Seafloor Map Generation for Autonomous Underwater Vehicle Navigation

Andrew E. Johnson and Martial Hebert

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

This research was supported by the National Science Foundation under contract DCR-8604199.

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. government.

Abstract

Elevation map generation is an essential component of any autonomous underwater vehicle designed to navigate close to the seafloor because elevation maps are used for obstacle avoidance, path planning and self localization. We present an algorithm for the reconstruction of elevation maps of the seafloor from side-scan sonar backscatter images and sparse bathymetric points co-registered within the image. Given the trajectory for the underwater vehicle, the reconstruction is corrected for the attitude of the side-scan sonar during the image generation process. To perform reconstruction, an arbitrary but computable scattering model is assumed for the seafloor backscatter. The algorithm uses the sparse bathymetric data to generate an initial estimate for the elevation map which is then iteratively refined to fit the backscatter image by minimizing a global error functional. Concurrently, the parameters of the scattering model are determined on a coarse grid in the image by fitting the assumed scattering model to the backscatter data. The reconstruction is corrected for the movement of the sensor by initially doing local reconstructions in sensor coordinates and then transforming the local reconstructions to a global coordinate system using vehicle attitude and performing the reconstruction again. We demonstrate the effectiveness of our algorithm on synthetic and real data sets. Our algorithm is shown to decrease the average elevation error when compared to real bathymetry from 4.6 meters for the initial surface estimate to 1.6 meters for the final surface estimate from a survey taken of the Juan de Fuca Ridge.

Keywords

autonomous underwater vehicles
underwater navigation
underwater mapping
elevation maps
acoustic backscatter
sidescan sonar
shape from shading
seafloor scattering models

Contact Author

Andrew E. Johnson
Smith Hall, Robotics
Carnegie Mellon University
Pittsburgh, PA 15213
Tel: (412) 268-6880
Fax: (412) 268-5571
Email: aej@ri.cmu.edu

List of Figures

regions of high and low albedo giving a segmentation of the seafloor based on material type (bottom).

FIGURE 13. Backscatter and elevation data with estimates of volume scattering model parameters. Elevation map (top left), backscatter data (top right), speed ratio parameter (middle left), density ratio parameter (middle right) and volume scattering albedo (bottom).

TABLE 1. Scattering models and descriptions of their parameters.

1 Introduction

Navigation of any autonomous vehicle requires a rich and accurate understanding of the vehicle's environment acquired using onboard sensors and sensing algorithms (Elfes 1987). A fundamental and useful representation of a vehicle's surroundings is an elevation map from which features can be extracted for vehicle localization, local and global path planning and obstacle avoidance, activities required for autonomous locomotion (Hebert 1989). The accuracy of vehicle localization and path planning from elevation maps is limited by the error on the estimate of the elevation, so algorithms and sensors for generating accurate elevation maps are needed.

The undersea world is an unfriendly environment for navigation sensors mainly because the sensors must sense through the water medium. Water severely attenuates and distorts light, so optical sensors produce poor results at far range. Sonar, on the other hand, is more suitable for sensing underwater. Water does not attenuate sound energy nearly as much as light energy, so sonar has a greater range. Sonar is also less susceptible to refraction. These benefits come at the price of lower resolution and increased noisiness of the sonar data. However, given the task of mapping and navigation underwater, sonar is the best alternative. To execute the precise trajectories needed for navigation close to the seafloor and manipulation of underwater objects the error on the seafloor elevation estimate must be much smaller than the size of the vehicle. For sonar sensors to be useful for navigation, algorithms must be developed for creating accurate seafloor elevation maps from sonar data.

Oeanographers and geologists have a strong desire to map the seafloor because the terrain and material composition of the seafloor hold clues to understanding the geological processes that transform the earth. The scattering characteristics of the seafloor are encoded as parameters of a scattering model which are related to actual physical properties of the seafloor. Therefore, measuring and mapping scattering model parameters will help geologists understand the physical makeup of the ocean floor. These physical properties also give clues to underlying materials in the seafloor which can be useful for mining and rare element acquisition.

In this paper we present and algorithm for generation of accurate global elevation maps and coarse maps of seafloor scattering model parameters from a side-scan sonar image, a common underwater sensor.

1.1 Side-Scan Sonar

A simple sonar sensor consists of one cylindrical source that creates a conic acoustic beam pattern that is symmetric around the axis of the source. This type of sonar will measure the range to the first surface it encounters within the cone and the intensity or echo of the return, however, the position of the surface cannot be localized within the cone. In general, the angle of the cone is too large to make range imaging with these type of simple sensors feasible. Side-scan sonars are arrays of sonars arranged in a straight line. When the sensor generates an acoustic pulse (pings), the pulses from the individual sonars are phased in such a way that the acoustic beam patterns from each individual sonar interfere to shrink the overall beam pattern of the array in the direction perpendicular to the line of the sensors. When the array is pinged, an acoustic pulse is generated that travels through the water. When the pulse reaches the ocean floor, part of the pulse is scattered back to the sensor. The sensor measures these echoes from the seafloor for a short time interval after the ping. Hence, given that the speed of sound in water is known, recording the measurements of the sensor will generate acoustic return as a function of range from the sensor.

Side-scan sonar are the sensors that are generally used for mapping of the seafloor. The sonar array is mounted on a torpedo shaped platform (sonar fish) that is towed through the water by a surface ship. An acoustic image is generated by pulling the platform through the water; the acoustic return vs. range function for each ping constitute the rows of the image which are stacked to form the image. If the fish is towed in a horizontal line, then the image is acoustic return as a function of range from the sensor and position of the sensor along the line.

1.2 Acoustic Backscatter

Acoustic Backscatter is the ratio of sound energy impinging on a surface to that leaving the surface. Generally, it is a function of incidence angle of the acoustic pulse and material properties of the surface. Because backscatter contains geometric information about the surface being imaged, it is the quantity desired when trying to reconstruct the seafloor elevation map from sonar images. Unfortunately, the acoustic backscatter from the surface is not directly measured by the sensor. Many other processes relating to the transmission of sound through the water and the equipment used to measure it affect the data collection. The sonar equation presented in Equation 1 details these processes (Urick 1983).

$$10\log S = RL - SL - \alpha R - TVG + 2TL - 10\log A - RBP - SBP \quad (1)$$

The backscatter *S* from the seafloor is a function of receiver level *RL*, source level *SL*, water column attenuation which is a function of range αR, the time varying gain of the sensor *TVG*, the transmission loss of the acoustic pulse due to spherical spreading TL, the equivalent ensonified area A, receiver beam pattern *RBP* and the source beam pattern *SBP*. If all quantities except for receiver level can be independently measured, then backscatter can be obtained from the receiver level measured by the sensor. An image of acoustic backscatter from a relatively flat seabed is given in Figure 1. Sidescan sonar calibration, which requires complex testing facilities and detailed modeling of sonar sensors, has been studied extensively (von Alt 1989, Mitchell & Somers 1989, Stewart 1989, Malik 1991, Stewart et al. 1994).

FIGURE 1. A typical unfiltered backscatter image of the seafloor.

1.3 Problem Statement

Side-scan sonar data in the form of a backscatter image does not directly convey the elevation of the seafloor. However, given knowledge of the physics of sonar scattering from the water-seafloor interface (i.e., a known scattering model) it is possible, if appropriate assumptions are made, to convert this backscatter image into an elevation map of the ocean bottom. The scattering model with respect to sound of the seafloor along with its correct parameters are inverted to obtain the incidence angle of the sonar pulse. These incidence angles are then converted to elevation derivatives which are integrated together to form the best estimate of the surface of the seafloor given the backscatter data. The problem our algorithm seeks to address is the following:

> Given an arbitrary but computable backscatter model, generate accurate and detailed global elevation and scattering model parameter maps of the seafloor using a sidescan sonar backscatter image, sparse (co-registered) bathymetric data and the trajectory of the sensor during the image generation process.

There has been a great deal of research conducted to determine a general and physically meaningful scattering model for the seafloor (Stanton 1984, Gensane 1989, Caruthers & Novarini 1993). To date the most detailed model for high frequency acoustic backscatter combines models for three different scattering processes from the seafloor (volume sediment, Kirchoff approximation and composite roughness) into one scattering model (Jackson et al. 1985, Mourad & Jackson 1989). The most detailed models do not assume that measurements of local surface shape of the

seafloor are available, so the models are presented as backscatter as a function of grazing angle, not incidence angle. However, some recent work combines known bathymetry to generate local slope estimates and backscatter to estimate scattering model parameters of the seafloor (de Moustier & Alexandrou 1991, Matsumoto et al. 1993, Clark 1994). These approaches use dense bathymetric images, and, as a result, are restricted to data gathered from bathymetric sidescan sonars. There also exist some automatic methods for estimation of seafloor scattering parameters using classification by neural networks (Stewart et al. 1992) and maximum likelihood estimation (Michlopoulou et al. 1994).

Existing methods of seafloor reconstruction have used a specific scattering model and somewhat arbitrarily selected scattering model parameters for the seafloor to determine incidence angles (Langer and Hebert 1991). If the model and parameters are not guessed correctly, then the integration will produce a surface that is not correct. Furthermore, if the integration does not have fixed starting and ending points, accumulation of errors will force the surface away from its real value. Our previous work (Johnson 1993) did not use arbitrarily set scattering parameters, but it did not account for the deviation of the path of the fish from a straight line; hence it could not generate accurate global elevation reconstructions. There exist methods for automatic correction of geometric distortions caused by movement of the sensor that do not rely on onboard navigation measurements (Cobra et al. 1992, Cervenka & de Moustier 1993), but, unlike our method, they do not reconstruct seafloor elevation.

2 Detailed Problem Statement

Ideally we would like to generate elevation maps of the seafloor using only sonar backscatter images. However, this is not possible because determining elevation from backscatter is a highly underconstrained problem. Echo intensity is related to the angle of incidence α of the sonar pulse with the surface through the acoustic scattering model σ at the interface between the water and the seafloor. This scattering model is a function of many scattering model parameters $\mathbf{p}$, the angle of incidence and the frequency of the sonar f. Furthermore, the angle of incidence of the surface is a function of the elevation z, its derivatives z_x and z_y and the attitude $\mathbf{P}(X,Y,Z,\theta,\varphi,\omega)$ of the sensor at each point. If we parameterize the surface in terms of its horizontal coordinates (x,y) we have the following relationship for all intensity values S:

$$S = \sigma(\alpha(z_x, z_y, z, \boldsymbol{P}), \boldsymbol{p}, f) \quad (2)$$

Thus, the scattering model has eleven or more variables, depending on the number of scattering model parameters. Determining the elevation at each point can be achieved by solving Equation 2, for z, but this cannot be done without knowledge of all the other variables. To make matters worse, we do not know the functional form of the scattering model with respect to its parameters. Assumptions must be made and additional data must be added to invert this equation.

To solve this problem, we have added two additional sets of data. First, we assume that we know the attitude for all backscatter values collected which can be obtained from position sensors on the sonar fish (pressure depth, altimeter, transponder nets, compass, etc.). The other assumption that we make is that we have an initial estimate of the shape of the seafloor. Usually this initial estimate will come from interpolation between points of known elevation in the backscatter image. One method for determining these points is the registration of the acoustic image with a coarser global elevation map (Rigaud and Marcé 1990) determined by some other method. Another method relies on the placement of acoustic transponders on the seafloor. As the fish follows a trajectory, the 3-D position of the transponders can be determined by fusing the fish pose estimates from on board sensors and the range estimates from the transponders using a Kalman filter (Leonard & Durrant-Whyte 1992). The initial estimate of the seafloor shape can also be extrapolated from backscatter images using simple geometric reasoning and assumptions (Cuschieri & Hebert 1990). To avoid local minimum in the reconstruction, the initial estimate of surface shape should be as close to the actual shape of the seafloor as possible. However, only three points of known elevation are needed to establish an initial shape estimate and proceed with the surface reconstruction.

We will assume that, for each data set being used, we know the functional form of the scattering model appropriate for the seafloor being imaged and the sonar sensor being used. The algorithm does not depend on a particular form for the scattering model; the only requirement is that the model and its derivatives with respect to its parameters have a computable form for all angles of incidence. The scattering model will have an explicit dependence on sonar frequency. However, since sonar frequency is usually known or can be measured independently, we will treat it as a constant in the scattering model. If the frequency is unknown and cannot be measured, it can be

treated as another parameter to be estimated in the scattering model (albeit one that does not vary locally, but is fixed for the entire image). In the scattering models shown in Table 1 the dependence of the models is not shown because it is implicit in the parameters of the model.

Given these additions and assumptions, we now know the exact form of the scattering model, where the data was collected, and we have a vague idea of the shape of the seafloor. However, we do not know any of the scattering model parameters and we do not know the elevation and its derivatives at each point; these are the unknowns that our algorithm will determine.

The algorithm uses the sparse elevation points (hereafter called control points) to generate an initial estimate for the elevation map and its derivatives. Then these initial estimates are iteratively adjusted to create a surface that fits the intensity data. At each iteration the parameters of the scattering model are estimated based on the current shape of the seafloor. The end results are elevation and scattering parameter maps which can be used for mapping and navigation.

The path of the fish as it is towed through the water affects the acoustic returns from the seafloor because the backscatter from the seafloor is a function of fish pose. Furthermore, deviations of the path from a straight line can cause the sensor to scan the same region of the seafloor from two different pings, thus causing overlap in the backscatter image. However, if the path of the fish is known, the pose of the fish for each ping can be used to correct for the movement of the fish during the acoustic image generation. As a result, the reconstruction will have a local phase where the fish is assumed to move in a straight line and the cylindrical geometry of the image is maintained, and a global phase that corrects for the movement of the fish by transforming the intensity data and the local elevation estimates to cartesian global coordinates.

The problem then will be to determine the elevation map and scattering parameter maps of the seafloor given a sidescan sonar image, sparse elevation control points, the attitude of the fish for all of the data collected and the functional form of the scattering model used.

3 Algorithmic Overview

Before going over the algorithm it is useful to define some coordinate systems used in the reconstruction.

- **Fish coordinates:** The coordinate system attached to the sensor for the current ping. This coordinate system is cylindrical in that its coordinates are (R,y,z). This coordinate system is different for each ping because the position of the fish changes for each ping.
- **Global coordinates:** The cartesian coordinate system to which all of the local reconstructions are transformed. It is fixed to the position of the fish at the first ping in the image.

FIGURE 2. Geometry of a side-scan sonar (left) and image generation coordinate systems (right).

A visualization of these coordinate systems is given in Figure 2. Assuming that locally the fish moves in a straight horizontal line, the elevation along each ping can be determined in fish coordinates for that ping independently of the position of the fish. Given this assumption, the reconstruction can proceed in two stages. In the first stage the surface is reconstructed locally for each ping only using data from nearby pings. This will give a correct reconstruction of the surface in fish coordinates for that ping. The second stage converts all the reconstructions to global coordinates, merges overlapping data and completes by performing a global reconstruction in cartesian coordinates. The following steps detail the stages in the algorithm.

- **Generate Initial Estimates:** To start the algorithm, an initial estimate of the elevation surface must be created from the sparse elevation control points. This initial estimate will ensure that the starting point for the surface is close to its actual value. If the assumption is made that locally the fish moves perpendicularly to the swath and in a horizontal straight line, then locally the derivatives of the elevation surface can be calculated by taking finite differences between consecutive rows and columns. The end result of this first interpolation are three maps embedded in the backscatter image that correspond to the elevation at each pixel $z(p,s)$, the derivative of the surface with respect to the range $z_s(p,s)$ and the derivative of the surface with respect to distance travelled $z_p(p,s)$.
- **Create Local Maps:** Once the initial estimates for the local derivatives have been determined local reconstructions can commence. For each ping, subimages of intensity, elevation and its derivatives are made from the parts of the images and maps that surround the current ping. These are the local maps for the current ping.

- **Estimate Local Scattering Parameters with Non-linear Minimization:** The next step in the local reconstruction is to determine the scattering model parameters at each point in the image. To do this, we make the justifiable assumption that pixels that are close to each other in the image have similar scattering parameters. Thus, given this grouping plus the fact that we know the elevation and its derivatives and hence that angle of incidence at each point (locally we are ignoring fish pose), we can do a non- linear fit of the scattering model to the data (S,α) to determine the scattering parameters.

- **Reconstruct Elevation Map:** Once the estimates of the scattering parameters have been determined, we are in a position to refine the initial estimate of the elevation and its derivatives given the dense intensity data. The idea behind this stage is to create a global error functional that measures the divergence of the surface from the data. Minimizing this functional will force the surface into a shape that is consistent with the intensity data. Procedures of this sort are very common in computer vision and are commonly called shape from shading techniques (Horn 1986). Numerical stability and integrability of the surface require the addition of smoothness terms to the functional that enforce integrability of the surface.

- **Iterate on Local Maps:** One simple pass through these procedures is not sufficient to determine the elevation or the scattering parameters because they are inheritly coupled. However, iterating by repeatedly estimating the parameters and then refining the surface will cause the surface to converge. This procedure will generate local elevation and scattering parameter maps. Creating these local maps for each ping will generate a good estimate of the elevation of the seafloor with respect to the local fish coordinates for that ping.

  Figure 3 shows the block diagram for the local algorithm. First, sparse bathymetric control points and a given scattering model are used to estimate the scattering model parameters and elevation map every where in the image. Then, dense backscatter data is used to iteratively refine the estimates of the scattering parameters and the elevation map. Iteration continues until the elevation values and scattering parameters have converged every where in the image.

- **Transform Local Reconstructions to Global Coordinates:** The local reconstructions are transformed to global coordinates using the position of the fish. It is common for multiple data points to fall into the same grid cell in the global map so special attention must be paid to com-

bining these points so that a minimal amount of information is lost. The general idea is to average the elevation values while keeping all of the intensity data around so that it can be used in the reconstruction.

- **Repeat Local Algorithm on Global Maps:** The global algorithm is essentially the same as the local algorithm. First initial estimates of the seafloor elevation and its derivatives with respect to x and y are determined. Next, the scattering parameters are estimated in regions using nonlinear fitting of the scattering model to all of the intensity data. Finally, the elevation surface is refined using all of the scattering data and the initial estimates for the elevation. The latter two stages are repeated until the surface changes very little.

FIGURE 3. Block diagram for map local generation.

4 Sensor and Data Specifications

We used real and synthetic data to test and debug our reconstruction algorithm. The real data was collected by members of the Deep Submergence Laboratory at Woods Hole Oceanographic Institution in 1991 during there multisensor survey of the Juan de Fuca Ridge (Stewart et al. 1994). The data sets we received came from two different geological provinces (a ridge flank and an axial valley). Simple sets of synthetic data were generated to test our algorithm for strengths and weaknesses.

4.1 Sensor

The particular sensor from which the data was acquired was a 120-kHz phase-difference sidescan sonar developed jointly by the Deep Submergence Laboratory at Woods Hole Oceanographic Institution, the Applied Physics Laboratory at the University of Washington, and Acoustic Marine Systems, Inc. of Redmond Washington. There are two sidescan sonar arrays on each side of the sonar fish. This sensor is able to generate acoustic images by averaging the acoustic receiver levels from both of the arrays. It is also able to generate dense co-registered bathymetry of the seafloor by taking the phase difference of the acoustic signals between the two arrays (Denbigh 1989). The phase difference can be converted directly into seafloor elevation in fish coordinates given the phase difference and the range to the surface. The geometry associated with this sensor is given in Figure 2. The incidence angle α, which is a key factor in determining the acoustic back-

scatter from the surface, is the angle between the vector from the sensor to the surface **R** and the surface normal **n**. The grazing angle θ is the angle between **R** and the horizontal and is calculated from the phase between the two arrays. Calibrated measurements of source level, transmission beam pattern, receiver beam pattern, system gains and transmission losses through the medium have been made for this sensor, so the receiver level at the sensor can be converted directly into acoustic backscatter from the seafloor. Also available are full six dimensional pose measurements for the fish which come from positional sensors on the fish and an acoustic transponder net.

This sensor provides the necessary information to test our algorithm on real data. The six dimensional positional information makes the transformation from fish coordinates to world coordinates possible. The calibration of the sensor makes available backscatter images of the seafloor. The bathymetric map provides the elevation control points as well as a ground truth for verifying the reconstructed elevation map. For the purposes of testing the algorithm on a sparse set of elevation values, only discrete elevation control points were used from the bathymetric image. It should be noted that sensors that generate dense co-registered bathymetric maps of the seafloor are complex and uncommon compared to the quantity of simple side-scan sonars in use, so algorithms for determining seafloor elevation from backscatter images are still extremely useful. Even if dense bathymetry is available from the sensor, it will be noisy. Our algorithm provides a method for fusing bathymetry and backscatter to provide less noisy estimates of seafloor shape and maps of scattering model parameters.

4.2 Filtering

The images generated using the AMS/DSL 120 sensor are extremely noisy due to multiple echoes, the statistical nature of acoustic backscatter and noise in the sensors, so filtering of the data was required. First, a 3x3 median filter was applied to the backscatter and bathymetric images to eliminate isolated spurious measurements. Then, a modified median filter with variable window size (Langer & Hebert 1991) that only eliminates dropouts in the data was applied. The thought behind this filtering procedure is that dropouts are much more likely to be due to noise while positive spikes are likely to correspond to interesting features in the image. The final filter applied was a graduated non convexity filter which smooths data while preserving discontinuities. Detailed description of the purpose and implementation of the graduated non-convexity filter is

given in Blake and Zisserman (1987). The combination of these filters provided smooth data yet did not eliminate any of the interesting features in the image. Before and after images of acoustic backscatter and bathymetry are given in Figure 4.

The pose of the fish also required filtering because of noise in the sensors and low resolution of some of the measurements (especially global X and Y). In general, a median filter of variable window size and then a gaussian filter were applied to all of the measurements independently. Filtering of the pose considering the dependence of the parameters was attempted, but was later deemed unnecessary.

FIGURE 4. Acoustic backscatter (right) and bathymetry (left) for a typical sonar ping before and after filtering

4.3 Applicable Scattering Models

The algorithm was tested using many different scattering models which are listed in Table 1. The models range from simple, as in the case of the Lambertian model which has one parameter, to complex, as in the case of the composite roughness model which has four parameters. The complex models give a more accurate description of back scatter from the seafloor at the expense of being difficult to compute and use. Even more complex models can be made by adding models together or blending them so that they are applied only at the range of angles for which they are physically feasible (Mourad & Jackson 1989). The lambertian scattering model is the simplest and most widely used scattering model in acoustics and computer vision. The Torrance and Sparrow model and the diffuse model are optical reflection models commonly used in computer vision. In these three scattering models, each parameter has a physical meaning: the albedo ρ determines the ability of the seafloor to absorb sound, C and ε are measurements of the roughness of the seafloor surface and σ is a constant that determines the specularity of the seafloor. The Kirchoff approximation, composite roughness and sediment volume models are common scattering models from underwater acoustics. The Kirchoff approximation is a well known model that is applicable near normal incidence ($\alpha = 0^{\circ}$) of the acoustic pulse on the surface. The sediment volume scattering model gives the functional form of the acoustic backscatter that originates from the acoustic pulse penetrating the seafloor surface and then being scattered back by a perfectly flat homogeneous surface. The composite roughness model explains the backscatter caused by the small scale roughness of the surface and it is mainly applicable at moderate incidence angles. The

parameters of these models, which are explained in Table 1, relate to the different physical properties of the materials being imaged by the sonar sensor.

TABLE 1. Scattering models and descriptions of their parameters.

Figure 5 shows the scattering models with the best fits to the data collected from an underwater ridge flank. The best parameters of the models were determined through the use of a genetic algorithm for determining functional parameters of complex and possibly non-differentiable functions (Baeck & Schwefel 1993). The scattering model with the least fit error is the complex model presented by Mourad and Jackson (1989) which is a continuous combination of the composite roughness, volume, and Kirchoff approximations models.

FIGURE 5. Best fit scattering models as given in Table 1 to the ridge flank data set.

4.4 Synthetic Data

In addition to the real data, sets of synthetic data were generated to test the algorithm in a controlled environment. This required simulation of the fish traveling on a well known path while scanning an object whose shape, scattering model and scattering model parameters were well known. Using data sets like this allowed us to test the efficacy of our algorithm in estimating the shape of the surface, estimating the scattering parameters given the correct scattering model, and merging multiple views of the same object. For most of the examples using synthetic data in this paper, the fish moves in a semi-circular path while scanning a cylindrical surface with a lambertian scattering model. This path and shape of the object are simple enough to make analysis of the results of the algorithm straightforward, yet at the same time complex enough to fully test all components of the algorithm.

5 Detailed Local Algorithm

The local reconstruction algorithm reconstructs the elevation profile of the current ping in fish coordinates. This reconstruction uses local information to generate an accurate reconstruction and assumes that locally the fish moves in a straight line

5.1 Initial Estimate of Local Maps

To start the algorithm, an initial estimate of the elevation surface must be created from the sparse elevation control points. This initial estimate ensures that the starting point for the surface is close to its actual value. The generation of the initial estimate in the general case is straightforward. The Delaunay triangulation of the *(x,y)* coordinates of the elevation points is determined using well known algorithms from computational geometry. Then the elevation values between control points are determined by linearly interpolating along the plane fit to each triangle in the triangulation. This interpolation is quick and sufficient for our algorithm. To simplify matters somewhat, we have assumed that the control points lie on a grid; hence, computation of the Delaunay triangulation is unnecessary because each corner in the grid determines a triangle in the triangulation.

The cylindrical nature of the sensor creates a sonar image whose columns correspond to acoustic samples which can be converted to range and whose rows correspond to the position of the fish for each ping. Thus the image can be parameterized in terms of the ping and the sample *(p,s)*. If the assumption is made that locally the fish moves perpendicularly to the swath and in a horizontal straight line, then locally the derivatives of the elevation surface can be calculated by taking finite differences between consecutive rows and columns. The finite differences along the columns will give the derivative of the surface with respect to range z_s, and the finite difference taken in the row direction will give the derivative of the surface with respect to the distance travelled along the assumed straight line fish path z_p. This first interpolation results in three maps embedded in the echo intensity image that correspond to the elevation at each pixel $z(p,s)$, the derivative of the surface with respect to the range $z_s(p,s)$ and the derivative of the surface with respect to distance travelled $z_p(p,s)$.

Once the initial estimates for the surface and its derivatives are generated, the local reconstruction for each ping can begin. Local maps of elevation and its derivatives are created from the rows of the image surrounding the current ping. The size of the image (i.e., number of rows included) can vary. The size of the local images should be chosen such that the overlap between pings in the image is minimal. Hence, small images should be used when the position of the fish changes greatly between pings; large images should be used when the fish basically moves in a straight line. Larger local images are better because the larger the image, the less likely boundary effects will affect the reconstruction. In all of the results shown in this paper we chose the images to be

40 rows wide. Once the local maps are created, the estimation of the parameters of the scattering parameters can proceed.

5.2 Local Scattering Parameter Estimation

The scattering model of a surface is a function of the incidence angle of the signal striking the surface and some set of parameters $\mathbf{p} = (p1,p2,...)$ that depend on the material properties of the surface and the sensor $\sigma = \sigma(\alpha,\mathbf{p})$. Even if the incidence angle of the surface is know at every point in the image, the expected intensity values predicted by the scattering model cannot be determined because the scattering parameters of the surface are not known. However, given the intensity image of the seafloor, the values of the scattering parameters can be predicted using least squares fitting of the scattering model to the backscatter data.

Because the seafloor is made up of regions of similar material type (which controls the values of the scattering model parameters), it is justifiable to assume that, in local neighborhoods on the seafloor, scattering parameters will have similar values. Using this assumption, we can break the backscatter image into regions where the scattering parameter values will be constant. For convenience we have made these constant regions square and of adjustable size. Because the scattering parameters are constant in square regions that are larger than one pixel they have a lower resolution than the elevation map. Each point in the image has an equation of the form $S_i = \sigma(\alpha_i,\mathbf{p})$, so increasing the size of the constant regions (hence the number of points where **p** is constant) will increase the number of equations which can be used to solve for the scattering parameters using non-linear minimization.

Under this assumption, in each region the parameters **p** that best fit the data S_i with error δ_i in the control point region can be found by minimizing the chi square merit function for fitting the scattering model to the data

$$\chi^2(\boldsymbol{p}) = \sum_{i=1}^{N} \left[\frac{S_i - \sigma\left(\alpha_i, \boldsymbol{p}\right)}{\delta_i} \right]^2 \tag{3}$$

where $\sigma(\alpha_i,\mathbf{p})$ is calculated at each pixel by substituting the current incidence angle and scattering parameters in the scattering model. This function was minimized with the Levenberg-Marquadt method for non-linear least squares minimization (Press et al. 1988). The resulting parameters

will be the best fit parameters to the data in the region given the scattering model and the intensity data.

A virtue of this fitting method is that different scattering models can be used simply by changing the function supplied to the minimization method. In this way different scattering models can be tested on the data to determine the one that is the most appropriate. The scattering model parameters are re-estimated each time a new estimate of the incidence angles is determined from the new elevation map. Convergence of the elevation map makes the incidence angles converge to their true values; as a result, the estimates of the scattering parameters should converge to their correct value. The scattering parameters are considered to be close to their true values when the value of χ^2 is close to 1.0. When this is true for all control point regions, the scattering parameters are considered to be converged for the surface. This is demonstrated by testing done on synthetic data where the true albedo is known.

A shortcoming of this method for parameter estimation is that it cannot always estimate all of the scattering parameters in the scattering model. Fortunately, this generally occurs when the parameter contributes little to the form of the scattering model. A specific example of this difficulty occurs when trying to estimate σ in the Torrance and Sparrow model at large values of α. This is precisely because, at large values of α, the specular component of scattering is negligible, so numerically a correct estimate of σ is not easy to obtain.

It is also difficult to estimate parameters correctly when there is a strong coupling between parameters. This coupling means there exist multiple sets of parameters that generate the same scattering model form. The data will fit the scattering model equally well for multiple sets of parameters, so the minimization method has no way to select the correct scattering parameters. However, the form of the scattering model will generally be estimated correctly. An example of this coupling occurs in the diffuse model where ρ and ε are multiplicatively coupled in the roughness term of the model.

Another difficulty in estimating the scattering parameters is that some parameters have definite constraints on the values that they can take. For instance all of the parameters cannot assume negative values. Occasionally, if the data is noisy or the surface estimate is very wrong, the estimate of the parameters creates a form for the scattering model that fits the data well and yet is generated

by parameters that are not physically feasible. To circumvent this problem, we add terms of the form $e^{p-p_{ub}}$ and $e^{p_{lb}-p}$ where p_{lb} and p_{ub} are the upper and lower bounds respectively on the parameter p to the summation in Equation 3. These terms become very large when the parameters go outside of their expected ranges. Thus when a parameter is outside its range, the value of χ^2 becomes very large. This will force the parameters to be adjusted in such a way as to decrease χ^2 and will occur only when the parameter that was outside of its range comes back into range.

Each time that the surface is reestimated, the scattering parameters are reestimated as well. In this way, the scattering parameters converge to their best fit form as long as the surface is converging to its true value. Figure 6 shows the convergence of the albedo parameter for a synthetic data generated from a cylinder. Initially the estimates of albedo for the quite different from there true value of 0.5. However, as the algorithm proceeds, the estimates of the albedo improve drastically. The isoplots show that the estimates of albedo have a much larger spread and are far from the actual value of albedo when the algorithm starts, but at the end of the algorithm, the estimates have a small spread and are close to the expected value for the albedo. The graphs show that initially the estimates of incidence angle are incorrect, but at the end of the iterations to improve the shape of the surface, the incidence angle are estimated correctly, so the albedo estimates are much more accurate. This test on synthetic data demonstrates that the algorithm for estimating scattering parameters is feasible.

FIGURE 6. Convergence of the albedo parameter for a synthetic data set following the lambertian scattering model. Initial albedo isoplot (top left), final albedo isoplot (top right), initial graph of backscatter vs. incidence angle (bottom left) and final graph of backscatter vs. incidence angle (bottom right).

5.3 Determining the Elevation Map

The backscatter image of the seafloor does not contain enough information about the shape of the seafloor to determine it exactly. If we had exact knowledge of the scattering model for the surface as well as its parameters, we would still find it impossible to determine the shape of the seafloor uniquely. Each pixel in the backscatter image can be directly related to the angle of incidence of the acoustic pulse with the surface, but the angle of incidence is a function of the elevation of the surface and its two surface derivatives. Thus with perfect scattering information for every pixel there will be only one equation for every 3 unknowns $S = \sigma(\alpha(z,z_s,z_p),\mathbf{p})$ In order to solve for the elevation map, some assumptions must be made.

Our approach is to solve for the surface by minimizing some global criteria that measure the goodness of fit to the data. Because the data can fit an infinite number of surfaces, constraints must be imposed on the shape of the surface to decrease the total number of possible surfaces to one. The procedure is to define an error functional that strikes a balance between the goodness of fit of the surface to the data and the need to constrain the surface so that a unique solution can be obtained. Minimization of such an error functional using shape from shading techniques appears frequently in computer vision. Originally we formulated the error functional in cylindrical coordinates and determined the elevation in a two stage relaxation process (Johnson 1993). However, after some further experimentation, we determined that the form given in Equation 4 using cartesian coordinates gave better results for the maps of z, z_s and z_p. This error functional is similar to that developed by Horn and Brooks (1988) for shape from shading with an integrability penalty term except that the seafloor backscatter model is a function of elevation as well as its derivatives.

$$\iint (S(s,p) - \sigma(z_s, z_p, z))^2 ds dp + \mu \iint \left(z_s - \frac{\partial z}{\partial s} \right)^2 + \left(z_p - \frac{\partial z}{\partial p} \right)^2 ds dp + \lambda \iint \left(z_{ss}^2 + z_{sp}^2 + z_{ps}^2 + z_{pp}^2 \right) ds dp \qquad (4)$$

The first integral in this functional ensures that the data matches the scattering model in the least squares sense and the second integral is a constraint term that tries to maintain the integrability of the surface. These two integrals are sufficient to uniquely determine the surface given the intensity scattering data; however, the solution obtained by using only these two terms will diverge if too much noise is present in the image. The third integral constrains the second derivatives of the surface (z_{ss}, z_{sp}, z_{ps} and z_{pp}) and ensures that the solution does not diverge by requiring that the reconstructed elevation map vary smoothly. The smoothness term will force the estimates of the surface to remain close together, thus preventing transient errors in the estimation from causing the surface estimates to diverge. The smoothness of the seafloor will vary from one geological province to the next, so the smoothness of the reconstructed surface can be adjusted by changing the weight λ on the smoothness error term. Similarly, the degree to which integrability is enforced on the surface is influenced by the value of μ in the functional. These values were empirically determined to give the best results when $\lambda = 0.1$ and $\mu = 0.01$.

The error functional is minimized by applying the calculus of variations to determine the Euler equations for the functional (Aleksandrov et al. 1964). Inserting finite differences estimates for the continuous derivatives will discretize the Euler equations (Horn 1986). These three discrete

Euler equations can then be solved for the three surface unknowns z, z_s and z_p to get the three iterative update equations given in Equation 5.

$$
\begin{aligned}
z_s^{n+1}(i,j) &= \frac{\mu}{\mu+\lambda}\frac{z^n(i,j+1)-z^n(i,j)}{\Delta s} + \frac{1}{\mu}\Big(S(i,j)-\sigma\Big(z_s^n, z_p^n, z^n\Big)\Big)\frac{\partial\sigma}{\partial z_s} + \frac{\lambda}{\mu}\overline{z_s^n}(i,j) \\
z_p^{n+1}(i,j) &= \frac{\mu}{\mu+\lambda}\frac{z^n(i+1,j)-z^n(i,j)}{\Delta p} + \frac{1}{\mu}\Big(S(i,j)-\sigma\Big(z_s^n, z_p^n, z^n\Big)\Big)\frac{\partial\sigma}{\partial z_p}\Big(z_s^n, z_p^n, z^n\Big) + \frac{\lambda}{\mu}\overline{z_p^n}(i,j) \\
z^{n+1}(i,j) &= \overline{z^n}(i,j) + \frac{z^n(i,j+1)-z^n(i,j)}{\Delta s} + \frac{z^n(i+1,j)-z^n(i,j)}{\Delta p} + \frac{1}{\mu}\Big(S(i,j)-\sigma\Big(z_s^n, z_p^n, z^n\Big)\Big)\frac{\partial\sigma}{\partial z}\Big(z_s^n, z_p^n, z^n\Big)
\end{aligned}
\tag{5}
$$

The bar represents an average of the four nearest neighbors of a point. The scattering model is a function of the z_s, z_p and z because, as is shown in the appendix, the incidence angle can be defined as given in Equation 6.

$$
\alpha = \operatorname{acos}\left(\frac{z - Rz_s}{\sqrt{R^2\left(z_s^2 + z_p^2 + 1\right) - 2Rzz_s}}\right). \tag{6}
$$

5.4 Algorithm

By iteratively adjusting the elevation map, the elevation derivative maps, and the scattering parameter maps until the estimates change very little, the algorithm will find the best estimate of the elevation in fish coordinates given the data. The complete local surface reconstruction algorithm is as follows.

For all pings in the acoustic backscatter image:

1. Linearly interpolate z between the known elevation control points.
2. Assuming that locally the fish moves in a straight line, create local derivative maps centered around the current ping by taking finite differences between rows and columns in the interpolated elevation image.
3. Estimate the scattering parameters by fitting the data from square regions in the local map to the scattering model using non-linear minimization.
4. Determine the new values for z_s, z_p and z using Equation 5.
5. Repeat steps 3 and 4 until the intensity predicted by the reconstructed surface is close to the intensity given by the data and the elevation estimate do not change very much.

5.5 Local Results

The convergence of the elevation estimate for a synthetically generated cylinder that employed the lambertian scattering model and had elevation control points spaced every 40 pixels can be seen in Figure 7. A cross-section of the three isoplots at row 25 shows the convergence of the estimate of the elevation *z*. For this experiment on synthetic data a fixed boundary condition is enforced to prevent the surface from diverging too greatly from its initial estimate. The fixed boundary will not affect the final reconstruction because for local reconstruction we are only concerned with the middle of the surface. Initially the surface is flat, but at the end of the iterations the surface takes on the shape of the curved cylinder. This test on synthetic data shows that the surface refinement algorithm can use backscatter data to improve the estimate the surface of the seafloor given by sparse elevation data points. The convergence of the local elevation estimate on a set of real data is shown in Figure 8. The initial estimate for the real surface was generated from points lying at the corners of a 250 m by 150 m rectangle. Each pixel in the images represents an area of 0.5 m by 0.5 m. The sequence of isoplots shows the initial flat estimate, the final reconstructed surface and the actual surface predicted from the bathymetry given by the sensor for a small section of the reconstructed surface. The real surface is not recovered exactly, but the final estimate is better than the initial estimate. The graph of a row from each of the three isoplots also shows that the final estimate of the surface is a better estimate of the surface that the initial estimate. Figure 9 shows the error histograms for difference between the true surface and the initial estimate and the true surface and the final estimate. The histograms show that the average error decreases from 4.6 m in the initial estimate to 1.6 m in the final estimate, a 2.875x improvement. Figure 9 also shows the average percentage decrease in the error between the actual surface and the estimates surface. The error decreases monotonically to a local minimum. The final elevation estimate does not completely agree with the actual surface because the scattering data does not

correspond exactly to that predicted by the surface and the surface generate must remain smooth, which will prevent any rapid changes in the curvature of the surface.

FIGURE 7. Convergence of local elevation for a synthetic image of a cylinder. Initial estimate (top left), final reconstruction (top right), actual surface (bottom left) and row 25 from each preceding isoplot (bottom right).

FIGURE 8. Convergence of the local elevation estimate for a small area of ridge flank. Initial estimate (top left), final reconstructed surface (top right), actual surface (bottom left) and row 55 from each of the three preceding isoplots (bottom right).

FIGURE 9. Initial and final elevation error histograms (left) for ridge flank and percent decrease in average error as a function of algorithm iterations (right).

In the tests we ran, the elevation estimates generally converge. In cases of divergence the values of λ and μ were increased and the test were ran again until convergence was achieved. Our algorithm finds the local minimum closest to the initial estimate that minimizes the functional given in Equation 4, so it is possible that the final estimate for the surface is not physically correct. To ensure that the global minimum for the surface is found, a more complicated, possibly stochastic, minimization procedure (e.g. simulated annealing) should be used.

A more sophisticated minimization procedure could also be used to decrease the number of iterations required for convergence. An iteration schedule that starts out with small values for λ and μ and the gradually increases there values as the number of iterations increases would initially fit the surface based solely on the backscatter data. Then, as the number of iterations increases, the values of λ and μ would increase and more smoothing would occur. This form of scheduling increases the speed of convergence because greater changes in surface shape are allowed during the initial iterations.

6 Local to Global

Acoustic images can be deceiving because they are created by stacking consecutive pings on top of each other. This stacking has the effect of tricking the viewer into thinking that the image was created by towing the fish in a horizontal straight line. In general, the fish's path will deviate from a straight line as it is towed through the water due to variable tension in the tether, surface turbulence and water currents. Since the pose of the fish influences the backscatter that the sensor detects (because the angle of incidence of the acoustic pulse is a function of the surface normal

and the position of the fish) the backscatter will have meaning only in the coordinates of the fish for the current ping. In the previous section, a method for reconstructing the elevation profile for the current ping in current fish coordinates was presented. However, these local reconstructions are useless unless they are transformed to some common coordinate system where true global proximity can be determined. Given this proximity, meaningful 2 1/2 D elevation maps of the sea-floor can be generated. Transformation of the elevation profiles for each ping to global coordinates is straightforward, while transformation of acoustic backscatter which is necessary for global elevation refinement is not.

6.1 Binning and the 3D Transformation

The global coordinates are arbitrarily chosen to be the coordinates of the fish for the first ping in the image. Given that the fish moves in a path that is close to a horizontal straight line, this assignment will make the transformation from local to global coordinates small, hence making the global surface appear similar to the local reconstructions. The global elevation surface is generated on a square grid whose origin corresponds to the origin of the global coordinate system. The rows of the elevation map correspond to the *y* global coordinate and the columns correspond to the *x* coordinate.

For each ping, the local elevation profiles which are in cylindrical fish coordinates are transformed into cartesian global coordinates using the pose of the fish as is shown in Equation 7.

$$\begin{bmatrix} x \\ y \\ z \end{bmatrix}_{global} = R_F^G(\theta, \varphi, \omega) \begin{bmatrix} \sqrt{R^2 - z^2} \\ y \\ z \end{bmatrix}_{local} + \begin{bmatrix} x \\ y \\ z \end{bmatrix}_{fish} \tag{7}$$

The x and y global coordinates of a point determine the grid cell in the global map in which to place the point. Depending on the resolution of the global map, multiple points may fall into the same grid cell. If this is the case, the global *z* values of all of the points are averaged to get the initial estimate of the global elevation for that grid cell. In general, the resolution of the global map is set such that, on average, each grid cell will have one local point fall into it. This resolution will strike a balance between the gaps and overlap that might appear in the global map. The transformation and binning of the local elevation estimates creates the initial global surface estimate for the that will be refined again using the backscatter data. The problem then is to find a way to

transform the backscatter data such that the data is still useful and little or none of its information is lost.

Transforming and then averaging of the local elevation points that fall in the same grid cell is appropriate because the addition of elevation values is meaningful. However, a meaningful way to combine backscatter is not apparent. The backscatter data is a function of the position of the sensor and can vary over the points that fall in the same grid cell. It is entirely possible for scattering values that fall in the same grid cell to originate from pings whose fish positions are very different. Simple averaging of these types of scattering values will be useless because the position of the fish of the averaged scattering value is undefined. The proposed solution to this problem is to create a bin data structure that stores all of the backscatter values and the corresponding position of the fish when the data was collected for points that fall in each grid cell. This solution ensures that none of the backscatter data is lost and that the backscatter data can be used because the position of the fish is known.

6.2 Importance of Local Reconstructions

The local reconstructions were necessary before the transformation to global coordinates because of the cylindrical nature of the data in fish coordinates. The estimate of local elevation determines the local x and z coordinates of the data point, so if the elevation estimate is bad, these coordinates will be bad; the data point could be transformed into the wrong global grid cell. Since the grid cell into which the point is transformed remains fixed after transformation, this error could mar the global reconstruction. The better the local elevation estimate is the better the global reconstruction will be. If the local reconstructions are perfect, then no global reconstruction needs to be done once the points are transformed.

7 Detailed Global Algorithm

The local estimates are transformed into global coordinates and averaged in bins to create the initial estimate for the global elevation surface. The estimates of the derivatives of the surface can then be derived by taking the finite differences of the elevation surface along the rows and columns of the map. These estimates will be refined into the final estimate for the shape of the seafloor in global coordinates given the backscatter data. These estimates will also be used to

determine the final scattering parameter maps. The global reconstruction algorithm is quite similar to the local one, but takes into account the cartesian form of the elevation map and the binning of the scattering data.

7.1 Global Scattering Parameter Estimation

Once the initial surface estimates are generated, estimates for the scattering parameters are generated. Just as in the local algorithm, the global map is broken into square regions where the scattering parameters are considered to be fixed. This assumption then allows solution for the scattering parameters given the data using non-linear minimization. The difference from the local algorithm is that the data in the regions might come from different places in the original image and that each grid cell could contain more than one backscatter data point. However, the position of the fish that produced each data point is stored in the bin data structure, so we can use this information when doing the minimization. In effect, all that has changed from the local algorithm is the calculation of the incidence angle for each data point and the fact that we must use all of the data points that fall in each grid cell. To get the best fit scattering parameters given the backscatter data, we minimize the following merit function.

(8)

$$\chi^2(\boldsymbol{p}) = \sum_{i=1}^{nbins} \sum_{j=1}^{nbins} \left[\frac{(S_{ij} - \sigma(\alpha(\boldsymbol{n_i}, \boldsymbol{P_{ij}}), \boldsymbol{p}))}{\delta_{ij}} \right]^2$$

In Equation 8 all of the data is considered because all of the bins in the region and all of the data in the bins are used in the summation. The surface normal **n** is fixed for each bin while the position of the fish **P** varies over the bins and the data in the bins, so the incidence angle α is predicted correctly given the surface estimate for each data point. The functional form of the scattering model does not change between local and global algorithms, so this merit function can be minimized using the Levenberg-Marquadt method as was done in the local algorithm. The estimates of the scattering parameters will be better because on average, there will be more than one backscatter data point in each bin and the surface shape has been corrected for the movement of the sensor.

The parameter estimation uses all of the scattering values in each bin to produce one consistent parameter vector for each bin.

7.2 Global Elevation Reconstruction

Given the estimates of the scattering parameters in the global maps, the next stage in the global algorithm is to adjust the estimates of the seafloor so that they fit the data according to some global error functional. As was done for the local algorithm, we have constructed an error functional that measures the goodness of fit of the data to the surface estimate while maintaining the integrability of the surface and the convergence of the algorithm. The functional given in Equation 9 is similar to the local error functional except that all of the data in each bin is considered. This functional is similar to those used in shape from shading techniques where images of the scene from multiple cameras are available (Horn 1986).

$$\iint \sum_{j=1}^{ndata} \left(S(x,y,j) - \sigma(\alpha(z_x, z_y, z, \boldsymbol{P_j}), \boldsymbol{p})\right)^2 dxdy + \mu \iint \left(z_x - \frac{\partial z}{\partial x}\right)^2 + \left(z_y - \frac{\partial z}{\partial y}\right)^2 dxdy + \lambda \iint \left(z_{xx}^2 + z_{xy}^2 + z_{yx}^2 + z_{yy}^2\right) dxdy \tag{9}$$

The first integral compares all of the data to the scattering model estimate which is a function of the current surface estimates as well as the position of the fish when the data was recorded. The second integral maintains integrability of the surface and the third integral keeps the surface estimates from diverging by constraining the second order surface derivatives z_{xx}, z_{xy}, z_{yx} and z_{yy}. Iterative update equations are generated using the calculus of variations and discrete forms for the derivatives in the exact same way as was done for the local algorithm.

The control of the global algorithm is identical to the local algorithm. Initial estimates are generated; then the estimates for scattering parameters, the surface and its derivatives are iteratively refined until the surfaces converge. These estimates are the best global estimates of the elevation map of the seafloor and its corresponding parameter maps given the acoustic backscatter image. The global algorithm for refinement of the seafloor uses all of the scattering data in each bin during the reconstruction, so no information is lost.

8 Global Results

The global reconstruction process was first tested on synthetic data to ensure its viability. Figure 10 shows various stages in the global reconstruction process on a set of synthetic data generated as the fish moved in a semi-circular path around a cylinder. The initial estimate for the surface does not resembles a cylinder because the shape of the cylinder is distorted by the curved path of the fish. The final local elevation map shows a surface with greater curvature, but the geometry of the surface is still extremely skewed. Not until the surface is transformed to global coordinates in the initial global elevation map, does the shape of the cylinder become apparent. The curvature generated by the local reconstructions is present in the initial global elevation map, but the true curvature of the cylinder is not realized until global reconstruction algorithm completes. The final global reconstruction has the same shape as the actual cylinder imaged except at the boundaries where the elevation was kept fixed to constrain the surfaces generated. As is shown in the density of overlap image, the path of the fish for this synthetic data set caused the highest density (white) of overlapping data to occur at the center of the fish's circular path. This synthetic data set clearly shows how a nonlinear fish path can change the geometry of the surface imaged, and it showed that our global algorithm works to reconstruct the actual geometry of the surface in a global sense given the path of the fish

FIGURE 10. Convergence of the global elevation estimate for a synthetic data set where the fish moves in a circle about a cylindrical surface. Initial local estimate (top left), final local estimate (top right), initial global estimate (middle left), final global estimate (middle right), actual surface and path of sonar fish (bottom left) and overlap of data (bottom right).

The result of the global reconstruction process performed on the ridge flank data set using the lambertian scattering model is given in Figure 11. The initial estimate of the surface was generated from a regular grid of elevation control points spaced every 40 meters. The reconstructed corresponds very well to the actual surface although some sharp features are lost in the reconstruction due to the smoothing term in the elevation error functional. The recovered albedo shows that the albedo increases perpendicularly to the direction of travel of the fish and that in general the albedo of the surface is very low which is in agreement with the soft materials that made up the ridge flank. Verification of the recovered albedo is not possible because there are no absolute measurements of scattering parameters. The overlap image (white corresponds to a high degree of

overlap) shows that the fish moved in a fairly straight line, however, there do exist places in the image where scan lines overlapped.

FIGURE 11. Global reconstruction on the ridge flank data set. Recovered elevation map (top left), actual global elevation surface given path of fish (top right), recovered albedo parameter (bottom left) and overlap of data in global image (bottom right).

The estimation of the scattering model parameters from the reconstructed elevation maps can be used to segment the seafloor based on scattering model parameters. Because scattering model parameters depend on the material properties of the seafloor, this segmentation essentially finds regions of the seafloor of similar material. Figure 12 depicts how the seafloor can be segmented based on material type. The data was generated by imaging an axial valley region of the seafloor. The reconstruction algorithm determined that there existed regions of high and low albedo in the image when the lambertian scattering model was applied. In the axial valley, these regions could come from the sandy bottom of the valley close to the fish (with a low albedo) and the rocky sides of the valley far from the fish. This result shows that qualitative segmentation of the seafloor is possible with our reconstruction and parameter estimation algorithm

FIGURE 12. Segmentation of data set by albedo parameter. Actual elevation of surface (top left), estimate albedo isoplot (top right) and estimated albedo image showing regions of high and low albedo giving a segmentation of the seafloor based on material type (bottom).

Estimation of the scattering model parameters using the volume scattering model was conducted on data set taken of the ridge flank at a different location from the one used in the previous figures. The three isoplots of the estimated parameters along with the corresponding actual surface and bathymetric data are given in Figure 13. The volume scattering albedo is generally small which is in agreement with the previous ridge flank data of Figure 4. The two density ratios are within the expected ratios (greater than one and generally less than three).

FIGURE 13. Backscatter and elevation data with estimates of volume scattering model parameters. Elevation map (top left), backscatter data (top right), speed ratio parameter (middle left), density ratio parameter (middle right) and volume scattering albedo (bottom).

9 Conclusion and Contributions

In this paper we have presented an algorithm that generates global maps of elevation and scattering model parameters from a side-scan sonar backscatter image, sparse bathymetric data of the seafloor and the trajectory of the sensor through the water. This algorithm works independently of

the scattering model used and, in regions bounded by known elevation values, determines the best set of scattering parameters given the scattering model using non-linear least squares fitting. The elevation map of the seafloor is determined by minimizing a global error functional that measures the error between the scattering data and the backscatter predicted by the elevation map while taking into account the deviations of the path of the fish from a straight line. This algorithm required the development of a sophisticated method for merging backscatter data in a global coordinate system that minimizes the loss of information.

Our algorithm has been shown to correctly predict elevation and scattering parameter values for synthetic and real data using different scattering models. This flexibility allows the algorithm to be used by many different types of sidescan sensors and on seafloors of differing material properties. The algorithm has also been shown to build accurate elevation maps and segment scattering parameters where appropriate on real data

Our algorithm has a computational complexity linear in the number of pixels in the backscatter image. The running time also depends linearly on the number of iterations at each pixel which is independent of image size. Since the elevation map and scattering parameters estimations are local computations, our algorithm can be parallelized by partitioning the image into local regions, each being processed on a separate processor. This will decrease the computational complexity to one that is linear in the number of pixels in each local region. The total number of iterations on all pixels can also be reduced by implementing our algorithm in a coarse to fine fashion. Initially elevation and scattering parameters are iteratively estimated on a coarse representation of seafloor. The final coarse estimates are then input as initial estimates into the next finer level of processing. The total number of computations at each pixel will be decreased because smaller images are being operated on and better initial estimates are input from coarse levels to fine levels, so less iterations need to take place. Given the increase in run time that a parallel and coarse to fine approach would provide, it is feasible that this algorithm could eventually be implemented as a real-time mapping and navigational package for an autonomous underwater vehicle.

Future work will center around extensive testing of the algorithm. For the algorithm to be truly useful for AUV navigation it will have to be tested on data taken using many different types of sensors and on data gathered from seafloor regions of varying material types. Comparison of the

elevation maps generated to real bathymetry will be essential for validation of the effectiveness of the algorithm.

10 References

Aleksandrov, A. D., Kolmogorov, A. N. and Lavrent'ev, M. A. 1964. *Mathematics: Its Content, Methods and Meaning*. MIT Press: Cambridge, MA.

Baeck, T. and Schwefel, H.P. 1993. An overview of evolutionary algorithms for parameter optimization. *Evolutionary Computation*, 1(1): 1-10.

Blake, A. and Zisserman, A. 1987. *Visual Reconstruction*. MIT Press: Cambridge, MA.

Caruthers, J. W. and Novarini, J. C. 1993. Modeling bistatic bottom scattering strength including a forward scatter lobe. *IEEE J. Oceanic Engineering*, 18(2): 100-107.

Cervenka, P. and de Moustier, C. 1993. Sidescan sonar image processing techniques. *IEEE J. Oceanic Engineering*, 18(2): 108-122.

Clarke, J. H. 1994 Toward remote seafloor classification using the angular response of acoustic backscatter: A case study from multiple overlapping GLORIA data. *IEEE J. Oceanic Engineering*, 19(1): 112-127.

Cobra, D. T., Oppenheim, A. V. and Jaffe, J. S. 1992. Geometric distortions in side-scan sonar images: a procedure for their estimation and correction. *IEEE J. Oceanic Engineering*, 17(3): 252-268.

Cuschieri, J. M. and Hebert, M. 1990. Three-dimensional map generation from side-scan sonar images. L. Energy Resources Technology, 112: 96-102, June.

Denbigh, P. N. 1989. Swath bathymetry: Principles of operation and analysis of errors. *IEEE J. Oceanic Engineering*, 14(4): 289-298.

Elfes, A. 1987. Sonar-based real world mapping and navigation. *IEEE J. Robotics and Automation*, RA-3(3): 249-265.

Gensane, M. 1989. A statistical study of acoustic signals backscattered from the sea bottom. *IEEE J. Oceanic Engineering*, 14(1): 84-93.

Hebert, M. 1989. Terrain modeling for autonomous underwater navigation. In *Proc. Unmanned Untethered Submersible Technology Conf.*, pp. 502-511.

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

Horn, B. K. P. and Brooks, M. J. 1988. The variational approach to shape from shading. *Computer Vision, Graphics and Image Processing*, 33(2): 174-208.

Jackson, D. R., Winebrenner, D. P., and Ishimaru, A. 1986. Application of the composite roughness model to high-frequency bottom backscattering. *J. Acoustical Society of America,* 79: 1410-1422, May.

Johnson, A. E. 1993. Incorporating different reflection models into surface reconstruction. In *Proc. Unmanned Untethered Submersible Technology Conf.*, pp. 446-459.

Langer, D. and Hebert, M. 1991. Building qualitative elevation maps from underwater sonar data for autonomous underwater navigation. In *Proc. IEEE Int. Conf. Robotics and Automation*, pp. 2478-2483.

Leonard, J. J. and Durrant-Whyte, H. F. 1992. Directed Sonar Sensing for Mobile Robot Navigation. Kluwer Academic: Norwell, MA.

Malik, S. 1991. *Quantitative seafloor backscatter characterization using an interferometric sidescan sonar.* Master's Thesis, U. Virginia.

Matsumoto, H., Dziak, R. P. and Fox, C. G. 1993. Estimation of seafloor microtopographic roughness through modeling of acoustic backscatter data recorded by multibeam sonar systems. *J. Acoustical Society of America,* 94: 2776-2787, November.

Mazel, C. 1985 *Side Scan Sonar Record Interpretation*. Klein Associates: Salem, NH.

Michalopoulou, Z., Alexandrou, D. and de Moustier, C. 1994. Application of a maximum likelihood processor to acoustic backscatter for the estimation of seafloor roughness parameters. *J. Acoustical Society of America*, 95: 2467-2477, May.

Mitchell, N. C. and Somers, M. L. 1989. Quantitative backscatter measurements with a long-range side-scan sonar. *IEEE J. Oceanic Eng*ineering,14(4): 368-374.

Mourad, P. D. and Jackson, D. R. High frequency sonar equation models for bottom backscatter and forward loss. In *Proc. IEEE Oceans 89 Conf.*, pp. 1163-1175.

de Moustier, C. and Alexandrou, D. 1991. Angular dependence of 12-kHz seafloor acoustic backscatter. *J. Acoustical Society of America*, 90: 522-531, July.

Oren, M. and Nayar, S. K. 1992. Diffuse scattering model for rough surfaces. Dept. of Computer Science Technical Report 057-92, Columbia University, New York, NY.

Press, W. H., Flannery, B. P., Teukolsky, S. A. and Vetterling, W. T. 1988. *Numerical Recipes in C*, Cambridge University Press: New York, NY.

Rigaud, V. and Marcé, L. 1990. Absolute location of underwater robotic vehicles by acoustic data fusion. In *Proc. IEEE Int. Conf. Robotics and Automation*, pp. 1310-1315.

Stanton, T. K. 1984. Sonar estimates of seafloor microroughness. *J. Acoustical Society of America*, 74: 809-818, March.

Stewart, W. K., Marra, M. and Jiang, M. 1992 A hierarchical approach to seafloor classification using neural networks. *IEEE Oceans 92 Conf.*, pp. 109-113, October.

Stewart, W. K. 1989. Three-dimensional modeling of seafloor backscatter from sidescan sonar for autonomous classification and navigation. In *Proc. Unmanned Untethered Submersible Technology Conf.*, pp. 372-392.

Stewart, W. K., Chu, D., Malik, S., Lerner, S. and Singh, H. 1994. Quantitative seafloor characterization using a bathymetric sidescan sonar. *IEEE J. Oceanic Engineering,* 19(4): 599-610.

Torrance, K. E. and Sparrow, E. M. Theory for off-specular scattering from roughened surfaces," *J. Optical Society of America*, 57: 1105-1114,

Urick, R. J. 1983. *Principles of Underwater Sound*. McGraw-Hill: New York, NY.

von Alt, C. 1989. A 200 kHz deep sea interferometric side scan sonar system. In *Proc. IEEE Oceans 89 Conf.*, pp. 1136-1141.

11 Acknowledgments

We would like to thank Ken Stewart of Woods Hole Oceanographic Institution for providing the sidescan sonar data and Hanumant Singh of Woods Hole and Dirk Langer of the Robotics Institute for answering many questions about accessing and using the data. We would also like to thank Chris Leger for supplying the genetic algorithm routines for fitting complex scattering models.

12 Appendix: Local Angle of Incidence

The surface **F** that is going to be reconstructed can be parameterized in R and y in the following coordinates.

$$\boldsymbol{F}(R,y) = \begin{bmatrix} x(R,y) \\ y(R,y) \\ z(R,y) \end{bmatrix} = \begin{bmatrix} \sqrt{R^2 - z^2(R,y)} \\ y \\ z(R,y)) \end{bmatrix} \tag{10}$$

Assuming that the surface is explicit, at any point (R,y) the normal to the surface is given by

$$\boldsymbol{n} = \frac{\partial \boldsymbol{F}}{\partial R} \times \frac{\partial \boldsymbol{F}}{\partial y}. \tag{11}$$

Differentiating Equation 10 with respect to R and y gives

$$\frac{\partial \boldsymbol{F}}{\partial R} = \begin{bmatrix} \frac{R - z z_R}{\sqrt{R^2 - z^2}} \\ 0 \\ z_R \end{bmatrix} \qquad \frac{\partial \boldsymbol{F}}{\partial y} = \begin{bmatrix} \frac{-z z_y}{\sqrt{R^2 - z^2}} \\ 1 \\ z_y \end{bmatrix} \tag{12}$$

where z_R and z_y are the derivatives of the surface z in the R and y directions respectively, so

$$\boldsymbol{n} = \begin{bmatrix} -z_R \\ \frac{-R z_y}{\sqrt{R^2 - z^2}} \\ \frac{R - z z_R}{\sqrt{R^2 - z^2}} \end{bmatrix} \tag{13}$$

Figure 2 shows that for a sidescan sonar

$$\boldsymbol{R} \cdot \boldsymbol{n} = \|\boldsymbol{R}\| \|\boldsymbol{n}\| \cos\alpha \quad . \tag{14}$$

Given

$$\boldsymbol{R} = \begin{bmatrix} \sqrt{R^2 - z^2(R, y)} \\ 0 \\ z(R, y) \end{bmatrix} \qquad \|\boldsymbol{R}\| = R \qquad \|\boldsymbol{n}\| = \sqrt{\frac{R^2 (z_R^2 + z_y^2 + 1) - 2 R z z_R}{R^2 - z^2}} \tag{15}$$

So Equation 14 can be solved for α to get

$$\alpha = \operatorname{acos}\left(\frac{z - R z_R}{\sqrt{R^2 (z_R^2 + z_y^2 + 1) - 2 R z z_R}} \right). \tag{16}$$

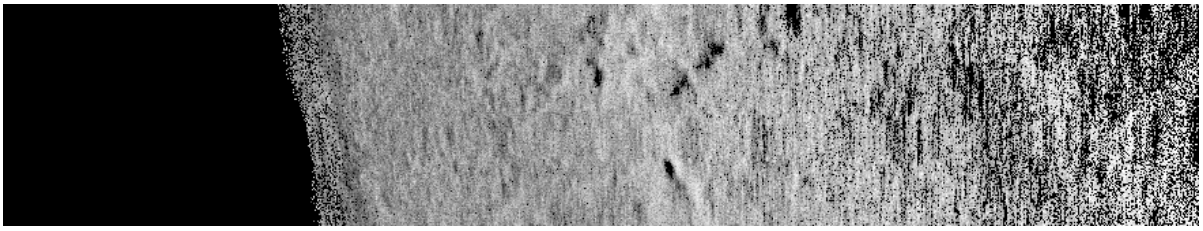

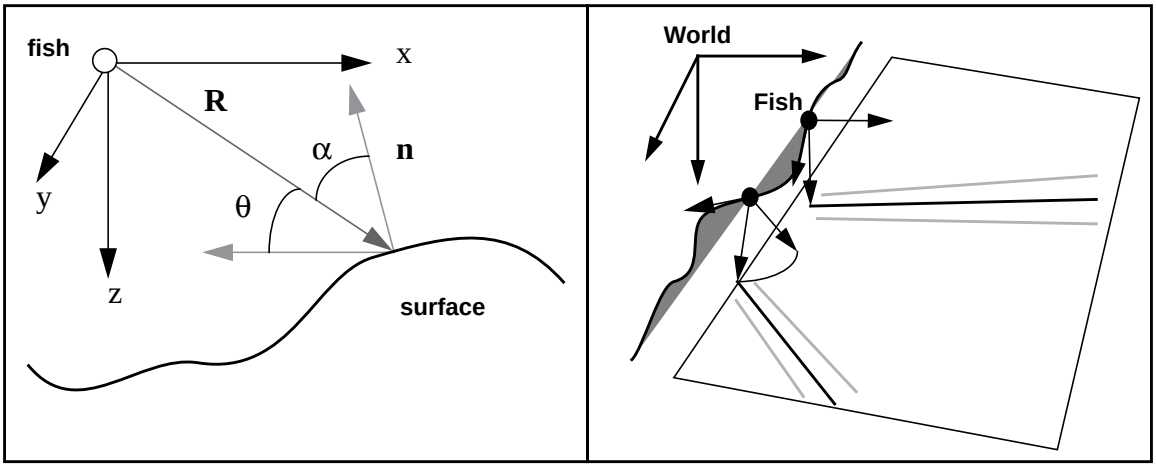
fish
x
R
α
n
y
θ
z
surface
World
Fish

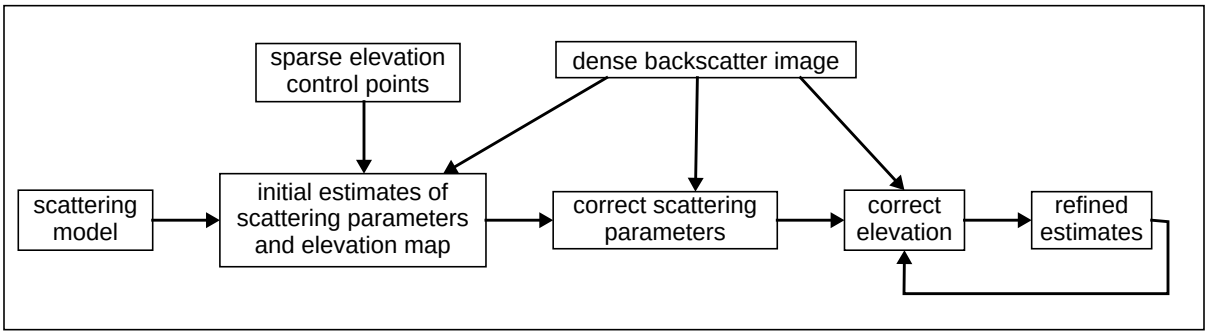
sparse elevation control points
dense backscatter image
scattering model
initial estimates of scattering parameters and elevation map
correct scattering parameters
correct elevation
refined estimates

Filtered and Unfiltered Backscatter

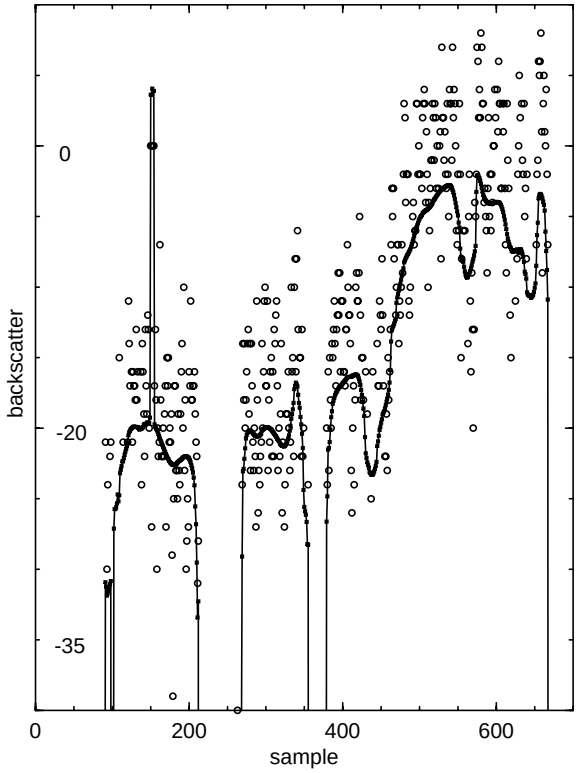

Filtered and Unfiltered Bathymetry

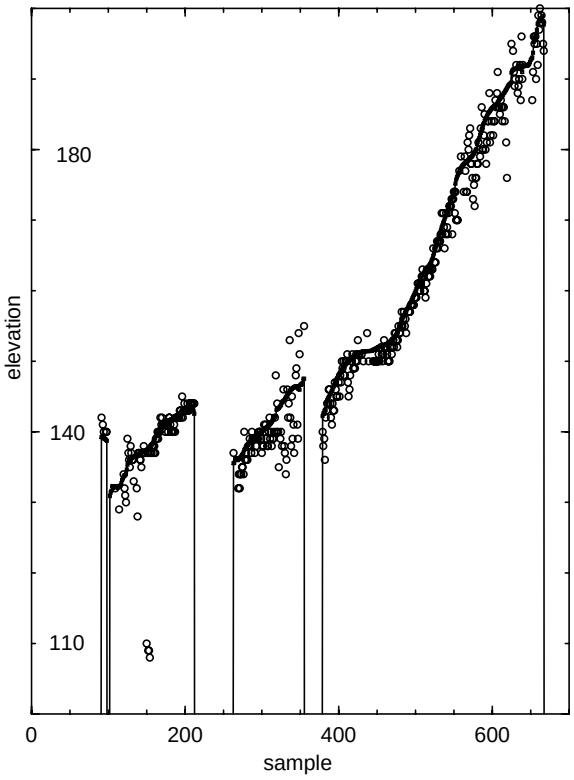

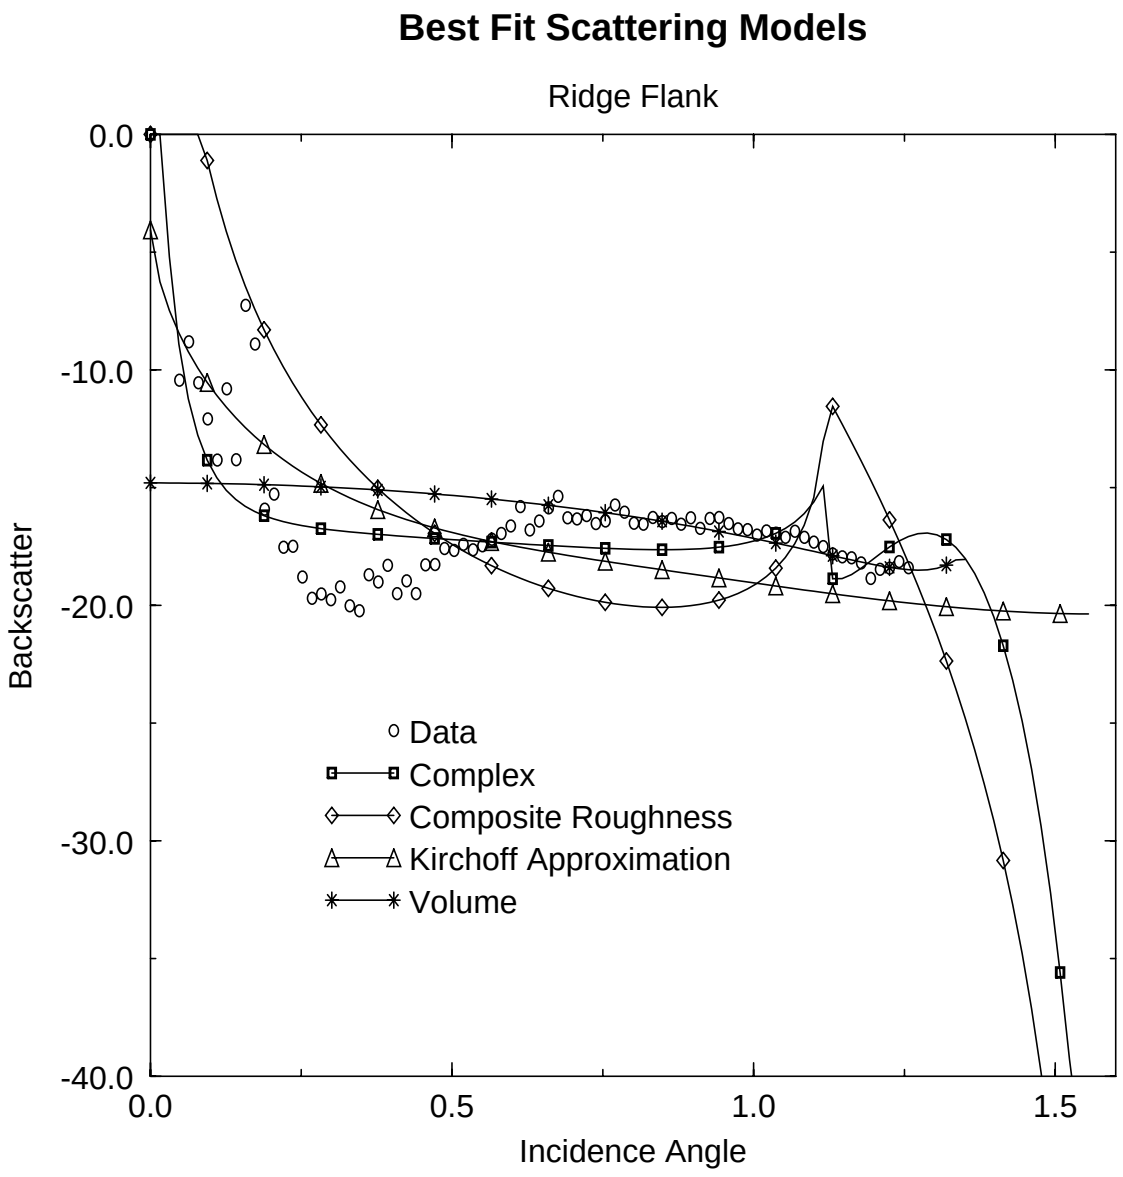
Best Fit Scattering Models
Ridge Flank
0.0
-10.0
-20.0
-30.0
-40.0
Backscatter
0.0
0.5
1.0
1.5
Incidence Angle
Data
Complex
Composite Roughness
Kirchoff Approximation
Volume

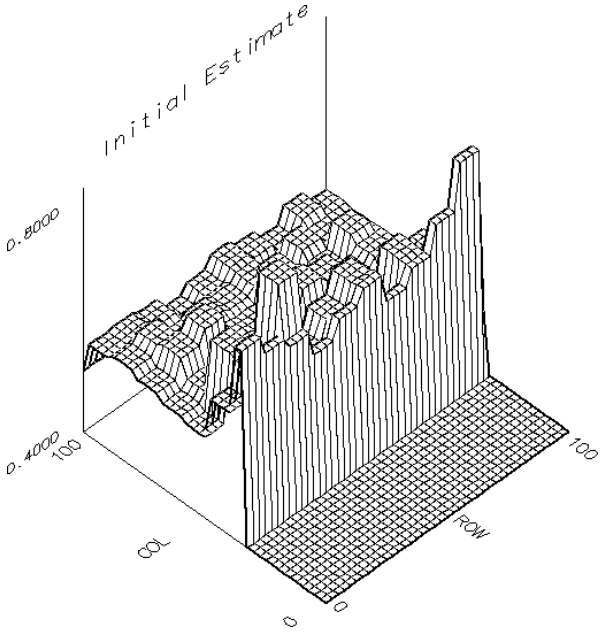

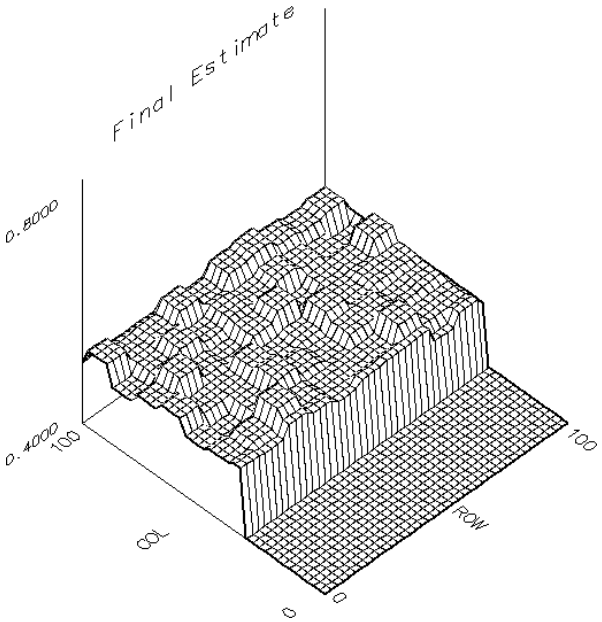

Initial Backscatter vs. Incidence Angle
with Fitted Scattering Models

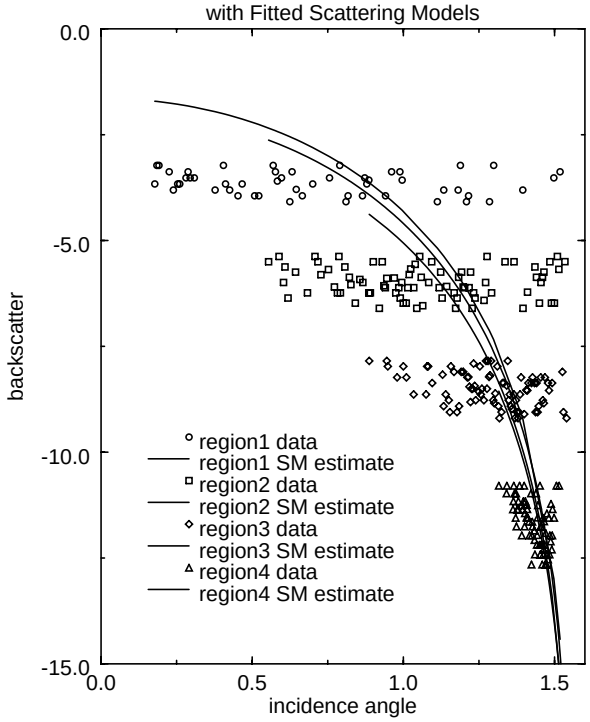

Final Backscatter vs. Incidence Angle
with Fitted Scattering Models

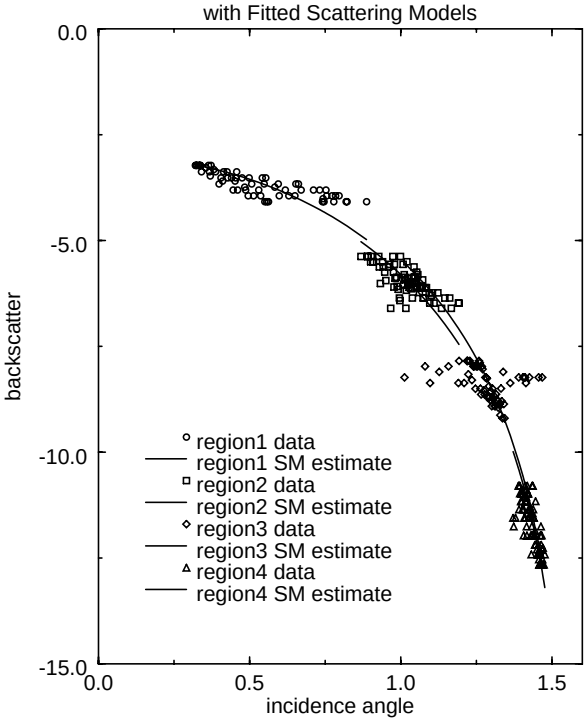

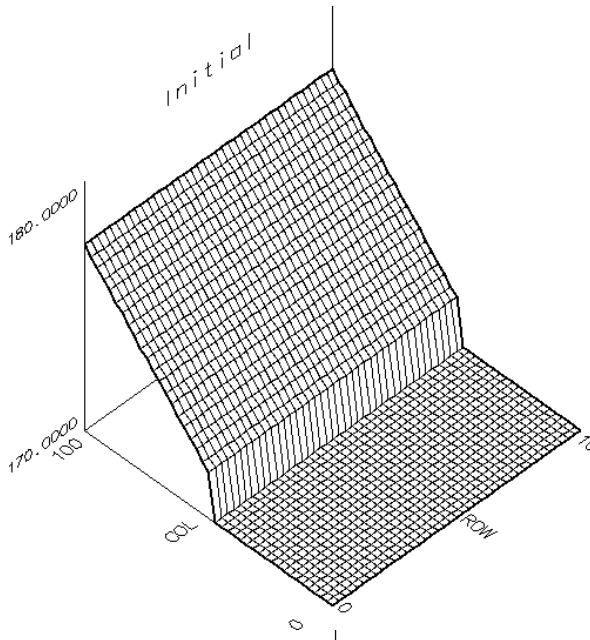
Initial
180.0000
170.0000
100
COL
ROW
0
0
100

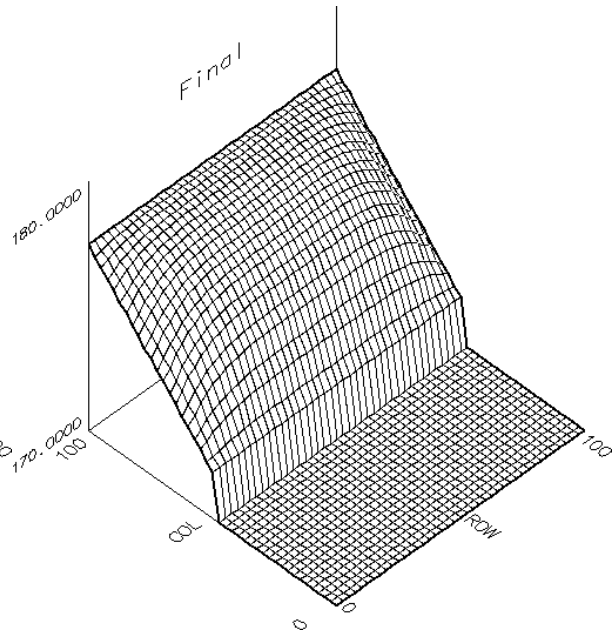
Final
180.0000
170.0000
100
COL
ROW
0
0
100

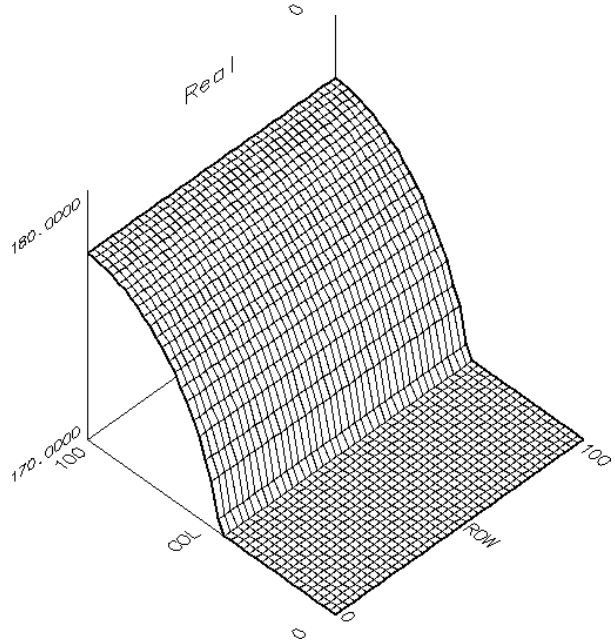
Real
180.0000
170.0000
100
COL
ROW
0
0
100

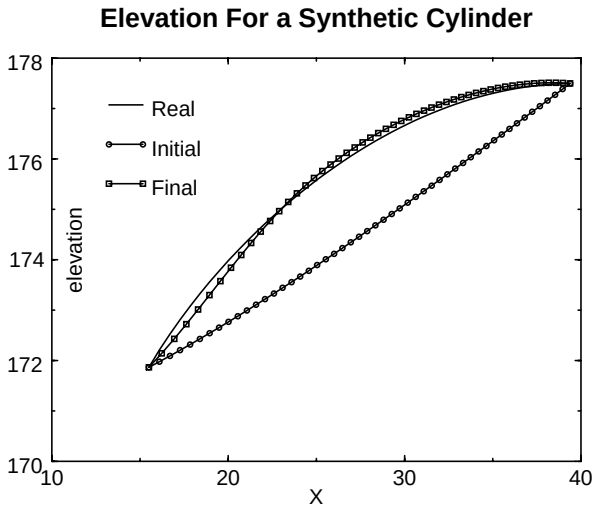
Elevation For a Synthetic Cylinder
178
176
174
172
170
elevation
Real
Initial
Final
10
20
30
40
X

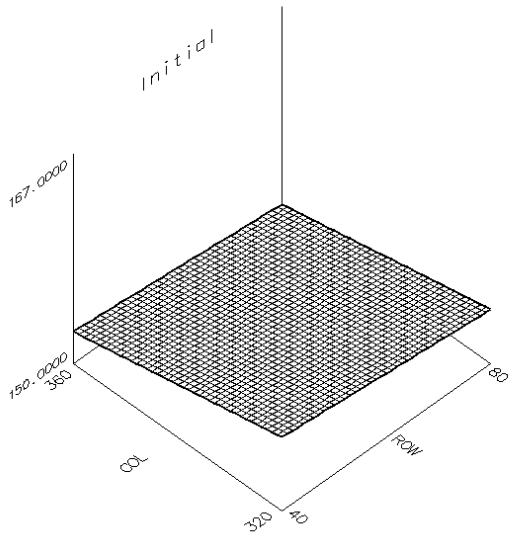
Initial
167.0000
150.0000
360
320
COL
40
ROW
80

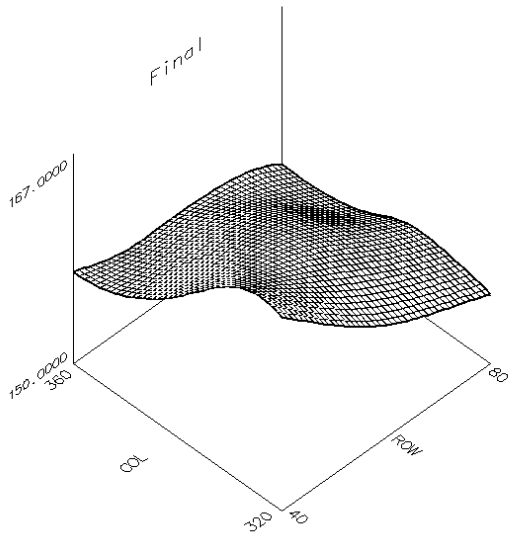
Final
167.0000
150.0000
360
320
COL
40
ROW
80

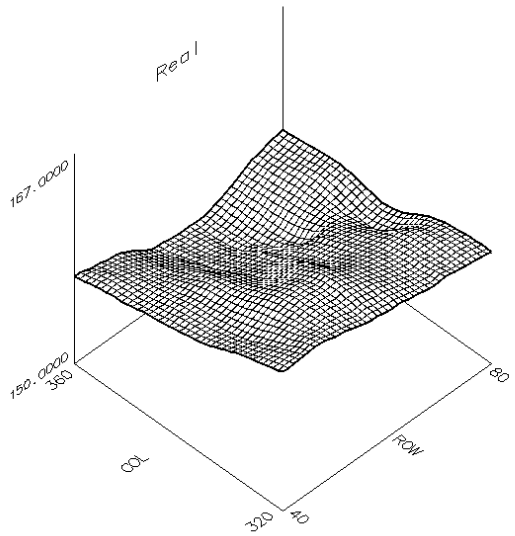
Real
167.0000
150.0000
360
320
COL
40
ROW
80

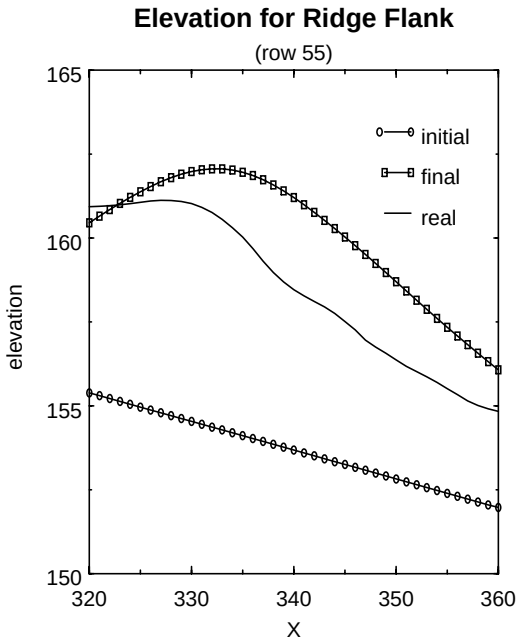
Elevation for Ridge Flank
(row 55)
165
160
155
150
elevation
320
330
340
350
360
X
initial
final
real

Local Elevation Error Histograms

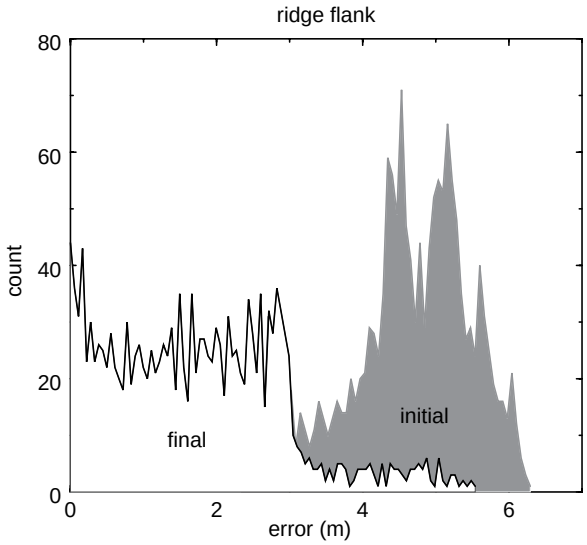

Average Elevation Error Convergence

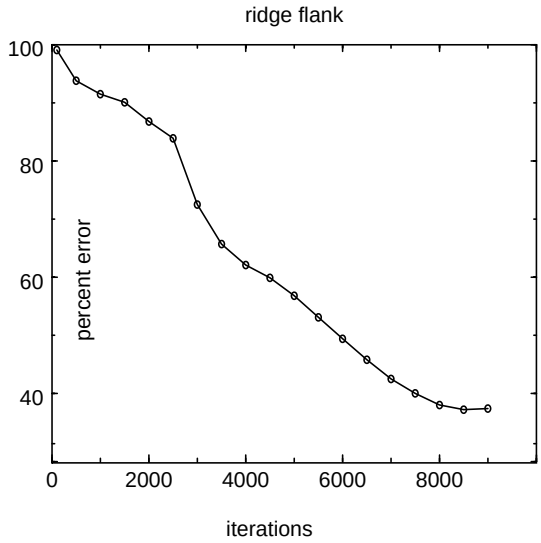

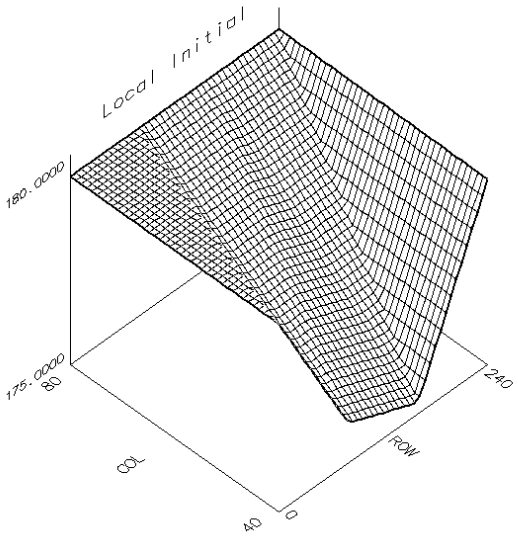

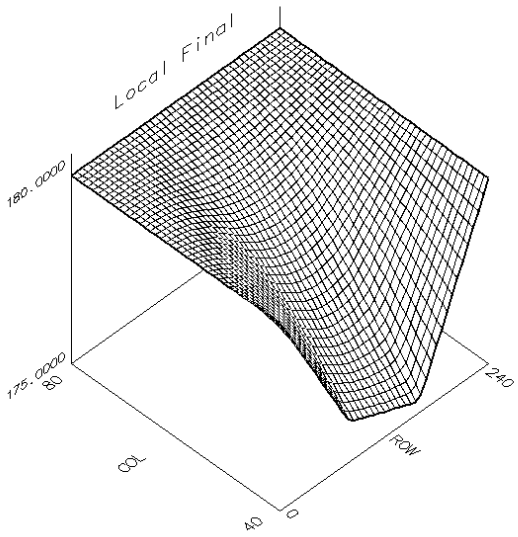

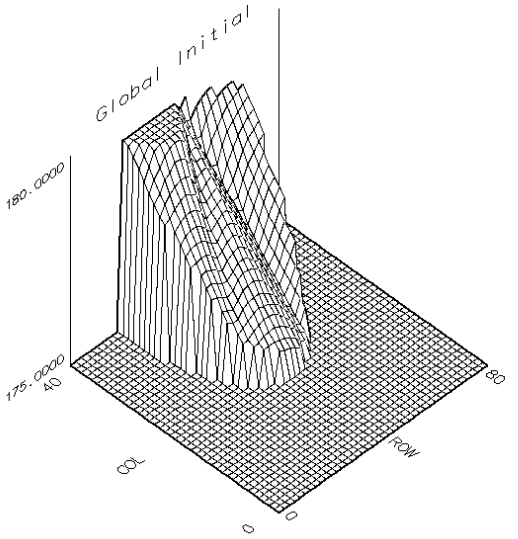

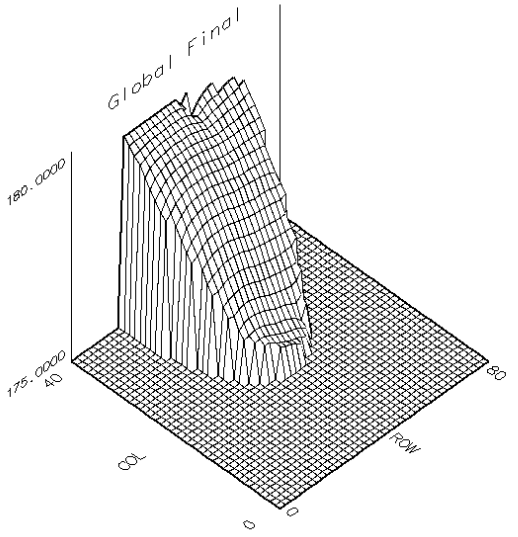

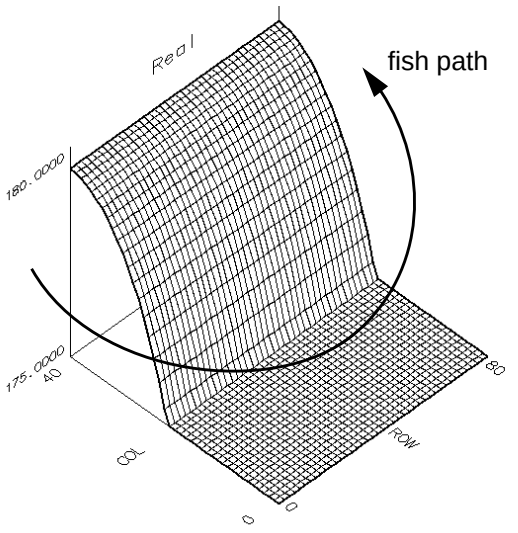

Overlap of Scans

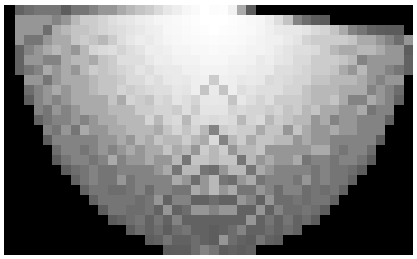

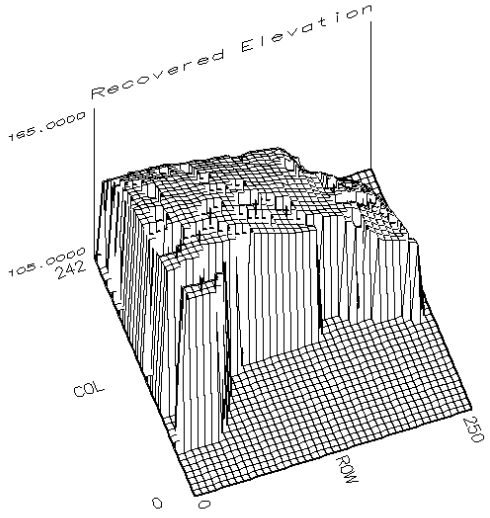

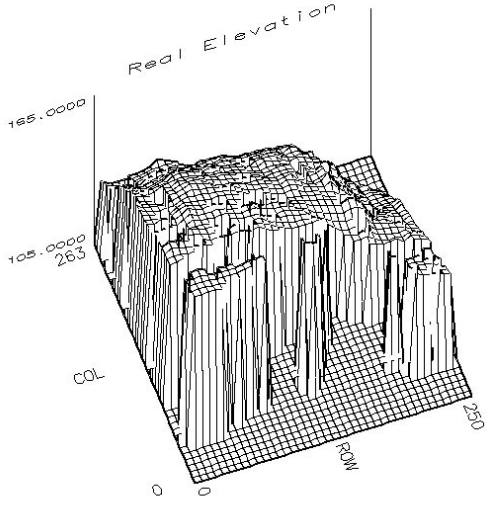

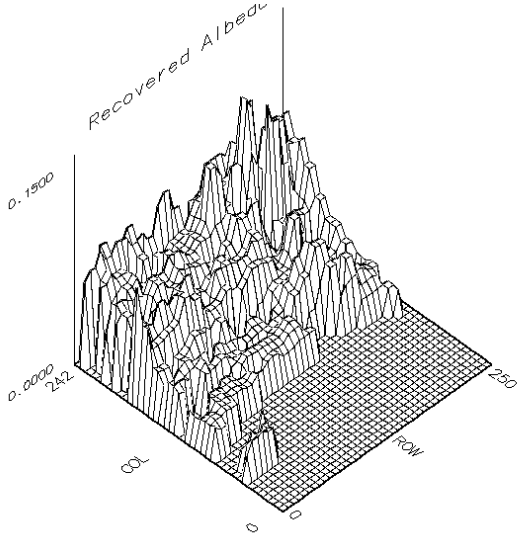

Overlap of Scans

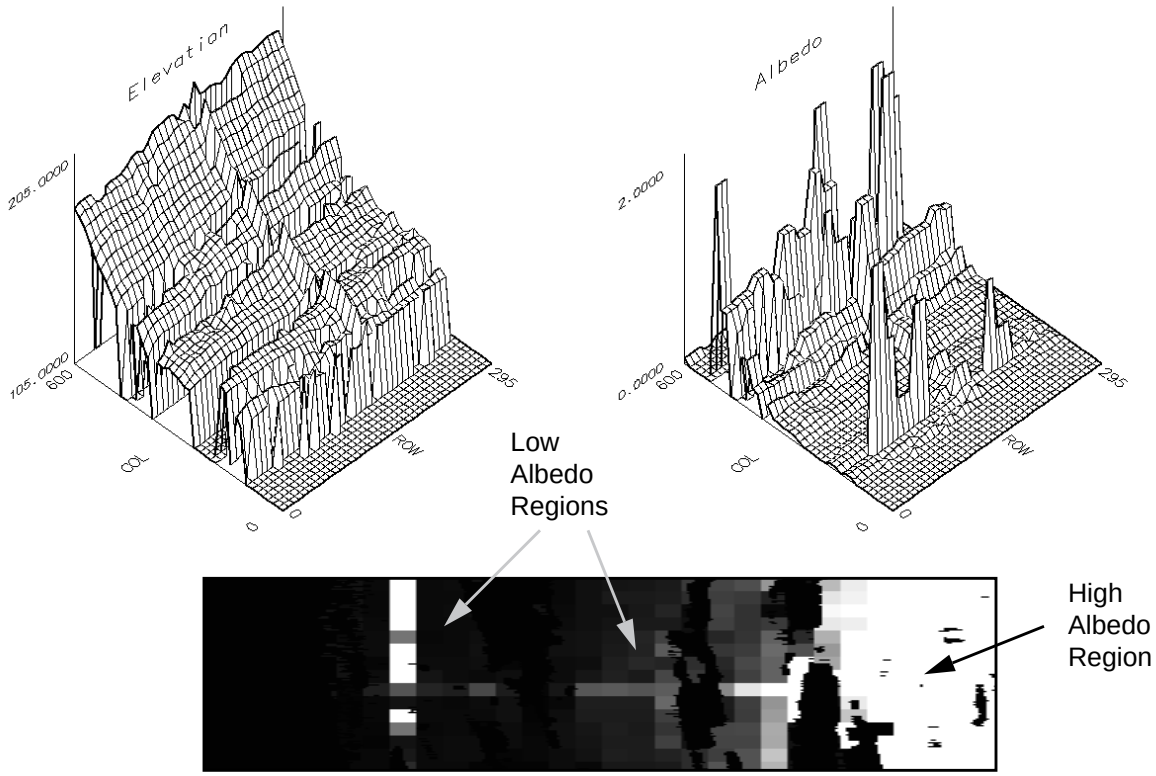
Elevation
205.0000
105.0000
600
COL
0
0
ROW
295
Albedo
2.0000
0.0000
600
COL
0
0
ROW
295
Low
Albedo
Regions
High
Albedo
Region

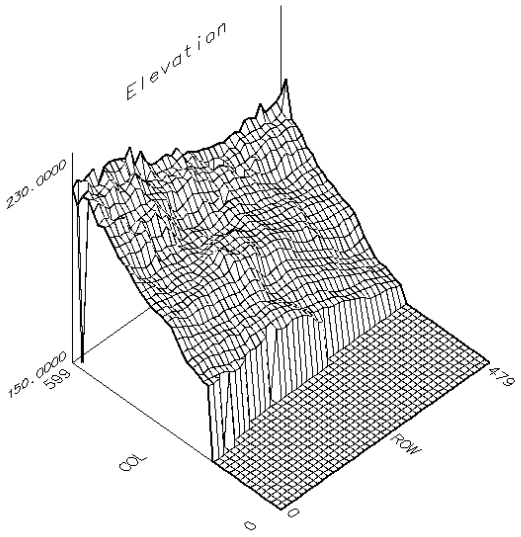
Elevation
230.0000
150.0000
599
COL
0
0
ROW
479

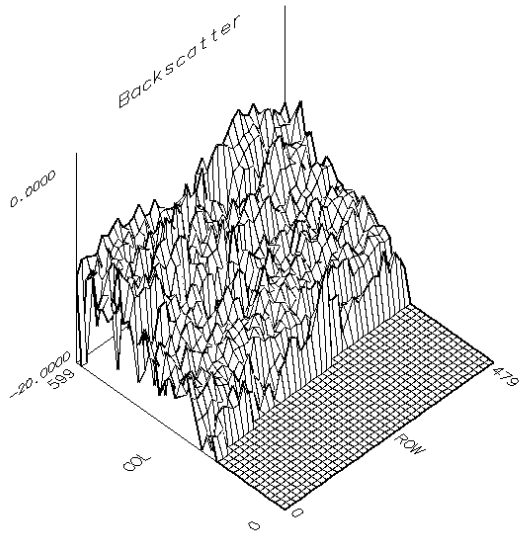
Backscatter
0.0000
-20.0000
599
COL
0
0
ROW
479

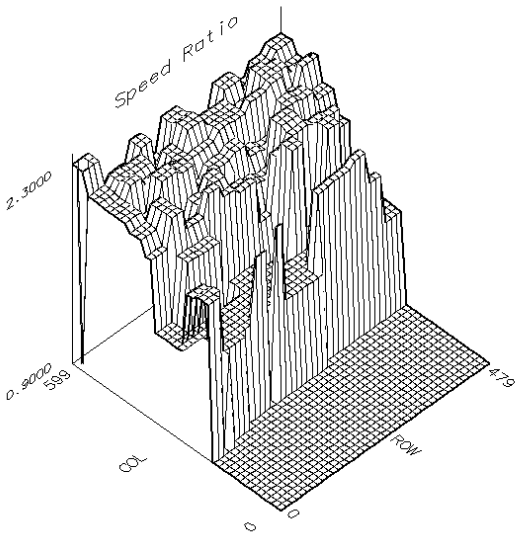
Speed Ratio
2.3000
0.9000
599
COL
0
0
ROW
479

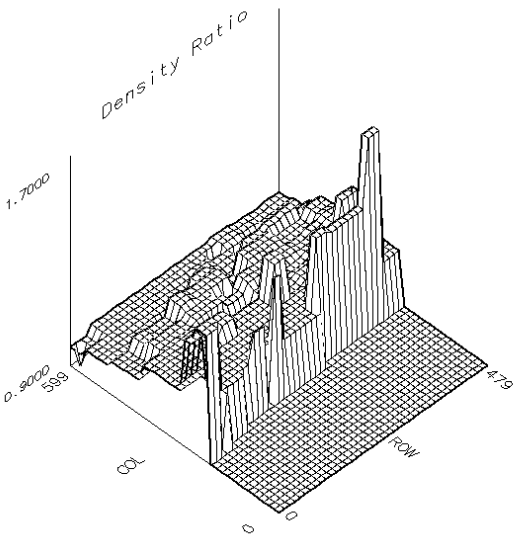
Density Ratio
1.7000
0.9000
599
COL
0
0
ROW
479

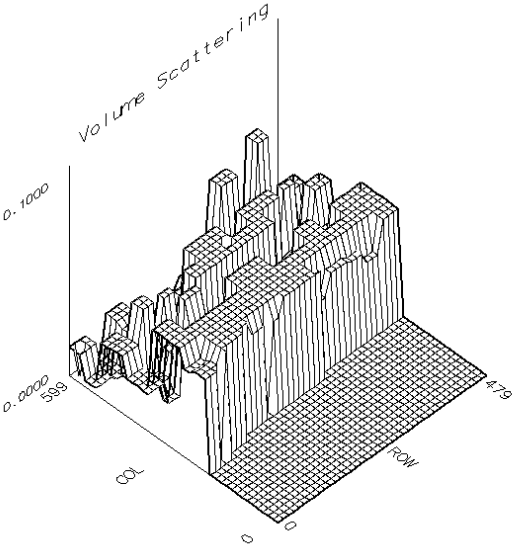
Volume Scattering
0.1000
0.0000
599
COL
0
0
ROW
479

Name	Form	Parameters
Lambertian	$\sigma = \rho \cos\alpha$	ρ = albedo
Diffuse (Oren & Nayar 1992)	$\sigma = \rho\left(\cos\alpha + \varepsilon \sin\alpha^2\right)$	ρ = albedo ε = roughness
Torrance & Sparrow (Torrance & Sparrow 1967)	$\sigma = \rho \cos\alpha\left((1-C) + 2Cexp\left(\frac{-\alpha^2}{2\sigma^2}\right)\right)$	ρ = albedo σ = rms surface roughness C = weight of body vs. surface scattering
Kirchoff Approximation (Jackson et. al. 1986)	$\sigma = \frac{bq_c R90^2}{8\pi\left(\sin\alpha^{4r} + \cos\alpha^4\right)^{\frac{1+r}{2r}}}$ $b = \frac{a^{\frac{r+1}{2r}}\Gamma\left(\frac{1}{r}\right)}{2r}$ $a = \frac{8r^2\Gamma\left(\frac{r+1}{2r}\right)}{\sqrt{\pi}\Gamma\left(\frac{1}{r}\right)\Gamma\left(\frac{1}{2r}\right)}$ $q_c = \frac{2\pi\beta\Gamma(2-r)2^{1-4r}}{r(1-r)\Gamma(1+r)}k_a^{2-2r}$ $R90 = \frac{\rho\upsilon - 1}{\rho\upsilon + 1}$	r = γ/2-1; γ = surface roughness spectrum exponent β = surface roughness spectral strength ρ =ratio of sediment mass density to water mass density ν = ratio of sediment sound speed to water sound speed k_a is the acoustic pulse wave number which can be adjusted for sensors of different frequencies.
Volume (Jackson et. al. 1986)	$\sigma(\alpha < \alpha_c) = \frac{5vol(1-g^2)^2\cos^2\alpha}{ln(10)sb}$ $g = \frac{\rho\upsilon\cos a - sb}{\rho\upsilon\cos\alpha + sb}$ $sb = \sqrt{1-\upsilon^2\sin^2\alpha}$	vol = volume scattering albedo ρ =ratio of sediment mass density to water mass density ν = ratio of sediment sound speed to water sound speed α_c = asin(1/ν) = surface critical angle
Composite Roughness (Jackson et. al. 1986)	$\sigma = 4k_a^4\cos^4\alpha FW$ $W = \beta(2k_a\sin\alpha)^{-2(r+1)}$ $F(\alpha < \alpha_c) = \frac{((\rho-1)^2\sin^2\alpha + \rho^2 + \upsilon^{-2})^2}{(\rho\cos\alpha + \sqrt{\upsilon^{-2} - \sin^2\alpha})^4}$ $F(\alpha \geq \alpha_c) = \frac{((\rho-1)^2\sin^2\alpha + \rho^2 + \upsilon^{-2})^2}{((1-\rho)^2\sin^2\alpha + \rho^2 + \upsilon^{-2})^2}$	r = γ/2-1; γ = surface roughness spectrum exponent β = surface roughness spectral strength ρ =ratio of sediment mass density to water mass density ν = ratio of sediment sound speed to water sound speed α_c = asin(1/ν) = surface critical angle