lambertian in reflectance. Since the interreflection model used by the recovery algorithm is based on the Lambertian assumption, errors in the shape and reflectance estimates are expected.
- For facets that are very close to one another and are oriented such that they exhibit strong interreflections, the discrete kernel does not provide a good approximation to the actual kernel. For this reason, errors in recovered shape are maximum in areas where the surface is discontinuous in orientation (e.g. edge in Figure 9b).

In our experiments, we have confined ourselves to surface orientations with tilt angles that are less that 70 degrees with respect to the viewing direction. As we pointed out earlier, large tilt angles would result in large facet areas for which the kernel given by equation 17 does not provide reasonable approximations.

7 Conclusion

- We have presented in this paper an algorithm for recovering the shape and reflectance of Lambertian surfaces in the presence of interreflections. The surfaces may be of arbitrary but continuous shape, and with possibly varying and unknown reflectance.
- The actual shape and reflectance are recovered from the pseudo shape and pseudo reflectance estimated by a local shape-from-intensity method. The recovery algorithm may therefore be used in conjunction with *any* shape-from-intensity method that is capable of producing accurate pseudo shape and pseudo reflectance estimates. Thus, the algorithm enhances the capability, and hence also the utility, of existing shape-from-intensity methods.
- We have analyzed in detail the convergence properties of the recovery algorithm for the special case of two planar facets. Qualitative arguments that support convergence for the general case are also given.
- We have included several simulation results as well as experimental results to demonstrate the robustness and accuracy of the algorithm.

In this paper, we have restricted ourselves to Lambertian surfaces. As the appearance of a Lambertian surface is vantage independent, we are able to model interreflections in a compact form. For surfaces with other reflectance properties, however, surface radiance is dependent on the viewing direction. In such cases, both the forward and the inverse interreflection problems are considerably more difficult to solve. As future research we are interested in extending our analysis to specular surfaces, and subsequently, to surfaces with more generic reflectance characteristics.

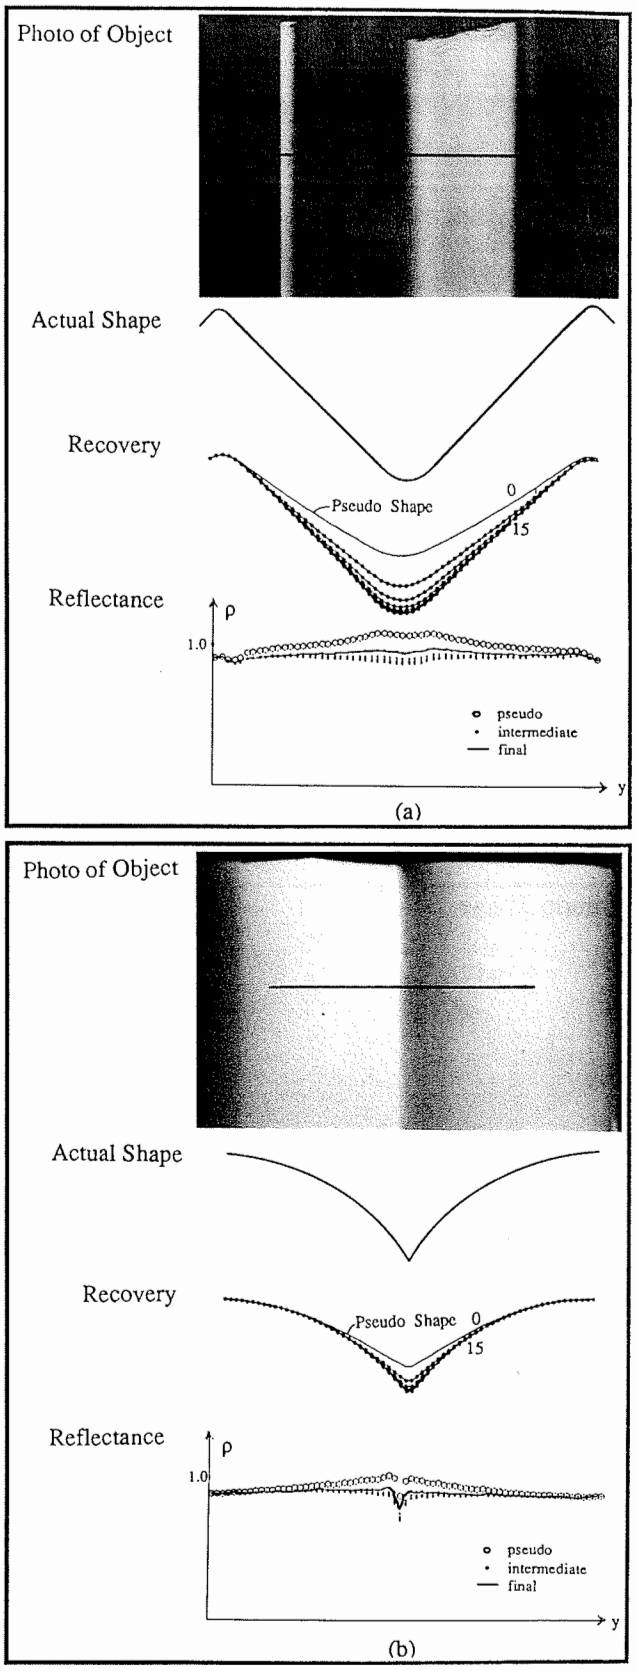

Figure 9: Experimental results for surfaces with translational symmetry in a single direction.

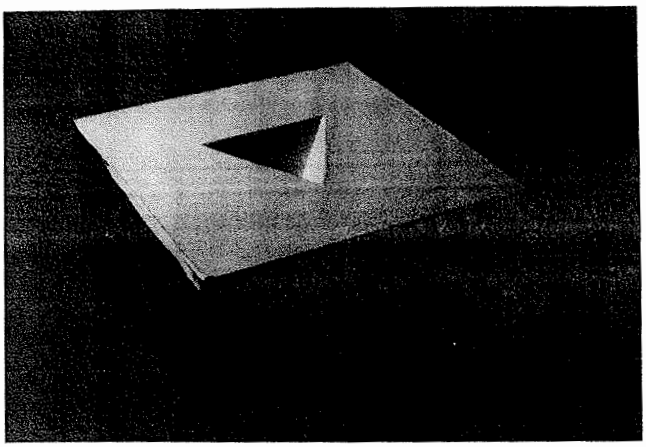

Figure 10: Photo of an inverted pyramid.

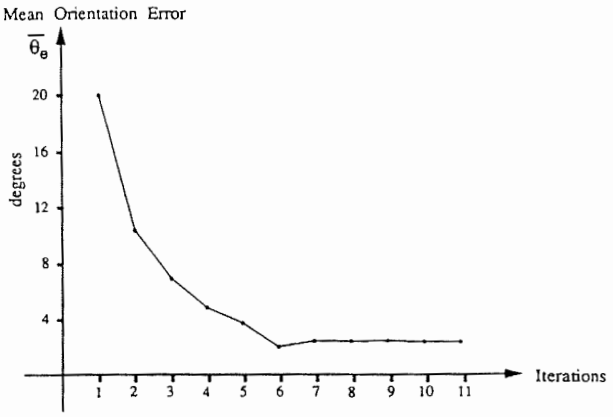

Figure 12: Convergence graph for the inverted pyramid shown in Figure 10.

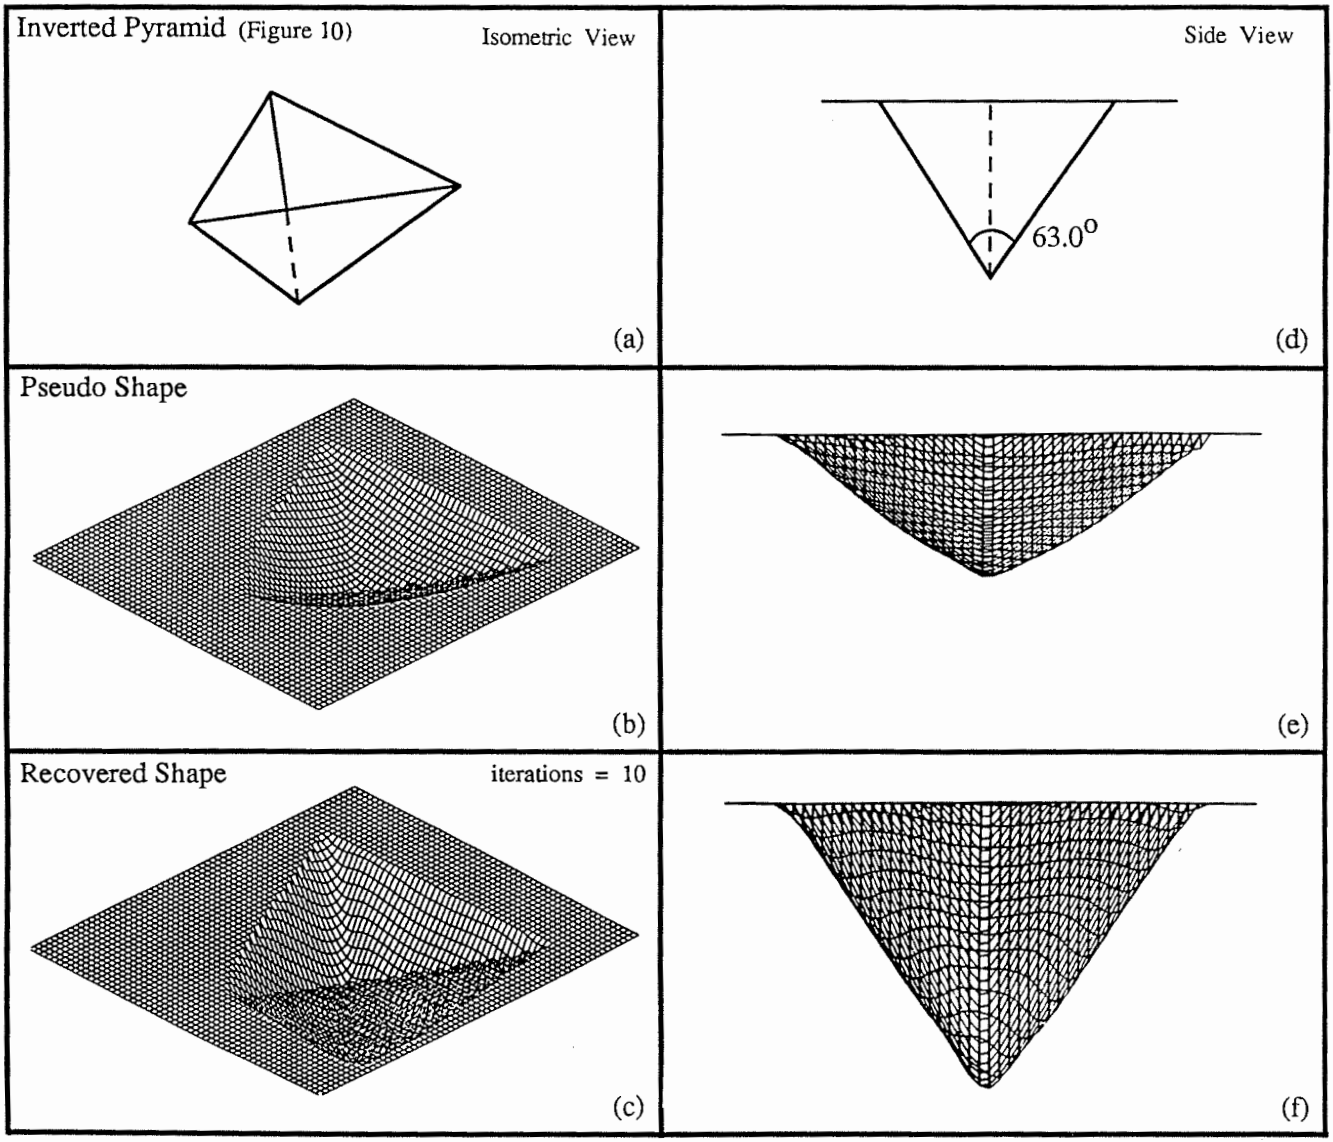

Figure 11: Shape recovery results for the inverted pyramid shown in Figure 10.

A Appendix

A.1 Radiometric Definitions

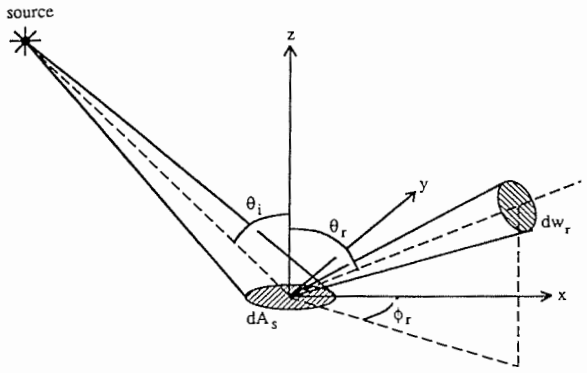

Figure 13: Basic geometry used to define radiometric terms.

We present definitions of radiometric terms that are useful in the analysis of interreflections. Detailed derivations and descriptions of these terms are given by Nicodemus et al. [11]. Figure 13 shows a surface element illuminated by a source of light. The *irradiance* E of the surface is defined as the incident flux density (W/m^{-2}):

$$E = \frac{d\Phi_i}{dA} \tag{28}$$

where $d\Phi_i$ is the flux incident on the area dA of the surface element. The *radiance* L of the surface is defined as the flux emitted per unit fore-shortened area per unit solid angle ($W/m^{-2}.sr^{-1}$). The surface radiance in the direction (θ_r, ϕ_r) is determined as:

$$L = \frac{d^2\Phi_r}{dA\cos\theta_r\, dw_r} \tag{29}$$

where $d^2\Phi_r$ is the flux radiated within the solid angle dw_r. The *Bi-Directional Reflectance Distribution Function* (BRDF) of a surface is a measure of how bright the surface appears when viewed from a given direction and illuminated from another given direction. The BRDF is determined as:

$$f = \frac{dL}{dE} \tag{30}$$

A.2 Kernel for Single Translational Symmetry Case

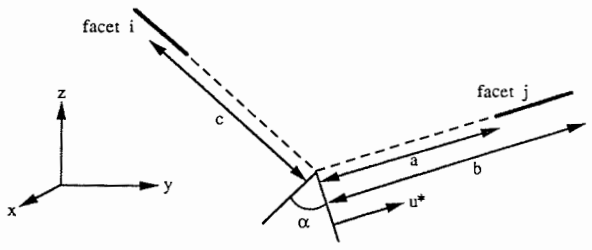

Figure 14: Cross-sectional view of two planar facets that are infinite in the x direction.

Forsyth and Zisserman [2] have derived the interreflection kernel for the special case two planar facets that have translational symmetry in a single direction. Figure 14 shows a cross-sectional view of two such facets that are infinite in the x direction. The kernel K_{ij} is derived [2] by integrating along the x and y directions the contribution of all points on facet j to the radiance of a point on the facet i:

$$K_{ij} = -\frac{\pi}{2}\left[\frac{c + u^* cos\alpha}{\left(c^2 + 2cu^* cos\alpha + u^{*2}\right)^{1/2}}\right]_{u^*=a}^{u^*=b} \tag{31}$$

where α is the angle between the surface normal vectors of the two facets and the parameter u^* respresents the cross-sectional length of the facet j. Since both facets are infinite in length, the same kernel is valid for all points on the facet i. Therefore, under the translation symmetry assumption, the kernel is two-dimensional in that it need only be evaluated for points along the cross-section of the surface. Note that the above kernel is valid only for surfaces that are infinite in the direction of symmetry. However, the kernel serves as a good approximation [2] for points that lie around the middle of a surface that is long though finite in the direction of symmetry.

Acknowledgements

The authors are grateful to Hideichi Sato, Carol Novak, Larry Matthies and Su-shing Chen for their valuable comments, and Gowthami Rajendran for reading earlier versions of this paper. The members of the VASC center at Carnegie Mellon University provided many useful suggestions. The experiments reported in this paper were conducted at the Calibrated Imaging Laboratory, Carnegie Mellon University. This research was supported in part by Westinghouse Electric Corporation and in part by DARPA under contract F33615-87-C-1499. Shree K. Nayar was supported by Westinghouse Electric Corporation.

References

[1] M. F. Cohen and D. P. Greenberg, *The Hemi-cube: A Radiosity Solution for Complex Environments*, SIGGRAPH 1985, 19, pp. 31-40, 1985.

[2] D. Forsyth and A. Zisserman, *Mutual Illumination*, Proceedings, CVPR, pp. 466-473, 1989.

[3] A. L. Gilchrist, *The Perception of Surface Blacks and Whites*, Scientific American, bf 240, pp. 112-124, 1979.

[4] B. K. P. Horn, *Mutual Illumination*, MIT AI Lab. Memo-335, August 1975.

[5] B. K. P. Horn, *Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View*, MIT Project MAC Internal Report TR-79 and MIT AI Laboratory Technical Report 232, November, 1970.

[6] J. A. Jacquez and H. F. Kuppenheim, *Theory of the integrating sphere*, Journal of Optical Society of America, Vol. 45, pp. 460-470, 1955.

[7] J. J. Koenderink and A. J. van Doorn, *Geometrical modes as a general method to treat diffuse interreflections in radiometry*, Journal of Optical Society of America, Vol. 73, No. 6, pp. 843-850, June, 1983.

[8] J. H. Lambert, *Photometria sive de mensura de gratibus luminis, colorum et umbrae*, Eberhard Klett, Augsburg, 1760.

[9] S. K. Nayar, K. Ikeuchi, T. Kanade, *Extracting Shape and Reflectance of Hybrid Surfaces by Photometric Sampling*, Proc. of Image Understanding Workshop, Palo Alto, May 1989.

[10] S. K. Nayar, K. Ikeuchi, T. Kanade, *Shape from Interreflections*, CMU-RI-TR-90-14, July, 1990.

[11] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis, *Geometrical Considerations and Nomenclature for Reflectance*, NBS Monograph 160, National Bureau of Standards, October 1977.

[12] R. J. Woodham, *Photometric stereo: A reflectance map technique for determining surface orientation from image intensity*, Proc. SPIE, Vol. 155, pp. 136-143, 1978.