ELSEVIER

Image and Vision Computing 15 (1997) 801–814

image AND vision COMPUTING

Iterative shape recovery from multiple images

Ruo Zhang, Mubarak Shah*

Computer Science Department, University of Central Florida, Orlando, FL 32816, USA

Received 30 October 1995; revised 9 January 1997; accepted 30 January 1997

Abstract

In general, shape from shading (SFS) involves the solution of an under-determined system, so it is difficult to always obtain a correct and unique solution. Since only one input image is used, in order to recover the shape of the object as completely as possible, the image has to be taken with careful light source placement in order to illuminate most of the object. Shape from photometric stereo avoids the under-determined problem by using three input images; however, the shape can only be recovered in the areas that are illuminated in all three images. In real life, when a larger sequence of images is available, it is possible to solve these problems.

In this paper, we present a method which recovers shape from a sequence of images taken with different illumination directions. The sequence can be of any length (at least three images), and given in any order, as long as the light source directions are not coplanar. The process can be viewed as cascading shape from photometric stereo, which is formulated in the framework of the linear Kalman filter in order to iteratively recover and refine the shape and the surface albedo. The algorithm can be repeated in multiple cycles over the same sequence of images, which results in an improvement of the recovered shape and albedo. It can also be initiated at any point, stopped at any point, and continued whenever new images arrive. By allowing a longer sequence of input images, the algorithm provides enough information to incrementally recover most of a scene with shadow areas. © 1997 Elsevier Science B.V.

Keywords: Physics-based vision; Shape from photometric stereo; Kalman filter

1. Introduction

Shape from shading (SFS) recovers the 3-D shape of an object from a single image. At each pixel, we know the brightness and the light source direction. The brightness can be described as a function of the surface shape and light source direction according to the Lambertian model. If the surface shape is described in terms of the surface normal, we have a linear equation with three unknowns. If the surface shape is described in terms of the surface gradient, we have a nonlinear equation with two unknowns. In either case, there are more unknowns than equations, therefore, finding a correct and unique solution to SFS is difficult. Another problem in SFS is that shape in shadow regions cannot be recovered due to the lack of shading information, so the image must be taken with careful light source placement in order to illuminate most of the object.

Shape from photometric stereo deals with shape recovery using shading information from *multiple* images illuminated from different directions. It was introduced by Woodham [1] in the early 1980s. In this method, Woodham solved for both the surface normal and the albedo (the ratio of the amount of reflected light over the amount of incoming light) of Lambertian surfaces using three images with non-coplanar light source directions. The use of three input images in photometric stereo reduces the nonlinear shape recovery problem into a linear problem; however, it does not solve the shadow problem. In fact, it makes the shadow problem even worse by enlarging the unrecovered shadow regions, since the shape is only recovered in the areas illuminated in all three images.

Zhang et al. [2,3] used a longer sequence of images as input, and iteratively refined the recovered shape by embedding SFS in the extended Kalman filter framework. The use of more images eliminated unrecovered shadow areas, and made shape refinement possible. However, the algorithm did not deal with varying albedo, and was affected by some shortcomings of the SFS algorithm embedded in it, such as the inaccuracy introduced by linearizing the reflectance function.

Lee and Kuo [4] used cascading photometric stereo on two images with the light source directions orthogonal to each other in order to guarantee overlapping of the gradient directions of one reflectance map with the tangent directions of the other. They applied SFS to the first image to obtain a shape estimate, then used this estimate as an initial value for SFS on the next image in order to improve the

* Corresponding author.

PII S0262-8856(97)00029-2

results. However, their approach did not deal with a long sequence of images, and had the same problems as Zhang et al. [2].

In this paper, we present a new method. Instead of applying cascading SFS, we apply cascading shape from photometric stereo. The process begins as soon as three images become available. When the next image arrives, the first of the three most recently processed images is replaced by the new image to create another trio of images for processing. The output from the previous iteration is used as the input for the next iteration and this process continues until all images are processed. It is not necessary that we always replace the first image in the trio by the new image: however, we do this for convenience (in picking up the next image) and consistency (so that all of the images, except the first two and the last two are processed three times). In our process, we use a linear Kalman filter to iteratively recover and refine the shape and the surface albedo. The estimation process can be performed in multiple cycles (a cycle is one pass over the entire sequence of input images) on the same image sequence by repeatedly processing the sequence to improve the results. This is equivalent to processing a longer sequence of images with each image occurring more than once.

There are several advantages to our approach:

- The use of shape from photometric stereo, rather than SFS, removes the assumption of constant albedo and enables us to easily recover the albedo, as well as the shape.
- The use of a linear Kalman filter guarantees an optimal solution for noisy images.
- By processing the sequence more than one cycle, the convergence of the Kalman filter can be improved, and the results enhanced, especially when the number of images is small.
- Since the images are processed in a cascading manner, the process can be initiated at any point, stopped at any point, and continued whenever new images arrive.
- The algorithm can process an input sequence of any length greater than two, in any order.
- Allowing a longer sequence of input images provides enough information to incrementally recover most of a scene with shadow areas.

Note that when only three images are processed our approach is equivalent to Woodham's photometric stereo algorithm [1]. This will be shown in Section 3.

2. Iterative shape recovery from multiple images

We assume Lambertian reflectance and infinite point source illumination. Let $\vec{L}=(L_x,L_y,L_z)$ be the unit light source direction, $\vec{N}=(N_x,N_y,N_z)$ be the unit surface normal, and ρ be the albedo. According to the Lambertian model, the brightness, I, is proportional to the dot product of the light source, $\vec{L}$, and the surface normal, $\vec{N}$:

$$I=\rho(\vec{N}\cdot\vec{L}) \tag{1}$$

For three images, we have three linear equations in the form of Eq. (1), and can formulate three functions f_1, f_2, f_3 as follows:

$$\begin{aligned} f_1&=\rho(\vec{N}\cdot\vec{L}_1)-I_1\\ f_2&=\rho(\vec{N}\cdot\vec{L})-I_2\\ f_3&=\rho(\vec{N}\cdot\vec{L})-I_3 \end{aligned} \tag{2}$$

where I_i (i = 1,2,3) is the brightness and $\vec{L}_i$ is the light source in the ith image. We can compute $\vec{X}=\rho\vec{N}$ from the above linear system. Since $\vec{N}$ is a unit vector, ρ will be the length of $\vec{X}$, and $\vec{N}$ will be the direction of $\vec{X}$ [1].

A new system of Eq. (2) is formed whenever a new image becomes available. We want to continually refine the solution of $\vec{X}$ by using the estimated $\vec{X}$ from the previous set of images as the initial value for the next set of images. This can be formulated as a linear, discrete time Kalman filtering problem [5].

The Kalman filter is a useful engineering tool that estimates one or more unknown parameters from a set of known measurements. The parameters are estimated and refined through the entire set of measurements based on how good the current measurement is, and how accurate the current estimate is. The estimate from the previous iteration is used together with the new measurement in the current iteration in order to gradually refine the estimate.

In our algorithm, the measurements consist of $\vec{L}_i$ and I_i ($i=1,2,3$). The parameters are represented by a vector $\vec{X}=\rho\vec{N}$, which can be computed recursively by the following Kalman filter:

$$\vec{X}^k=\vec{X}^{k-1}+K(\vec{Y}-M\vec{X}^{k-1}) \tag{3}$$

$$K=S^{k-1}M(W+MS^{k-1}M^T)^{-1} \tag{4}$$

$$S^k=(I-KM)S^{k-1} \tag{5}$$

$$\vec{Y}=\frac{\partial\vec{f}}{\partial\vec{X}}\vec{X}^{k-1}-\vec{f}(\vec{I},L,\vec{X}^{k-1}) \tag{6}$$

$$M=\frac{\partial\vec{f}}{\partial\vec{X}} \tag{7}$$

where S is the 3 by 3 covariance matrix of the estimation error for $\vec{X}$, and

$$\begin{aligned} \vec{f}&=(f_1,f_2,f_3)^T\\ \vec{I}&=(I_1,I_2,I_3)^T\\ L&=(L_1,L_2,L_3)^T\\ W&=\frac{\partial\vec{f}}{\partial\vec{I}}\Lambda\frac{\partial\vec{f}^T}{\partial\vec{I}} \end{aligned}$$

Λ is a 3 by 3 matrix which indicates the covariance of the measurement.

In order to obtain an optimal solution, the Kalman filter requires that noise in the input data be Gaussian with zero mean (white noise). We assume that the noise in real images can be approximated by white noise.

3. Relationship with photometric stereo

In this section, we will show that our method reduces to standard photometric stereo when processing three images.

It can be shown that the Kalman filter obtains a final estimate $\vec{X}$, which minimizes the following criterion [5]:

$$C=(\vec{X}^0-\vec{X})^T S^{0-1}(\vec{X}^0-\vec{X}) + \sum_{i=1}^{3}(Y_i-M_i\vec{X})^T W^{-1}(Y_i-M_i\vec{X}) \quad (8)$$

Here Y_i is the ith component of Y, and M_i is the ith row of the 3 by 3 matrices M. As the covariance matrix of the estimation error S^{0-1} (Eq. (5)) is initialized to zero, and W^{-1} is a positive definite matrix, Eq. (8) is minimized when $Y=M\vec{X}$.

If we consider three input images and an initial estimate $\vec{X}^0$, from Eqs. (2), and (6) we have

$$\begin{aligned} Y &= \frac{\partial \vec{f}}{\partial \vec{X}}\vec{X}^0 - \vec{f}(\vec{I},L,\vec{X}^0) \\ &= S\vec{X}^0 - \vec{f}(\vec{I},L,\vec{0}) \\ &= S\vec{X}^0 - (S\vec{X}^0 - I) \\ &= \vec{I} \end{aligned} \quad (9)$$

From Eqs. (2), and (7) we also have

$$\begin{aligned} M\vec{X} &= \frac{\partial \vec{f}}{\partial \vec{X}}\vec{X} \\ &= S\vec{X} \\ &= S(\rho\vec{N}) \end{aligned} \quad (10)$$

Combining Eq. (9) with Eq. (10), Eq. (8) is minimal when

$$\vec{I}=S(\rho\vec{N}) \quad (11)$$

This shows that our algorithm is equivalent to Woodham's photometric stereo algorithm when the input sequence consists of three images.

4. Results

In this section we will show the input images for each sequence, shaded output from photometric stereo and our algorithm using light source $< 0,0,1 >$, and 3-D plots of the reconstructed shape from photometric stereo, our algorithm, and SFS for comparison. The photometric stereo output was obtained by choosing the three images most suitable for photometric stereo (the light sources are at equal angles from each other). The SFS output was generated by two different algorithms, Bichsel and Pentland's [6] and Lee and Kuo's [4]. These two algorithms were used for comparison simply because according to a recent study they produce the best results among the most popular SFS algorithms [7] for the images used in this paper. The initial condition for Bichsel and Pentland's algorithm is singular points, which are detected automatically in the algorithm as the brightest points. There is no initial condition for Lee and Kuo's algorithm. Note that both SFS algorithms recover depth directly, while in the case of photometric stereo and our algorithm the recovered surface normals are integrated to obtain a depth map, using the algorithm presented in [8].

To further understand the integration, we should look at the algorithm [8] used here. This integration algorithm starts at one surface point $p_0 = (x_0,y_0)$ (we arbitrarily pick the center of the image), and assigns it a fixed depth value. Since the line connecting p_0 and $p_1 = (x_0 + 1,y_0)$ is approximately perpendicular to the average normal between these two points, the dot product of the slope of this line and the average normal should be zero. As the slope is the difference in depth between p_0 and p_1, and we know the normals at these two points, the depth at p_1 can be derived easily. This allows us to propagate the depth along the x-axis. Similarly, we can propagate the depth along the y-axis.

4.1. Synthetic images

The results shown here are for two synthetic Lambertian images: vase and Mozart. The images were generated from range data.

The average recovered albedo error (defined as the absolute difference between the true albedo and the recovered albedo per pixel) for both photometric stereo and our algorithm will be given in the text. The recovered surface normal errors are given in Table 1.

4.1.1. Without noise

Fig. 1 shows the results for a sequence of nine Mozart images. The sequence was generated using light sources with a constant slant of 50° and variable tilt (tilt was changed in steps of 40°, starting at 0). The images were generated with a constant albedo of 0.75. Our algorithm had an average albedo error of 1.2257×10^{-14}; photometric stereo had an average albedo error of 2.5482×10^{-17}.

Fig. 2 shows the results, including the recovered albedo map, for a sequence of eight vase images. The sequence was generated using light sources with a constant slant of 60° and the variable tilt (the tilt was changed in steps of 45°, starting at 0). The images were generated with a constant albedo of 0.75. Our algorithm had an average albedo error of 5.8845×10^{-17}; photometric stereo had an average albedo error of 1.5388×10^{-17}. We can see that both photometric

Table 1
Surface normal error for synthetic images

	Total error		Partial error	
	Ours	Photometric stereo	Ours	Photometric stereo
Vase				
No noise	2.0381×10^{-16}	0.27	4.0137×10^{-16}	2.1721×10^{-16}
$\sigma = 10^{-2}$				
1 cycle	2.5535×10^{-2}	0.30	6.9167×10^{-2}	7.5252×10^{-2}
10 cycles	5.7597×10^{-3}		1.4873×10^{-2}	
20 cycles	4.0285×10^{-3}		1.0317×10^{-2}	
Mozart				
No noise	1.1411×10^{-14}	0.11	1.6514×10^{-13}	5.3128×10^{-14}
$\sigma = 10^{-2}$				
1 cycle	3.8326×10^{-2}	0.14	0.3030	0.3087
10 cycles	1.345×10^{-3}		1.2226×10^{-2}	
20 cycles	1.1493×10^{-3}		8.7178×10^{-3}	

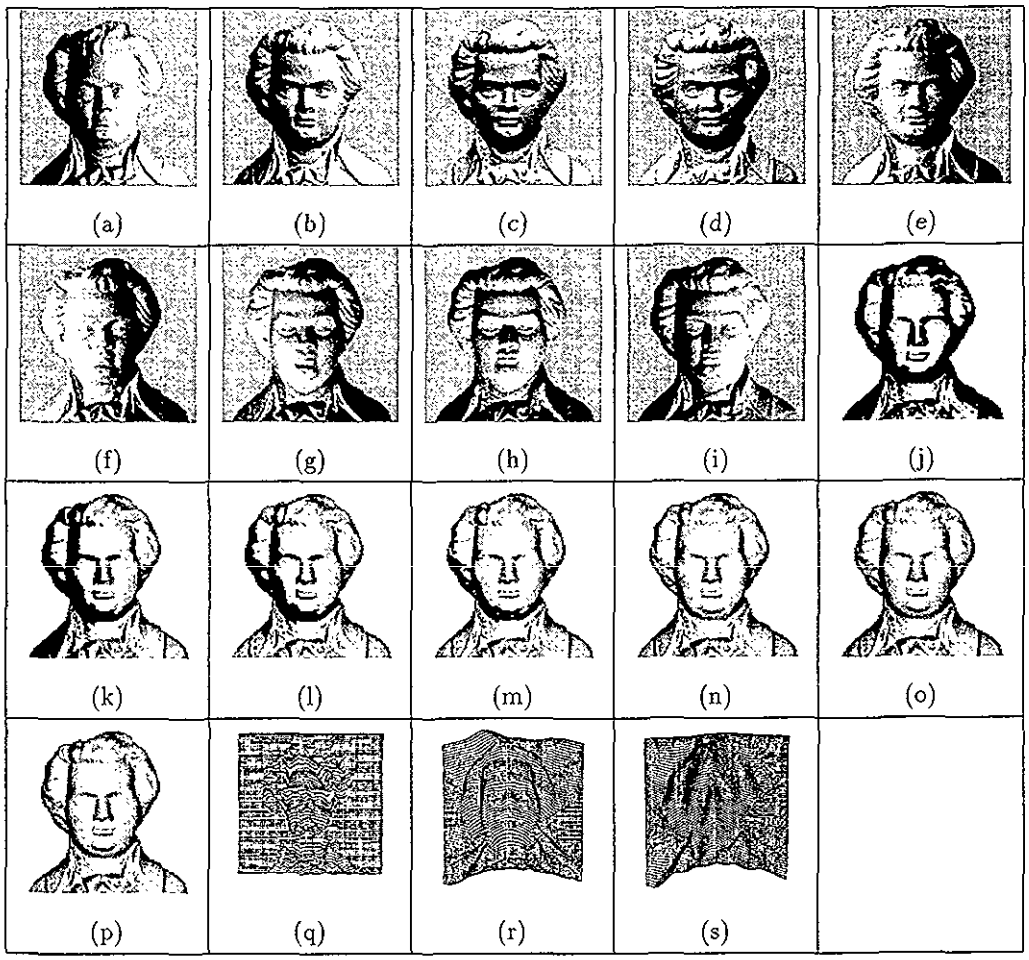

Fig. 1. Results for Mozart without noise: (a)–(i) nine input images; (j) shaded output of photometric stereo using input images (a), (d) and (g); (k)–(o) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (p) final shaded output of our algorithm after processing all images; (q) depth map (integrated from surface normals) from photometric stereo; (r) depth map (integrated from surface normals) from our algorithm; (s) reconstructed depth map, using SFS by Lee and Kuo [4], from (a).

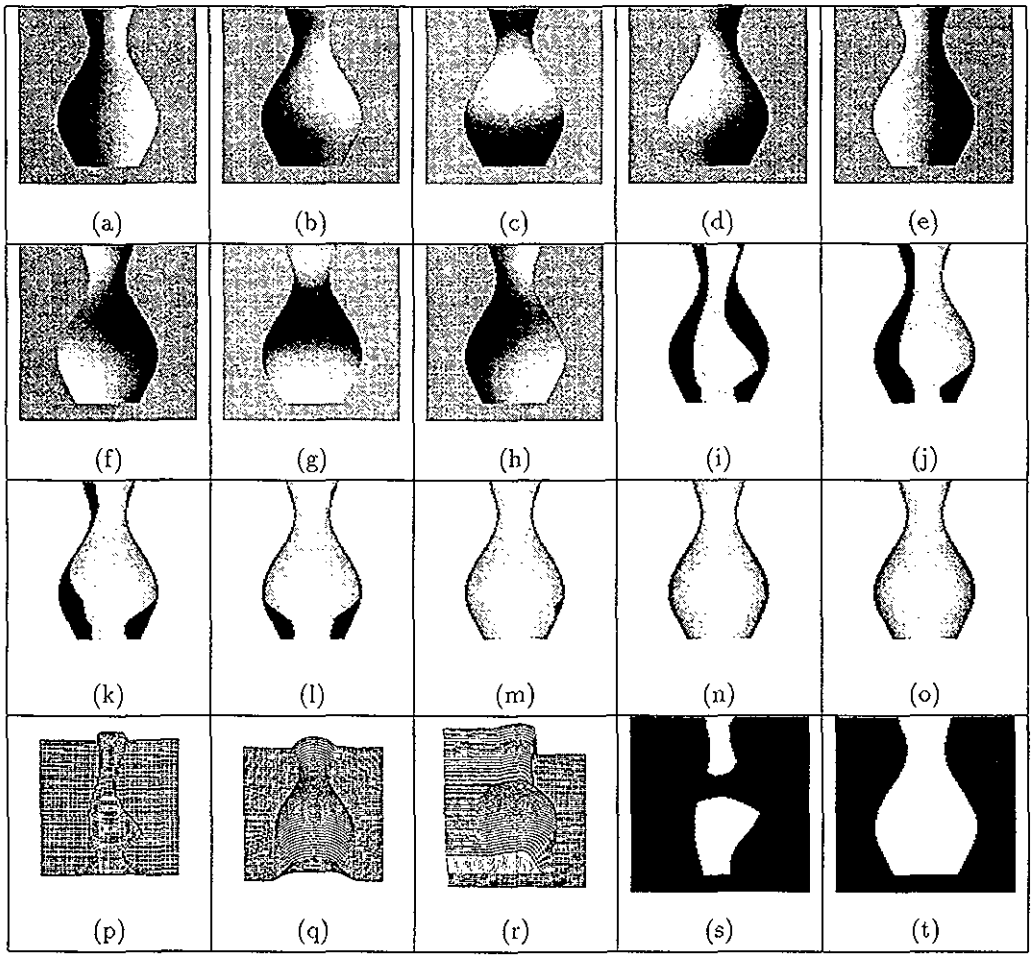

Fig. 2. Results for vase without noise: (a)–(h) eight input images; (i) shaded output of photometric stereo using input (a), (c) and (f); (j)–(n) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (o) final shaded output of our algorithm after processing all images; (p) depth map (integrated from surface normals) from photometric stereo; (q) depth map (integrated from surface normals) from our algorithm; (r) reconstructed depth map, using SFS by Bichsel and Pentland [6], from (a), choosing all points with maximal intensity as singular points; (s) recovered albedo map from photometric stereo; (t) recovered albedo map from our algorithm.

stereo and our algorithm recovered a smooth and accurate shape.

We also ran our algorithm on images generated with varying albedo. Fig. 3 shows the results of a vase sequence generated by two different albedo values (0.5 for the top half of the vase and 0.75 for the bottom half). Fig. 4 shows a similar vase sequence, but the albedo is gradually changed from 0 (at the top of the vase) to 1 (at the bottom of the vase) by setting the albedo to the row number divided by the total number of rows. As we expected, the algorithm can handle varying albedo surfaces.

4.1.2. With noise

Gaussian noise with zero mean and 10^{-2} variance was added to the input sequences in order to test the sensitivity to noise. Since our normalized intensity range is 0 to 1, the corresponding variance for the intensity range 0 to 255 would be 2.55.

The shape of the objects in the synthetic images without noise was recovered well after only one cycle, so multiple cycles were not used. However, it is worthwhile to run more cycles on noisy images to stabilize the results and provide a better recovery. Therefore, for the sequences with added noise we will also show the results of our algorithm after 10 cycles. Fig. 5 shows results for the Mozart sequence. The images were generated with a constant albedo of 0.75. Our algorithm had an average albedo error of 0.039847 after 1 cycle and 0.001699 after 10 cycles; photometric stereo had an average albedo error of 0.036416.

Fig. 6 shows results, including the recovered albedo map, for the vase sequence. The images were generated with a constant albedo of 0.75. Our algorithm had an average albedo error of 0.008581 after 1 cycle and 0.001913 after 10 cycles; photometric stereo had an average albedo error of 0.009193. Notice that the boundaries of the recovered regions from photometric stereo have more error due to the noise. However, our results at the boundary do not

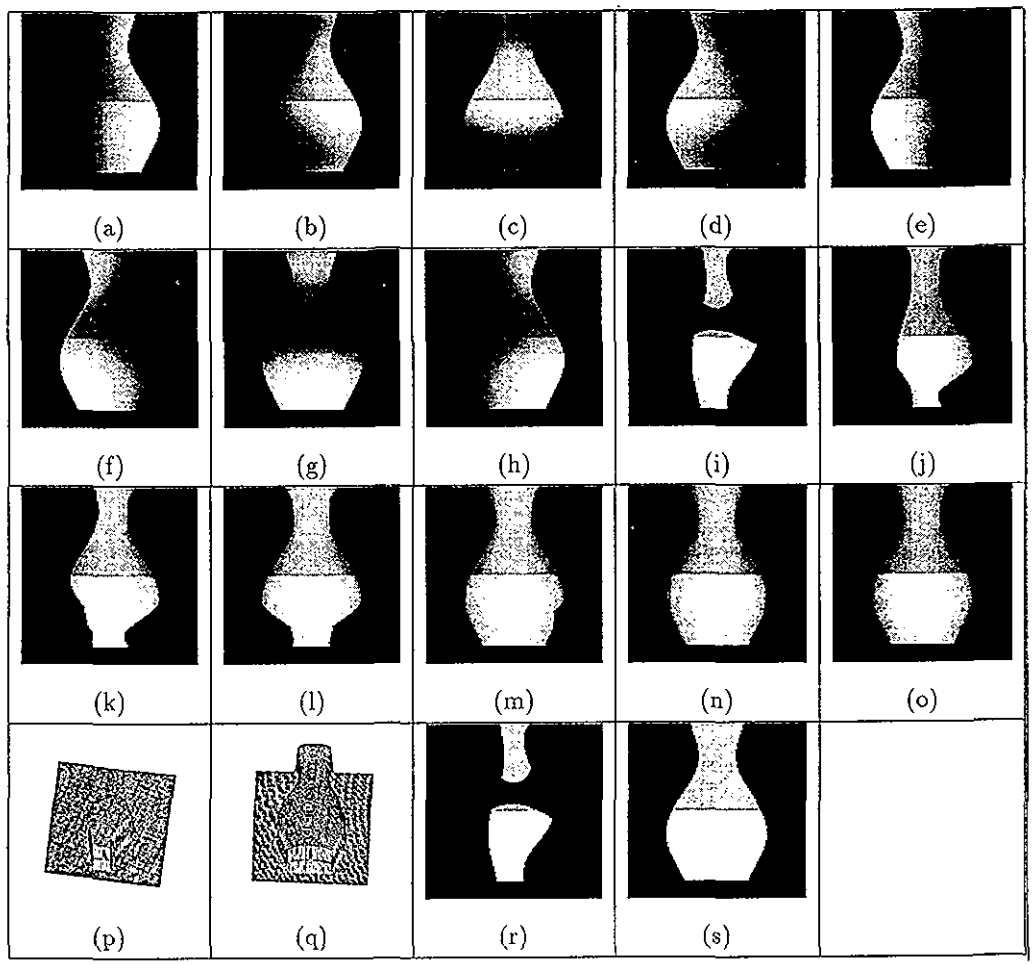

Fig. 3. Results for vase sequence with two albedo values and no noise: (a)–(h) eight input images; (i) shaded output of photometric stereo using input images (a), (d) and (g); (j)–(n) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5: images 4, 5, 6; and images 5, 6, 7; (o) final shaded output of our algorithm after processing all images; (p) depth map (integrated from surface normals) from photometric stereo; (q) depth map (integrated from surface normals) from our algorithm; (r) recovered albedo map from photometric stereo; (s) recovered albedo map from our algorithm.

have this problem since successive refinements were performed using the Kalman filter.

Fig. 7 shows the results of a vase sequence generated by two different albedo values. Fig. 8 shows a similar vase sequence with variable albedo.

4.1.3. Error analysis

In order to compare our results with those of photometric stereo, the total and partial errors in the recovered surface normal were computed. The total error was the error in the estimated surface normal over the entire image, which was computed by

$$\frac{\sum\sum\{|N_x-\hat{N}_x|+|N_y-\hat{N}_y|+|N_z-\hat{N}_z|\}}{NP} \tag{12}$$

Here NP is number of pixels, (N_x,N_y,N_z) is the true surface normal, $(\hat{N}_x,\hat{N}_y,\hat{N}_z)$ is the recovered surface normal, and the summation was taken over the entire image. The partial error was the error in the estimated surface normal over the regions recovered by photometric stereo. The errors are given in Table 1 for the sphere, vase, and Mozart images. For sequences with noise, the results of our algorithm are shown after 1, 10, and 20 cycles to illustrate the improvement gained by using multiple cycles.

From the error tables, we can see that the more cycles we run, the better results our algorithm provides. We achieved significant improvement in both the partial and the total error over photometric stereo after 10 cycles.

In order to show the improvements through multiple cycles, we ran our algorithm for 20 cycles, then computed the average albedo error, and the partial and total surface normal errors, for each cycle. The error plots are shown in Fig. 9. We can see that all three errors drop quickly before 10 cycles, and become stable around 20 cycles.

We also ran our algorithm on images generated with variable albedo and different noise levels. Fig. 10 shows the partial and total surface normal errors and average albedo error for different levels of noise from the variable albedo vase sequence. Our algorithm provides reasonable results

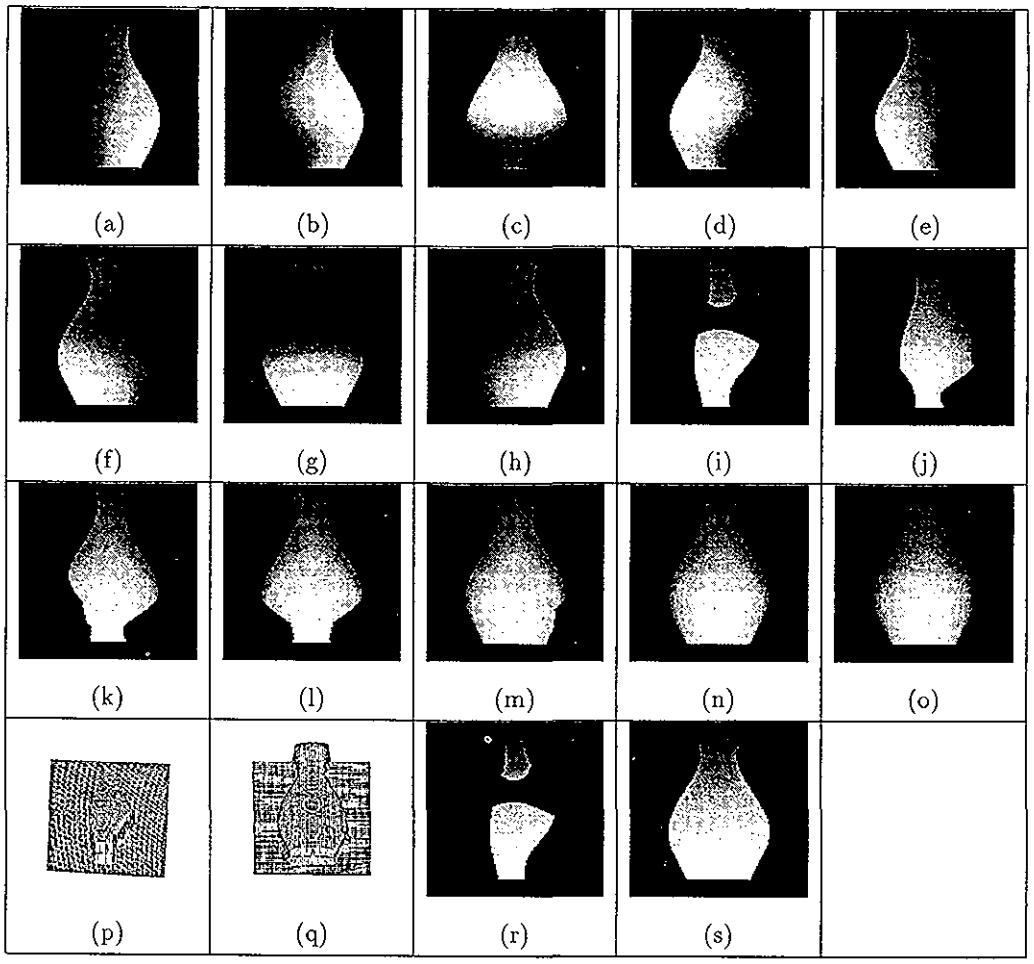

Fig. 4. Results for vase sequence with variable albedo values and no noise: (a)–(h) eight input images; (i) shaded output of photometric stereo using input images (a), (d) and (g); (j)–(n) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (o) final shaded output of our algorithm after processing all images; (p) depth map (integrated from surface normals) from photometric stereo; (q) depth map (integrated from surface normals) from our algorithm; (r) recovered albedo map from photometric stereo; (s) recovered albedo map from our algorithm.

for noise levels up to $\sigma = 10^{-2}$ (which is equivalent to 2.55 for the intensity range 0 to 255).

4.2. Real images

In this section we present results for two sequences of real images. In each figure, the 3-D plots from the algorithms are shown from different views that best represent the performance of the respective algorithm based on our visual observations.

Fig. 11 shows the results for a sequence of nine images of a mask. The sequence of light sources used had a slant of approximately 45.4° and the tilt changed every 40°, starting at 10°. Our algorithm was run for 10 cycles. It can be seen that the newly recovered areas are gradually merged into the areas already recovered. Note that the 3-D plot of the integrated surface shape from photometric stereo does not contain the entire reconstructed surface shown in Fig. 11(j). This is due to the discontinuity of the recovered surface. For a surface with discontinuous patches, since we don't know relative depths between the different patches, the integration of depth will only work for each patch separately. Consequently, we only show the plot for the center patch, which contains the starting point.

Fig. 12 shows the results for a sequence of seven images of a column. The sequence of light sources used had a slant of approximately 47.6° and tilts of 130, 210, 290, 10, 90, 170 and 50°. We show our output after 1 cycle and 5 cycles. Similar to the mask sequence, the results were improved after several cycles.

In each figure, the 3-D plots are shown from the view that best represents the performance of the respective algorithm based on our visual observation.

5. Conclusions

We have presented an algorithm which gradually

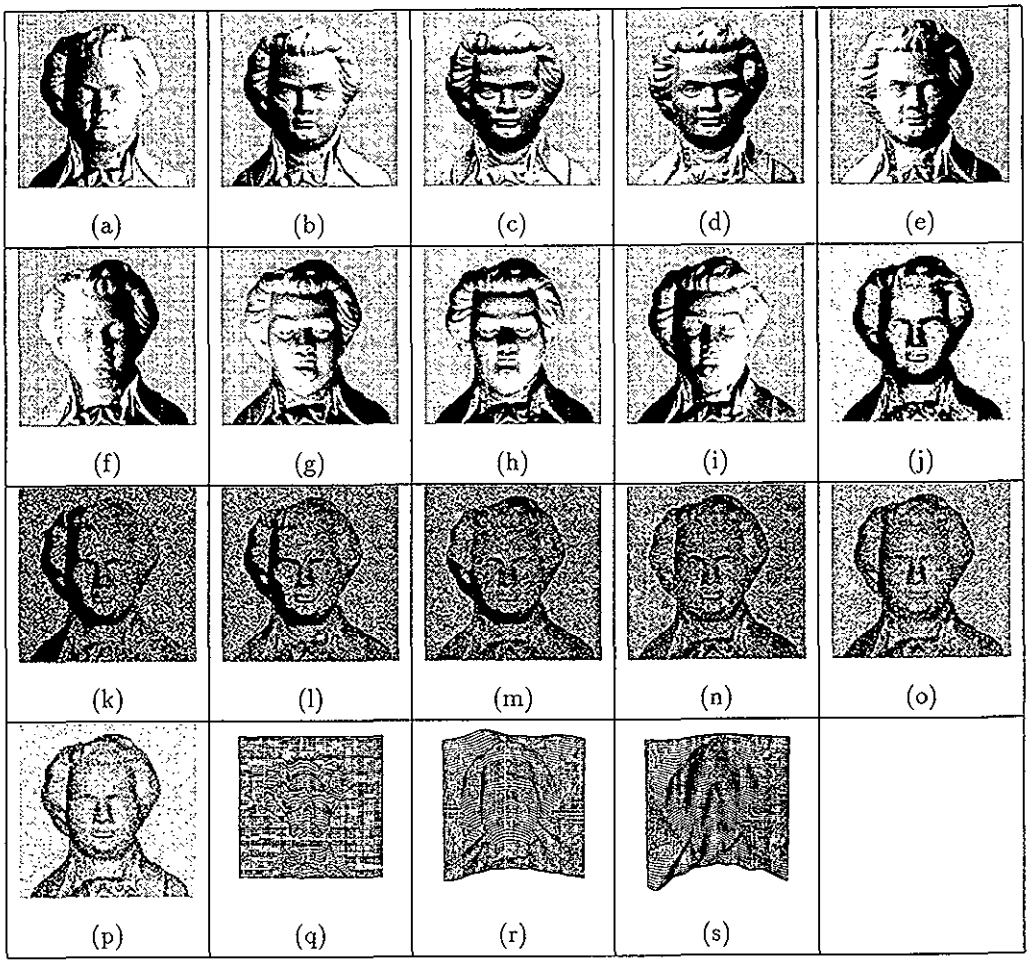

Fig. 5. Results for Mozart with noise: (a)–(i) nine input images; (j) shaded output of photometric stereo using input images (a), (d) and (g); (k)–(o) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (p) final shaded output of our algorithm after processing all images; (q) depth map (integrated from surface normals) from photometric stereo; (r) depth map (integrated from surface normal) from our algorithm; (s) reconstructed depth map, using SFS by Lee and Kuo [4], from (a).

recovers and refines the shape and albedo of an object from a sequence of images taken with different illumination directions. The process can be viewed as cascading shape from photometric stereo, formulated in the framework of a linear Kalman filter in order to iteratively recover and refine the shape and the surface albedo. The results have shown that the proposed algorithm has an order of magnitude of improvement over photometric stereo on noisy images even if we only consider the partial area recovered by photometric stereo (Table 1). If we consider the entire object area, the improvement is even greater since this algorithm can recover larger portions of the object through the use of more images and multiple cycles. The same is true for the recovered albedo results. In conclusion, when a longer sequence of images is available, this algorithm provides effective shape recovery from noisy images, especially when the algorithm is run for multiple cycles. It also should be pointed out that this algorithm is an extension of photometric stereo to a longer sequence of images, not a replacement of photometric stereo. It is equivalent to photometric stereo in the case of three images.

When there are at least six images available, this algorithm can be extended to recover both the shape and the light sources (direction and strength) by applying a least squares fit to recover the light sources first, then to recover the shape from the estimated light sources and known intensities [9]. However, the technique of using look-up tables to recover shape from general reflectance maps [10] will not be easily integrated into our framework, since the Kalman filter requires an analytical expression for the reflectance map, and nonlinear formulas have to be linearized before applying the Kalman filter.

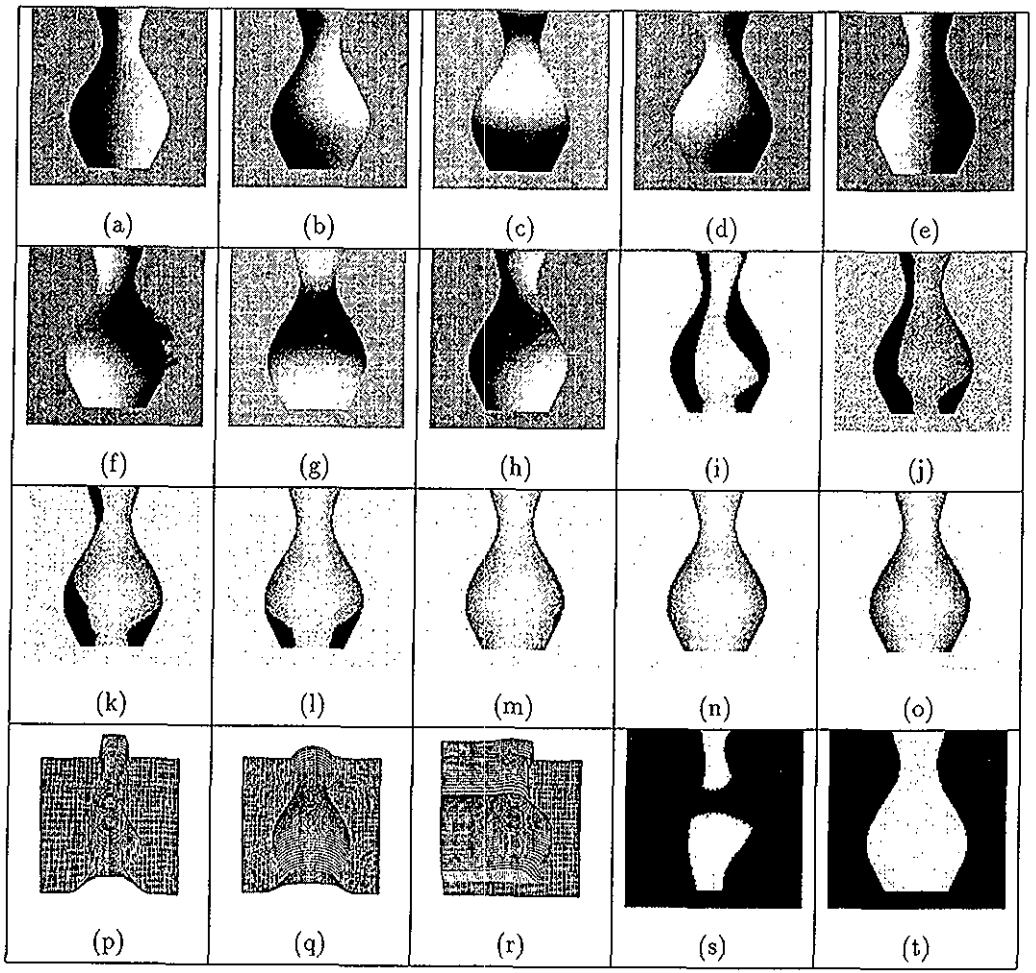

Fig. 6. Results for vase with noise: (a)–(h) eight input images; (i) shaded output of photometric stereo using input images (a), (c) and (f); (j)–(n) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (o) final shaded output of our algorithm after processing all images; (p) depth map (integrated from surface normals) from photometric stereo; (q) depth map (integrated from surface normals) from our algorithm; (r) reconstructed depth map, using SFS by Bichsel and Pentland [6], from (a), choosing all points with maximal intensity as singular points; (s)

Fig. 7. Results for vase sequence with two albedo values and with noise: (a)–(h) eight input images; (i) shaded output of photometric stereo using input images (a), (d) and (g); (j)–(n) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (o) final shaded output of our algorithm after processing all images; (p) depth map (integrated from surface normals) from photometric stereo; (q) depth map (integrated from surface normals) from our algorithm; (r) recovered albedo map from photometric stereo; (s) recovered albedo map from our algorithm.

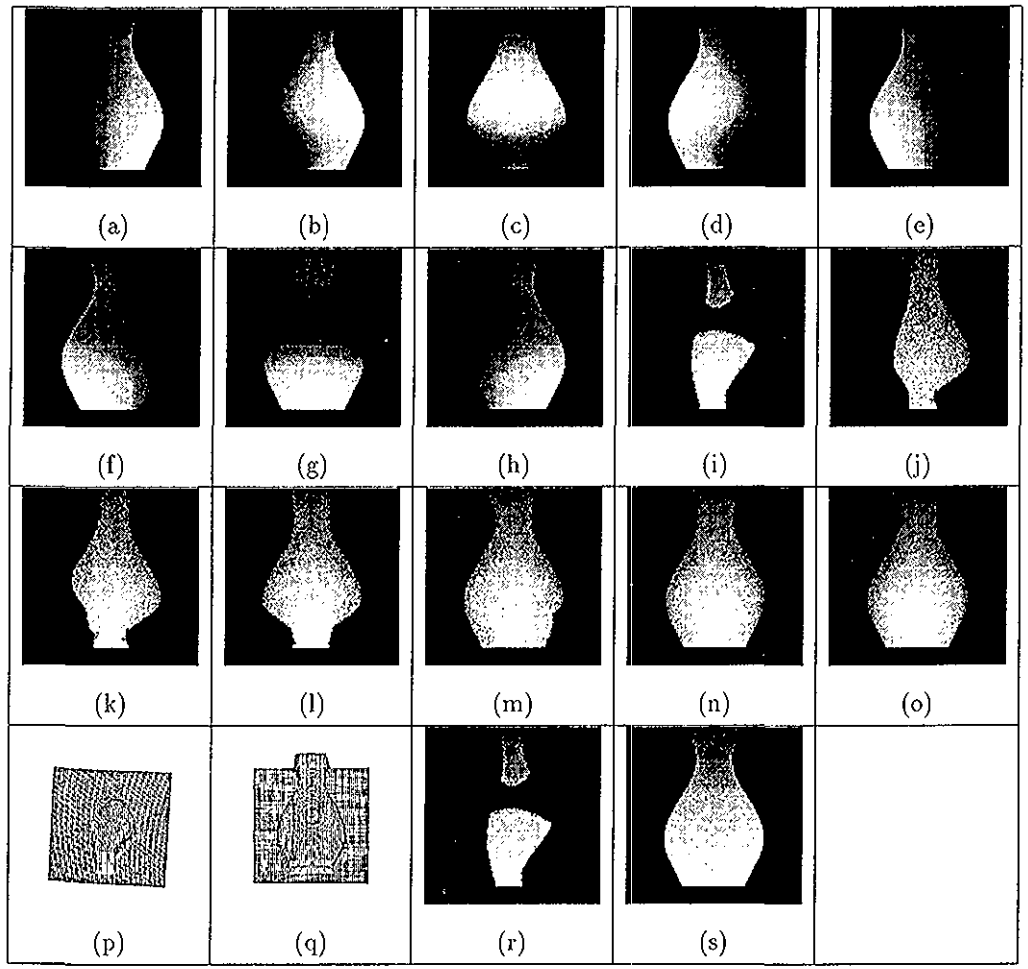

Fig. 8. Results for vase sequence with variable albedo values and with noise: (a)–(h) eight input images; (i) shaded output of photometric stereo using input images (a), (d) and (g); (j)–(n) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (o) final shaded output of our algorithm after processing all of the images; (p) depth map (integrated from surface normals) from photometric stereo; (q) depth map (integrated from surface normals) from our algorithm; (r) recovered albedo map from photometric stereo; (s) recovered albedo map from our algorithm.

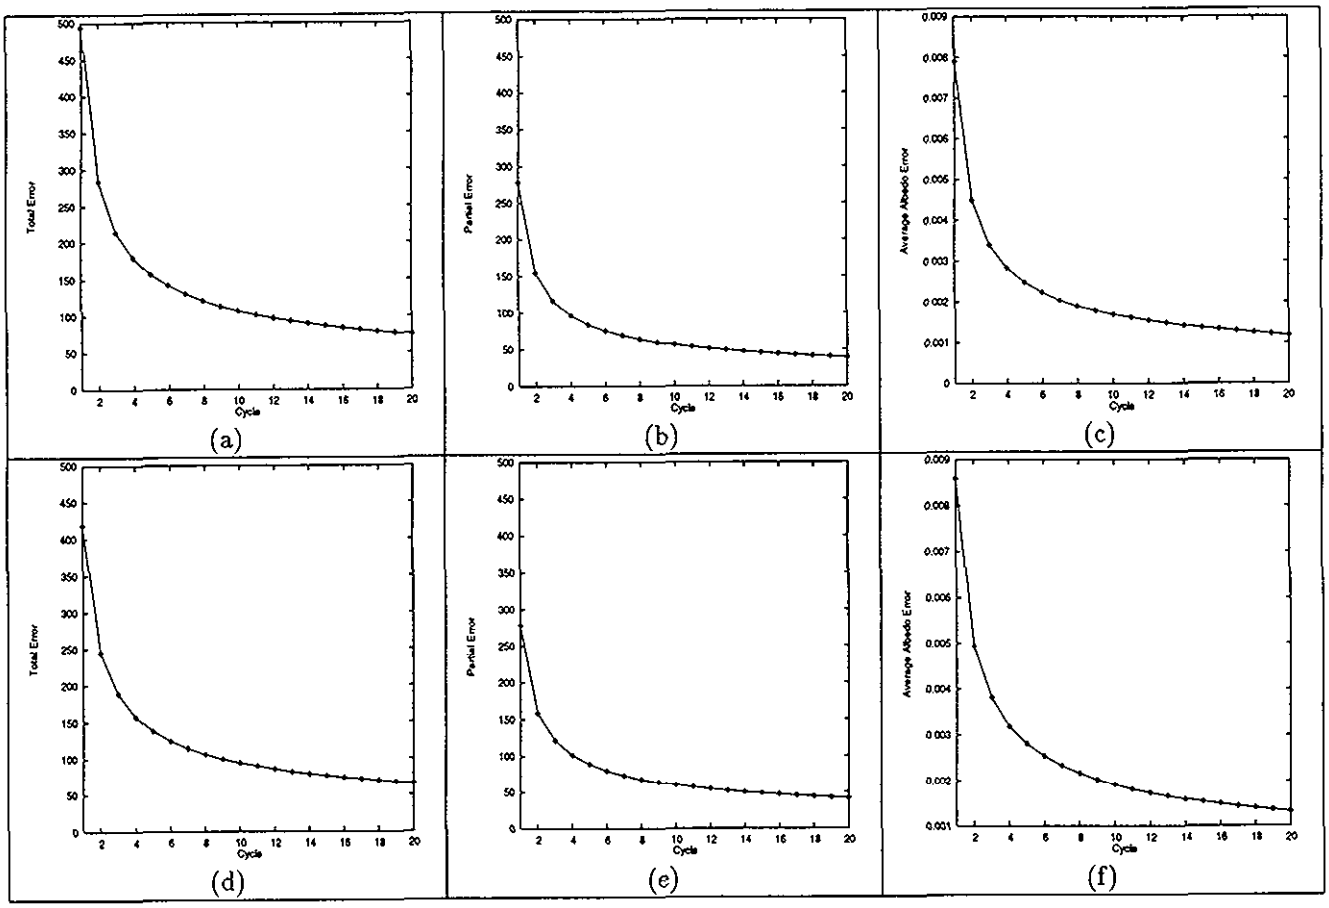

Fig. 9. Error plots, for noisy synthetic images, for our algorithm: total surface normal error, partial surface normal error and average albedo error, (a)–(c) for Mozart; (d)–(f) for vase.

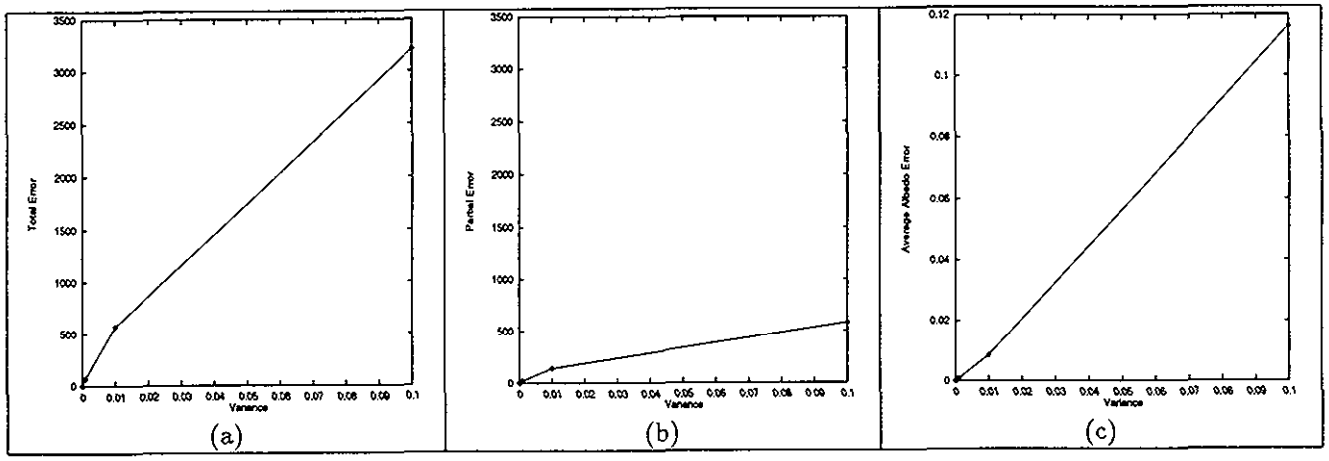

Fig. 10. Error plots, for variable albedo vase sequence with different noise levels, for our algorithm: (a) total surface normal error; (b) partial surface normal error; (c) average albedo error.

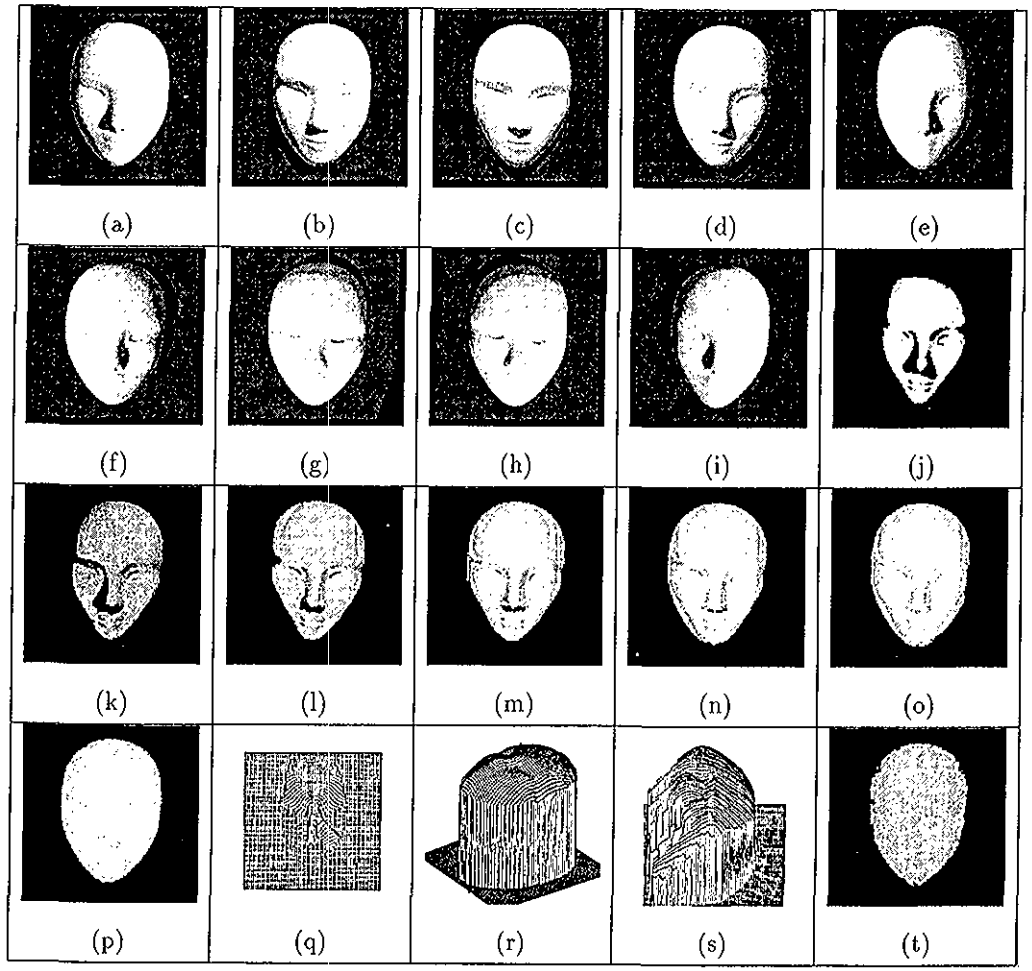

Fig. 11. Results for the real images of a mask: (a)–(i) nine input images; (j) shaded output of photometric stereo using input images (a), (d) and (g); (k)–(o) shaded output of our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (p) final shaded output of our algorithm after processing all images; (q) depth map (integrated from surface normals) from photometric stereo; (r) depth map (integrated from surface normals) from our algorithm; (s) reconstructed depth map, using SFS by Bichsel and Pentland [6], from (a), choosing all points with maximal intensity as singular points; (t) recovered albedo map from our algorithm.

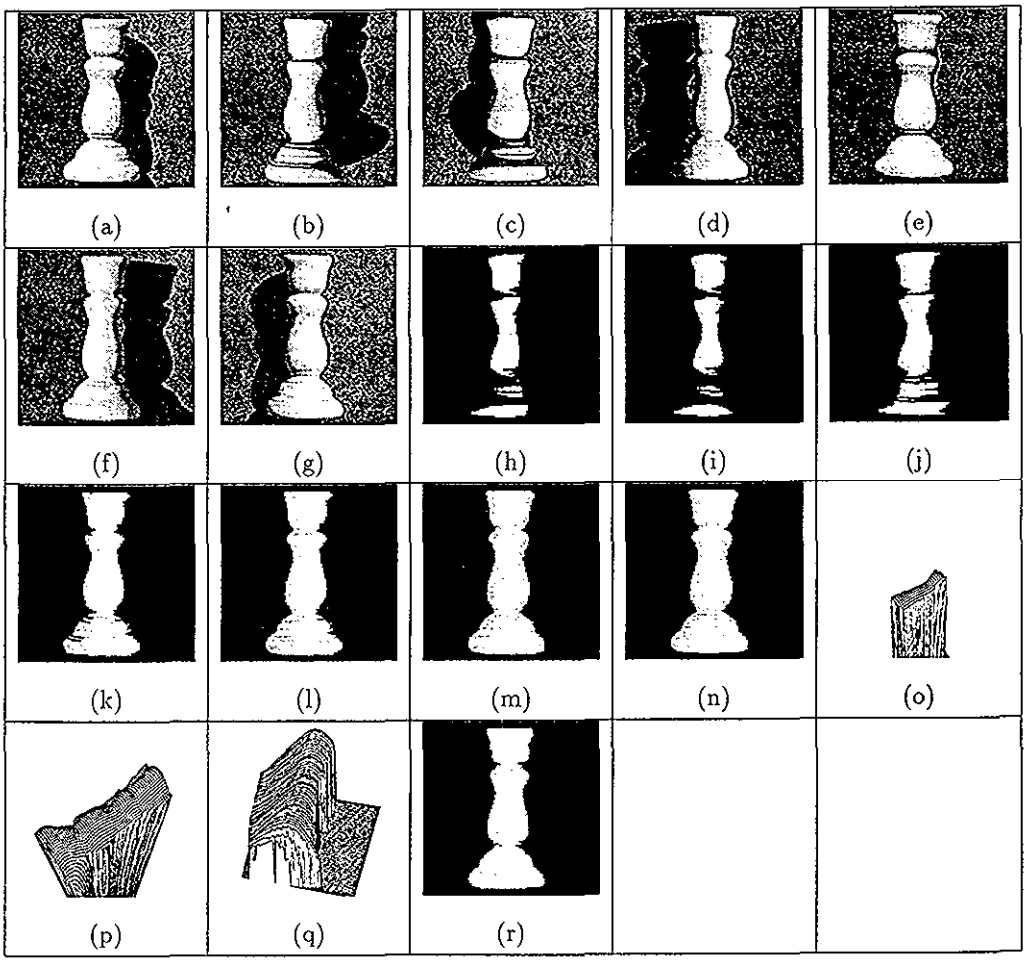

Fig. 12. Results for the real images of a column: (a)–(g) seven input images; (h) shaded output from photometric stereo using input images (a), (c) and (f); (i)–(m) shaded output from our algorithm after processing images 1, 2, 3; images 2, 3, 4; images 3, 4, 5; images 4, 5, 6; and images 5, 6, 7; (n) final shaded output from our algorithm after 5 cycles; (o) depth map (integrated from surface normals) from photometric stereo; (p) depth map (integrated from surface normals) from our algorithm; (q) reconstructed depth map, using SFS by Bichsel and Pentland [6], from (a), choosing all points with maximal intensity as singular points; (r) recovered albedo map from our algorithm.

Acknowledgements

The research reported here was supported by the National Science Foundation grants CDA-9122006 and IRI-9220768.

References

[1] R.J. Woodham, Photometric method for determining surface orientation from multiple images, Optical Engineering 19 (1) (1980) 139–144.
[2] R. Zhang, P.S. Tsai, M. Shah, Shape from photomotion, Proc. Conf. Computer Vision and Pattern Recognition, IEEE, 1993, pp. 140–141.
[3] R. Zhang, P. Tsai, M. Shah, Photomotion, Computer Vision and Image Understanding 63 (2) (1996).
[4] K.M. Lee, C.C.J. Kuo, Shape reconstruction from photometric stereo, Proc. Conf. Computer Vision and Pattern Recognition, IEEE, 1992, pp. 479–484.
[5] N. Ayache, Artificial Vision for Mobile Robots: Stereo Vision and Multisensory Perception, MIT Press, Cambridge, MA, 1991.
[6] M. Bichsel, A.P. Pentland, A simple algorithm for shape from shading, in: Proc. Conf. on Computer Vision and Pattern Recognition, 1992, pp. 459–465.
[7] R. Zhang, P.S. Tsai, J. Cryer, M. Shah, Analysis of shape from shading techniques, Proc. Conf. on Computer Vision and Pattern Recognition, IEEE, 1994, pp. 377–384.
[8] E.N. Coleman Jr., R. Jain, Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry, Computer Vision, Graphics and Image Processing 18 (1982) 309–328.
[9] R.J. Woodham, Y. Iwahori, R.A. Barman, Photometric stereo: Lambertian reflectance and light sources with unknown direction and strength, Technical Report, 91-18, 1991.
[10] R.J. Woodham, Surface curvature from photometric stereo, Technical Report, 90-20, 1990.