

APPLICATION OF SHAPE-FROM-SHADING TO SYNTHETIC APERTURE
RADAR

By

GLENN WILLIAM POPE

B.Sc., The University of British Columbia, 1984

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF
MASTER OF SCIENCE

in

THE FACULTY OF GRADUATE STUDIES
(DEPARTMENT OF COMPUTER SCIENCE)

We accept this thesis as conforming
to the required standard

THE UNIVERSITY OF BRITISH COLUMBIA

March 1990

© Glenn William Pope, 1990

In presenting this thesis in partial fulfilment of the requirements for an advanced degree at the University of British Columbia, I agree that the Library shall make it freely available for reference and study. I further agree that permission for extensive copying of this thesis for scholarly purposes may be granted by the head of my department or by his or her representatives. It is understood that copying or publication of this thesis for financial gain shall not be allowed without my written permission.

Department of Computer Science

The University of British Columbia
Vancouver, Canada

Date March 26th, 1990

Abstract

This thesis investigates the viability of applying a shape-from-shading technique to SAR imagery. A shape-from-shading algorithm is derived and tested on a single site for which both a Seasat SAR image and Digital Elevation Model (DEM) were available.

The shape-from-shading technique used in this thesis follows an approach proposed by Frankot and Chellappa for processing slant range SAR imagery. The algorithm incorporates a one-step technique for projecting non-integrable surface orientation estimates onto an integrable set in the frequency domain along with the iterative convergent shape-from-shading algorithm of Brooks and Horn. The significant issues and choices made in implementing the shape-from-shading algorithm and in preparing the SAR data and DEM are discussed.

The shape-from-shading algorithm was applied to both the test site SAR image and images synthesized from the DEM. Reflectance models were derived from the SAR image and DEM. By quantitatively comparing the shape-from-shading results with the initial conditions used for the experiments, it was found that the algorithm produced substantially better results when applied to the synthesized images; however, when applied to the SAR image, there was no significant improvement over the initial conditions.

Contents

Abstract	ii
Contents	iii
List of Tables	v
List of Figures	vi
Acknowledgement	viii
1 Introduction	1
2 Review of Literature	3
2.1 Energy Minimization	4
2.2 Integrability	6
2.3 Synthetic Aperture Radar	8
3 Theory of Shape-from-Shading Algorithm	10
3.1 Problem Framework	10
3.1.1 Imaging Geometry	12
3.1.2 Surface Orientation	14
3.1.3 Reflectance Function	15
3.2 Relaxation Algorithm	17
3.2.1 Continuous Form	17
3.2.2 Discrete Form	19
3.3 Integrability	22
3.3.1 Projection Technique	23
3.3.2 Fourier Expansion	25
3.4 Adaptation for SAR	27
3.4.1 Coordinate System	30

3.4.2	Reflectance Function	31
3.5	Summary	32
4	Software Implementation	35
4.1	Relaxation Algorithm	36
4.2	Integrability Projection	41
4.3	Analysis Tools	44
5	Data Preparation	47
5.1	Data Sources	47
5.2	DEM Registration	50
5.3	DEM Coordinate Transformation	55
6	Results	61
6.1	Simulated Data	61
6.2	Reflectance Functions for the Test Site	67
6.3	Experiments on the Test Site	72
7	Observations	83
8	Conclusions	86
	Bibliography	88

List of Tables

6.1	Gradients of Simulated Surfaces in Ground and Slant Coordinates	66
6.2	Results from Synthesized Images After 100 Iterations	76
6.3	Results from SAR Image After 100 Iterations	77
6.4	Results from Fine Resolution SAR Image After 100 Iterations	81

List of Figures

3.1	Imaging Geometry	11
3.2	SAR Imaging Geometry	28
4.1	Relaxation Algorithm Surface Orientation Estimates Data Flow	37
5.1	Slant Range SAR Image	49
5.2	DEM Registered to Slant Range SAR Image	53
5.3	Rectified SAR Image Registered to Slant Range SAR Image	54
5.4	DEM Rotation Geometry	59
6.1	Simulated Surface and Associated Image	63
6.2	Simulated Surface Rotated 45 Degrees and Associated Image	63
6.3	Profile of Surface Model (dotted line) and SFS Result (solid line) in Ground Coordinates	65
6.4	Profile of Surface Model (dotted line) and SFS Result (solid line) in Slant Coordinates	65
6.5	DEM of Test Site in Ground Coordinates	68
6.6	Histogram of Incident Angles Across Test Site	68
6.7	SAR Mean Reflectance Function: Mean SAR Pixel Value Per Degree of Incident Angle	69
6.8	Linear Regression Reflectance Function and Keydel Reflectance Function	69
6.9	Histograms of SAR Image and Synthesized Images	71
6.10	Surface Represented by Initial Conditions	73
6.11	SFS Results on Keydel image, 100 iterations	76
6.12	SFS Results on SAR Mean reflectance image, 100 iterations	77
6.13	SFS Results on SAR image using Keydel Reflectance Function	78
6.14	Initial Conditions: height difference and surface orientation error with respect to DEM	80
6.15	Synthesized Keydel results: height difference and surface orientation error with respect to DEM	80

6.16 SAR results: height difference and surface orientation error with respect to DEM	82
--	----

Acknowledgement

I would like to thank Dr. Bob Woodham for suggesting this topic, his insightful guidance, and his immense patience in helping me bring this to completion. A special thanks also to several people at MacDonald Dettwiler and Associates: Dr. Frank Wong, who handled many SAR questions; Bill Chong, who got me just the right SAR image; and Rob Reid, for access to the computing facilities. Last but not least, I would like to thank Lynn, for all her support and encouragement throughout; and Matthew, for the inspiration to get it done!

Chapter 1

Introduction

Shape-from-shading is the problem of recovering the shape of a smooth surface from shading information in a single image. Shape-from-shading techniques use the visible points of a surface in an image to compute variations in surface depth from variations in image brightness.

Frankot and Chellappa [FRAN87] have proposed applying shape-from-shading to Synthetic Aperture Radar (SAR) imagery. They have performed tests on synthesized SAR images and Seasat SAR images with apparent success. However, their published results have only been evaluated in terms of visual appearance; no quantitative comparison has been made with topographical data of the study sites.

The purpose of this thesis is to investigate the viability of applying shape-from-shading to SAR imagery. This is done by quantitatively comparing the results of experiments on a Seasat SAR image with a Digital Elevation Model (DEM). A single test site containing mountainous terrain is used for the experiments.

The algorithm used for the experiments is an implementation of the one outlined by Frankot and Chellappa in their two papers [FRAN87, FRAN88]. Their method incorporates a one-step technique for projecting non-integrable surface orientation estimates onto an integrable set in the frequency domain with the iterative convergent algorithm of Brooks and Horn [BROO85]. The significant issues and choices that were faced in implementing the shape-from-shading algorithm for this application, and in preparing the test data for the experiments are presented in this thesis, along with the results of the experiments.

Chapter 2

Review of Literature

Horn [HORN75] was the first to formulate a general form for the shape-from-shading problem, equating image irradiance to scene radiance, to give the non-linear partial differential equation, which is now commonly called the image irradiance equation:

$$I(x, y) = R(z_x, z_y) \tag{2.1}$$

The image irradiance equation makes explicit the difficult nature of the shape-from-shading problem: how to find the two components of surface orientation z_x and z_y given only a single image brightness value $I(x, y)$.

Horn [HORN77] also introduced the reflectance map, an important shape-from-shading tool, used to represent the surface radiance function $R(z_x, z_y)$. The reflectance map describes surface radiance values for all possible combinations of surface orientation given a particular imaging geometry and type of surface material. The reflectance map is graphically depicted using contour lines to represent points of surface orientation with equal

radiance values.

2.1 Energy Minimization

Ikeuchi and Horn [IKEU81] originated the idea of using a energy minimization framework for solving the shape-from-shading problem. Instead of strictly enforcing the equality between the two sides of the image irradiance equation (Equation 2.1), they posed the problem in terms of minimizing the square of the difference of the two sides of the equation across the whole image with the added constraint that the solution must be acceptably smooth:

$$\text{minimize } \int_x \int_y (I(x, y) - R(z_x, z_y))^2 + \lambda S(z_x, z_y) \quad (2.2)$$

Here $S(z_x, z_y)$ represents a measure of surface smoothness as a function of surface orientation, and λ is a regularization parameter, which is used to trade off surface smoothness against the accuracy of the estimated surface orientations as applied to the image irradiance model. The smoothness constraint is believed to select (it has not been formally proved) a unique set of surface orientations out of all possible sets of surface orientations which satisfy the image irradiance equation to a certain degree of accuracy. Altering the value of the regularization parameter selects a new set of surface orientations which optimally meet the constraints across the image. Ikeuchi and Horn used the sum of the squares of the partial derivatives of the surface orientation estimates for their smoothing

function:

$$S(z_x, z_y) = f_x^2 + f_y^2 + g_x^2 + g_y^2 \quad (2.3)$$

To derive their shape-from-shading algorithm, Ikeuchi and Horn applied the method of calculus of variations to find a function which minimizes Equation 2.2. The function being sought represents an extrema of the integral equation over the set of possible functions, and so must satisfy a set of associated Euler equations. From the Euler equations an iterative algorithm is obtained. The algorithm converges upon a solution given initial boundary conditions from which surface orientation values may propagate.

Ikeuchi and Horn derived their algorithm using stereographic coordinates (f, g) to represent surface orientation because they make it possible to include the surface orientations of occluding boundaries, which can be deduced from the edges in the image.

The algorithm can be adapted to the quality of the image using the regularization parameter. Increasing the weight of the smoothness component in the energy minimization equation offsets large errors or noise in an image. A high quality image, on the other hand, can be accurately processed with very little smoothing by minimizing the value of the regularization parameter. However, the regularization parameter introduces a new problem in that there is no clear means of deriving an appropriate value for it. Furthermore, the presence of the smoothness constraint may actually prevent the algorithm from converging to the correct solution and can even cause the algorithm to walk away from

the correct solution [HORN89]. Therefore, the Ikeuchi and Horn algorithm computes a smoothed approximation to the shape-from-shading problem as opposed to maintaining complete fidelity with the data.

A methodology for deriving an iterative algorithm from a energy minimization function using the calculus of variations was presented by Horn and Brooks in [HORN86b].

2.2 Integrability

For a surface with height $z(x, y)$, integrability is defined as

$$z_{xy}(x, y) = z_{yx}(x, y). \quad (2.4)$$

To obtain consistency between the surface orientation estimates z_x and z_y in shape-from-shading results, it is desirable to include an integrability constraint. The constraint also implies that the result of reconstructing surface height from surface orientation estimates will be independent of the integration path chosen. Without this constraint in the shape-from-shading problem, it is most likely that the surface orientation estimates will be non-integrable.

Ikeuchi and Horn proposed, as reported in [HORN86b], that an integrable surface could be derived by finding the best fit of an integrable surface to the possibly non-integrable results of their iterative shape-from-shading algorithm. In a manner similar to that in which they derived their shape-from-shading algorithm, they pose the surface

reconstruction problem as one of minimizing the square of the difference between the results of the shape-from-shading algorithm and the surface orientations of an integrable surface. Once again, the calculus of variations is used, Euler equations are found, and an iterative convergent algorithm is obtained. Boundary conditions for the algorithm are obtained by enforcing the integrable surface orientations along the same boundary curve used for initial constraints in the shape-from-shading algorithm.

Horn and Brooks [HORN86b] attempted to improve upon previous work on this problem by directly including integrability as a constraint in the original shape-from-shading problem. They derived a procedure for developing iterative shape-from-shading algorithms using the calculus of variations and applied their methodology to a energy minimization function which included the integrability constraint. Unfortunately, their effort yielded non-linear Euler equations for which they were unable to find a convergent iterative algorithm.

As an alternative, Horn and Brooks proposed to use the integrability constraint in place of the surface smoothness constraint used by Ikeuchi and Horn in the energy minimization equation (Equation 2.1). In this role, the integrability constraint serves as a penalty function to push the shape-from-shading solution towards integrability. With this approach, there is no guarantee of an integrable solution when the convergence stops since integrability is not strictly enforced. Only an approximately integrable solution is obtained.

Frankot and Chellappa [FRAN88] have proposed a technique that strictly enforces integrability. Possibly non-integrable surface orientation estimates in gradient-space are transformed into the frequency domain and projected onto a set of integrable orthonormal basis functions using a least-squares fit. The projection recovers surface height, as represented in the frequency domain, which may then be transformed back into the spatial domain to produce the surface, or used to derive new gradient estimates.

Frankot and Chellappa's method is similar to Ikeuchi and Horn's technique in that a projection is used to map non-integrable surface orientations to an integrable surface. The former approach restricts the set of integrable surfaces to those that can be constructed from the set of basis functions whereas the latter does not. On the other hand, Frankot and Chellappa's projection can be performed in a single operation, unlike Ikeuchi and Horn's technique which must iteratively converge to an acceptable solution.

Frankot and Chellappa have implemented their projection technique as an additional step in the iterative algorithm of Ikeuchi and Horn. In a further paper, [FRAN87], they adapted the algorithm for processing SAR data. It is this particular algorithm which is the focus of investigation of this thesis.

2.3 Synthetic Aperture Radar

Synthetic Aperture Radar (SAR) imaging differs from conventional imaging in that it employs an active sensing mechanism and uses the distance measuring property of radar.

It is a valuable alternative to optical imaging since the radar used for SAR imaging can reveal surface features not visible at optical wavelengths and can also penetrate cloud cover. Being an active imaging system, the image quality does not depend upon external illumination conditions that may vary with time of day, or season, like the position of the sun.

As an application area for shape-from-shading, SAR imagery is believed to be well suited since SAR data exhibits a high dependence of image pixel intensity on the surface topography as opposed to variations in the ground cover [WILD86, GUIN89]. The quality of the correlation between pixel intensity and surface orientation is a function of the incident angle [LILL87] with small angles (0 to 30 degrees) having a strong correlation, moderate angles (30 to 70 degrees) being more affected by surface roughness, and large angles (over 70 degrees) mostly falling into shadow.

SAR images have a lower signal-to-noise ratio than conventional images. The most dominant form of noise is speckle or clutter, which is caused by interference between the return signals from several sources on a surface [SWOR86]. Clutter noise has a high standard deviation and is most evident as bright speckles in the SAR image.


Chapter 3

Theory of Shape-from-Shading Algorithm

This chapter describes the derivation of the shape-from-shading algorithm implemented in this thesis. The theory behind the algorithm is presented in three stages. First, the basic algorithm based upon Horn and Brook's variational approach [HORN86b] is introduced. Second, the integrability technique of Frankot and Chellappa [FRAN88] is included. Last of all, the algorithm is adapted to the particular requirements of processing SAR imagery. To derive the algorithm it is first necessary to clearly establish the domain of the shape-from-shading problem.

3.1 Problem Framework

The shape-from-shading problem is formulated as the inverse of the image formation process, see Figure 3.1. Mathematically, this process is represented by the image irradiance equation, which equates surface radiance, emitted in the viewing direction, to the



where

\vec{i} is the illumination vector

\vec{n} is the surface normal

\vec{e} is the emittance vector

Figure 3.1: Imaging Geometry

brightness values registered in the image:

$$I(x, y) = R(\vec{i}, \vec{n}, \vec{e}, \rho) \quad (3.1)$$

The function $I(x, y)$ represents the brightness value at each point, (x, y) , in the image. The radiance function, R , is dependent upon four sets of parameters. Vectors \vec{i} and \vec{e} represent the direction from which the surface is illuminated and the direction in which energy is emitted from the surface to the image, respectively. The surface normal \vec{n} specifies surface orientation and is aligned perpendicular to a plane tangent to a point on the surface. The albedo, ρ , represents the reflectance properties of the surface material as determined by its composition, micro-structure, and the wavelength of the illuminating energy.

The first step to be taken in analyzing the shape-from-shading problem is to simplify the image irradiance equation as much as possible. The following subsections outline the simplifying assumptions that will be used.

3.1.1 Imaging Geometry

It is convenient to formulate the shape-from-shading problem using a camera-centred coordinate system for both the object and the image, as shown in Figure 3.1. The optical axis extends perpendicularly from the image plane, $I(x, y)$, and passes through the object. The optical axis is labelled as the z -axis in the (x, y, z) coordinate system. The origin of the coordinate system is placed on the opposite side of the surface from the

image plane, with the positive z -axis extending towards the image plane. The position of the origin with respect to the surface along the z -axis is otherwise set arbitrarily. The z -axis serves to measure surface height, so that surface points closer to the image plane will be considered higher than those closer to the origin of the z -axis.

The first assumption to be made in the imaging model is that the projection of surface points onto the image is orthographic. For convenience, the size of the object is considered to be scaled to the image so that a distance in xy coordinates on the image is the same distance in xy coordinates over the object. This means that for each point (x, y, z) on the surface, it will appear in the image at location (x, y) . Essentially, this is the same as specifying that the object to image distance is very large in comparison to the variations in surface height of the object, so that there are no perspective distortions in the image due to changes in surface relief. The orthographic projection provides the simplest mapping of surface locations to image locations. Equally important though, it implies that every ray of energy emitted from the surface that reaches the image travels in the same direction, parallel to the z -axis. In terms of the image irradiance equation, this means that the value of \vec{e} is constant across the whole surface.

A similar assumption is also made for the illuminating energy, that, for all points on the imaged surface, the lighting conditions, as represented by \vec{l} , are constant. Two interpretations of this condition are possible, either:

- a) point source illumination is being used, where the light sources are so distant from the surface that the incoming rays of light are collimated and uniform across the surface,

or,

- b) a hemispherical lighting condition exists, so that all points on the surface are receiving the same amount of incoming energy from any one particular direction.

Another assumption which is implied by specifying a constant illumination \vec{r} is that there are no multiple reflections of light off of the surface reaching the image, or that such occurrences have a negligible effect in the image. If this were not so, some surface points would have to be considered as having more than one illumination direction, considerably complicating the imaging model.

3.1.2 Surface Orientation

The shape-from-shading algorithm derived in this thesis follows the choice of Frankot and Chellappa [FRAN88] in using surface gradients to represent surface orientation. The reason they chose to use gradient space in their shape-from-shading algorithm is that it is the representation used as input and output to their integrability projection.

For a surface $z(x, y)$, the gradients are p and q , where,

$$p = z_x = \frac{\delta z}{\delta x} \quad \text{and} \quad q = z_y = \frac{\delta z}{\delta y}. \quad (3.2)$$

The gradient is simple to compute but is limited in representing surface slope. Any surface point that faces 90 degrees or more away from the viewer cannot be represented, which means that occluding boundaries cannot be represented. Occluding boundaries provide valuable constraints in solving shape-from-shading, as mentioned in Section 2.1, and cannot be precisely incorporated in a solution if the algorithm uses gradient-space. This problem can be overcome in part by using very large gradient values to approximate boundary values.

3.1.3 Reflectance Function

With the assumptions that have been made thus far, and the additional assumption that the surface material of the object is homogeneous, so that the albedo factor ρ can be considered a constant, it is possible to rewrite the surface radiance function in terms of just the two parameters of surface orientation:

$$R(\vec{r}, \vec{n}, \vec{e}, \rho) \implies R(p, q) \quad (3.3)$$

This simplified expression, $R(p, q)$, is called the reflectance function, or reflectance map. It gives an image irradiance (or image brightness) value for each possible surface orientation for a particular configuration of imaging geometry and type of surface material. The reflectance function may be derived empirically using a sample of the material and a controlled lighting environment to measure reflectance values at different surface orientations.

An alternate approach is to use a mathematical model for the reflectance function. Surfaces can generally be characterized as containing two different components of reflectance: specular and diffuse [WOOD84]. An ideal specular surface is a mirror, which reflects all incoming light in one direction as determined by the incident angle i , so that the emergent angle e is equal to it and lies in the same plane. A diffuse surface radiates incoming light in all directions so that the surface appears to have a flat finish as opposed to a glossy one. An ideal diffuse surface is called a Lambertian surface. Such a surface appears equally bright when viewed from any direction, as long as the illumination conditions are constant. For a distant point source of illumination, the reflection from a Lambertian surface can be modelled as the cosine of the incident angle i alone, since the emergent angle e determines the viewing direction. The incident angle i can be derived from the gradients of a surface patch by expressing the gradients in vector form and computing their normalized dot product to get the cosine of the angle of incidence. Therefore, in this case, the Lambertian function is just the normalized dot product of two vectors. Given a point source of illumination in direction $\vec{i} = (p_s, q_s, 1)^T$ and that the orientation of a surface patch is represented by the surface normal $\vec{n} = (p_n, q_n, 1)^T$, then the model for a Lambertian surface is

$$R(p, q) = \cos(i) = \left(\frac{\vec{i} \cdot \vec{n}}{\|\vec{i}\| \|\vec{n}\|} \right) = \left(\frac{1 + p_s p_n + q_s q_n}{\sqrt{1 + p_n^2 + q_n^2} \sqrt{1 + p_s^2 + q_s^2}} \right). \quad (3.4)$$

3.2 Relaxation Algorithm

With the assumptions made in the previous section, the general image irradiance equation (Equation 3.1) can be rewritten as

$$I(x, y) = R(p, q). \quad (3.5)$$

The problem of recovering surface orientation from $I(x, y)$ given a reflectance function $R(p, q)$ is ill-posed since there is generally no unique mapping of an image irradiance value to p and q . Instead, there are an infinite number of possible solutions for surface orientation. Further constraints must be employed to select a single surface orientation, preferably the correct one, from the set of possible solutions for a brightness value at a particular location in an image.

The steps in deriving the relaxation algorithm were formulated by Horn and Brooks [HORN86b] and are used here in presenting the basic shape-from-shading algorithm used in this thesis.

3.2.1 Continuous Form

The relaxation algorithm is derived by posing the shape-from-shading problem in a form that allows the “goodness” of a solution to be measured. The solution being sought is one in which the estimated surface best fits the given data according to a measure of smoothness imposed upon it. The equation which expresses this idea provides a

non-negative measure of the departure of the estimated surface from the ideal solution. By solving the equation for its minimum, the best solution is obtained. Frankot and Chellappa [FRAN88] use the following energy minimization equation:

$$\text{minimize } \int \int (I(x, y) - R(p, q))^2 + \lambda(p_x^2 + p_y^2 + q_x^2 + q_y^2) dx dy. \quad (3.6)$$

The first term of the integral provides the constraint of the image irradiance equation in the form of a mean-squared intensity error to account for noise in the image and errors in modelling the reflectance function of the surface. The second term is a quadratic measure of variation in surface slope. The factor λ is a Lagrangian multiplier that is used to adjust the relative influence of the two constraint terms. A large value of λ means that the smoothing constraint will play a more dominant role in the final solution than conformity to the image irradiance constraint.

When a minimization problem is posed in the form,

$$\int \int F(p, q, p_x, p_y, q_x, q_y), \quad (3.7)$$

the calculus of variations can be used. The calculus of variations is a technique for searching for the extrema of an expression in which the variable is a function (in which case it is called a functional). The functional in this problem is the field of gradients over the space of the surface. The conditions necessary to find an extrema of such an expression can be found with Euler equations. Since the integral is a combination of two unknown functions p and q , there will be two Euler equations, one for each function. For

Equation 3.7, which is a second-order partial differential equation, the form of the Euler equations are,

$$\begin{aligned} F_p - \frac{\delta}{\delta x} F_{p_x} - \frac{\delta}{\delta y} F_{p_y} &= 0, \\ F_q - \frac{\delta}{\delta x} F_{q_x} - \frac{\delta}{\delta y} F_{q_y} &= 0. \end{aligned} \quad (3.8)$$

When applied to Equation 3.6 they produce,

$$\begin{aligned} (I(x, y) - R(p, q))R(p, q)_p + \lambda \nabla^2 p &= 0, \\ (I(x, y) - R(p, q))R(p, q)_q + \lambda \nabla^2 q &= 0. \end{aligned} \quad (3.9)$$

The ∇^2 operator is the Laplacian operator where,

$$\nabla^2 = \frac{\delta^2}{\delta x^2} + \frac{\delta^2}{\delta y^2}. \quad (3.10)$$

Since there is no upperbound on the energy minimization equation, the Euler equations will only find an extrema which is also a minimum.

3.2.2 Discrete Form

The energy minimization equation, Equation 3.6, can be expressed in discrete form by considering how an element (i, j) of image I is used. The first part of the equation relating image irradiance to scene radiance is,

$$r_{ij} = [I_{ij} - R(p_{ij}, q_{ij})]^2. \quad (3.11)$$

The second part of the equation is a quadratic measure of the smoothness of the surface slopes. A measure of the partial derivatives of surface slope can be obtained using a finite difference approximation:

$$s_{ij} = \left[\left(\frac{p_{i+1j} - p_{i-1j}}{2} \right)^2 + \left(\frac{p_{ij+1} - p_{ij-1}}{2} \right)^2 + \left(\frac{q_{i+1j} - q_{i-1j}}{2} \right)^2 + \left(\frac{q_{ij+1} - q_{ij-1}}{2} \right)^2 \right] \quad (3.12)$$

Combined together, they form the discrete energy minimization equation:

$$e = \sum_i \sum_j (r_{ij} + \lambda s_{ij}). \quad (3.13)$$

To find the ideal minimal solution to the equation it is differentiated with respect to the two unknowns p_{ij} and q_{ij} ,

$$\begin{aligned} \frac{\delta}{\delta p_{ij}} e &= 2\lambda \kappa(\bar{p}_{ij} - p_{ij}) - 2[I_{ij} - R(p_{ij}, q_{ij})] \frac{\delta}{\delta p_{ij}} R(p_{ij}, q_{ij}), \\ \frac{\delta}{\delta q_{ij}} e &= 2\lambda \kappa(\bar{q}_{ij} - q_{ij}) - 2[I_{ij} - R(p_{ij}, q_{ij})] \frac{\delta}{\delta q_{ij}} R(p_{ij}, q_{ij}). \end{aligned} \quad (3.14)$$

The differentiated form of the smoothing term s_{ij} results in the terms $\kappa(\bar{p}_{ij} - p_{ij})$ and $\kappa(\bar{q}_{ij} - q_{ij})$ which are discrete approximations of the Laplacian operator ∇^2 . Here, \bar{p}_{ij} and \bar{q}_{ij} are the local averages surrounding the points p_{ij} and q_{ij} respectively. From the four term finite difference measure of surface slope smoothness a five-point approximation is obtained where,

$$\begin{aligned} \bar{p}_{ij} &= \frac{1}{4} [p_{i+1j} + p_{ij+1} + p_{i-1j} + p_{ij-1}], \\ \bar{q}_{ij} &= \frac{1}{4} [q_{i+1j} + q_{ij+1} + q_{i-1j} + q_{ij-1}]. \end{aligned} \quad (3.15)$$

Including adjacent corner elements in the discrete approximation of the Laplacian produces a nine-point approximation, whose lowest-order error term is of a higher order than that of the five-point approximation [HORN89]:

$$\begin{aligned}\bar{p}_{ij} &= \frac{1}{5} [p_{i+1j} + p_{ij+1} + p_{i-1j} + p_{ij-1}] + \frac{1}{20} [p_{i+1j+1} + p_{i-1j-1} + p_{i+1j-1} + p_{i-1j+1}], \\ \bar{q}_{ij} &= \frac{1}{5} [q_{i+1j} + q_{ij+1} + q_{i-1j} + q_{ij-1}] + \frac{1}{20} [q_{i+1j+1} + q_{i-1j-1} + q_{i+1j-1} + q_{i-1j+1}].\end{aligned}\quad (3.16)$$

A drawback of the nine-point approximation is that more computation is required. For the five-point Laplacian, κ equals 4, whereas for the nine-point Laplacian κ equals $\frac{10}{3}$.

To find the minima of the discrete energy minimization equation, its partial derivatives, as found in Equation 3.14, are set to zero. Rearranging this equation gives,

$$\begin{aligned}\kappa\lambda p_{ij} &= \kappa\lambda\bar{p}_{ij} + [I_{ij} - R(p_{ij}, q_{ij})] \frac{\delta}{\delta p_{ij}} R(p_{ij}, q_{ij}), \\ \kappa\lambda q_{ij} &= \kappa\lambda\bar{q}_{ij} + [I_{ij} - R(p_{ij}, q_{ij})] \frac{\delta}{\delta q_{ij}} R(p_{ij}, q_{ij}).\end{aligned}\quad (3.17)$$

From this set of equations an iterative algorithm is obtained:

$$\begin{aligned}p_{ij}^{n+1} &= \bar{p}_{ij}^n + \frac{1}{\kappa\lambda} [I(i, j) - R(p_{ij}^n, q_{ij}^n)] R_p(p_{ij}^n, q_{ij}^n), \\ q_{ij}^{n+1} &= \bar{q}_{ij}^n + \frac{1}{\kappa\lambda} [I(i, j) - R(p_{ij}^n, q_{ij}^n)] R_q(p_{ij}^n, q_{ij}^n).\end{aligned}\quad (3.18)$$

Here, the superscripts n and $n + 1$ refer to the iteration number of the surface gradient estimates.

The algorithm may be unstable in this form since two values of the gradient estimate, \bar{p}_{ij} and p_{ij} or \bar{q}_{ij} and q_{ij} are used in each equation. If the two values lie on opposite sides

of the actual gradient value to which they are converging, the algorithm will force the new gradient estimate away from it. As a result the shape-from-shading algorithm will oscillate around the solution or possibly diverge from it. To fix this problem, the average surface gradient can be used in all parts of the right-hand side of the equation. The final form of the basic shape-from-shading algorithm is,

$$\begin{aligned} p_{ij}^{n+1} &= \bar{p}_{ij}^n + \frac{1}{\kappa\lambda}[I(i, j) - R(\bar{p}_{ij}^n, \bar{q}_{ij}^n)]R_p(\bar{p}_{ij}^n, \bar{q}_{ij}^n), \\ q_{ij}^{n+1} &= \bar{q}_{ij}^n + \frac{1}{\kappa\lambda}[I(i, j) - R(\bar{p}_{ij}^n, \bar{q}_{ij}^n)]R_q(\bar{p}_{ij}^n, \bar{q}_{ij}^n). \end{aligned} \quad (3.19)$$

The iterative algorithm requires an initial set of boundary conditions where p_{ij} and q_{ij} are known in order to work. The algorithm can be run by repeating it for a fixed number of iterations or by running it until the energy minimization function (Equation 3.13) becomes smaller than a chosen value of ϵ . The value of λ chosen will also determine the success of the algorithm. If λ is too small, the algorithm will be unstable and unable to converge to a solution. If λ is very large, the errors in the estimated gradient values from the correct surface gradients may be very large.

3.3 Integrability

Integrability in surface slopes can be expressed as the constraint that,

$$\frac{\delta}{\delta y} z_x = \frac{\delta}{\delta x} z_y. \quad (3.20)$$

The surface slope estimates produced by the relaxation shape-from-shading algorithm generally do not meet this constraint, and so are non-integrable.

3.3.1 Projection Technique

Frankot and Chellappa [FRAN88] have proposed a technique to project possibly non-integrable surface slope estimates (\hat{z}_x, \hat{z}_y) onto a set of integrable surface slope estimates $(\tilde{z}_x, \tilde{z}_y)$. Their approach minimizes across the whole surface the distance between the integrable and non-integrable set of surface slopes according to the measure,

$$d\{(\hat{z}_x, \hat{z}_y), (\tilde{z}_x, \tilde{z}_y)\} = \int \int |\tilde{z}_x - \hat{z}_x|^2 + |\tilde{z}_y - \hat{z}_y|^2 dx dy. \quad (3.21)$$

The least-squares distance measure ensures that the projection of non-integrable surface slopes onto an integrable set of surface slopes is an orthogonal projection for all slopes.

To solve for this projection, a linear combination of a finite set of integrable basis functions $\phi(x, y, \omega)$ are used to represent the surface $z(x, y)$ so that,

$$z(x, y) = \sum_{\omega \in \Omega} C(\omega) \phi(x, y, \omega). \quad (3.22)$$

The set of expansion coefficients $C(\omega)$ spans a finite space Ω in two dimensions (ω_1, ω_2) . The components of surface slope are represented by the partial derivatives of Equation 3.22,

$$\begin{aligned} z_x(x, y) &= \sum_{\omega \in \Omega} C(\omega) \phi_x(x, y, \omega), \\ z_y(x, y) &= \sum_{\omega \in \Omega} C(\omega) \phi_y(x, y, \omega). \end{aligned} \quad (3.23)$$

However, if the basis functions $\phi_x(x, y, \omega)$ and $\phi_y(x, y, \omega)$ are assumed to be mutually orthogonal, then each set of surface slopes z_x and z_y can be represented by an independent set of expansion coefficients,

$$\begin{aligned} z_x(x, y) &= \sum_{\omega \in \Omega} C_1(\omega) \phi_x(x, y, \omega), \\ z_y(x, y) &= \sum_{\omega \in \Omega} C_2(\omega) \phi_y(x, y, \omega). \end{aligned} \quad (3.24)$$

The projection problem can now be expressed in terms of finding the set of expansion coefficients $\tilde{C}(\omega)$ of the integrable surface given the expansion coefficients for the non-integrable surface slopes $\hat{C}_1(\omega)$ and $\hat{C}_2(\omega)$. The final projection equation is derived by substituting in the expressions for the basis function representations of the surface into the distance minimization expression in Equation 3.21,

$$d\{(\hat{z}_x, \hat{z}_y), (\tilde{z}_x, \tilde{z}_y)\} = \int \int \left| \sum_{\omega \in \Omega} \tilde{C} \phi_x - \sum_{\omega \in \Omega} \hat{C}_1 \phi_x \right|^2 + \left| \sum_{\omega \in \Omega} \tilde{C} \phi_y - \sum_{\omega \in \Omega} \hat{C}_2 \phi_y \right|^2 dx dy. \quad (3.25)$$

This can be simplified to

$$d\{(\hat{z}_x, \hat{z}_y), (\tilde{z}_x, \tilde{z}_y)\} = \int \int \sum_{\omega \in \Omega} |\tilde{C} - \hat{C}_1|^2 |\phi_x|^2 + \sum_{\omega \in \Omega} |\tilde{C} - \hat{C}_2|^2 |\phi_y|^2 dx dy. \quad (3.26)$$

Interchanging the order of integration and summation produces

$$d\{(\hat{z}_x, \hat{z}_y), (\tilde{z}_x, \tilde{z}_y)\} = \sum_{\omega \in \Omega} |\tilde{C}(\omega) - \hat{C}_1(\omega)|^2 P_x(\omega) + |\tilde{C}(\omega) - \hat{C}_2(\omega)|^2 P_y(\omega), \quad (3.27)$$

where $P_x(\omega) = \int \int |\phi_x(x, y, \omega)|^2 dx dy$ and $P_y(\omega) = \int \int |\phi_y(x, y, \omega)|^2 dx dy$.

To find the set of expansion coefficients $\tilde{C}(\omega)$ that minimizes the distance equation, the above expression is differentiated with respect to $\tilde{C}(\omega)$. The terms in the summation expression can be minimized individually, giving,

$$\frac{\delta}{\delta \tilde{C}(\omega)} d\{(\hat{z}_x, \hat{z}_y), (\tilde{z}_x, \tilde{z}_y)\} = P_x(\omega)2[\tilde{C}(\omega) - \hat{C}_1(\omega)] + P_y(\omega)2[\tilde{C}(\omega) - \hat{C}_2(\omega)]. \quad (3.28)$$

By equating the right side of Equation 3.28 with zero in order to find the minimum, and by rearranging the equation to solve for $\tilde{C}(\omega)$, the final form of the projection equation is obtained,

$$\tilde{C}(\omega) = \frac{\hat{C}_1(\omega)P_x(\omega) + \hat{C}_2(\omega)P_y(\omega)}{P_x(\omega) + P_y(\omega)}. \quad (3.29)$$

3.3.2 Fourier Expansion

Frankot and Chellappa [FRAN88] chose to use the Fourier series as the set of basis functions for the integrability projection so that $\phi(x, y, \omega) = \exp^{j(\omega_1 x + \omega_2 y)}$. The Fourier representation of a measurable property over a finite two dimensional space uses two finite sets of sinusoidal basis functions, $(\omega_1, \omega_2) \in \Omega$, with wavefronts in orthogonal directions, to capture the full information content in the source space. Using Equation 3.22, an integrable surface can be represented in the frequency domain as,

$$\tilde{z}(x, y) = \sum_{\omega \in \Omega} \tilde{C}(\omega) \exp^{j\omega \cdot (x, y)}, \quad (3.30)$$

where $\tilde{C}(\omega)$ represents the coefficients of the Fourier series expansion that indicate the magnitude and phase of the sinusoidal spatial frequencies which make up the surface.

A convenient property of the frequency domain is that the derivatives of the surface they represent can be computed by a multiplication:

$$\phi_x = j\omega_x \phi, \quad (3.31)$$

$$\phi_y = j\omega_y \phi,$$

where ω_x and ω_y are the frequency domain representations of the derivative operator for the ω_1 and ω_2 dimensions respectively.

Using this property, the integrability projection as expressed in Equation 3.29 can be rewritten in terms of a set of Fourier basis functions. First, let the terms \hat{C}_x and \hat{C}_y represent the non-integrable surface orientation estimates z_x and z_y , then the mutually orthogonal sets of expansion coefficients can be computed as,

$$\begin{aligned} \tilde{C}_1(\omega) &= \frac{\hat{C}_x(\omega)}{j\omega_x} \\ \tilde{C}_2(\omega) &= \frac{\hat{C}_y(\omega)}{j\omega_y} \end{aligned} \quad (3.32)$$

The terms P_x and P_y are equivalent to the terms ω_x^2 and ω_y^2 for an expansion term in the Fourier series. Substituting terms, the integrability projection as expressed with the Fourier series is,

$$\tilde{C}(\omega) = \frac{-j\omega_x \hat{C}_x(\omega) - j\omega_y \hat{C}_y(\omega)}{\omega_x^2 + \omega_y^2}. \quad (3.33)$$

The new integrable surface slopes \tilde{z}_x and \tilde{z}_y can be computed in terms of their frequency coefficients $\tilde{C}_x(\omega)$ and $\tilde{C}_y(\omega)$ by taking the derivative of the surface $z(x, y)$ expansion

coefficients $\tilde{C}(\omega)$,


$$\tilde{C}_x(\omega) = j\omega_x \tilde{C}(\omega) \quad (3.34)$$

$$\tilde{C}_y(\omega) = j\omega_y \tilde{C}(\omega)$$

The only point at which the above equations are invalid is at $\omega = (0,0)$. This point corresponds to the information describing the average value of the surface represented, or the DC component in the frequency domain. For the coefficients of the integrable surface slopes \tilde{C}_x and \tilde{C}_y , the value of the coefficient for the expansion term $(0,0)$ can be acquired from the coefficients of the non-integrable surface slopes \hat{C}_x and \hat{C}_y .

3.4 Adaptation for SAR

A SAR sensor scans the terrain along a line perpendicular to its direction of flight (see Figure 3.2). The direction of flight is known as the azimuth direction of the sensor and it corresponds to the y -axis direction in the figure. Perpendicular to the azimuth along the ground is the range direction of the sensor which corresponds to the x -axis. A radar pulse signal is emitted from the sensor out towards the ground in the range direction. The returned signal is partitioned by time into cells that correspond to pixels along a single strip in the range direction. By repeating this process as the sensor travels in the azimuth direction, the adjacent strips can be used to construct a 2-dimensional image. The slant range resolution of a pixel can be computed from the duration of the radar



where

- d is the depression angle
- l is the look angle
- S is the slant range resolution
- G is the ground range resolution

Figure 3.2: SAR Imaging Geometry

pulse signal, the actual method used depends upon the type of SAR processing involved. For an introduction to SAR image formation and SAR processing see [LILL87], for a detailed description see [WONG84].

The imaging geometry for SAR is simplified considerably by the fact that the illumination source and viewing apparatus are in the same location. As a result, the incident angle i of the radar pulse to the surface equals the emittance angle e of the reflected pulse back to the sensor and the phase angle g between the two angles is zero. Therefore, only half as much information is required to establish the geometry for SAR shape-from-shading as for optical applications.

Geometric distortion in SAR imagery due to variations in terrain can be a significant problem in rugged areas [WONG84, LILL87]. Relief displacement occurs in the range direction and is caused by points of higher ground returning the radar pulse earlier than would occur if they were lower, since increased ground elevation reduces the slant range distance to the sensor. As a result, slopes facing the sensor in the range direction are foreshortened at a rate depending upon the angle of the slope. If the slope exceeds a critical limit then pixel layover occurs. This shows up in a SAR image when a target further away from the sensor along the ground appears to be closer to the radar in the image than other targets which are closer on the ground. Radar shadow occurs on slopes facing away from the sensor in the range direction. The extent of relief displacement and shadowing depends upon the severity of the terrain and the look angle of the sensor.

With a large look angle, relief displacement is minimized and radar shadows are large.

A small look angle will give greater relief displacement and minimal radar shadow.

3.4.1 Coordinate System

Frankot and Chellappa [FRAN87] devised their shape-from-shading algorithm to work on SAR strip maps which measure the range distance along the slant travelled by the radar signal as opposed to along the ground. The relaxation form of the shape-from-shading algorithm works on the principle that the projection of surface points to image points is orthographic (see Section 3.1.1). In order to maintain this property, it is therefore necessary to introduce a new coordinate system (r, y, u) to process the SAR imagery. The axes r , y , and u are mutually orthogonal with the r -axis running parallel to the direction of the incident and reflected radar signal and the y -axis extending in the azimuth direction. The (r, y, u) coordinate system is rotated from the ground coordinate system (x, y, z) about the y -axis.

In introducing a new coordinate system two assumptions are made. First, that the slant range direction is constant across the range of the surface imaged. Second, that the relation between the ground range and slant range coordinate systems can be modelled by a rotation. In reality, the slant range direction of the SAR signal does change along the range of the image in a non-linear fashion. However, given a large distance between the sensor and the ground imaged this effect is reduced, and over a small range distance in

an image is probably minimal. The correction of slant range to ground range actually is more appropriately modelled by a hyperbolic function in order to make spacing of image points proportional to horizontal ground features [LILL87]. For large range distances the earth's curvature is another factor which must be considered [WONG84].

3.4.2 Reflectance Function

The interaction between the radar signal and a surface is modelled by the radar cross section σ . The radar cross section is a function which measures the energy directed from a target back to the sensor [SWOR86]. It depends on target size, shape, material, direction of view and illumination, radar frequency and polarisation. The reflectance function of a radar signal σ_o is obtained by normalizing the radar cross section σ by the average area illuminated by the radar A .

$$\sigma_o(\alpha) = \frac{\sigma(\alpha)}{A(\alpha)} \quad (3.35)$$

Both σ and A are functions of the incident angle α .

Frankot and Chellappa [FRAN87] chose to use a semi-empirical model for the radar cross section called a generalized Lambert model. Keydel [KEYD82] describes this model as a good estimator for the average σ_o of rough surfaces,

$$\sigma_o(\alpha) = \begin{cases} \gamma \frac{\cos^\mu \alpha_g}{\sin^\nu \alpha_g}, & \text{for } 0^\circ \leq \alpha \leq \alpha_g \\ \gamma \frac{\cos^\mu \alpha}{\sin^\nu \alpha}, & \text{for } \alpha_g < \alpha \leq 90^\circ \end{cases} \quad (3.36)$$

where μ , ν , γ , and α_g (normally, $\alpha_g \leq 10^\circ$) are arbitrary constants that are chosen to fit the data based on measurements obtained for the surface type. Frankot and Chellappa

chose a simplified form of this function, with $\mu = 2$, $\nu = 1$, $\gamma = 1$, and α_g not used but replaced with a small constant $\delta \sim 10^{-4}$ in the denominator which is used as a bias to prevent σ_o from becoming infinite when $\alpha \rightarrow 0$.

Surprisingly, Barrick [BARR70] in his discussion of scattering models for rough surfaces claims that this particular form of the generalized Lambert model correlates poorly with measured results for most terrain and natural surfaces. The basic assumption underlying this model is that the terrain is a diffuse scatterer, not specular. It works best when the incident beam is somewhat away from the surface normal, so that the incident angle α is not too close to zero.

3.5 Summary

Based on the material presented thus far in this chapter, the shape-from-shading algorithm used in this thesis is summarized here. The algorithm uses the (r, y, u) coordinate system for processing SAR imagery where r represents the range direction, y represents the azimuth direction, and u represents the surface height relative to the (r, y) plane. The relaxation algorithm originally proposed by Ikeuchi and Horn [IKEU81], modified to use surface gradients, is combined with the integrability projection of Frankot and Chellappa [FRAN88] to create the algorithm used by Frankot and Chellappa [FRAN87] for processing SAR imagery. The steps are as follows:

1. Set the boundary conditions in the surface slope estimates u_r and u_y .

2. Use the relaxation technique to propagate the initial boundary conditions and converge towards a solution,

$$\begin{aligned} u_r^{n+1} &= \bar{u}_r^n + \frac{1}{\kappa\lambda} [I(r, y) - R(\bar{u}_r^n, \bar{u}_y^n)] R_r(\bar{u}_r^n, \bar{u}_y^n), \\ u_y^{n+1} &= \bar{u}_y^n + \frac{1}{\kappa\lambda} [I(r, y) - R(\bar{u}_r^n, \bar{u}_y^n)] R_y(\bar{u}_r^n, \bar{u}_y^n). \end{aligned} \quad (3.37)$$

3. Apply the Fourier transform to the surface slope estimates u_r^{n+1} and u_y^{n+1} to get the frequency domain expansion coefficients $\hat{C}_r(\omega)$ and $\hat{C}_y(\omega)$ of the non-integrable surface slopes.
4. Apply the integrability projection to get the frequency coefficients of the estimated surface,

$$\tilde{C}(\omega) = \frac{-j\omega_r \hat{C}_r(\omega) - j\omega_y \hat{C}_y(\omega)}{\omega_r^2 + \omega_y^2}. \quad (3.38)$$

5. Compute the derivative of the surface in the frequency domain to get a set of integrable surface slopes $\tilde{C}_r(\omega)$ and $\tilde{C}_y(\omega)$,

$$\tilde{C}_r(\omega) = j\omega_r \tilde{C}(\omega) \quad (3.39)$$

$$\tilde{C}_y(\omega) = j\omega_y \tilde{C}(\omega)$$

6. Apply the inverse Fourier transform to the frequency coefficients $\tilde{C}_r(\omega)$ and $\tilde{C}_y(\omega)$ to get the integrable surface slope estimates \tilde{u}_r^{n+1} and \tilde{u}_y^{n+1} .
7. Repeat the previous six steps, using \tilde{u}_r^{n+1} and \tilde{u}_y^{n+1} as input to the next iteration.

The algorithm is either run for a fixed number of iterations or until an acceptable degree of convergence is indicated by the energy minimization equation,

$$\int \int (I(r, y) - R(u_r, u_y))^2 + \lambda(u_{rr}^2 + u_{ry}^2 + u_{yr}^2 + u_{yy}^2) dr dy. \quad (3.40)$$

The final form of the surface $u(r, y)$ can be constructed by applying the inverse Fourier transform to the frequency coefficients of the surface $\tilde{C}(\omega)$.

Chapter 4

Software Implementation

The important choices and significant details of implementing the shape-from-shading algorithm, as derived in Chapter 3, are presented in this chapter. The software was developed in three stages:

1. The relaxation algorithm of Ikeuchi and Horn [IKEU81] was implemented, using surface gradients instead of stereographic coordinates.
2. The integrability projection of Frankot and Chellappa [FRAN88] was added.
3. The software tools required to evaluate the shape-from-shading results were developed.

At the end of each development stage the software was tested on simulated data to verify its correctness. The software is written in Common LISP.

Four basic data types are handled by the shape-from-shading software: surface heights, surface derivatives, frequency coefficients, and image pixels. The 2-dimensional arrays for

storing intermediary values and results for these data types use fixed point representations. All operations on the data use floating point arithmetic to maximize the precision of the calculations. Both types of surface data as well as the frequency coefficients are stored as signed 32-bit values. Image pixels are 8-bit unsigned values.

4.1 Relaxation Algorithm

As explained in Section 3.1.2, gradient space was chosen to represent surface orientation in the relaxation algorithm. Gradient values are computed from the surface model either for providing initial boundary constraints, or for generating a simulated image of the surface. Numerical differentiation is done using finite-difference calculations. The central difference expression,

$$p = \frac{z(x+1, y) - z(x-1, y)}{2b} \quad \text{and} \quad q = \frac{z(x, y+1) - z(x, y-1)}{2c}, \quad (4.1)$$

is used except for along array boundaries where either the forward-difference or backward-difference expression is required. The parameters b and c represent the units used for pixels along the x and y dimension, respectively.

This implementation of the relaxation algorithm uses three sets of 2-dimensional array structures, with coordinates (x, y) for simulated scenes and (r, y) for SAR imagery, to hold surface orientation values p and q through the processing stages, see Figure 4.1. The first set of array structures contains the initial surface orientation estimates which,

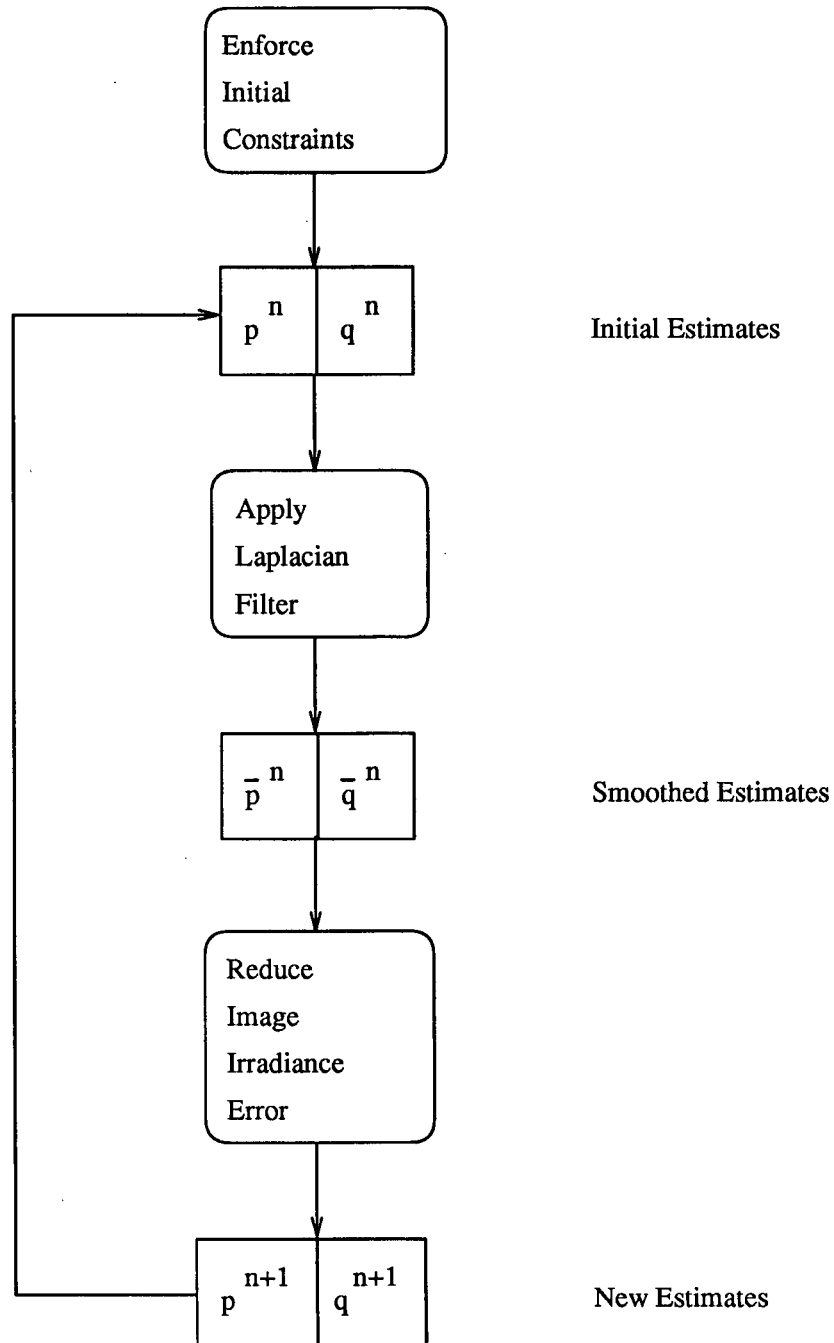


Figure 4.1: Relaxation Algorithm Surface Orientation Estimates Data Flow

except for the first iteration, are produced by the previous iteration of the algorithm. On the first iteration, the surface orientation estimates are arbitrarily set to values which represent a horizontal surface as an initial guess for the surface shape. For ground range processing in (x, y, z) coordinates, this value is zero for both p and q . For slant range processing in (r, y, u) coordinates, q is still zero but p becomes set to the tangent of the sensor depression angle.

For initial constraints specific to the problem domain, points of known surface orientation, obtained from surface boundaries, singular points or a priori knowledge of the surface shape, are then included, overriding existing estimates. An interpolation of the surface from these special points could be done prior to starting the shape-from-shading algorithm in order to start off with a more reasonable guess of the shape. This was not done in this implementation. Frankot and Chellappa's approach for initially setting the field of surface orientations is not known.

The initial domain specific constraints can be treated either as initial guesses about the shape of the surface, or taken as factual knowledge about the surface. In this implementation, the latter situation is assumed, so that the points of known surface orientation are enforced at the beginning of each iteration of the shape-from-shading algorithm. If the initial constraints are just conjectures of the surface shape, then it is probably more appropriate to phase out their influence on the computed estimates as further iterations of the algorithm are run. In this situation, the initial constraints serve to push the

shape-from-shading algorithm towards a solution; however, they should be removed as the solution converges so that there is not a bias towards weakly supported knowledge. Frankot and Chellappa did use initial surface gradient constraints in some of their experiments [FRAN88]; however, it is not clear whether they enforced them on every iteration or not.

The second set of array structures contains the results of smoothing the initial surface orientation estimates. A nine-point smoothing filter is convolved with the initial estimates using the 3 by 3 kernel,

$$k \cdot \begin{pmatrix} 1 & 4 & 1 \\ 4 & 0 & 4 \\ 1 & 4 & 1 \end{pmatrix}, \quad (4.2)$$

where k equals $1/20$, except for when processing along array boundaries where k equals $1/9$ for array corners and $1/14$ array edges. This is the same smoothing filter used by Frankot and Chellappa.

The final processing step uses the smoothed surface orientation estimates to compute the next iteration of estimates (see Equation 3.19 or Equation 3.37). The value of the reflectance function in this calculation is based upon the incident angle of illuminating energy. As mentioned in Section 3.1.3, the incident angle can be derived from the inverse cosine of the normalized dot product of the vectors for the surface normal \vec{n} and the direction of illumination \vec{r} . For processing SAR data in the (r, y, u) coordinate system, the fact that the sensor's illuminating energy travels parallel to the r -axis means that

the illumination direction can be expressed as $\vec{i} = (-1, 0, 0)^T$. The surface normal is computed from the cross product of the gradient vectors to give $\vec{n} = (-u_r, -u_y, 1)$. The equation used to calculate the incident angle α is,

$$\alpha = \cos^{-1} \left(\frac{u_r}{\sqrt{u_r^2 + u_y^2 + 1}} \right). \quad (4.3)$$

This is the technique used by Frankot and Chellappa [FRAN87].

Derivatives of the reflectance function $R(p, q)$ are calculated using the central finite-difference expression,

$$R_p(p, q) = \frac{R(p + \Delta, q) - R(p - \Delta, q)}{2\Delta} \quad \text{and} \quad R_q(p, q) = \frac{R(p, q + \Delta) - R(p, q - \Delta)}{2\Delta}, \quad (4.4)$$

where Δ is set to a very small value. Frankot and Chellappa used numerical derivatives for the reflectance function, but the form they used is not known.

Frankot and Chellappa [FRAN88] made a provision in their relaxation algorithm for handling shaded regions. By identifying the shaded areas in their imagery prior to processing, they were able to skip over processing these areas so that the only possible propagation of surface orientation estimates across these regions was by the smoothing filter. The idea behind their approach is to prevent image areas which lack information about the surface from influencing the shape-from-shading result. This particular feature was not included in the relaxation algorithm used here in order to see what the effects of radar shadows would be in the final result.

4.2 Integrability Projection

The integrability projection is performed in the frequency domain. The non-integrable surface slope estimates \hat{p} and \hat{q} produced by the relaxation algorithm are transformed using the discrete Fourier transform,

$$F(\omega_1, \omega_2) = \sum_{x=1}^N \sum_{y=1}^N f(x, y) \exp \frac{2\pi j}{N} (\omega_1 x + \omega_2 y), \quad (4.5)$$

for an $N \times N$ array of data, to get the frequency coefficients \hat{C}_x and \hat{C}_y . After projecting the estimates and finding the derivatives \tilde{C}_x and \tilde{C}_y of the integrable surface, the inverse discrete Fourier transform is applied,

$$f(x, y) = \frac{1}{n^2} \sum_{\omega_1=1}^N \sum_{\omega_2=1}^N F(\omega_1, \omega_2) \exp \frac{2\pi j}{N} (\omega_1 x + \omega_2 y), \quad (4.6)$$

to get the integrable surface slope estimates \tilde{p} and \tilde{q} . The fast Fourier transform (FFT) algorithm is used to compute the transformation. An implicit assumption in using the FFT is that the surface slopes are periodic along both x and y axes. No measures were taken in the software to achieve this condition, so abrupt discontinuities in slope values may occur along array boundaries. Frankot and Chellappa also used the periodic form of the FFT in their work.

Derivative operators $a_x(\omega_1)$ and $a_y(\omega_2)$ are needed in computing the integrability projection (see Equation 3.33) and the integrable slope estimates \tilde{C}_x and \tilde{C}_y (see Equation 3.34) in the frequency domain. The Fourier transform of the central difference

operator (Equation 4.1) is taken to obtain the frequency coefficients of the operator so that,

$$\begin{aligned} a_x(\omega_1) &= \frac{1}{2b} \exp^{j\omega_1} - \frac{1}{2b} \exp^{-j\omega_1} = \frac{1}{2b} j \sin(\omega_1), \\ a_y(\omega_2) &= \frac{1}{2c} \exp^{j\omega_2} - \frac{1}{2c} \exp^{-j\omega_2} = \frac{1}{2c} j \sin(\omega_2). \end{aligned} \quad (4.7)$$

The indices ω_1 and ω_2 are integer multiples of $\frac{2\pi n}{N/2}$ where $n \in (0, 1, \dots, (N/2) - 1)$ for an $N \times N$ sized image. Essentially this means that there are two oscillations of the sine wave along either the ω_1 or ω_2 dimension. As with the surface gradients, the parameters b and c are set to the pixel units for the derivative operator in the frequency domain. Frankot and Chellappa appear to have used a slightly different formulation of the derivative operator [FRAN87], with only a single sinusoidal oscillation along either the ω_1 or ω_2 dimension and no apparent units or scaling factor. Substituting the values for $a_x(\omega_1)$ and $a_y(\omega_2)$ into Equation 3.33 gives the implemented form of the integrability projection,

$$\tilde{C}(\omega) = \frac{-a_x(\omega_1)\hat{C}_x(\omega) - a_y(\omega_2)\hat{C}_y(\omega)}{|a_x(\omega_1)|^2 + |a_y(\omega_2)|^2}, \quad (4.8)$$

and from Equation 3.34 the form of the derivative operation,

$$\tilde{C}_x(\omega) = a_x(\omega_1)\tilde{C}(\omega), \quad (4.9)$$

$$\tilde{C}_y(\omega) = a_y(\omega_2)\tilde{C}(\omega).$$

When either ω_1 or ω_2 is an integer multiple of π then the value of the corresponding derivative operator ω_x or ω_y is zero. If both coefficients are zero simultaneously then the

integrability projection cannot be evaluated. When this occurs, $\tilde{C}(\omega_1, \omega_2)$ is set to zero since the divisor term of the integrability projection serves just to normalize the result of the projection.

The integrability projection is not valid at the point $\omega = (0, 0)$. As a result, the DC component of the slope information is lost in the projection. For primarily horizontal surfaces this is not a problem since the DC component of the surface slope in the x or y direction is zero. However, when processing SAR imagery, a horizontal surface in the (x, y, z) coordinate system becomes a slanted surface along the range axis of the (r, y, u) coordinate system. Therefore, the DC component of the range slope estimates is non-zero and will be lost in the projection. In order to recover a correct set of integrable surface slopes the DC component must be retained. In this implementation, the DC components of the slope estimates are preserved independently of the projection technique so that,

$$\hat{C}_x(0, 0) \text{ is assigned to } \tilde{C}_x(0, 0) \quad \text{and} \quad \hat{C}_y(0, 0) \text{ is assigned to } \tilde{C}_y(0, 0). \quad (4.10)$$

Initial constraints can be introduced in the frequency domain as well as the spatial domain. Low frequency coefficients of the integrable surface $\tilde{C}(\omega)$ can be replaced with low frequency coefficients from the Fourier transform of an existing low resolution model of the surface. When working in the ground range (x, y, z) coordinate system, it is sufficient to take the FFT of the surface heights $z(x, y)$ and copy the low frequency coefficients into the integrable surface $\tilde{C}(\omega)$. In the slant range this will not work. An

abrupt shift in surface height $u(r, y)$ between the two ends of the r -axis is typical for a surface rotated into the slant range coordinate system. The FFT of the slant range surface $u(r, y)$ will include information describing the abrupt shift in surface heights. This information is not required when working with surface orientations only, and in fact will be misleading if used as a constraint. To overcome this problem, surface gradients of $u(r, y)$ are computed, transformed to the frequency domain, and with the integrability projection, used to establish the frequency coefficients for the surface. Frankot and Chellappa do use low frequency constraints in some of their experiments; however, they did not indicate that any special considerations were necessary for slant range processing.

The projection constraint can be used either as a part of the iterative shape-from-shading algorithm or as the final processing step after running the relaxation algorithm. In this implementation the projection constraint was included in the iterative algorithm in order to see its effects on convergence to a solution, as was done by Frankot and Chellappa.

4.3 Analysis Tools

Several analysis tools were developed to assist in evaluating the performance of the shape-from-shading algorithm. The energy minimization equation is useful for measuring the

convergence of the algorithm. Its discrete form is,

$$e = \sum_{i=1}^N \sum_{j=1}^N (I(i, j) - R(p(i, j), q(i, j)))^2 + \lambda(p_x(i, j)^2 + p_y(i, j)^2 + q_x(i, j)^2 + q_y(i, j)^2), \quad (4.11)$$

where the partial derivatives p_x , p_y , q_x , and q_y are computed using the central, forward or backward finite-difference equations.

A second tool was developed to obtain a more accurate picture of the change in surface orientation estimates between iterations. The error in the slope estimates p_{est} and q_{est} was computed by comparing them with the gradients of the original surface p_{real} and q_{real} . The angle of error was measured by taking the inverse cosine of the normalized vector dot product of the surface slope estimates $\vec{n}_{est} = (p_{est}, q_{est}, -1)^T$ and the actual slope values $\vec{n}_{real} = (p_{real}, q_{real}, -1)^T$,

$$\theta_e = \cos^{-1} \left(\frac{\vec{n}_{est} \cdot \vec{n}_{real}}{\|\vec{n}_{est}\| \|\vec{n}_{real}\|} \right). \quad (4.12)$$

Using this measure, the mean and standard deviation in surface orientation error θ_e is computed across the whole surface. Frankot and Chellappa [FRAN88] use this type of measure themselves, but have not indicated how they compute it.

A surface model of the shape-from-shading surface orientation estimates can be constructed by applying the inverse Fourier transform to the projected surface $\tilde{C}(\omega)$. To make the average surface heights equal, the DC component of the original surface must be copied to the normalized integrable surface. For recovering a surface in the slant range coordinate system, information describing the shift between the two ends of the

r -axis must be included. This information can be obtained by taking the FFT of a low resolution surface height model $u(r, y)$ and copying across the frequency coefficients with no azimuth component, that is for $\omega_2 = 0$. If this is not done an approximately flat surface will be recovered as opposed to a surface slanted along the r -axis. Frankot and Chellappa make no mention of this problem in their work.

Chapter 5

Data Preparation

This chapter describes how the test site data was prepared for the shape-from-shading experiments. To execute and evaluate the experimental results a SAR image and surface model are required in slant range coordinates. The SAR imagery is originally in this coordinate space, the surface model is not. The task of transforming the surface model into a form that is aligned with the SAR image and is in slant range coordinates is the subject of the following sections.

5.1 Data Sources

The test site is an 11.1 km² area on Vancouver Island, British Columbia, west of Shawnigan Lake centred at geographic coordinates W 123:50:00, N 48:35:30. The area is mountainous terrain, ranging from 140 to 880 metres in elevation. The landscape is primarily forest; however, a sizeable portion has been harvested so there are areas of old growth forest, second growth forest and recent logging, as indicated in previous studies of the

area [WOOD85].

The SAR image of the site is from the Seasat-1 satellite that was launched June 27, 1978 and operated only 99 days before failing. The satellite orbited the earth at an altitude of 800 km and used an L-band radar (23.5 cm wavelength) with HH polarization, imaging swaths 100 km across track. The actual image used for this thesis is from orbit 193. The spacecraft was travelling approximately North-North-West relative to the site, and scanning North-East-East in the range direction. The imagery was processed to a slant range resolution of 12.5 metres in both the range and azimuth direction for a sensor look-angle of 20.5 degrees and is shown in Figure 5.1. In the figure, the sensor is to the left of the image with the range direction extending to the right of the page and the azimuth direction towards the top of the page, both axes parallel to the image boundaries. The bright lines in the image correspond to slopes facing the sensor.

A Digital Elevation Model (DEM) generated for another study involving this area [WOOD85] was used for elevation data. The DEM was manually digitized from a 1:50 000 Class A1 map sheet 92B/12 (Shawnigan Lake) from the Canadian National Topographic Series (NTS). The DEM is on a regular grid with 60 metre spacing and is aligned to the Universal Transverse Mercator (UTM) projection of the source map.




Figure 5.1: Slant Range SAR Image

5.2 DEM Registration

The SAR image is not in a UTM map projection. It must be kept in the slant range projection in order to preserve the orthogonality of the range and azimuth axes as well as the alignment of the image grid to these axes. These conditions are required to maintain a constant direction of illumination from the sensor across the image and to simplify the calculation of the incident angle for the reflectance function. Therefore, to compare the shape-from-shading results from the SAR image with the DEM, the most convenient way is to match the DEM spatially to the SAR image.

A spatial transformation from the UTM map projection to the slant range projection of a SAR image is difficult to model due to the terrain dependent effects of relief distortion in SAR imagery (see Section 3.4). The non-linearity between ground range coordinates and slant range coordinates is another distortion to be considered as well. Instead of using complex models to derive a transformation, it was decided to create a simple transformation based on the registration of a few DEM points to their corresponding locations in the slant range SAR image. Accurately identifying corresponding points between the DEM and SAR image is difficult. A version of the SAR image that had been rectified to the DEM in the UTM map projection, taking into account relief displacement effects, was obtained from [WONG84] and used to facilitate this process. By scaling the DEM to the rectified SAR image, the transformation used to register the rectified SAR

image to the slant range SAR image was used to register the DEM as well.

The procedure used to register the DEM to the slant range SAR image is as follows:

1. The DEM was scaled from a 60 metre grid spacing to the 12.5 metre grid spacing of the rectified SAR image.
2. Both the DEM and the rectified SAR image were rotated 26 degrees clockwise from their North up alignment. This was done to give both the rectified SAR image and the slant range SAR image the same orientation so that registration points could be marked accurately.
3. An affine transformation from the rotated rectified SAR image to the slant range SAR image was defined and applied to the DEM. The affine transformation from an (x, y) grid to a (u, v) grid has six coefficients,

$$u = a_0 + a_1x + a_2y \tag{5.1}$$

$$v = b_0 + b_1x + b_2y,$$

which can be completely defined using three registration points. Sharp curves along lake shorelines were used to accurately identify and match the registration points. The points were chosen to be on lakes as close to each other in elevation as possible to minimize distortion due to relief displacement and at the same time to be as far apart as possible horizontally to prevent large errors from accumulating outside the

triangle of the three points at the outer edges of the DEM.

The result of applying the transformation to the DEM and the rectified SAR image is shown in Figure 5.2 and Figure 5.3 respectively. Comparing Figure 5.1 and Figure 5.3 it is apparent that the bright lines of the slant range SAR image are much thicker in the rectified image. This “smearing” occurs only in the range direction and is an artifact of the correction process used to compensate for the slope foreshortening effect in the slant range image. Essentially, wherever there was a mapping of one pixel to many locations, no radiometric correction was performed to distribute the energy of the pixel amongst the rectified pixels. Instead the intensity of the slant range pixel was copied to all the rectified pixels it mapped to.

The accuracy of the transformation was verified by comparing the range position of several reference points between the registered rectified SAR image and the slant range SAR image. The transformation was considered to be sufficiently accurate only if, for each reference point, the deviation in range between the two images was less than the change in distance along the ground that could be attributed to relief displacement. Relief displacement is a function of the sensor’s look-angle as well as range distance across the swath (see [WONG84]). It provides a measure of the displacement distance of a target along the ground in an image given the height of the target above the base ground level at which there are no displacement effects. For ground range displacement, this relation



Figure 5.2: DEM Registered to Slant Range SAR Image



Figure 5.3: Rectified SAR Image Registered to Slant Range SAR Image

can be approximated by a constant of proportionality equal to the tangent of the sensor depression angle θ . Therefore, given the change in elevation Δz between two reference points, the relief displacement Δd along the ground in the range direction is,

$$\Delta d = \Delta z \tan \theta. \quad (5.2)$$

For example, a reference point on Wild Deer Lake, elevation 395 metres, has a range displacement of 562.5 metres between the two SAR images. Comparing this point against the registration point on Lois Lake, elevation 680 metres, the possible error due to relief displacement is,

$$\Delta d = \Delta z \tan \theta = (680 - 395) \tan 69.5^\circ = 762.3 \text{ metres.}$$

This shows that the correction for relief displacement in the rectified SAR image accounts for the error in the reference point in the range direction.

Error in the azimuth direction was found to be minimal, as expected, since relief displacement has no effect in this direction.

5.3 DEM Coordinate Transformation

As mentioned in Section 3.4.1, the shape-from-shading algorithm must use the slant range coordinate system (r, y, u) to process the SAR imagery. In addition to using the DEM to evaluate the shape-from-shading results, it is also needed to provide initial constraints, in either the spatial domain or frequency domain (see Section 4.1 and Section 4.2), in

order to run the algorithm. Therefore, the DEM must be transformed from the ground range coordinate system (x, y, z) to the slant range coordinate system (r, y, u) . This transformation is modelled as a rotation about the y -axis, maintaining orthogonality between the axes. Mathematically, it is expressed as,

$$\begin{pmatrix} r \\ y \\ u \end{pmatrix} = \begin{pmatrix} r_s \\ y_s \\ u_s \end{pmatrix} \begin{pmatrix} \cos \theta & 0 & -\sin \theta \\ 0 & 1 & 0 \\ \sin \theta & 0 & \cos \theta \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} r_o \\ y_o \\ u_o \end{pmatrix}, \quad (5.3)$$

where θ is the sensor's depression angle (69.5 degrees for Seasat); r_s , y_s , and u_s are scaling factors; and r_o , y_o , and u_o are bias factors.

The equation requires six coefficients to be derived. Since the rotation is about the y -axis, there is no change to the azimuth component of the DEM so the y -axis coefficients can be set as $y_o = 0$ and $y_s = 1$. The base "elevation" of the u -axis can be set arbitrarily to $u_o = 0$ since the shape-from-shading algorithm solves for surface gradients, not surface depth. Both the x and z axes use equal scaling (metres) and this proportionality must be maintained between the r and u axes, so $u_s = r_s$. What remains to be solved for now is r_o and r_s , which are used in the equation,

$$r = r_s(x \cos \theta - z \sin \theta) + r_o \quad (5.4)$$

Values for x , z and r are known at the points used for computing and verifying the transformation from the rectified SAR image to the slant range SAR image. By examining a few of these sets of points, an educated guess can be made for r_o . From these sets of points, the values for x and z can be substituted into the expression $(x \cos 69.5 - z \sin 69.5)$

to form the matrix A and the values for r can be combined with r_o to form the vector $\vec{b} = \vec{r} - R_o$. This gives an overdetermined system of linear equations,

$$\vec{b} = Ar_s, \quad (5.5)$$

which can be rearranged to find a least-squares solution for r_s ,

$$r_s = \frac{A^T \vec{b}}{A^T A}. \quad (5.6)$$

To verify the accuracy of the rotation, the result of applying Equation 5.4 to the data points is compared with the errors attributed to relief displacement. For the slant range, the relief displacement is proportional to the sine of the sensor depression angle θ . Residuals are computed for each data point $\Delta \vec{r} = \vec{b} - Ar_s$ and used to compare data points i and j against one another to see if,

$$\Delta d = |\Delta \vec{r}_i - \Delta \vec{r}_j| \leq |z_i - z_j| \sin \theta. \quad (5.7)$$

For example, the reference point at Lois Lake has an elevation of 680 metres and a residual of 69 metres and the reference point at Wild Deer Lake has an elevation of 395 metres and a residual of -81 metres. Comparing the two points shows that,

$$\Delta d = (69 + 81) = 150 \text{ metres} < (680 - 395) \sin 69.5 = 267 \text{ metres}.$$

Using this method, coefficients r_o and r_s were found that gave residuals \vec{r} that were within the tolerance allowed by error due to relief displacement.

The same set of coefficients can be used to compute the rotation back from (r, y, u) coordinates to (x, y, z) coordinates using the equation,


$$\begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix} \left[\begin{pmatrix} 1/r_s \\ 1/y_s \\ 1/u_s \end{pmatrix} \begin{pmatrix} r - r_o \\ y - y_o \\ u - u_o \end{pmatrix} \right]. \quad (5.8)$$

The units used for processing slant range data, as required in Equation 4.1 and Equation 4.7, must be carefully considered in the slant range. The transformation from (x, y, z) coordinates to (r, y, u) coordinate introduces scale changes to the r and u components but not to the y component. To maintain proportionality between all three axes, the y units should be set to $r_s \cos \theta$.

The software implementation of the rotation involves two steps. The first is to compute values for r and u on the (x, y) grid so that two structures represented by $r = g(x, y, z)$ and $u = h(x, y, z)$ are obtained. The next step is to determine the appropriate value for u on a regular (r, y) grid by using $g(x, y, z)$ to interpolate from $h(x, y, z)$. Interpolation is necessary along the r -axis but not the y -axis. Usually, an integer value of r will fall between two (x, y) grid points, so that $g(x_1, y, z) < r < g(x_2, y, z)$ where $x_1 < x_2$. In which case the value of $u(r, y)$ can be linearly interpolated using,

$$u(r, y) = \left[\frac{r - g(x_1, y, z)}{g(x_2, y, z) - g(x_1, y, z)} \right] \cdot [h(x_2, y, z) - h(x_1, y, z)]. \quad (5.9)$$

Figure 5.4 shows the relation between the ground range x and z axes and the slant range r and u axes. From the figure it can be seen that a horizontal surface in the (x, y, z) coordinate system (shown as a dotted line) has a constant slope in the (r, y, u)



where

d is the depression angle

l is the look angle

Figure 5.4: DEM Rotation Geometry

coordinate system equal to the sensor depression angle d . It is also apparent that any slope facing the sensor in the range direction which is greater than the look angle l in the (x, y, z) coordinate space cannot be represented in the (r, y, u) coordinate space. For Seasat SAR, this means that any slope in the DEM greater than 20.5 degrees facing the sensor will be lost. Furthermore, slopes exceeding this limit will cause decreasing values of r to be generated for increasing values of x in the $g(x, y, z)$ structure. The interpolation algorithm handles this situation by skipping over regions of decreasing r in $g(x, y, z)$ until a value larger than the next integer value of r is obtained. The upper and lower boundaries of the (r, y, u) grid may not have corresponding $g(x, y, z)$ values for r depending upon the surface elevation z at the x boundary. The u values for r must then be extrapolated from $g(x, y, z)$ and $h(x, y, z)$ in order to provide a reasonable surface height estimate.

Chapter 6

Results

In this chapter, the experiments to evaluate the viability of applying shape-from-shading to SAR data are described and their results given. The experimental work is presented in three stages. First a brief look is taken at tests performed on a simulated scene; second, reflectance functions are derived for the test site from the DEM and SAR image; and finally, the results of running the shape-from-shading algorithm on data from the test site is presented.

6.1 Simulated Data

The correctness of the software developed was verified on a simulated scene. A surface model was constructed in the form of a partial sphere protruding from a horizontal plane. The surface is smooth with the exception of a discontinuity in the smoothness at the joint between the two objects. This is the same type of surface as Frankot and Chellappa used for some of their tests in [FRAN88]. An image was generated from the surface using

a Lambertian reflectance model with point source illumination at 45 degrees above the horizon, directly over the x -axis. Figure 6.1 shows the surface and the image of it that was produced.

A second form of this surface was derived by rotating it into a slant range coordinate system 45 degrees from the original. An image was then generated, maintaining the same Lambertian reflectance function and the same position of the illumination source, so that the illumination vector was now parallel to the r -axis. Figure 6.2 shows the rotated surface and its associated image. The image appears different in the slant range than in the ground range because the viewing direction has rotated 45 degrees away from the illumination source, from the z -axis to the u -axis, so that illuminated side of the sphere covers less area in the image.

Similar experiments were run on both the ground range and slant range representations of the surface model. Experiments were run for only 25 iterations, as this was found to be a sufficient number to verify the correctness of the algorithm on this simple surface. The experiments tested the shape-from-shading algorithm under various values for λ and varying amounts of initial constraints in both the spatial and frequency domains. Overall, it was found that the fit to data $(I - R)^2$ was eight to ten times worse in the slant range than in the ground range. Also, the mean orientation error was about three degrees worse in the slant range than the ground range. The ratio between the smoothness measure in the ground range and slant range had a much greater variation in




Figure 6.1: Simulated Surface and Associated Image




Figure 6.2: Simulated Surface Rotated 45 Degrees and Associated Image

values; however, all were considerably less than the ratio of the smoothness of the slant range surface model to the ground range surface model: 22 to 1.

From these experiments the question arises as to why there is a difference between the accuracy of the results for a surface processed in ground coordinates as opposed to slant range coordinates. There are some differences between the two sets of experiments which must be considered. First, in this case in changing the coordinate system to slant range the viewing point was changed as well so that the darker region of the sphere occupies more of the image. This means that when the shape-from-shading algorithm is run, different regions of the reflectance map will be used to compute image irradiance values.

A second reason as to why the results are different can be gained by examining profiles of the surfaces generated in the two different coordinate systems. A comparison of surface height between the surface model and the test results, made by taking a slice through the surfaces, is shown in Figures 6.3 and 6.4 for the ground and slant coordinate systems respectively. Here it is evident that in the slant range, the shape-from-shading algorithm considerably smooths over the transition from plane to sphere in the lower range of the r -axis.

Third, the vast difference in the measures of surface smoothness between the surface model in the slant range and ground range coordinate systems is due to the non-linearity of gradient space. When the points on a unit sphere are mapped to gradient space,

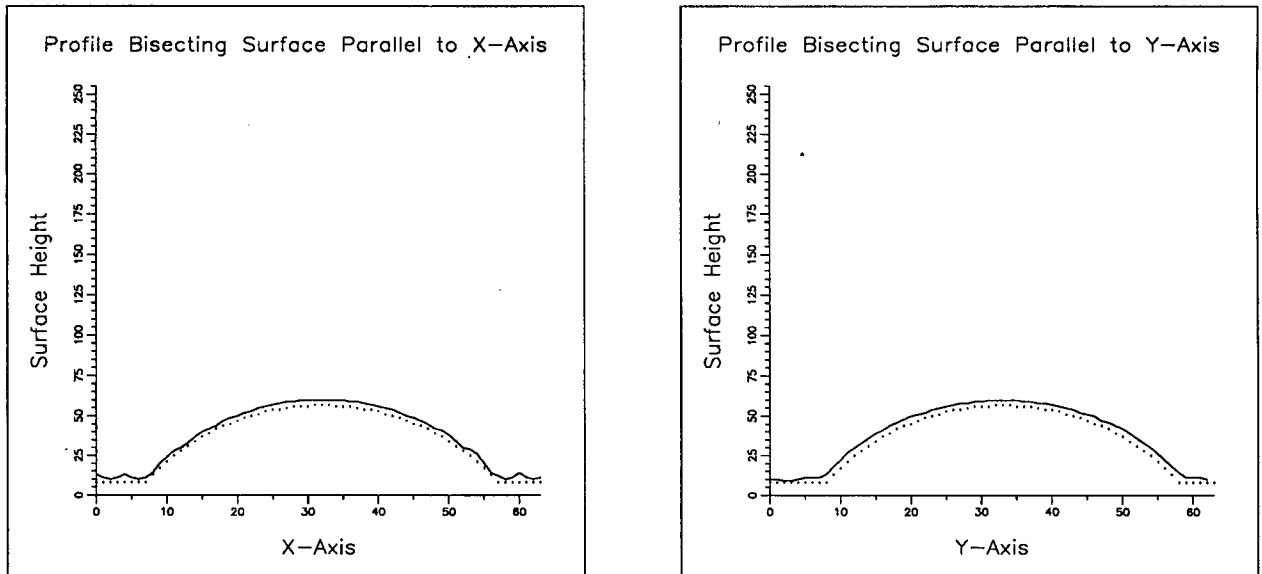


Figure 6.3: Profile of Surface Model (dotted line) and SFS Result (solid line) in Ground Coordinates

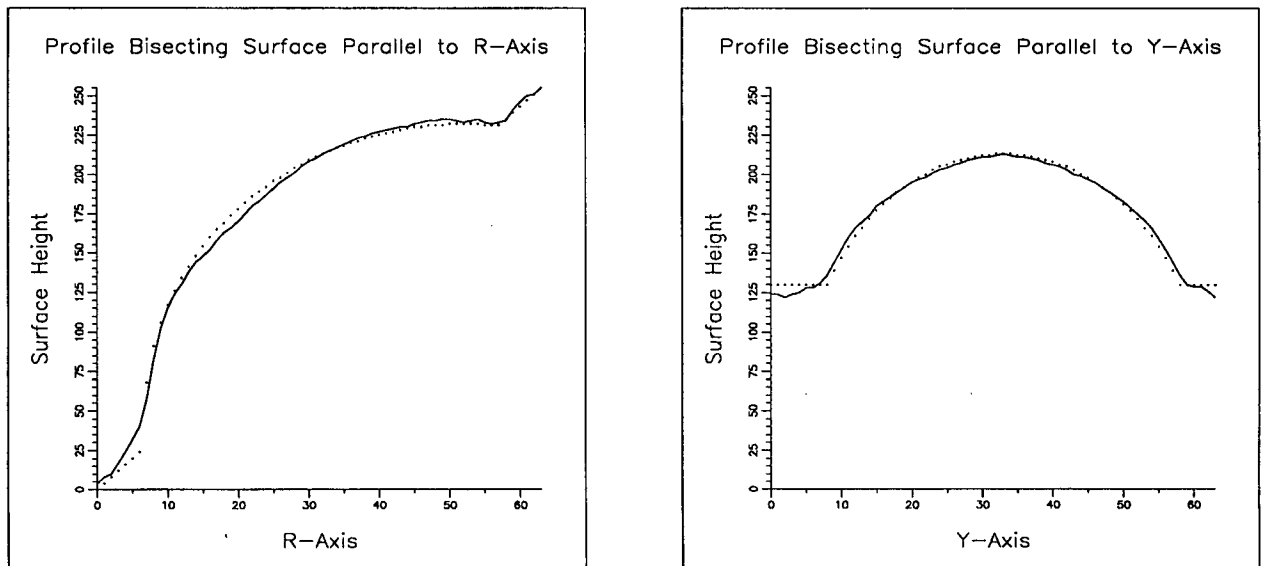


Figure 6.4: Profile of Surface Model (dotted line) and SFS Result (solid line) in Slant Coordinates

Coordinates	Gradient	Max	Min	Mean	Std. Dev.
Ground	p	1.124	-1.124	0.000	0.349
Ground	q	1.124	-1.124	0.000	0.349
Slant 45 deg	p	0.077	-9.219	-1.000	0.939
Slant 45 deg	q	5.695	-5.695	0.000	0.808

Table 6.1: Gradients of Simulated Surfaces in Ground and Slant Coordinates

distances between points are not correctly preserved. This means that when the surface model is converted from ground range to slant range, the relative size of gradient values for the surface model is not preserved. Table 6.1 shows that the gradient values of the surface model in the ground coordinates are centred around the origin of gradient space, whereas in the slant range coordinates the values are no longer centred along the p -axis and the range in gradient values is much greater. Therefore, surface smoothness, as calculated in the energy minimization function, is a coordinate system dependent measure.

6.2 Reflectance Functions for the Test Site

The DEM and SAR image of the test site were reduced to 162.5 metre square pixels to fit a 64 by 64 array for testing. The DEM was rotated 69.5 degrees, the depression angle of the Seasat SAR sensor, to place it into the slant range coordinate system. Figure 6.5 shows a plot of the DEM from the southeast in the ground, (x, y, z) , coordinate system.

The incident angle at each point in the array was computed and a histogram of these values was generated to examine the distribution, see Figure 6.6. The peak of the histogram occurs at an angle of 22 degrees, which corresponds to flat ground in the slant range coordinate system. However, the distribution is clearly asymmetric, the majority of the angles being greater than 22 degrees. The distribution mean is actually 32 degrees.

Three different reflectance functions were created. The first one was derived by computing the mean pixel value in the SAR image for each degree of incident angle obtained by applying the viewing geometry to the DEM. The results are shown in Figure 6.7 and will be referred to as the SAR Mean reflectance function.

A second reflectance curve was obtained from the SAR Mean reflectance function by performing a non-weighted linear regression on the data. The data was split into two sets at the angle corresponding to the peak of the distribution of angles, 22 degrees, to produce two lines of differing slope. The resulting curve is shown in Figure 6.8 overlaying the SAR Mean reflectance function.




Figure 6.5: DEM of Test Site in Ground Coordinates




Figure 6.6: Histogram of Incident Angles Across Test Site




Figure 6.7: SAR Mean Reflectance Function: Mean SAR Pixel Value Per Degree of Incident Angle





Figure 6.8: Linear Regression Reflectance Function and Keydel Reflectance Function

For the third reflectance curve, the model proposed by Keydel [KEYD82] and used by Frankot and Chellappa [FRAN87] was chosen. A bias term β was added to the model to provide a better fit of to the empirical data. The final formula used is,

$$\beta + \gamma \frac{\cos^2 \alpha}{\sin \alpha}. \quad (6.1)$$

Since the majority of the incident angles fall in the range of 15 to 50 degrees (as evident from Figure 6.6), values for β and γ were derived that would provide a good fit along this stretch of the curve. To establish values for β and γ , two pairs of $(\alpha, R(\alpha))$ were required. Values for $R(\alpha)$ were estimated by taking the mean of several points along the curve of the empirical reflectance function. The resulting reflectance curve is shown in Figure 6.8, overlaying the SAR Mean reflectance curve.

Images of the test site were generated using these three reflectance functions. The synthesized images were found to have a standard deviation in pixel value of only 4 to 5 whereas the SAR image had a standard deviation of 15. The correlation coefficient r was computed for each synthesized image with respect to the SAR image. The proportion of variance that each synthesized image and SAR image have in common can be estimated from the value of r^2 . For all the synthesized images, this was found to be about 91 percent. The histograms of the images are shown in Figure 6.9.







Figure 6.9: Histograms of SAR Image and Synthesized Images

6.3 Experiments on the Test Site

The initial conditions for the experiments run on the test site were provided in the frequency domain using coefficients copied from the Fourier transform of the surface model. The highest frequency of coefficients used were equivalent to $\frac{1}{2}$ cycle over 64 pixels, or 1 cycle per 20.8 km, in both the r and y axes. A plot of the surface represented by the initial conditions is shown in Figure 6.10.

Comparing the surface represented by the initial conditions with the surface model revealed that the mean surface orientation error of the initial conditions was 19.6 degrees and that the correlation between the surface heights was 0.9984. The smoothness of the two surfaces was calculated using the smoothness equation of the energy minimization function. The initial conditions produced a measure of 0.0045λ indicating a very smooth surface and the surface model produced a measure of 3.472λ , indicating a much rougher surface.

The first set of experiments performed used images generated from the surface model for each of the three reflectance functions. Each experiment was run for 100 iterations to ensure that changes could propagate throughout the 64 by 64 arrays of surface slope estimates. A scheme for reducing the effects of the smoothness constraint was setup whereby λ was decremented by a constant amount after each iteration of the shape-from-shading algorithm. Several experiments were run to find a fairly optimal choice




Figure 6.10: Surface Represented by Initial Conditions

for λ and the decrement value so that the smoothness constraint did not inhibit the convergence of the algorithm and yet was not too small to cause instability. Table 6.2 shows the results of the experiments from the final choices for λ .

The first column of the table indicates what reflectance function was used to generate the image for the experiment and to compute surface radiance for the shape-from-shading algorithm. The next two columns give the values for the two parts of the energy minimization function. The total value of the energy minimization function is not shown since it is dependent upon the value for λ . The values for the fit to data measure, $(I - R)^2$, are shown on a per pixel basis as opposed to across the whole image. For these experiments, a lower fit to data value is most likely (but not necessarily) an indication of a better result. The orientation error, measured in degrees, is the difference in surface orientation between the shape-from-shading results and the surface orientation derived from the surface model, averaged across the whole test site. The height correlation is simply the correlation value computed from the surface model and the reconstructed heights of the shape-from-shading results.

The table shows that both the Keydel and Linear images behave similarly. Their mean orientation error is lower than that of the initial conditions, their surface height correlation is higher than that of the initial conditions, and they are not as smooth as the initial conditions. In other words, a better estimate of the surface slopes has been computed. The plot in Figure 6.11 of the of the surface computed from the Keydel

image, confirms this. The table also shows that when the shape-from-shading algorithm was applied to the SAR Mean reflectance image only marginally better results were generated than what was given in the initial conditions. Figure 6.12 shows the surface.

These same experiments were then run on the SAR image, maintaining the same set of initial conditions, initial value for λ , and decrement as used for the corresponding experiments with the synthesized images. It was found that there was not enough of the smoothness constraint, the surfaces were overly rough, the mean orientation was much larger than that for the initial conditions, and the correlation value was less.

Further experiments were run with larger values of λ . Table 6.3 shows the best set of results obtained after 100 iterations from tests over varying ranges of λ . Although these results may not be the best possible results attainable for 100 iterations of the algorithm, several trials over the space of possible values for λ seem to indicate that they are close. As can be seen, the tests with the Keydel and Linear reflectance functions saw only marginal improvement in the algorithm results over the initial conditions, as measured in terms of the mean orientation error and correlation value. None of the tests with the SAR Mean reflectance function could improve upon the initial conditions. The surface reconstructed from the results using the Keydel reflectance function are shown in Figure 6.13.

The large values for the fit to data measure reflect the fact that none of the images derived from the surface model using the reflectance functions closely matched the SAR

Image/ Ref Fnc	Min. Function		Orientation Error	Height Correlation
	Fit to Data	Smoothness		
Keydel	8.47	0.26λ	15.4	0.9989
Linear	7.74	0.17λ	16.5	0.9988
SAR Mean	21.45	0.21λ	19.5	0.9986

Table 6.2: Results from Synthesized Images After 100 Iterations




Figure 6.11: SFS Results on Keydel image, 100 iterations




Figure 6.12: SFS Results on SAR Mean reflectance image, 100 iterations

Reflectance Function	Min. Function		Orientation Error	Height Correlation
	Fit to Data	Smoothness		
Keydel	192.96	0.02λ	18.8	0.9985
Linear	213.77	0.01λ	18.6	0.9985
SAR Mean	219.82	0.03λ	20.0	0.9983

Table 6.3: Results from SAR Image After 100 Iterations




Figure 6.13: SFS Results on SAR image using Keydel Reflectance Function

image. The fit to data measurement was applied to the synthesized images with respect to the SAR image. For the images derived using the Keydel, Linear and SAR Mean reflectance functions it was found that the fit to data values were 223.1, 224.0, and 231.0 respectively.

To gain a better understanding of the nature of the results, two types of images were generated from them, see Figure 6.14. The height difference image measures the difference between the estimated surface and the DEM. Dark areas represent where the DEM is much lower than the estimated surface, light areas represent where the DEM is much higher than the estimated surface, areas of medium brightness represent equality. The surface orientation error image indicates the difference between the estimated surface orientation values and the surface orientation values derived from the DEM. Dark areas indicate low error, bright areas indicate high error.

Figures 6.14, 6.15, and 6.16 show the result of computing these images with the initial conditions provided in the frequency domain, the shape-from-shading results from using the synthesized Keydel image, and the shape-from-shading results from using the SAR image with the Keydel reflectance function respectively. With the SAR shape-from-shading results there is very little noticeable change from the initial conditions. However, with the results from the synthesized Keydel image there is a noticeable blurring in the height difference image and a darkening in the surface orientation error image, both of which indicate an improvement in the results. These images confirm what the statistical





Figure 6.14: Initial Conditions: height difference and surface orientation error with respect to DEM



Figure 6.15: Synthesized Keydel results: height difference and surface orientation error with respect to DEM

Reflectance Function	Min. Function		Orientation Error	Height Correlation
	Fit to Data	Smoothness		
Keydel	254.31	0.01λ	21.1	0.99806
Linear	270.01	0.01λ	20.9	0.99805
SAR Mean	293.51	0.14λ	23.5	0.99879

Table 6.4: Results from Fine Resolution SAR Image After 100 Iterations

measures have already indicated about the shape-from-shading results.

The set of experiments were repeated with a finer resolution grid of the test site, 128 by 128 pixels, to determine if the algorithm would produce the same quality of results with less information in the initial constraints. The highest frequency of coefficients in the constraints were equivalent to $\frac{1}{2}$ cycle over 128 pixels, in both the r and y axes. The mean surface orientation error of the initial conditions from the surface model was 21.9 degrees. The correlation between the surface model and the surface reconstructed from the initial constraints was 0.99802. The smoothness of the surface produced by the initial conditions was 0.0006λ , whereas the surface model was 5.291λ . The experiments were run for 100 iterations with the same set of reflectance functions and the same choices for λ and its decrement on each iteration. The results listed in Table 6.4 show that the relative convergence from the initial conditions was no greater in the case of the 128 by 128 grid than for the 64 by 64 grid.




Figure 6.16: SAR results: height difference and surface orientation error with respect to DEM

Chapter 7

Observations

The reason for the poor performance of the shape-from-shading algorithm on the SAR data does not appear to have anything to do with the actual algorithm. Instead, the experimental results show that the algorithm can work well if the reflectance function accurately models image irradiance and the reflectance function is a simple curve; as was the case when the algorithm was run on images synthesized from the Keydel and Linear reflectance functions.

The primary reason for the lack of success in the experiments on the test site is apparently due to the poor modelling of the reflectance properties of the terrain. The much larger standard deviation in the SAR image than in the synthesized images seems to indicate that there are phenomena that have not been accounted for by the reflectance functions. These phenomena may be linked to variations in ground cover within the test site, noise in the image, sensor distortions, or other causes.

The question arises as to whether 100 iterations of processing is sufficient to reach

any conclusions about the SAR experiments. From the experiments performed using the images synthesized from the DEM, it is apparent that 100 iterations were quite sufficient to obtain a measurably significant improvement in the surface orientation estimates for the set of initial conditions that were given.

It was observed that the mean surface orientation error was always three to four times worse for gradients in the azimuth direction as opposed to the range direction for the SAR experiments. This was true for both the experiments performed on images synthesized from the DEM and for the SAR image. This is probably because the portion of the reflectance map used for these experiments has much greater variation in image irradiance values for a small change in the range surface orientation values than for an equal change in value in the azimuth direction. This means that the shape-from-shading technique will then generally perform better at distinguishing variation in surface orientation along a range scan as opposed to across the azimuth.

In addition to the need for an improved model of the reflectance characteristics of the terrain, a method to calibrate the model to the image is required. The calibration technique would be best to take advantage of the data used for the initial conditions of the shape-from-shading algorithm.

The setup of the shape-from-shading algorithm to work in the slant range coordinate space presents problems. If a surface slope facing the sensor exceeds an angle greater than the look angle from horizontal in the ground coordinate space, then it cannot be

represented in the slant range space. For the Seasat SAR image of the test site this angle was fairly small, for other sensors it may not be as great a problem.

It was also noted that a gradient space measure of surface smoothness is dependent upon the coordinate system. This makes it difficult to compare quantitatively the quality of results between processing in the slant range and ground range.

One other drawback with the shape-from-shading algorithm is the problem of choosing λ and a method to reduce λ as the algorithm converges and requires less of the stability provided by the smoothing constraint. It was observed that this parameter is sensitive to both the content of the image and the choice of reflectance function.

Chapter 8

Conclusions

The shape-from-shading algorithm of Frankot and Chellappa was found to produce substantially better results for images synthesized from a DEM; however, when applied to a SAR image, the algorithm could not produce a surface any more accurate with respect to the DEM than that supplied by the initial conditions of the experiment.

The results suggest that a better reflectance model is needed for the SAR image. A possible area for future investigation is to determine whether modelling variations in ground cover, image noise, or other phenomena over the space of the image can be used to improve the reflectance model and the shape-from-shading result.

It was also observed that the shape-from-shading results from both the synthesized images and the SAR image were considerably more accurate along range than azimuth. Therefore, another possible area of investigation is to see if the overall quality of the shape-from-shading results can be enhanced with a better technique for deriving surface orientation values in the azimuth direction.

An alternate approach to the reconstruction of surface shape would be to see whether applying a shape-from-shading technique to a very noisy estimate of a surface can provide a more accurate surface model than simply using a smoothing filter on the surface model.

Bibliography

- [BARR70] Barrick, D.E. (1970) "Rough surfaces," Chapter 9 in *Radar Cross Section Handbook*, G.T. Ruck (ed.), Plenum Press, New York, NY, pp. 671-772.
- [BROO85] Brooks, M.J. & Horn, B.K.P. (1985) "Shape and source from shading," *Proceedings of the National Conference on Artificial Intelligence*, Los Angeles, CA, August 18-23, pp. 932-936.
- [FRAN87] Frankot, R.T. & Chellappa, R. (1987) "Application of a shape from shading technique to SAR imagery," *Proceedings IEEE International Geoscience Remote Sensing Symposium*, Ann Arbor, MI, pp. 1323-1329.
- [FRAN88] Frankot, R.T. & Chellappa, R. (1988) "A method for enforcing integrability in shape from shading algorithms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 10, No. 4, pp. 439-451.
- [GUIN89] Guindon, B. (1989) "Development of a shape-from-shading technique for the extraction of topographic models from individual spaceborne SAR images," *Proceedings of the 1989 International Geoscience and Remote Sensing Symposium (IGARSS)*, Vancouver, Canada, July 10-14, pp. 597-602.
- [HORN75] Horn, B.K.P. (1975) "Obtaining shape from shading information," Chapter 4 in *The Psychology of Computer Vision*, P.H. Winston (ed.), McGraw-Hill, New York, NY, pp. 115-155.
- [HORN77] Horn, B.K.P. (1977) "Understanding image intensities," *Artificial Intelligence*, Vol. 8, No. 2, pp. 201-231.
- [HORN86a] Horn, B.K.P. (1986) *Robot Vision*, MIT Press, Cambridge, MA & McGraw-Hill, New York, NY.
- [HORN86b] Horn, B.K.P. & Brooks, M.J. (1986) "The variational approach to shape from shading," *Computer vision, Graphics and Image Processing*, Vol. 33, No. 2, pp. 174-208.

- [HORN89] Horn, B.K.P. (1989) "Height and gradient from shading," Memo 1105, Artificial Intelligence Laboratory, MIT, Cambridge, MA, May.
- [IKEU81] Ikeuchi, K. & Horn, B.K.P. (1981) "Numerical shape from shading and occluding boundaries" *Artificial Intelligence*, Vol. 17, No. 1-3, pp. 141-184.
- [KEYD82] Keydel, W. (1982) "Application and experimental verification of an empirical backscattering cross section model for the earth's surface," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 20, No. 1, pp. 67-71.
- [LILL87] Lillesand, T.M. and Kiefer, R.W. (1987) *Remote Sensing and Image Interpretation*, 2nd edition, Wiley, New York, NY.
- [RIND66] Rindfleisch, T. (1966) "Photometric method for lunar topography," *Photogrammetric Engineering*, Vol. 32, No. 2, pp. 262-276.
- [SWOR86] Swords, S.S. (1986) *Technical History of the Beginnings of Radar*, Peter Pergrinus Ltd., London, U.K.
- [THOM89] Thomas, J., Kober, W., and Leberl, F. "Multiple-image SAR shape from shading," *Proceedings of the 1989 International Geoscience and Remote Sensing Symposium (IGARSS)*, Vancouver, Canada, July 10-14, pp. 592-596.
- [WILD86] Wildey, R.L. (1986) "Radarclinometry for the Venus Radar Mapper," *Photogrammetric Engineering and Remote Sensing*, Vol. 52, No. 1, pp. 41-50.
- [WONG84] Wong, F.H. (1984) "A unified approach to the geometric rectification of remotely sensed imagery," *Phd. Thesis*, University of British Columbia, May 1984.
- [WOOD84] Woodham, R.J. (1984) "Photometric method for determining shape from shading," Chapter 4, *Image Understanding*, Ullman and Richards, Ablex Publishing Corpl, Norwood, NJ.
- [WOOD85] Woodham, R.J., Catanzariti, E. and Mackworth, A. K. (1985) "Analysis by synthesis in computational vision with application to remote sensing," *Computational Intelligence*, Vol. 1, No. 2, pp. 71-79.