

Recovering Shape in the Presence of Interreflections

Shree K. Nayar

Department of Computer Science, Columbia University, New York, N.Y. 10027

Katsushi Ikeuchi and Takeo Kanade

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

Abstract

All shape-from-intensity methods assume that points in a scene are only illuminated by the sources of light. This assumption is valid only when the scene consists of a single convex surface. Most scenes consist of concave surfaces where points reflect light between themselves. In the presence of these interreflections, shape-from-intensity methods produce erroneous (pseudo) estimates of shape and reflectance. We show that for Lambertian surfaces the pseudo shape and reflectance are unique, and can be mathematically related to the actual shape and reflectance of the surface. An iterative algorithm is presented that simultaneously recovers the actual shape and reflectance from the pseudo estimates. Experiments were conducted to demonstrate the accuracy and robustness of the algorithm.

1 The Interreflection Problem

Points in the scene, when illuminated, reflect light not only in the direction of the sensor but also between themselves. This is always true with the exception of scenes that consist of only a single convex surface, in which case, no two points on the surface are visible to one another. In general, however, scenes include concave surfaces and/or concavities resulting from multiple objects. In such cases, points in the scene reflect light between themselves. These *interreflections* (also referred to as mutual illuminations) can appreciably alter the appearance of the scene. Existing vision algorithms do not account for the effects of interreflections and hence often produce erroneous results.

A class of vision algorithms that are directly affected by interreflections are the shape-from-intensity algorithms, such as, shape-from-shading [4], photometric stereo [11], and photometric sampling [7]. 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, these shape-from-intensity methods produce erroneous results when applied to concave surfaces. 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 been a serious limitation of shape-from-intensity methods.




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

1.1 Forward and Inverse Problems

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 [5] discussed the changes in image intensities due to interreflections caused by polyhedral surfaces that are Lambertian in reflectance. Koenderink and van Doorn [6] 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 [2] 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. More recently, Forsyth and Zisserman [3] used a similar numerical solution to the forward problem to compare predicted and measured image intensities.

Here, our goal 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 "chicken and egg" problem. 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?

1.2 A Solution to the Inverse Problem

We have developed an algorithm that recovers accurate shape information in the presence of interreflections. This solution to the inverse problem is valid for Lambertian surfaces of arbitrary (but continuous) shape, with possibly varying but unknown reflectance (albedo). First, an existing shape-from-intensity method is applied to the concave surface to obtain pseudo (erroneous) estimates of shape and reflectance. We have shown that the pseudo shape and reflectance, though erroneous, can be mathematically related to the actual shape and reflectance of the surface. Further, the pseudo shape has certain interesting properties; it is unique for a given actual shape, and is less concave than the actual shape. A recovery algorithm uses these properties to iteratively recover the actual shape and reflectance from the pseudo shape and reflectance. The algorithm uses a physical interreflection model to compute interreflections in each iteration. Experiments have been conducted to demonstrate the robustness, accuracy, and practical feasibility of the recovery algorithm.

The results presented here demonstrate two points that we feel are vital to vision research. First, interreflection effects can cause shape recovery methods to produce unacceptably erroneous results. Hence, interreflections must not be ignored. Second, interreflection problems (certainly some, if not all) are tractable and solvable.

2 A Physical Interreflection Model

Our solution to the inverse interreflection problem is based on the solution to the forward problem, i.e. modeling interreflections for a surface of given shape and reflectance. The interreflection model described here is primarily based on the formulation proposed by Koenderink and van Doorn [6]. 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 (ρ) may vary from one surface point to the next.

Consider the concave surface shown in Figure 2. The surface is divided into m infinitesimal facets. Let \mathbf{x}_i and $d\mathbf{x}_i$ represent the three-dimensional coordinates and the surface area of the i^{th} facet, respectively. The radiance (brightness) and albedo (reflectance) values of each facet 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 $\rho_i = \rho(\mathbf{x}_i)$. Consider the two facets i and j . The radiance of the facet i due to the radiance of the facet j is determined using basic radiometric definitions [10] as:

$$L_i = \frac{\rho_i}{\pi} K_{ij} L_j \quad (1)$$

where the factor K_{ij} is given by:

$$K_{ij} = \frac{[\mathbf{n}_i \cdot \mathbf{r}_{ij}] [\mathbf{n}_j \cdot \mathbf{r}_{ji}]}{[\mathbf{r}_{ij} \cdot \mathbf{r}_{ij}]^2} \cdot d\mathbf{x}_j \quad (2)$$

K_{ij} is a function of the relative positions and the orientations of the two facets; it determines the interreflections between i and j from a purely geometrical perspective. It is referred to as the interreflection kernel.




Figure 2: Modeling the surface as a collection of facets, each with its own radiance and albedo values.

Now, let us consider the entire surface shown in Figure 2. Assume the surface to be illuminated by a point source of light. Rays of light that impinge upon the surface are reflected between the facets. Since the albedo of each facet is less than unity, the rays of light lose a fraction of their energy with each bounce. Eventually¹, the radiance at each point on the surface converges to a final steady-state value. Hence, the radiance of the facet i may be expressed as a sum of the radiance due to direct illumination from the source and the radiance due to the final radiance values of other facets on the surface:

$$L_i = L_{s_i} + \frac{\rho_i}{\pi} \sum_{j \neq i} L_j K_{ij} \quad (3)$$

where L_{s_i} is the radiance due to direct illumination from the source and the summation term corresponds to the radiance due to mutual illumination. This is the *interreflection equation*. Note that the radiance values L_i and L_j are assumed to be constants in the above equation. It is important to note that this assumption is valid only for Lambertian surfaces; the radiance of a Lambertian surface element is independent of the viewing direction.

The interreflection equation for the complete surface can be written using vector notation. We define the *facet radiance vector* as $\mathbf{L} = [L_1, L_2, \dots, L_m]^T$ and the *source contribution vector* as $\mathbf{L}_s = [L_{s_1}, L_{s_2}, \dots, L_{s_m}]^T$. We also define the *albedo matrix* \mathbf{P} and the *kernel matrix* \mathbf{K} as:

$$\mathbf{P} = \frac{1}{\pi} \begin{bmatrix} \rho_1 & 0 & \dots & 0 \\ 0 & \rho_2 & 0 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & \rho_m \end{bmatrix} \quad \mathbf{K} = \begin{bmatrix} 0 & K_{12} & \dots & \dots & \dots \\ K_{21} & 0 & \dots & \dots & \dots \\ \dots & 0 & \dots & \dots & \dots \\ \dots & \dots & 0 & \dots & \dots \\ \dots & \dots & \dots & 0 & \dots \\ \dots & \dots & \dots & \dots & 0 \end{bmatrix} \quad (4)$$

Then, equation 3 may be written as:

$$\mathbf{L} = \mathbf{L}_s + \mathbf{PKL} \quad (5)$$

or:

$$\mathbf{L} = (\mathbf{I} - \mathbf{PK})^{-1} \mathbf{L}_s \quad (6)$$

¹After possibly an infinite number of bounces.

where \mathbf{I} is the identity matrix. Thus, we have obtained 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 compute the source contribution vector \mathbf{L}_s . Then the radiance of the surface facets, \mathbf{L} , can be determined using the above equation.

3 The Extracted Pseudo Shape

From equation 6 it is clear that surface radiance values are affected by 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, 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 local methods that extract both shape (orientation) and reflectance (albedo) information. Photometric stereo [11] and photometric sampling [9] are examples of such shape-from-intensity methods². In the presence of interreflections, these shape-from-intensity 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 2). The i^{th} facet may be mathematically represented as:

$$\mathbf{N}_i = \frac{\rho_i}{\pi} \mathbf{n}_i \quad (7)$$

where $\mathbf{n}_i = [n_{xi}, n_{yi}, n_{zi}]^T$ is the unit surface normal and ρ_i is the albedo value for the facet. Therefore, the term "facet" represents both local orientation as well as local reflectance information. The complete surface is then defined by the *facet matrix* $\mathbf{F} = [\mathbf{N}_1, \mathbf{N}_2, \dots, \mathbf{N}_m]^T$. Consider, once again, the interreflection equation given by equation 6. 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:

$$\mathbf{L}_s = \mathbf{F} \mathbf{s} \quad (8)$$

Hence, we obtain:

$$\mathbf{L} = (\mathbf{I} - \mathbf{PK})^{-1} \mathbf{F} \mathbf{s} \quad (9)$$

We define the matrix \mathbf{F}_p as:

$$\mathbf{F}_p = (\mathbf{I} - \mathbf{PK})^{-1} \mathbf{F} \quad (10)$$

We see that \mathbf{F}_p has the same dimensions as the facet matrix \mathbf{F} . In fact, in the absence of interreflections, \mathbf{K} is a null matrix and hence $\mathbf{F}_p = \mathbf{F}$. In the presence of interreflections, \mathbf{F}_p may be viewed as representing another Lambertian surface whose shape and reflectance differ from those of \mathbf{F} . Therefore, if a local shape-from-intensity method is applied to the concave surface,

²We 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.

the extracted shape and reflectance is \mathbf{F}_p and not the actual shape and reflectance given by \mathbf{F} . We refer to \mathbf{F}_p as the *pseudo facet matrix*; it represents the pseudo shape and pseudo reflectance that are extracted in the presence of interreflections.

In practice, the pseudo facet matrix may be computed by using a local shape-from-intensity method. Let us consider, for instance, the photometric stereo method. Three different source directions, \mathbf{s}_1 , \mathbf{s}_2 , and \mathbf{s}_3 , are used sequentially to illuminate the surface. The three resulting surface radiance vectors \mathbf{L}_1 , \mathbf{L}_2 , and \mathbf{L}_3 may be expressed as:

$$[\mathbf{L}_1, \mathbf{L}_2, \mathbf{L}_3] = \mathbf{F}_p \cdot [\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3] \quad (11)$$

The pseudo facet matrix is computed as:

$$\mathbf{F}_p = [\mathbf{L}_1, \mathbf{L}_2, \mathbf{L}_3] \cdot [\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3]^{-1} \quad (12)$$

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

$$\mathbf{N}_{p_i} = \frac{\rho_{p_i}}{\pi} \mathbf{n}_{p_i} \quad (13)$$

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 conclude this section by highlighting three important properties of the pseudo shape and reflectance:

- The pseudo shape and reflectance are *illumination invariant*. In equation 10, note that the albedo matrix \mathbf{P} , the kernel matrix \mathbf{K} , and the actual facet matrix \mathbf{F} are all invariant to the direction and intensity of the illumination. As a result, the matrix \mathbf{F}_p is also illumination invariant. It is independent of sources directions used by the shape-from-intensity method to illuminate the surface!
- The pseudo shape and reflectance are *unique*. From equation 10 we see that the pseudo facet matrix \mathbf{F}_p is dependent on the actual facet matrix \mathbf{F} , the albedo matrix \mathbf{P} , and the kernel matrix \mathbf{K} . Note that \mathbf{P} and \mathbf{K} are in turn determined by \mathbf{F} . Hence, \mathbf{F}_p is dependent only on \mathbf{F} . In other words, there exists only a single pseudo shape and pseudo reflectance corresponding to a given actual shape and reflectance.
- The pseudo shape tends to be *less concave* than the actual shape of the surface. A more formal proof of this property is provided in [9]. Figure 3 illustrates this property through a few examples of actual shapes and pseudo shapes. All the surfaces are assumed to have a constant albedo value, $\rho = 0.95$. The pseudo shapes are computed using equation 10 and are seen to be less concave than the actual shapes.

4 Recovering Actual Shape and Reflectance

We now present an algorithm that *simultaneously* recovers both actual shape and actual reflectance of the surface from the extracted pseudo shape and reflectance. The algorithm is based on the three properties of the pseudo shape and reflectance described in the previous section.

Figure 4 illustrates the flow of the recovery 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




Figure 3: A few actual shapes (with $\rho = 0.95$) and their pseudo shapes.

shape and reflectance are simply the actual ones. However, if the scene consists of concavities, the pseudo shape and reflectance differ from the actual ones. As we showed in the previous section, the pseudo shape is a shallower (less concave) version of the actual shape. Hence, the algorithm uses the pseudo shape and reflectance as conservative initial estimates of the actual shape and reflectance, to model the interreflections and produce estimates for the albedo matrix \mathbf{P} and the kernel matrix \mathbf{K} . The computed \mathbf{P} , \mathbf{K} , and the pseudo facets \mathbf{F}_p are then inserted in equation 10 to obtain the next estimate of the actual facets. This estimate of the surface is expected to be more concave than the previous one and is used in the next iteration to obtain an even better estimate. The algorithm may hence be written as:

$$\mathbf{F}^{k+1} = (\mathbf{I} - \mathbf{P}^k \mathbf{K}^k) \mathbf{F}_p \quad (14)$$

where $\mathbf{F}^0 = \mathbf{F}_p$

In the above equation, $\mathbf{P}^k = \mathbf{P}(\mathbf{F}^k)$ and $\mathbf{K}^k = \mathbf{K}(\mathbf{F}^k)$. Note that each estimate of \mathbf{F} provides estimates of *both* shape and reflectance. With each iteration, more accurate estimates of shape and reflectance are obtained and the result finally converges to the actual shape and reflectance. The convergence properties of the algorithm are discussed in detail in [9]. The following are a few assumptions and observations related to the above algorithm.

- The surface is assumed to be continuous. Note that the inter-reflection kernel depends not only on the orientations of individual facets but also on their relative positions. Therefore, a depth map of the scene must be reconstructed (by integration) from the orientation map computed in each iteration

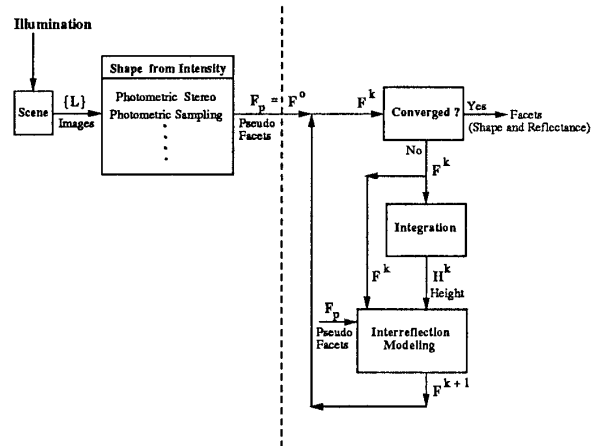


Figure 4: The shape and reflectance recovery algorithm.

of the algorithm. The continuity assumption is necessary to ensure integrability of the orientation maps.

- The proposed recovery algorithm may be used in conjunction with existing shape-from-intensity methods. The shape-from-intensity method used must be capable of computing accurate estimates of both pseudo shape and pseudo reflectance. 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 4.
- 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 $O(Mn^2)$ 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.

5 Experimental Results

Several simulations as well as laboratory experiments have been conducted [9] to demonstrate the accuracy and practical feasibility of the shape and reflectance recovery algorithm. Here, we will present one of our experimental results to illustrate the performance of the algorithm on real surfaces.

Figure 5 shows a photograph of an inverted pyramid. The surface is painted and has a matte (Lambertian-like) finish. Incandescent light sources were used to illuminate the surface from three different directions, and a 512×480 CCD camera was used to obtain images of the surface. The pseudo shape and reflectance of the surface were computed by photometric stereo. The recovery algorithm was used to obtain the actual shape and reflectance from the pseudo estimates. Figures 7a and 7f illustrate isometric and front views of the structure of the inverted pyramid in Figure 5. Figures 7b and 7g show the isometric and front views of the

pseudo shape of the inverted pyramid extracted by photometric stereo. The pseudo shapes are followed by a few intermediate estimates of the shape produced by the recovery algorithm. The convergence graph for the inverted pyramid is shown in Figure 6. The shape estimate converges in about 6 iterations with a mean orientation error³ $\bar{\theta}_e \approx 3$ degrees.

6 Conclusion

We have presented 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 (e.g. photometric stereo). Thus, the algorithm enhances the performance and the utility of existing shape-from-intensity methods.

From the results reported here, two observations can be made that are pertinent to machine vision: (a) interreflections can cause vision algorithms to produce unacceptably erroneous results and hence should not be ignored; (b) interreflection problems (at least some) are tractable and solvable.

In this paper, we have restricted ourselves to Lambertian surfaces. As future research we are interested in extending our interreflection analysis to specular surfaces, and subsequently, to surfaces with more general reflectance characteristics.

Acknowledgements

The authors are grateful to Berthold Horn for his valuable comments, and Gowthami Rajendran for reading earlier versions of this paper. This research was conducted at the VASC center, Carnegie Mellon University.

References

- [1] R. Bajcsy, S. W. Lee, and A. Leonardis "Image Segmentation with Detection of Highlights and Interreflections Using Color," GRASP LAB 182 MS-CIS-89-39, University of Pennsylvania, Dept. of Computer and Info. Science, June 1989.
- [2] M. F. Cohen and D. P. Greenberg, "The Hemi-cube: A Radiosity Solution for Complex Environments," *SIGGRAPH 1985*, 19, pp. 31-40, 1985.
- [3] D. Forsyth and A. Zisserman, "Mutual Illumination," *Proc. of CVPR*, pp. 466-473, 1989.
- [4] 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.
- [5] B. K. P. Horn, "Image Intensity Understanding," MIT AI Lab. Memo-335, August 1975.

³The orientation error at a surface point is defined as the angle (in degrees) between the actual normal vector and the estimated normal vector.

- [6] 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.
- [7] S. K. Nayar, K. Ikeuchi, T. Kanade, "Determining Shape and Reflectance of Hybrid Surfaces by Photometric Sampling," *IEEE Trans. on Robotics and Automation*, Vol. 6, No. 4, pp. 418-431, August, 1990.
- [8] S. K. Nayar, K. Ikeuchi, T. Kanade, "Shape from Interreflections," *Proc. of Third International Conference on Computer Vision*, pp. 2-11, December, 1990.
- [9] S. K. Nayar, "Shape Recovery using Physical Models of Reflection and Interreflection," Ph.D. Dissertation, Department of Electrical and Computer Engineering, Carnegie Mellon University, December, 1990.
- [10] 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.
- [11] R. J. Woodham, "Photometric stereo: A reflectance map technique for determining surface orientation from image intensity," *Proc. SPIE*, Vol. 155, pp. 136-143, 1978.




Figure 5: Photo of an inverted pyramid.




Figure 6: Convergence graph for the inverted pyramid shown in Figure 5.

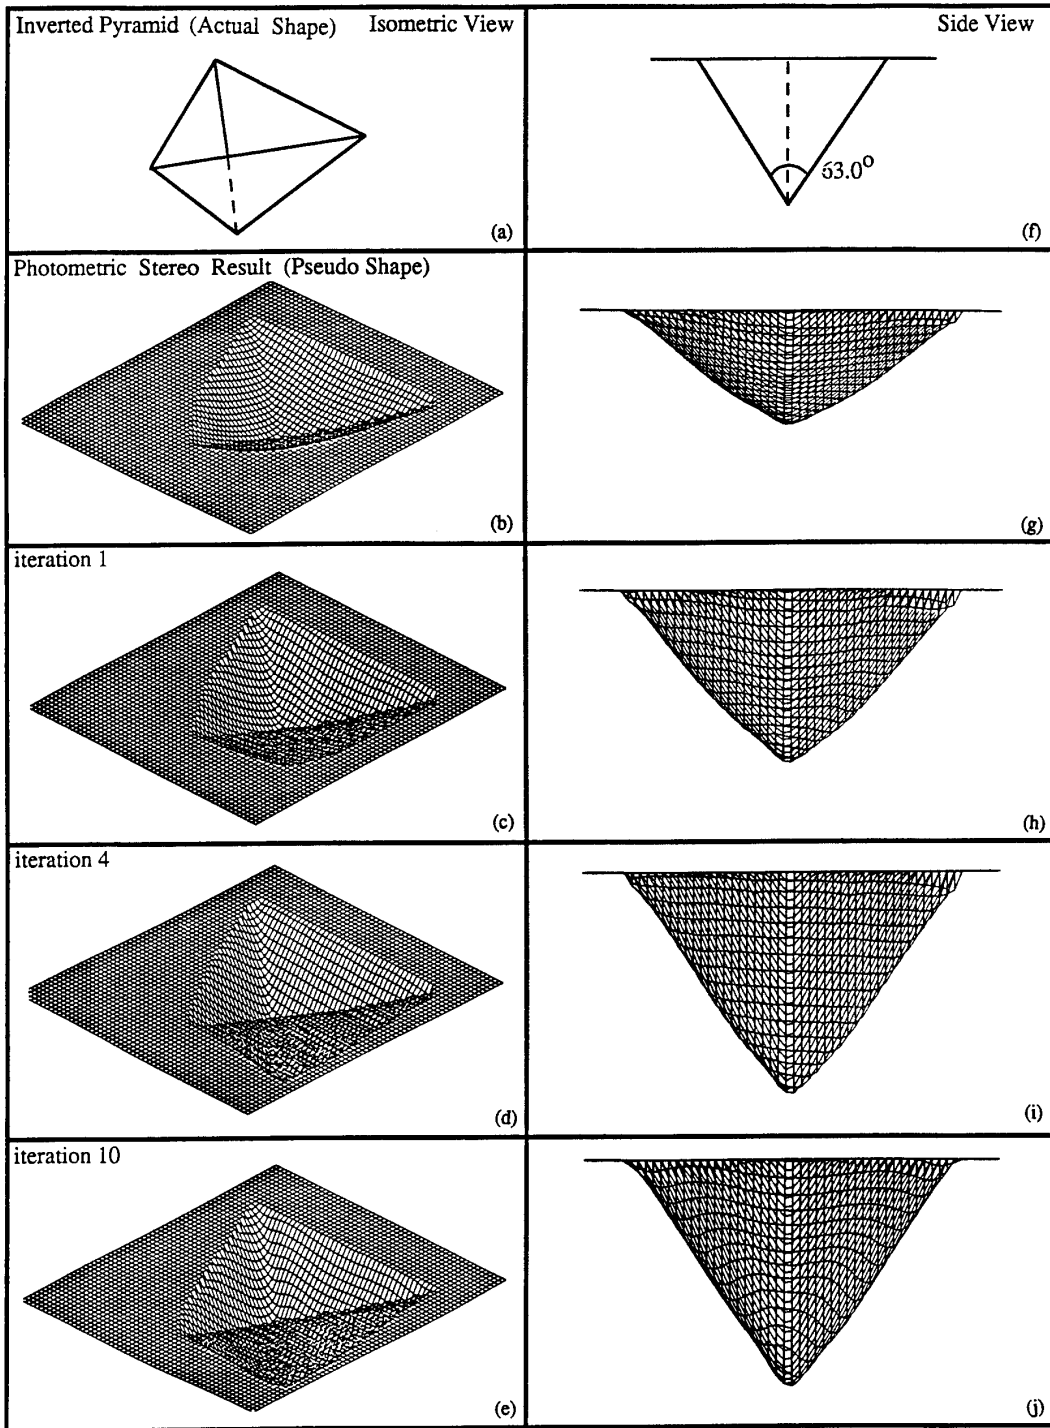


Figure 7: Shape recovery results for the inverted pyramid shown in Figure 5.