AN ITERATIVE MULTI-RESOLUTION SHAPE-FROM-SHADING ALGORITHM AND ITS APPLICATION TO PLANETARY MAPPING

P L Van Hove
Massachusetts Institute of Technology
Room 36-647, Cambridge, MA 02139

M J Carlotto
The Analytic Sciences Corporation
1, Jacob Way, Reading, MA 01867

ABSTRACT

Experimentation with an algorithm for reconstructing the 3-D structure of a scene from a single image is reported. Such algorithms are referred to as shape-from-shading in the image understanding community, photoclinometry in the planetary science community. The algorithm presented here is the synthesis of a two-step reconstruction of elevation maps with multi-resolution iterative algorithms. Performance of the algorithm is assessed on synthetic test imagery. Convergence to metric accuracy is obtained only after large numbers of iterations, and is affected by image size and illumination conditions. Initial reconstructions are presented for an impact crater and a mesa on Mars from Viking orbiter imaging.

Keywords: Shape from shading, Photoclinometry, Topographic mapping, Planetary remote sensing, Image processing.

1. INTRODUCTION

The reconstruction of terrain surfaces from remotely sensed imagery is an important activity in planetary exploration. Using visible imagery, topography can be reconstructed using three basic techniques, namely by an analysis of shadows, by stereo matching, and by shape-from-shading or photoclinometry. Shadow analysis involves relating lengths of shadows to heights of the objects casting the shadows and is only able to reconstruct the silhouette of the shadow-casting objects. Stereo techniques involve matching features in one image of a stereo pair to the corresponding features of the other image, in order to determine the heights of these features. The Shape-from-shading method attempts to reconstruct the shape of a terrain being imaged by relating shading information to surface orientations. Stereo matching techniques are preferred in the presence of high densities of distinct surface features. In case of a lack of distinct surface features, the primary source of shape information is shading so that if the surface material composition is homogeneous, shape-from-shading methods should be preferred. Finally, shadow analysis provides only very partial shape reconstruction and is applicable only for nearly grazing illumination.

This work was conducted in part while P. Van Hove was working at The Analytic Sciences Corporation.

The purpose of this paper is to present a practical algorithm for the reconstruction of relief from shading information. The algorithm is based on a variational formulation (Ref.1) and implemented using an iterative multi-resolution approach (Ref.2). Section 2 summarizes previous work on shape-from-shading, leading to the implementation described in this paper. Section 3 discusses implementation issues associated with the particular algorithm. Experimental results are presented in section 4; they include the reconstruction of known objects from synthetic images and imagery of Martian terrain from Viking orbiter. Section 5 summarizes the contributions of the paper.

2. PREVIOUS WORK ON SHAPE-FROM-SHADING

This section reviews some of the work performed in the machine vision community on extracting shape information from brightness images. Reference (3) contains a good tutorial review as well as a list of references. The shape-from-shading problem consists of estimating an elevation map of the terrain, given a gray scale image of this terrain, characteristics of the surface materials and the distribution of light sources.

Gray levels in an image measure the irradiances $E(x,y)$ in the image plane, which are proportional to the radiances $L(x,y)$ at the corresponding points in the scene in the direction of the camera. With the assumption that both light sources and observer are distant from the scene, $E(x,y)$ depends for one surface material, and for one imaging device, only on the surface orientations at the corresponding points in the scene.

$$E(x,y) = \alpha L(x,y) = \alpha\rho E_0 R(p(x,y), q(x,y)) \quad (1)$$

where $R(p,q)$ is the reflectance map (Ref.3), for a given distribution of light sources. Its dependence on surface orientation is expressed through the gradients (slopes) $p = \partial z/\partial x$ and $q = \partial z/\partial y$. The constant α represents effects of the imaging system, ρ surface albedo and E_0 illumination irradiance; the product $\alpha\rho E_0$ will be made equal to 1 by scaling of $E(x,y)$. Coordinates x,y are used both for points in the image plane and corresponding points in the base plane of the scene.

Equation (1) provides a direct solution to shading synthesis, i.e. the determination of image irradiances $E(x,y)$ for a given scene relief $z(x,y)$. It is clear that, in general, this equation can not directly be inverted for p and q or for z in terms of $E(x,y)$.

Proceedings of IGARSS' 86 Symposium, Zürich, 8–11 Sept. 1986. Ref. ESA SP-254 (published by ESA Publications Division, August 1986).

This inverse problem leads to a non-linear partial differential equation for the elevation map $z(x,y)$ in terms of the image brightness map $E(x,y)$. As problems arise in the numerical solution of this equation, variational formulations have been developed for the shape-from-shading problem.

2.1 Differential Formulation of Shape-from-Shading

One solution to extracting relief information from shading information is to consider equation (1) as a differential equation for $z(x,y)$, namely

$$R\left(\partial z(x,y)/\partial x, \partial z(x,y)/\partial y\right) = E(x,y) \qquad (2)$$

which is in general, a non-linear partial differential equation. This differential approach to the problem was phrased in the more general case of perspective projection by Horn (Ref.4) and is related to photoclinometry (Ref.5). Solutions to the differential equation in (2) can be obtained by integration along characteristic strips. The major drawback of this solution method is that the paths of the base characteristics are data dependent. As a result, first it is not possible to control the spacing between nearby characteristics. Second, it is difficult to predict where initial conditions are necessary to construct the solution on a given domain. Third, in the presence of noise, base characteristics may intersect with incompatible values for z, p, q. Finally, the method is limited by the difficulty of exploiting boundary conditions on a closed contour for only z or only a combination of p and q.

2.2 Variational Formulation of Shape-from-Shading

A variational equivalent to equation (1) is given by the optimization problem

$$\min_{p,q} \mathbf{F} = \qquad (3)$$

$$\min_{p(x,y),q(x,y)} \iint \left| E(x,y) - R(p(x,y), q(x,y)) \right|^2 dxdy$$

This formulation takes into account inaccuracies in the data by attempting to reduce the departure from equation (1) instead of imposing strict equality everywhere. A large variety of additional sources of information can be incorporated into this formulation by adding constraints with Lagrange multipliers or penalty factors.

As equation (1) is underdetermined when p and q are independent unknown functions, it is clear that the problem in equation (3) is underdetermined as well. Several methods have been proposed to resolve this ambiguity, such as regularization, using constraints on p and q, or directly reconstructing $z(x,y)$. We have chosen to use the constraint that the gradients p and q must be curl-free, and to exploit this constraint by adding a penalty term in the functional. The optimization problem is then

$$\min_{p(x,y),q(x,y)} \mathbf{F} + \lambda \iint \left| p_y - q_x \right|^2 dxdy \qquad (4)$$

where λ is a constant chosen a-priori. This formulation has several advantages. First, it leads to a simple iterative solution on a discrete grid which will be discussed later in the text. Second, the exact gradients solve equation (4) for noise-free data. However, since the curl-free constraint for the gradients is not exploited as a strict constraint, the reconstructed gradients are in general incompatible so that the elevations can not be reconstructed by a simple integration.

Boundary conditions for the problem are now considered. The analysis of singular points of the differential equation (2) reveals that the shape-from-shading problem has a unique solution in some cases without the need for fixing any boundary conditions. In general however, it is necessary to specify boundary conditions for p, q, or a combination of p and q on a closed contour. In the field of machine vision, the major source of boundary conditions is the occluding boundaries of smooth objects, see references (6, 7). In our application of the shape reconstruction method, flat terrain was assumed over part or all of the boundary of the region being processed.

2.3 Numerical Solution

A numerical solution method for the optimization problem in equation (4) is now discussed. In a first step, a solution is attempted on a closed domain, with known values for the gradients p and q along the boundary of the region. This is the case for example when estimating the shape of a relief feature emerging over a flat terrain. The formal solution to the minization problem is given by the Euler equations of variational calculus, which provide the following necessary condition for the optimal gradient fields p and q inside the domain,

$$\begin{cases} \left| E(x,y) - R(p,q) \right| R_p(p,q) + \lambda \left(p_{yy} - q_{xy} \right) = 0 \\ \left| E(x,y) - R(p,q) \right| R_q(p,q) + \lambda \left(q_{xx} - p_{yx} \right) = 0 \end{cases} \qquad (5)$$

which is a set of coupled second order partial differential equation for p and q. This type of equations can be solved by iterative methods. The solution is obtained on a rectangular grid by first sampling the functions in equation (5) and approximating derivatives by finite differences. In the resulting discrete equations, the term corresponding to the central point is isolated and the remainder of the equation is considered as an update expression for the value at the central location. The resulting update formula for p corresponding to equation (5) is given by

$$p_{ij}^{NEW} = \bar{p}_{ij} - \tilde{q}_{ij} + \left| E_{ij} - R(p_{ij}, q_{ij}) \right| R_p(p_{ij}, q_{ij}) \qquad (6)$$

where p_{ij}, q_{ij} are the sample values of p, q at the grid point (i,j), and $\bar{p}_{ij}$, $\tilde{q}_{ij}$ are linear combinations of the values of p, q at points around (i,j); a similar expression applies for the update of q.

The iterative method is started from initial values for p, q, in our case zero gradients. Each iteration consists of updating p at each point interior to the domain by equation (6) and q by the corresponding formula. In principle, the update may be done recursively or non-recursively. Recursive updates were chosen in our implementation for better stability.

3. IMPLEMENTATION ISSUES

In this section, implementation issues are discussed and general observations made on preliminary tests are reported.

It has been experimentally observed that, with synthetic data, the iterative algorithm generally converges to a gradient field which closely matches the original data. Stability of the iteration is good when the image data is consistent with a real continuous surface and when the reflectance map used in the reconstruction process closely matches the reflectance map used in synthesizing the image. However, stability degrades when large amounts of noise are present in the input data or when the reflectance map used in the reconstruction is very different from the actual reflectance map. Indeed, this type of degradations may make the data (image and reflectance map) strongly inconsistent. The iteration also becomes unstable when the value chosen for the penalty factor λ is too small. In that case, the reduction of ambiguity in equation (3) by the penalty term in equation (4) becomes too weak. Convergence of the iterative algorithm is generally slow after a fast initial decrease of the error. The slow convergence is especially marked for grid sizes in excess of 25x25 pixels. This slow convergence can be attributed to the update in equation (6), which propagates the constraints only one grid spacing at a time. This convergence rate issue can be alleviated by applying multi-resolution methods.

Multi-resolution methods have been proposed to solve partial differential equations in general (Ref.8), and surface reconstruction problems in particular (Ref.2). The basic principle consists of solving a given problem simultaneously for different samplings of the same domain. Updates such as defined in equation (6) propagate the constraints much faster on the coarser grids, whereas accuracy is maintained by iterations on the finer grids. In our implementation, the measured brightness image is repeatedly smoothed and subsampled by factors of 2 until the grid size is approximately 10x10. The reconstruction is first performed on this coarse grid, for which convergence is fast. The resulting gradients and elevations are then interpolated to provide initial conditions for the iteration at the next finer level. The sequence of interpolations and iterations is continued until the gradients are reconstructed on the finest grid, see Figure 1.

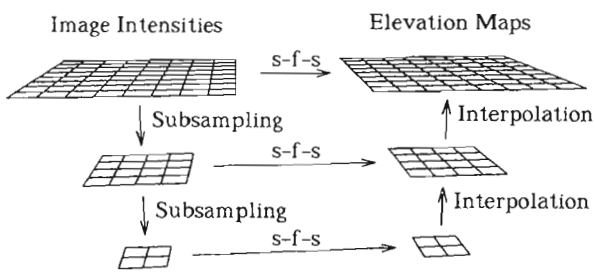

Figure 1. Multi-Resolution Reconstruction Method

With this method, reasonable convergence can be attained on large grid sizes with a limited total number of iterations. Multi-resolution iterative solutions have been suggested for the shape-from-shading problem in (Ref.7) and and example reported in (Ref.2) for a regularized version of equation (3). However, performance of the multi-resolution algorithm for shape-from-shading reconstruction had not been much investigated.

In real applications, it is often desirable to apply the shape-from-shading algorithm to image regions of arbitrary size and shape. Our implementation has been designed specifically to address this question, and to produce correct subsamplings and interpolations when coarser grids are not aligned with the region boundary.

In addition, the simple boundary condition of a flat terrain may not be reasonable everywhere along the region outline. For part of the region outline, it may not be possible to estimate the slopes a-priori, so that free boundary conditions must be implemented. For points on free boundaries, the update in equation (6) has to be modified because some of the neighbors required in the computation of $\bar{p}_{ij}$, $\tilde{q}_{ij}$ are outside the region. Modified expressions have been obtained in a coherent manner by discretizing the minimization problem in equation (4) and determining for each boundary point, which terms are present in the expressions for $\bar{p}_{ij}$, $\tilde{q}_{ij}$. This coherent treatment of boundary conditions is presented for other problems in (Ref.2) and is referred to as molecular inhibition. Special software has been developed to implement the free boundary conditions on regions with arbitrary shape.

Performance of the algorithm in the presence of shadows is now discussed. As the reflectance map in equation (1) is a local relation, it does not correctly model shadows. As a result, errors can be expected in the reconstructed shape in the presence of shadows. In the case of a point source, shadow areas correspond to zero image brightness. Although the shape of the shadowed objects can obviously not be reconstructed, it turns out that the algorithm reconstructs over the shadowed area, a surface which follows the light rays grazing the shadow-casting objects; see Figure 2.

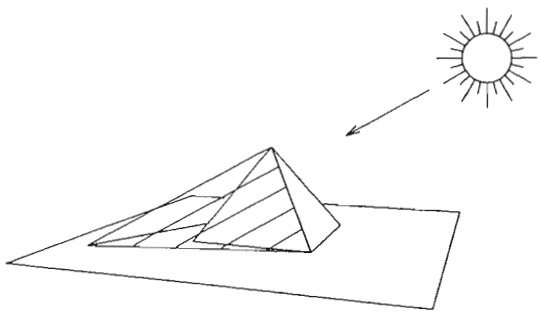

Figure 2. Surface of Grazing Rays Reconstructed in a Shadow Area.

In the case of a point source, the shadowed areas can be detected from the brightness data and their presence does not adversely affect the reconstructed shape elsewhere in the image. On the contrary, the grazing rays reconstructed by the algorithm enforce consistency between lengths of the shadows and heights of the shadow-casting objects. As a result, the shape-from-shading reconstruction implicitly includes in the reconstructed shape, the information that could otherwise be obtained by shadow analysis.

The algorithm described so far reconstructs only the gradient fields p and q. As the curl-free constraint for p and q was applied only as a penalty term, the resulting fields are not integrable in general. An elevation map which best fits these inconsistent gradients is obtained as the solution of

$$\min_{z(x,y)} \iint \left| (z_x - p)^2 + (z_y - q)^2 \right| dxdy \qquad (7)$$

This problem is well behaved when z is fixed at one point at least. Multi-resolution iterative solution methods similar to those described in the previous section have been developed for this problem.

4. EXPERIMENTAL RESULTS

The iterative multi-resolution shape-from-shading algorithm was initially tested on synthetic images of simple geometric objects such as pyramids, hemispheres, and raised cosine surfaces, see Figure 3. A Lambertian reflectance map was used to synthesize the test images and to reconstruct the elevation maps displayed in Figure 3. Qualitatively, the reconstructed surfaces closely resemble the original objects. The accuracy of the algorithm was then examined as a function of the number of iterations performed , the size of the object being reconstructed, and the incidence angle of the illumination.

In a first experiment, the influence of the number of iterations was investigated. An image of a pyramid with base 32 x 32 and unit height was synthesized for a point source illumination with an azimuth of 70° and an incidence angle of 30°. The shape was reconstructed with the multi-resolution algorithm discussed in section 3, applied to grid sizes of 32 x 32, 16 x 16 and 8 x 8 pixels. The resulting accuracy is given below as a function of the number of iterations performed at each level.

Number of Iterations			RMS Error	Peak Error
@full res.	@1/2 res.	@1/4 res.		
25	50	100	0.09	- 17%
12	25	50	0.13	- 26%
6	12	25	0.21	- 46%
3	6	12	0.30	- 69%

Although the general shape of the surface is recovered after relatively few iterations, a large number of iterations are required for good metric accuracy.

Next, the effect of varying the incidence angle of the illumination was investigated. A test image of a 32 x 32 x 1 pyramid was generated at an azimuth of 70° degrees and various incidence angles. The results after 40, 20 and 10 iterations at 1/4, 1/2 and full resolution, respectively are shown below for various incidence angles.

Incidence Angle	RMS Error	Peak Error
5°	0.34	- 81%
10°	0.23	- 53%
20°	0.16	- 33%
40°	0.14	- 28%
60°	0.14	- 27%

As the incidence angle approaches 0° (Zenith), the solutions become increasingly poor due to the lack of shading information. The best performance was achieved at relatively high incidence angles where shading effects are most pronounced.

Finally, the accuracy of the solution as a function of the size of the object relative to the sampling grid was assessed. Images of unit pyramids of different base dimensions were generated with the illumination at an azimuth angle of 70° and an incidence angle of 40°. The results are presented below for the shape reconstructed after 40, 20 and 10 iterations at 1/4, 1/2 and full resolution respectively.

Object Size (Base Dimension)	RMS Error	Peak Error
64 x 64	0.14	- 28 %
32 x 32	0.036	- 18 %
16 x 16	0.026	- 23 %
8 x 8	0.023	- 40 %

Initially, as the size of the object decreases, the error decreases since information needs to be propagated over shorter distances in the grid. The result is that the algorithm converges at a faster rate. If the object size is further decreased, the error begins to rise because the object becomes so small that it cannot be fully resolved at coarser resolutions. As the incidence angle, or the object size, or the number of iterations performed was decreased, a "ringing" in the solution surface was also noted, particularly along the base of the pyramid.

To summarize the results of our experiments on synthetic data, general surface shapes are well reconstructed with the algorithm, although numerical accuracy is relatively poor. This suggests the combined use of shading and stereo matching for the estimation of terrain elevations, for example by using stereo matching algorithms to compute a coarse but accurate grid of heights that is subsequently "filled in" using information from a shape-from-shading algorithm.

Finally, the algorithm was tested on real terrain imagery. Figure 4 presents examples of reconstructed elevation maps for Martian terrain. Elevations reconstructed with the shape-from-shading algorithm are encoded into gray scales in figure 4, with dark levels corresponding to low elevations. It has been noted that Mars is better suited for photoclinometry or the shape-from-shading methods than the Moon (Ref.5). On the Moon, ejecta from the impact of meteors tend to form rays of material with different albedo. On Mars, winds redeposit materials so that the surface material properties are fairly homogeneous locally within the same geological province. For the initial results contained in this paper, a Lambertian reflectance map was used. Current work involves incorporating photometric measurements of Mars by O'Leary and Jackel (Ref.9).

5. SUMMARY

An algorithm for reconstructing the 3-D shape of surfaces from a single image was described. The algorithm is based on a variational formulation of the problem and on its solution with a multi-resolution iterative method. Experimental results show that the qualitative shape of the scene is reconstructed after a small number iterations. However, many iterations are required to obtain metric stability. In addition, divergence from hypotheses of homogeneous surface material and of the surface material reflectance function introduce quantitative errors which are difficult to estimate. In addition to exploiting the shading information, it has been verified that the algorithm also exploits shadows in reconstructing the height of the shadow-casting objects, so that the relief reconstructed with shape-from-shading implicitly includes information that could be obtained independently with shadow analysis. The relatively poor quality of the algorithm in estimating relief elevations suggests the combined use of shading and stereo matching for the automatic estimation of dense and accurate elevation maps. The relief reconstruction algorithm is currently being used to reconstruct the shape of several unusual Martian surface features from Viking orbiter imagery.

6. REFERENCES

[1] T.M. Strat, *A Numerical Method for Shape form Shading from a Single Image*. M.S. Thesis, Department of Electrical Engineering and Computer Science, MIT, 1979.

[2] D. Terzopoulos, "Computing Visible-Surface Representations," Memo AIM-800, MIT Artificial Intelligence Laboratory, March 1985.

[3] B.K.P. Horn, *Robot Vision*. Cambridge, MA: MIT Press, 1986.

[4] B.K.P. Horn, "Obtaining Shape from Shading information, in," in *The Psychology of Computer Vision*, ed. P.H. Winston, Mc Graw Hill, 1975.

[5] L. Wildey, "Generalized Photoclinometry for Mariner 9," *ICARUS*, vol. 25, pp. 613-626, 1975.

[6] K. Ikeuchi and B.K.P. Horn, "Numerical Shape from Shading and Occluding Boundaries," *Artificial Intelligence*, vol. 17, pp. 141-185, 1981.

[7] B.K.P. Horn and M.J. Brooks, "The Variatonal Approach to Shape from Shading," Memo AIM-813, MIT Artificial Intelligence Laboratory, March 1985.

[8] A. Brandt, "Multi-Level Adaptive Solutions to Boundary Value Problems," *Math. Comp.*, vol. 31, pp. 333-390, 1977.

[9] B. O'Leary and L. Jackel, "The 1969 Opposition Effect of Mars," *ICARUS*, vol. 13, pp. 437-448, 1970.

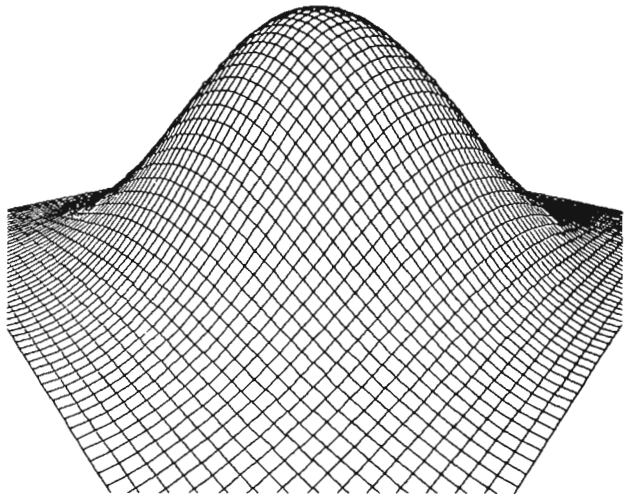

a) Raised Cosine.

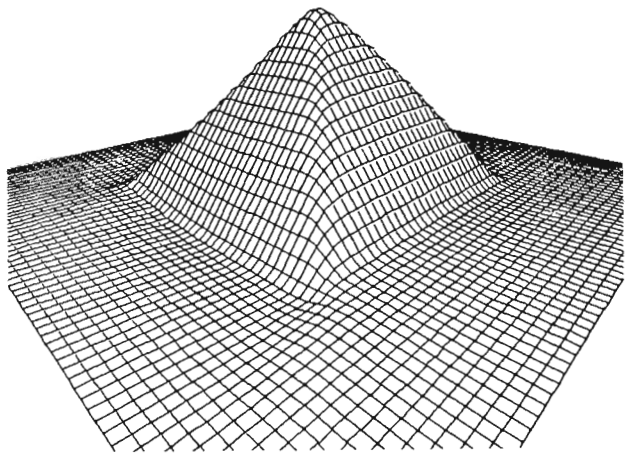

b) Pyramid.

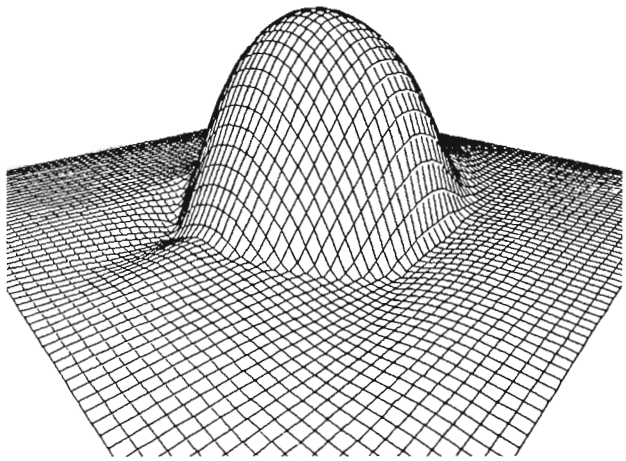

c) Hemisphere (Exaggerated Elevations).

Figure 3. Perspective Plots of Shapes Reconstructed from Synthetic Images.

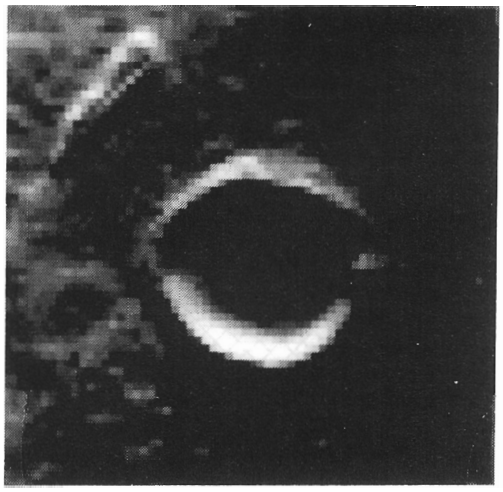

a) Image of an Impact Crater

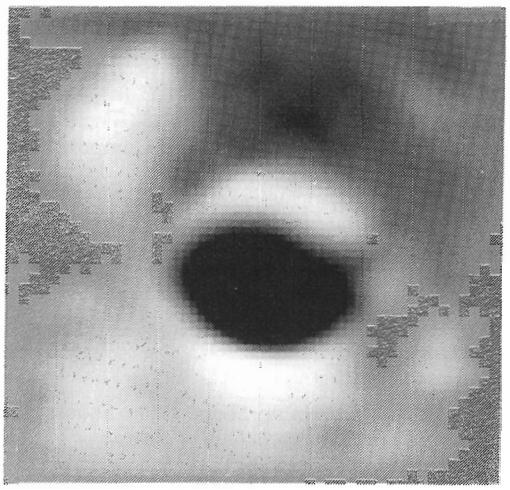

b) Half-Tone Coded Elevations Reconstructed from a)

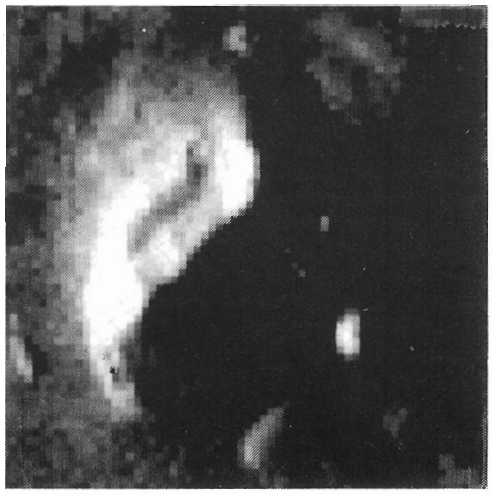

c) Image of a Mesa

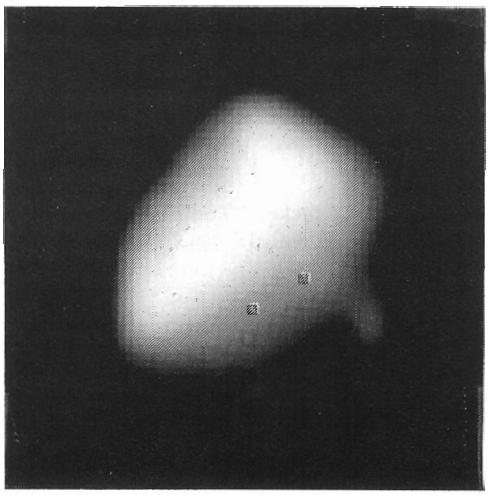

d) Half-Tone Coded Elevations Reconstructed from c)

Figure 4. Reconstructions of Mars Terrain