

SIGNAL PROCESSING ON DIGITIZED LADAR WAVEFORMS FOR
ENHANCED RESOLUTION ON SURFACE EDGES

by

Kevin D. Neilsen

A thesis submitted in partial fulfillment
of the requirements for the degree

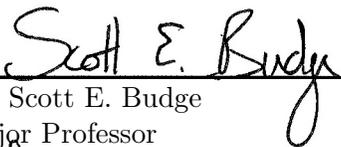
of

MASTER OF SCIENCE

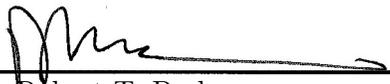
in

Electrical Engineering

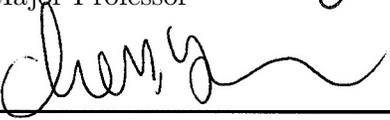
Approved:



Dr. Scott E. Budge
Major Professor



Dr. Robert T. Pack
Committee Member



Dr. YangQuan Chen
Committee Member

Dr. Byron R. Burnham
Dean of Graduate Studies

UTAH STATE UNIVERSITY
Logan, Utah

2011

Copyright © Kevin D. Nielsen 2011

All Rights Reserved

Abstract

Signal Processing on Digitized Ladar Waveforms for Enhanced Resolution on Surface
Edges

by

Kevin D. Neilsen, Master of Science

Utah State University, 2011

Major Professor: Dr. Scott E. Budge
Department: Electrical and Computer Engineering

Automatic target recognition (ATR) relies on images from various sensors including 3-D imaging ladar. The accuracy of recognizing a target is highly dependent on the number of points on the target. The highest spatial frequencies of a target are located on edges. Therefore, a higher sampling density is desirable at these locations. A ladar receiver captures information on edges by detecting two surfaces when the beam lands partially on one surface and partially on another if the distance between the surfaces is greater than the temporal pulse width of the laser.

In recent years, the ability to digitize the intensity of the light seen at the ladar receiver has led to digitized ladar waveforms that can be post-processed. Post-processing the data allows signal processing techniques to be implemented on stored waveforms. The digitized waveform provides more information than simply a range from the sensor to the target and the intensity of received light. Complex surfaces change the shape of the return.

This thesis exploits this information to enhance the resolution on the edges of targets in the 3-D image or point cloud. First, increased range resolution is obtained by means of deconvolution. This allows two surfaces to be detected even if the distance between them

is less than the width of the transmitted pulse. Second, the locations of multiple returns within the ladar beam footprint are computed.

Using deconvolution on the received waveform, an increase from 30 cm to 14 cm in range resolution is reported. Error on these measurements has a 2 cm standard deviation. A method for estimating the width of a 19 cm slot was reported to have a standard deviation of 3.44 cm. A method for angle estimation from a single waveform was developed. This method showed a 1.4° standard deviation on a 75° surface. Processed point clouds show sharper edges than the originals.

The processing method presented in this thesis enhances the resolution on the edges of targets where it is needed. As a result, the high spatial frequency content of edges is better represented. While ATR applications may benefit from this thesis, other applications such as 3-D object modeling may benefit from better representation of edges as well.

(103 pages)

Acknowledgments

I would like to thank my committee—Dr. Scott Budge, major professor; Dr. Robert Pack, committee member; and Dr. YangQuan Chen, committee member—for the time they have put into this thesis. I also appreciate Dr. Todd Moon and Roger West for their valuable input and ideas. I want to thank NAVAIR for funding this work and for the use of equipment to collect data. Finally, I want to thank my wife, Martilyn, and daughter, Rachel, for their support in completing this thesis.

Kevin D. Neilsen

Contents

	Page
Abstract	iii
Acknowledgments	v
List of Tables	viii
List of Figures	ix
Acronyms	xii
1 Introduction	1
1.1 Prior Work	2
1.1.1 Surface Estimation	2
1.1.2 Spatial Relationship	4
1.2 Problem Statement	5
1.3 Overview	5
1.4 Thesis Contribution	6
2 Background	7
2.1 Three-dimensional Imaging Ladar (Laser Detection and Ranging)	7
2.2 VISSTA ELT	12
2.3 LadarSIM	15
3 Signal Processing to Obtain the Surface Response	19
3.1 Data Processing Model	19
3.2 Low-Pass (Matched) Filtering	20
3.3 Interpolation	21
3.4 Deconvolution	26
3.5 Range Discrimination	29
4 Surface Response Estimation	31
4.1 Surface Response Derivation	31
4.1.1 Surface Response of Two Offset Surfaces	32
4.1.2 Surface Response of a Planar Surface	34
4.2 Regularization	36
4.3 Surface Response Results	42
4.3.1 Range Error on a Single Surface	42
4.3.2 Resolution of Two Surfaces	44
4.3.3 Effect of SNR on Range Resolution	47
4.3.4 Angle Estimation	50
4.3.5 Slot Width Estimation	55

5 Estimation of Spatial Location Within the Footprint	57
5.1 Method for Locating Returns Within the Footprint	61
5.2 Results of Locating Returns Within the Footprint	68
5.2.1 Slotted Surface	68
5.2.2 Angled Surface	73
5.2.3 Complex Objects	75
6 Conclusion	86
6.1 Summary	86
6.2 Further Research	87
References	89

List of Tables

Table		Page
4.1	Range error at 493 meters.	43
4.2	Angle estimate.	53
4.3	Estimate of slot width.	56
5.1	Error of points from angled surface fitted to a plane.	73

List of Figures

Figure	Page
2.1 Point cloud of a car with trees in the background.	8
2.2 Example of walk error due to pulse amplitude.	10
2.3 Digitized ladar waveform.	11
2.4 Digitized ladar waveform from a surface angled at 75°	12
2.5 The VISSTA van.	13
2.6 ELT transmitted pulse.	16
2.7 Spectrum of a received waveform from the ELT.	16
2.8 Main user interface of LadarSIM.	18
3.1 Data acquisition model for the received waveform.	19
3.2 Processing method used in the VISSTA ELT.	20
3.3 Waveform of a shot at a sufficient range to show high noise levels.	22
3.4 Plot of the power spectrum of the received signal and the low-pass filter.	23
3.5 Waveform of interpolated pulse.	24
3.6 The process of interpolation.	25
3.7 Plot of the impulse response $\tilde{h}(t)$	28
4.1 The footprint covering a discontinuous surface that gives two returns.	33
4.2 Separation of two Gaussian pulses.	35
4.3 Planar surface at an angle of ϕ	35
4.4 Convolution of two signals with a system in the time domain.	38
4.5 Difference between the convolution between the system and each signal.	39
4.6 Convolution of two signals with a system in the frequency domain.	40

4.7	Under-regularized surface response resulting from hitting a planar surface at 75°	42
4.8	Simulation measurements from returns hitting two surfaces as the distance of separation between the surfaces varied.	45
4.9	The target used for testing in these experiments.	47
4.10	ELT measurements from returns hitting two surfaces as the distance of separation between the surfaces varied.	48
4.11	The ELT received waveform processed with the different deconvolution methods.	49
4.12	Simulation measurements from returns hitting two surfaces at 20 cm as the SNR varied.	49
4.13	ELT measurements from returns hitting two surfaces at 20 cm as the SNR varied.	51
4.14	Side profile of spatial beam footprint showing areas and boundaries of peaks.	52
4.15	Relation of peaks to the spatial locations within the footprint.	54
4.16	Test board used in experiments.	55
5.1	Incorrect point estimation when the center of the footprint is assumed as the location of points.	58
5.2	Side view depiction of an interpolated discontinuous surface that contains a multiple return.	60
5.3	Choosing the orientation of the surface to promote smooth surfaces.	62
5.4	Nearest neighbor selection.	65
5.5	Pixel spacing of interpolated points.	65
5.6	Shot on planar surface before adjusting the locations of the points.	69
5.7	Shot on planar surface after solving for ϕ	69
5.8	Shot on planar surface after rotating by θ	70
5.9	Side view of ELT point cloud of slotted surface.	71
5.10	Side view of ELT point cloud of slotted surface with a practical shot spacing. The dotted line shows the true surface. The solid line shows the surface estimate by interpolation.	72

5.11 Top view of ELT point cloud of angled surface.	74
5.12 Estimate of θ for $\phi = 60^\circ$	76
5.13 Estimate of θ for $\phi = 75^\circ$	77
5.14 Testboard scanned by the ELT.	79
5.15 Strap on the testboard. First returns are shown in red. Other returns are shown in blue.	80
5.16 Cube scanned by the ELT.	81
5.17 Top view of cube scanned by the ELT.	82
5.18 Scan of shed with an I-beam.	83
5.19 I-beam scanned by the ELT.	84
5.20 ELT waveform with a low detection threshold.	85
5.21 Surface reponse of ELT waveform with a low detection threshold.	85

Acronyms

APD	Avalanche Photo-Diode
ATR	Automatic Target Recognition
BRDF	Bi-Directional Reflectance Distribution Function
ELT	Eye-Safe Ladar Testbed
EM	Expectation Maximization
FPA	Focal Plane Array
FWHM	Full Width at Half Maximum
lidar	Light Detection and Ranging
ladar	Laser Detection and Ranging
LM	Levenburg-Marquardt
NNLS	Non-Negative Least Squares
RJMCMC	Reversible Jump Markov Chain Monte Carlo
RL	Richardson-Lucy
SNR	Signal-to-Noise Ratio
SQNR	Signal-to-Quantization Noise Ratio
VISSTA	Vehicle Integrated Sensor Suite for Targeting Applications

Chapter 1

Introduction

Automatic target recognition (ATR) is a machine vision task that uses sensors to collect data and attempts to identify targets of interest. To increase the ability to accurately detect and identify targets, multiple imaging sensors may be used. One of these sensors is 3-D imaging lidar. Due to its short optical wavelength, lidar is capable of producing high resolution images of the target compared to radar. This is advantageous because the number of shots on the target improves the fidelity of ATR. Although lidar can provide high spatial resolution images, there is a need for improvement with regards to the range resolution.

Range resolution in a lidar system is the ability to detect two surfaces that lie within the field of view (FOV) of a single detector. This is limited by the temporal pulse width of the laser. For example, if the laser beam hits two surfaces that are separated by more than the temporal pulse width, the receiver can detect multiple returns in the waveform. However, if two surfaces are closer than the pulse width, the receiver detects only one return.

Lidar receivers have been built that digitize the intensity of returning light. A digital waveform can be used to improve the range resolution of the lidar. Signal processing techniques can be applied to the digitized waveform to obtain information about the shape of the targeted surface within the detector FOV. Knowing the shape of the surface allows for interpolating shots on the target for increased resolution in both range and spatial dimensions.

This thesis explores signal processing techniques that can be applied to improve the range resolution associated with a given lidar waveform. Much research has been conducted for the application of terrain mapping from airborne sensors; however, the literature does not contain much research on applying signal processing to digitized lidar waveforms to

improve the resolution of complex, man-made and natural targets. This thesis focuses on representing surfaces from these types of targets.

1.1 Prior Work

1.1.1 Surface Estimation

A common approach to estimate the shape of the surface is by means of modeling the waveform as the sum of individual pulses [1]. Several authors use a Gaussian pulse profile [1–5]. The assumption of a Gaussian pulse may not always be accurate. Another pulse profile used by Chauve et al. is a generalized Gaussian [6]. Once the pulse profile that best fits the data has been determined by the user, the signal is decomposed into individual returns. The most common methods for decomposition are the expectation maximization (EM) algorithm and the Levenburg-Marquardt (LM) algorithm [7].

After decomposing the received waveform into individual returns, features are extracted from the parameters of the fitted profiles. The simplest interpretation of a feature is that the range to the surface is determined by the mean of the pulse. Another interpretation used by Neuenschwander et al. in the application of terrain mapping is that the width of the first return can be interpreted as the width of the tree canopy [5]. Other application-specific interpretations of features exist and are usually left to the end user of the data to interpret.

A drawback of the above method is that the number of pulses contained in the estimated waveform must be known *a priori*. Jung and Crawford use a greedy EM algorithm proposed by Vlassis and Likas to overcome this drawback [2, 8]. This approach begins by assuming that the return is made up of one pulse. It then increases the number of assumed pulses in the decomposed waveform until the residual in the estimate begins to increase. Thus, the number of pulses assumed in the waveform is the number that leads to the smallest error.

Hernandez et al. use the reversible jump Markov chain Monte Carlo (RJMCMC) technique for pulse decomposition [9]. This work does not use the generalized Gaussian pulse, but rather uses a piecewise defined function that models the pulse in their system. A more

complex approach to surface classification that uses a marked point process and the RJMCMC algorithm is given by Mallet et al. [10]. This work argues that the lidar return is not accurately modeled by the generalized Gaussian pulse shape or any other single shape. Therefore, the waveform is modeled with a library of functions from different surfaces. The closest fit indicates the surface that was most likely hit.

A different method for surface estimation is deconvolution. The received waveform can be modeled as the response of the system convolved with the surface response of the surface. Deconvolution is an operation that obtains the surface response given the received waveform and the system response.

For deconvolution, Jutzi and Stilla applied the Wiener filter to lidar pulses [11]. The Wiener filter is an optimal linear filter used to minimize the mean squared error between the expected surface response and the estimated surface response [12]. Using the Wiener filter, Jutzi and Stilla showed that it is possible to detect two surfaces that differ in range by 15 cm with a 1 GHz bandwidth receiver that samples a 5 ns pulse at 20 GHz [11]. This experiment used a specular surface at a range of 100 m. The conclusion of this work is that under ideal conditions, i.e. high sample rate and high signal-to-noise ratio (SNR), obtaining the surface response significantly improves the ability to detect two offset surfaces.

Harsdorf and Reuter compared the Wiener filter with two other methods on a simulated lidar waveform [13]. One of these methods is the Richardson-Lucy (RL) algorithm. This method is an iterative technique and has the constraint that the surface response must be positive [14]. This constraint is applicable to the lidar waveform because light returning to the receiver cannot be negative. The other method compared is the non-negative least squares (NNLS) algorithm. This method computes the least squares solution for the surface response with the constraint that the surface response is positive.

Harsdorf and Reuter's work looked at two possible surface responses—a delta function and a square pulse. They concluded that the NNLS algorithm works very well at recovering the delta function surface response, but not at recovering the square pulse; however, the RL algorithm can recover both surface response shapes well with 100 iterations. The Wiener

filter method did not perform as well as the other two methods.

1.1.2 Spatial Relationship

Each lidar shot is independent one from another; however, spatial correlation exists and can be used to increase the quality of the point cloud. One way that this correlation has been used is to increase the SNR by means of averaging waveforms that are spatially proximate. In terrain mapping, this has been used by Magruder and Neuenschwander to detect faint returns along the ground [15]. In an urban environment, the same was performed by Stilla et al. on roof tops [16]. This method is a trade off between SNR and spatial resolution. Nonetheless, in areas where shots are highly correlated, it is an effective means not only of extracting more information out of the signal for detection, but also of obtaining a better range estimate of the surface due to the increase in SNR.

The matched filter is commonly used in pulse detection. It correlates the received waveform with the transmitted waveform to increase the SNR. An obstacle to using the matched filter with lidar waveforms is that an estimate of the surface response must be known. To overcome this obstacle, Kirchhof et al. use surrounding points in the point cloud to estimate the surface response [17]. This is done in an iterative procedure that begins by estimating the surface response to be a delta function. After the initial point cloud is created, the points are used to update the surface response. At each iteration, the new surface response is used. This procedure is continued until the rate of improvement slows.

In urban mapping, Jutzi et al. use shots along a roof top to obtain better spatial resolution along the edge of the roof [18]. When the beam footprint lands partially on the roof and partially on the ground, the energy in each return can be computed. Assuming that the ground and roof reflectivity is the same, the distance between the edge of the roof and the middle of the footprint can be computed by integration over the footprint. They report an increase of sub-pixel resolution by a factor of 10.

1.2 Problem Statement

The quality of the 3-D ladar image, also known as a point cloud, is influenced by the spatial and range resolutions. This is especially important on edges of targets where the surface is complex. Without signal processing, the complex surface may be blurred causing edge details to be lost.

The problem to be resolved by this thesis is how to apply signal processing methods to the digitized ladar waveform and extract information about the shape of the surface seen by a single detector. However, the orientation of the surface cannot be determined using the waveform from a single detector. Therefore, a second problem to be resolved is how to use neighboring points in the point cloud to determine the orientation of the surface. Solving these problems allows for improved resolution in the point cloud by interpolating the estimated surface.

1.3 Overview

The method used to solve the problem is done in two parts. First, deconvolution is used on the waveform to increase the range resolution. Then, neighboring points in the point cloud are used to find the orientation of the surface.

Background information on ladar and the instruments used for data collection are given in Chapter 2. Chapter 3 describes the signal processing required to obtain the surface response from the output of the Eye-safe Ladar Testbed (ELT) developed by the Navy. This is a state-of-the-art waveform digitizing ladar. Chapter 4 derives and tests the surface response for two types of surfaces of interest. These are two discontinuous surfaces within the footprint that are separated by less than the pulse width in range and have angled surfaces. Chapter 5 describes the method used to incorporate information from proximate points. This solves the problem of not knowing the orientation of the surface from the surface response. It also obtains a better estimate of the spatial location of each return. Conclusions are given in Chapter 6.

1.4 Thesis Contribution

A list of contributions from this thesis are given below.

- A signal processing model is presented for a state-of-the-art, waveform digitizing lidar.
- Several existing issues for lidar waveform processing are addressed including the matched filter, interpolation, and deconvolution.
- The three deconvolution methods studied by Harsdorf and Reuter on simulated data [13], were applied for real data from a lidar that digitizes the waveform at a practical sample rate.
- The surface response is derived for two important surface types.
- A new method for treating the surface response of angled surfaces is given.
- A new method is presented that utilizes information from neighboring points to determine the orientation of the estimated surface.
- The presented methods result in a technique to obtain enhanced resolution in a lidar point cloud.

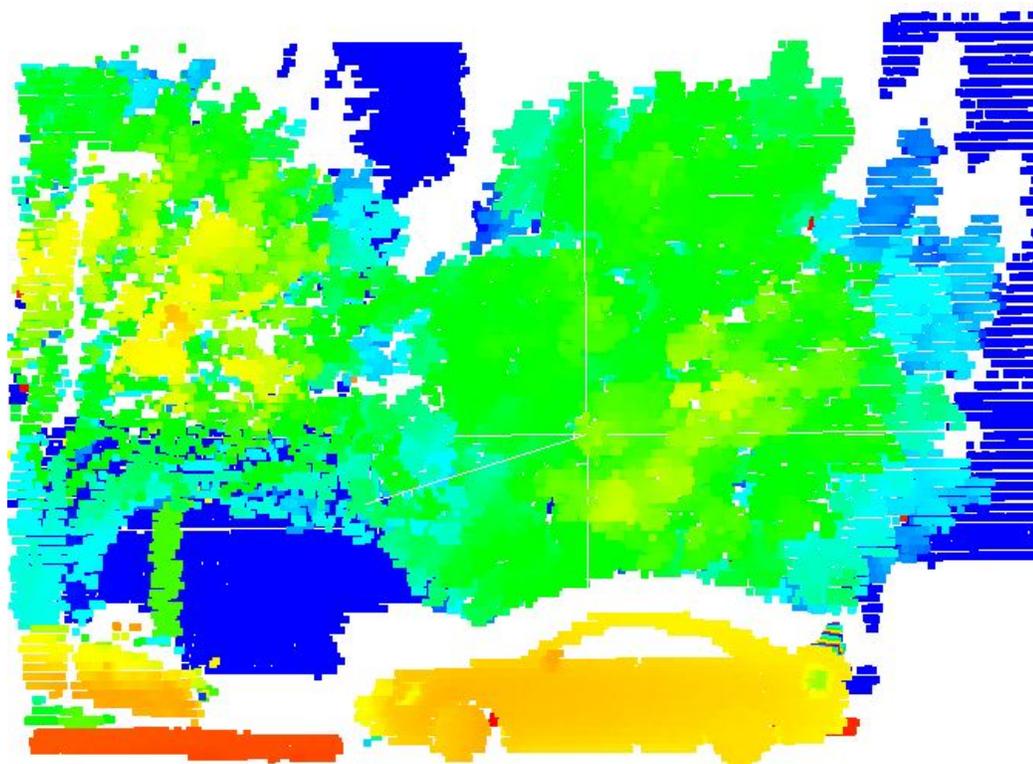
Chapter 2

Background

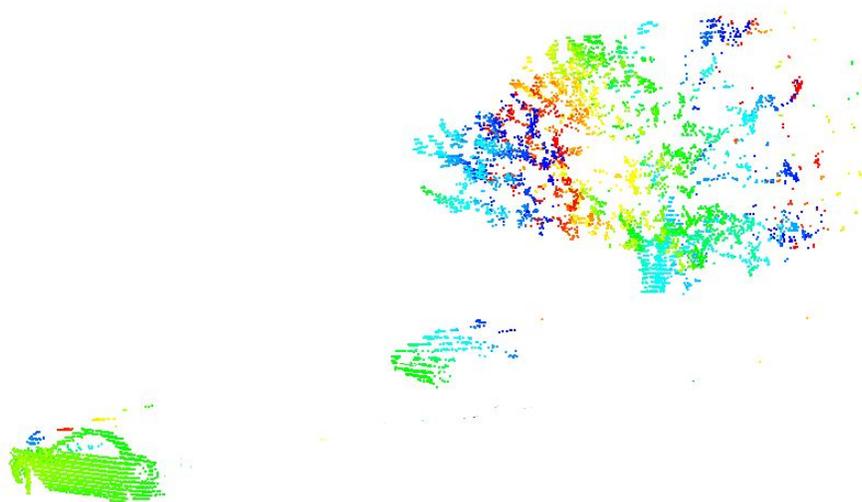
2.1 Three-dimensional Imaging Ladar (Laser Detection and Ranging)

Three-dimensional imaging ladar is a technology commonly used for terrain mapping, object recognition, and autonomous navigation. The term lidar (light detection and ranging) refers to the same basic technology, and is used almost interchangeably with ladar in the literature. Images obtained by ladar sensors contain information about the range from the sensor to the target and are typically called point clouds. Figure 2.1 shows a point cloud taken with the ladar sensor used in this thesis. Figure 2.1(a) shows a car in front of two large trees and a building at the back. With red being the closest and blue being the farthest, colors show ranges from the sensor to the target. Although both trees in Fig. 2.1(a) are marked with green, the tree on the left is closest. To show more detail, colors cycle five times for this image. Figure 2.1(b) demonstrates the capability of ladar point clouds to be rotated to the side to show the 3-D nature of the data. This shows the front part of the scene which includes the car and the tree on the left of Fig. 2.1(a). Notice that only the front side of the car can be displayed whereas the tree shows full 3-D detail. With only one scan, there is no way of knowing what is behind a solid target, thus creating a shadow. Some applications require that the shadow be removed by fusing multiple scans from different views. Alternatively, ladar can be used to create a range image. A range image contains the same information as a point cloud, but differs in the way it is displayed. Rather than displaying a 3-D space filled with points, it appears as a 2-D image where each pixel contains the range to the target.

Ladar is an active form of sensing that requires transmission of a laser pulse. For this reason, it is not currently possible to have large arrays comparable to digital cameras and



(a) Front view of the entire scene.



(b) Side view of the car and the first tree in the scene.

Fig. 2.1: Point cloud of a car with trees in the background.

preserve the ability to light each element with sufficient optical power. Common sensor arrays vary from as little as one element to at most tens of thousands of elements. To compensate for small arrays, a mechanical scanner is often used to change the array's line of sight. Scanning is frequently accomplished using low mass mirrors. Knowing the scanner's position and the range to the target for thousands of points allows the point cloud to be created.

The range from the sensor to the target can be found using different types of sensors; however, this thesis focuses on pulse time of flight lidar. The pulse time of flight principle is based on firing a short laser pulse (on the order of a few nanoseconds) toward a target and recording the time of flight. Knowing the speed of light and the time of travel, the range from the sensor to the target is found by

$$R = \frac{c}{2}t, \quad (2.1)$$

where c is the speed of light and t is the round trip time for the transmitted pulse to return to the sensor.

Different techniques exist for discriminating the start and stop time of the pulse. The simplest of these techniques uses a threshold. When the intensity of the light seen at the receiver exceeds the threshold, a detection is declared. A problem with this method is that it is susceptible to walk error. Walk error is defined as error in the range measurement due to a change in pulse amplitude [19]. Pulse amplitude is highly dependent on the reflectivity and distance of the target. If a target contains dark and bright areas, the amplitude of the returns over the dark area would be lower than the amplitude of the returns over the bright area due to the reflectivity. Example pulses detected using the leading edge discriminator are given in Fig. 2.2. Although, the returns have the same mean, the larger return which came from the bright area, is estimated to be closer than the smaller return.

A more accurate discrimination method uses a high-pass filter as a differentiator [19]. The detection is declared when two conditions are met—the differentiated signal crosses zero and the pulse is above a threshold. This method is commonly known as the crossover

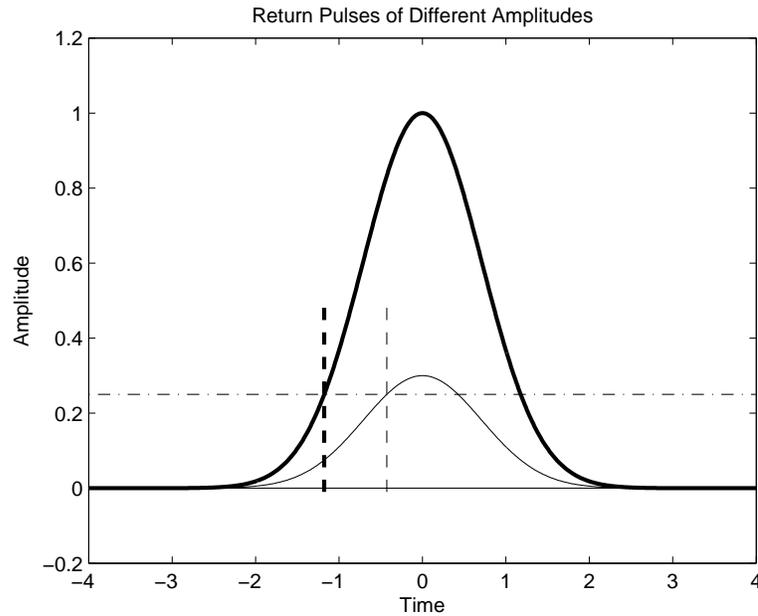


Fig. 2.2: Example of walk error due to pulse amplitude.

method. Because this method uses a differentiator, the timing point does not depend on the amplitude of the pulse.

Another approach to discriminate a pulse is the constant fraction method [19]. The constant fraction method is implemented in hardware by creating a scaled, delayed copy of the received signal. The difference between the original and the delayed signal is computed. When the difference is equal to zero, a threshold is checked. If the signal is above the threshold, a detection is declared at the crossover point. Like the derivative crossover method, the timing point does not depend on the amplitude of the pulse.

In the past, lidar systems have been subject to analog methods of detection including the methods just discussed. Since 2004, however, commercial systems have been built that digitize and store the intensity of the light returning to the sensor [7]. This allows digital signal processing techniques to be performed on the waveform to increase the quality of the point cloud. Processing the waveform can be done in real-time, or can be post-processed for implementation of more complex algorithms. Figure 2.3 shows a digitized waveform from the lidar used in this thesis.

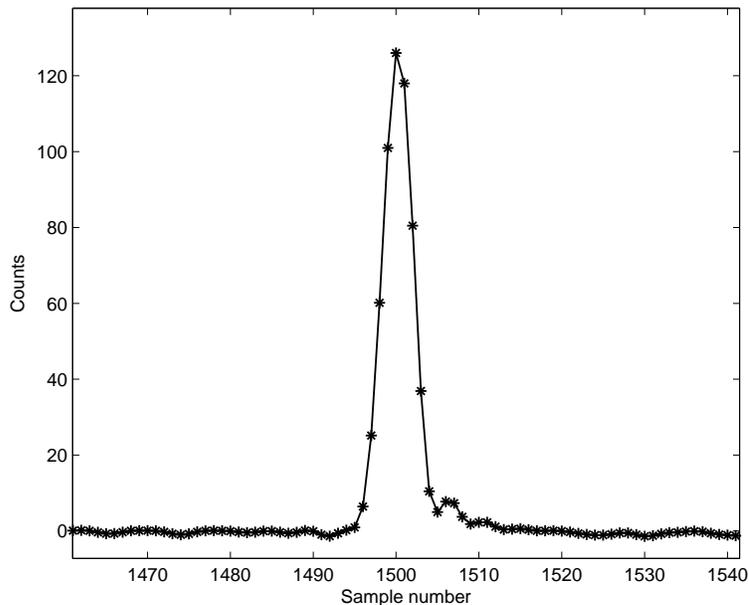


Fig. 2.3: Digitized ladar waveform.

An advantage of waveform processing comes from the fact that the laser beam diverges as it propagates, creating a beam footprint. Due to the divergence, the beam footprint may fall onto more than one target. This causes multiple returns in the ladar waveform. Analog systems typically report the first or last return. However, waveform digitizing systems store the information to find each return in the waveform. Evidence of this can be seen in the trees in Fig. 2.1(b) where returns were detected on multiple leaves from each ladar shot. Not only can waveform digitizing ladar systems detect multiple returns better than analog systems, but also can provide more information about the surface that was hit. For example, when the laser hits a surface at an angle, the light does not return to the receiver all at the same time. Instead of showing a delayed copy of the transmitted pulse, the digitized waveform shows a blurred version of that pulse [20]. A pulse that hit a surface at 75° is shown in Fig. 2.4 and can be compared to the pulse in Fig. 2.3 which was obtained from hitting a surface straight on.

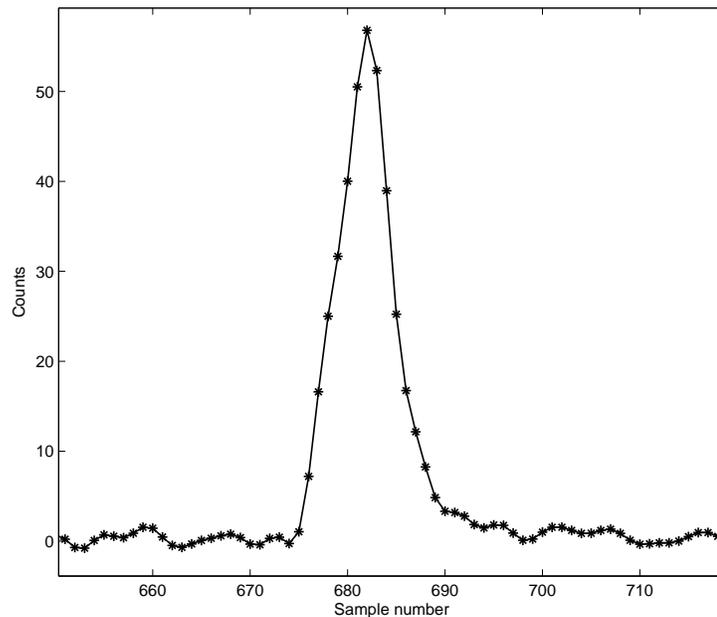


Fig. 2.4: Digitized lidar waveform from a surface angled at 75° .

2.2 VISSTA ELT

The Vehicle Integrated Sensor Suite for Targeting Applications Eye-safe Ladar Testbed (VISSTA ELT), shown in Fig. 2.5, is the lidar system that was used for collecting data for this thesis [21]. The VISSTA van was developed as an experimental lidar to provide a mobile sensor platform designed for research. A target area of research is in the understanding and exploitation of digitized lidar waveforms. The van contains the ELT and an IR camera housed in a movable turret. Inside the cabin, multiple computers run specific components of the system. These components include the fast scanning mirror, pointing controller, visible camera, data recorder, and host computer.

The VISSTA ELT contains an eye-safe wavelength laser ($1.5 \mu\text{m}$) and a visible wavelength camera co-boresighted to observe the same field of view. This allows the operator to accurately point the lidar as well as capture visible images and lidar point clouds simultaneously. The lidar and visible image can be used to create 3-D texel images which contain range data fused with textured data from a visible camera [22].



(a) The ELT aperture can be seen at the right side in the turret behind the cab.



(b) Displays used for system control and display of sensor data.

Fig. 2.5: The VISSTA van.

The receiver in the ELT includes a single sensor avalanche photo-diode (APD). This requires a complex scanner capable of scanning in both the vertical and horizontal directions in order to form point clouds. The scanner utilizes a fast scanning mirror to produce various scanning patterns. The most common pattern is a square wave in the horizontal (azimuth) direction while constantly zig-zagging in the vertical (elevation) direction. The scanner can also reverse the vertical and horizontal scan roles, be fixed with no motion, or create a mosaic of scans. Shots are generally collected at 20 kHz, but can be collected at up to 100 kHz. Because the van is experimental, there is still work needed to improve the scanning patterns. Point clouds created by the ELT contain high levels of noise in the scanning measurement relative to the noise in the range measurement.

The ELT is able to produce a point cloud in real-time as well as digitize and store lidar waveform data for off-line processing. It samples data at 2 GHz, or equivalently, every .5 ns providing digital waveforms. According to (2.1), this corresponds to 7.5 cm in range. Due to the sampling, the real-time system quantizes the range into 7.5 cm bins. The ELT is capable of producing point clouds with range error on the order of millimeters; thus, binning at 7.5 cm compromises the quality of the point cloud without post processing interpolation [21]. At this sample rate, the ELT can sustain data rates high enough to capture up to 500 m of range data per shot at the 20 kHz shot rate.

The waveform data are quantized to signed 8-bit integers after being amplified with linear amplifiers. The ELT has an adjustable gain setting that is set by the user. If the target of interest is either far from the sensor or has a low reflectivity, the gain can be easily increased so that the signal spans the range of the digitizer. Adjusting the gain of the receiver to the amplitude of the pulse takes advantage of the dynamic range and increases the signal-to-quantization noise ratio (SQNR). This is effective when surfaces of similar reflectivities are targeted, but produces a lower SQNR when the intensities of the returns vary significantly. The quantization noise, σ_q^2 , is $\frac{1}{12}$ counts, and the maximum signal amplitude, A_{max} , is 127 counts for the ELT pulse [23]. If the SQNR is defined as

$$SQNR = \frac{A}{\sigma_q}, \quad (2.2)$$

then by plugging in the numbers, $SQNR = 3.5A$ and $SQNR_{max} = 440$ counts. Noise from the electronics is determined by the gain settings in the receiver and typically dominates the quantization noise.

The ELT has a pulse width of 1.5 ns defined by the full width at half maximum (FWHM) and an example pulse is shown in Fig. 2.6. This pulse was measured by the ELT laser manufacturer. The bandwidth of the receiver that processes the pulse is estimated to be 500 MHz. This bandwidth is given as the width of the spectrum for which the SNR is greater than one. The power spectrum of a pulse after passing through the receiver electronics is shown in Fig. 2.7.

2.3 LadarSIM

LadarSIM is another tool utilized for the experiments in this thesis. LadarSIM is a Matlab-based ladar simulation that has been shown to accurately model the performance of the VISSTA ELT [21]. Aimed at the engineering design process, LadarSIM was designed to provide an interface that allows the user to change parameters to the setup and quickly view the outcome. Simulation fidelity options are included in the simulation so that the user can decide the trade-off between simulation time and accuracy. To maintain a user friendly interface, the Matlab GUI shown in Fig. 2.8 is used to configure the simulation.

The entire data acquisition process is modeled from the hardware setup to the movement of targets in the scene. The user first describes the ladar to be simulated. The main hardware components are the focal plane array (FPA), scanner, transmitter, optics, and receiver electronics. Then the user places the ladar into an environment and describes the path and sensor pointing criteria of the sensor. The environment may contain multiple targets that may contain motion during data acquisition.

The simulation occurs in two stages. The first stage focuses on the geometry of the setup. The sensor and targets are able to move and scanner error enters in, but radiometric

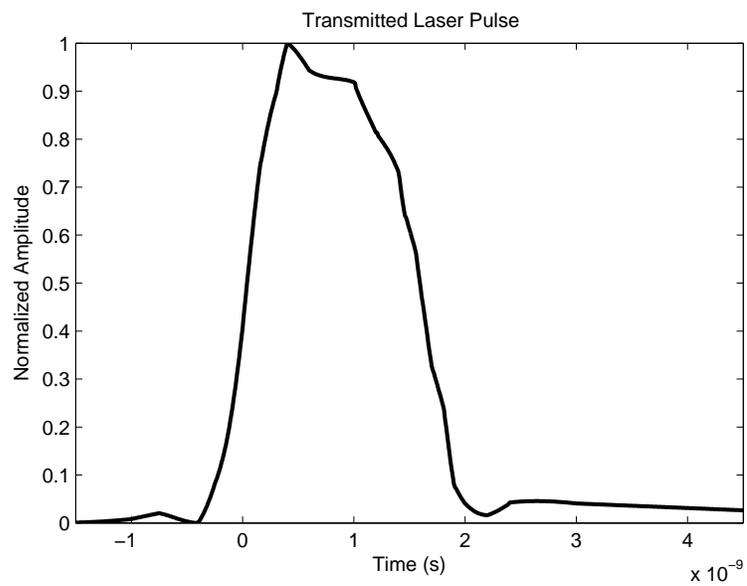


Fig. 2.6: ELT transmitted pulse.

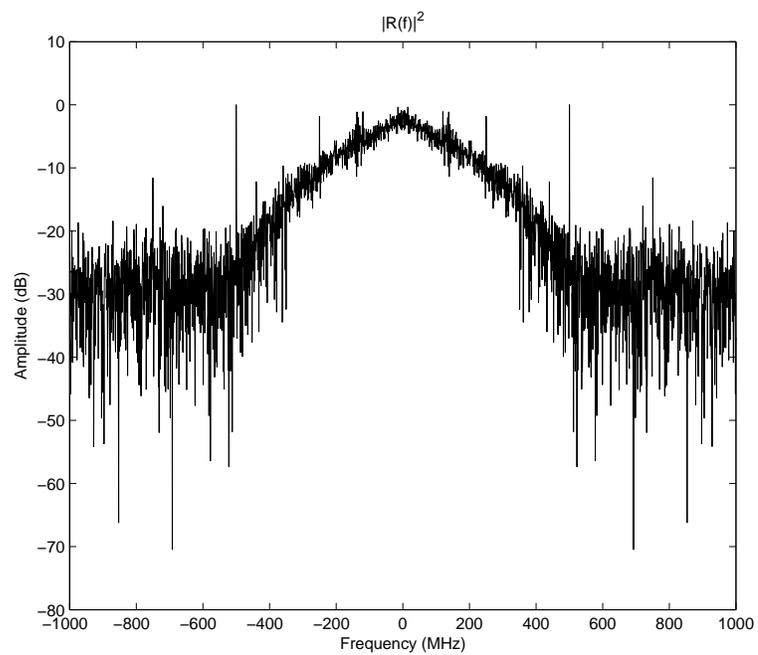


Fig. 2.7: Spectrum of a received waveform from the ELT.

errors are not computed at this point in the simulation. This provides the user with the true range to the targets. Once a geometry scan is completed, the user begins a radiometry scan. The radiometry scan takes the geometry scan output and factors in the radiometric errors. Radiometric errors come from noise sources such as solar background noise, backscatter noise, shot noise, and noise in the electronics.

LadarSIM effectively models complex surfaces when the user chooses to perform waveform processing [24]. The ladar footprint is sampled by a specified number of beamlets. This allows for sub-detector modeling of the surface. The true range depends on the surface and is found at the location of each beamlet. Next, a waveform is created for each beamlet with the corresponding range. The optical waveform seen at the detector is then modeled as the superposition of the waveforms of the beamlets. Finally, the optical waveform is filtered by the electronics and the signal is processed for range discrimination.

After a noiseless signal has been processed, errors are added to the simulation. These errors are computed statistically from the noise in the system. Three sources of error due to signal discrimination are dropouts, false alarms, and range error [25]. Dropouts are points that should have been on a surface but are not detected. False alarms are points that did not come from a surface, but were detected as such due to noise. Range error is measurement error on a surface. It is defined as the standard deviation of the range measurements. After errors are added, the point cloud is created.

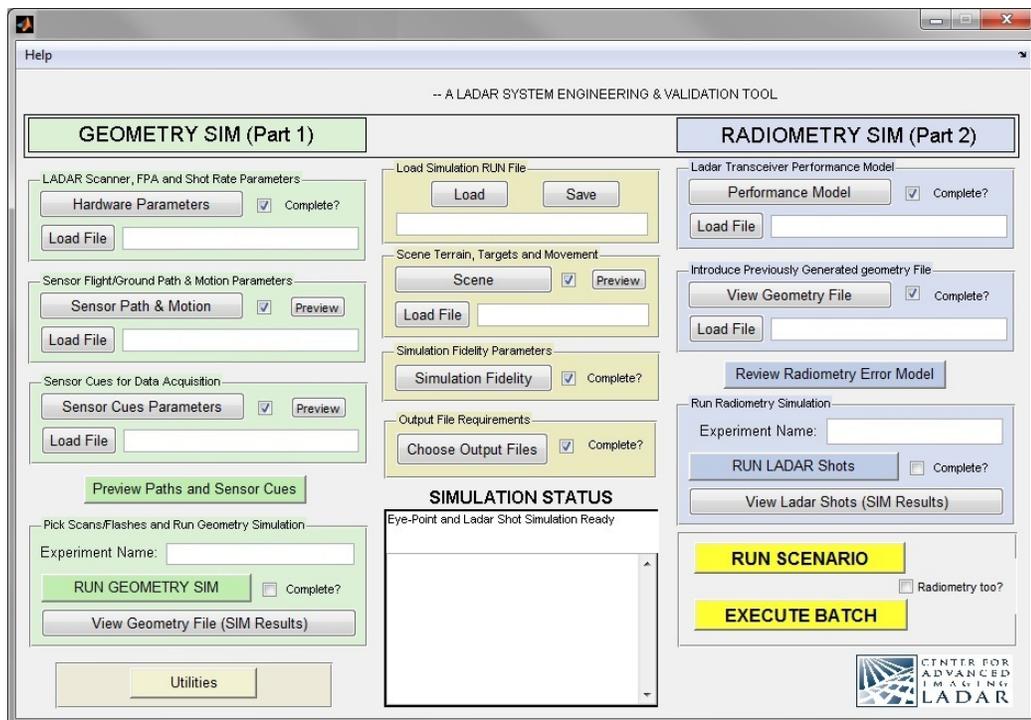


Fig. 2.8: Main user interface of LadarSIM.

Chapter 3

Signal Processing to Obtain the Surface Response

3.1 Data Processing Model

The data acquisition model used for signal processing the ladar waveform is given in Fig. 3.1.

The input to the system is the Dirac delta function $\delta(t)$, representing an ideal transmitted pulse shape. The input is then passed through the laser pulse transfer function $P(s)$ to create a transmitted laser pulse waveform. An estimate of $P(s)$ for the ELT is shown in Fig. 2.6. The pulse then travels into the environment where it interacts with the target's surface. This is the convolution of the transmitted pulse waveform with the target's surface response, $S(s)$. The surface response will be derived in Chapter 4 for the surfaces studied in this thesis. Next, backscatter, solar background noise, dark current noise, and noise in the receiver electronics are lumped together as additive white noise $n_t(t)$. Amplitude dependent (multiplicative) shot noise from the avalanche photo diode (APD) is added at the input of the electronics. In the design of the VISSTA ELT, separate electronics are used to detect the transmitted pulse as it exits the sensor and the return pulse from the target; however, they have the same design. Different electronics transfer functions, $H(s)$, are used

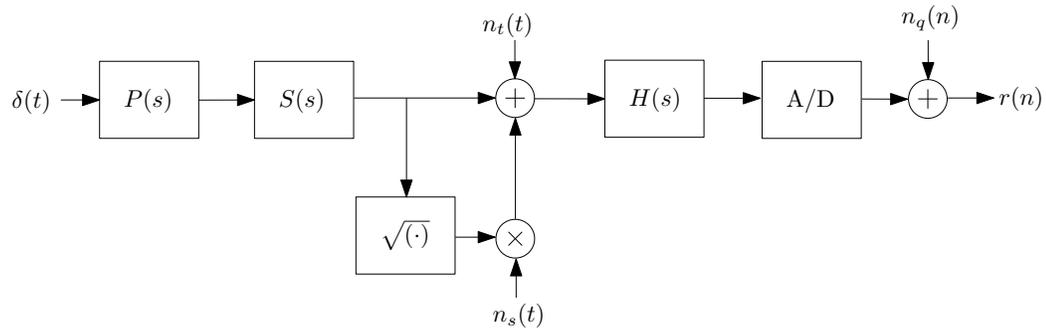


Fig. 3.1: Data acquisition model for the received waveform.

in modeling the digitized exiting and returning pulses due to slight mismatches between the received and transmitted pulse electronics. The transmitted pulse electronics have a slightly lower bandwidth than the received pulse electronics. Finally, the signal is quantized where quantization noise $n_q(n)$ in the A/D converter is added to obtain the sampled waveform $r(n)$.

Due to the difficulty in obtaining an estimate of the signal-dependent shot noise, the model was simplified by assuming it can be modeled as additive white noise with variance proportional to the average value of the pulse waveform. This leads to the frequency-domain sampled signal

$$R(f) = P(f)S(f)H(f) + N(f)H(f) + N_q(f), \quad (3.1)$$

where

$$N(f) = N_t(f) + \sqrt{\overline{p_s(t)}}N_s(f), \quad (3.2)$$

$$p_s(t) = p(t) * s(t), \quad (3.3)$$

$\overline{p_s(t)}$ is the time average of $p_s(t)$ during the duration of the pulse, f is the normalized discrete frequency, $N_t(f)$ and $N_s(f)$ are white noise, and $p(t)$ and $s(t)$ are the inverse Laplace transforms of $P(s)$ and $S(s)$, respectively.

The goal of processing the data described in this chapter is to recover the discrete-time surface response $s(n)$ from the received waveform $r(n)$ and to use it for signal discrimination. For the ELT waveform, this is achieved by filtering, interpolation, deconvolution, and range discrimination as shown in Fig. 3.2.

3.2 Low-Pass (Matched) Filtering

To reduce unnecessary computation, it is desirable to be able to detect potential return

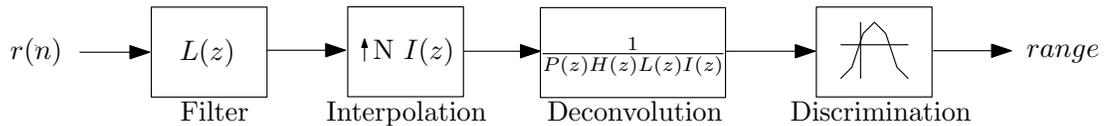


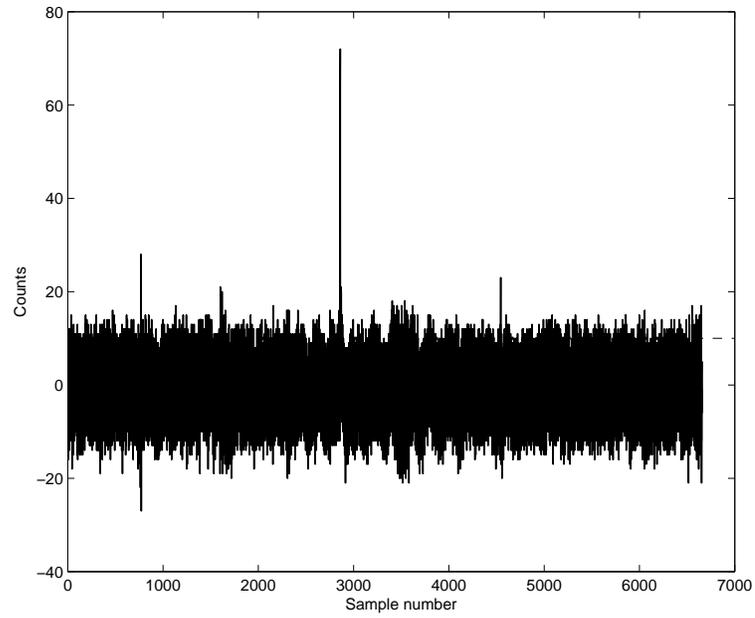
Fig. 3.2: Processing method used in the VISSTA ELT.

pulses and deconvolve only those parts of the signal. The ability to detect pulses is improved if the signal is filtered by a matched filter $L(z)$ that reduces false alarms and improves the SNR. As an approximation to a matched filter, a low-pass filter with a band edge located at the highest frequency in the transmitted pulse spectrum can be used. A true matched filter is not used because a matched filter assumes that the shape of the return is known. Due to the surface response, the received signal takes on many different possible forms which are not known until after the signal is received. It is significantly easier to detect return pulses in the resulting low passed signal. An example of the improvement is illustrated in Fig. 3.3. The data used to create this waveform were collected at a significant range to show high noise levels. Contrasting the figures shows a significant decrease in the noise floor of the signal. In this case, if a threshold level is set at ten counts, the number of false returns in this waveform would be reduced from hundreds to under five.

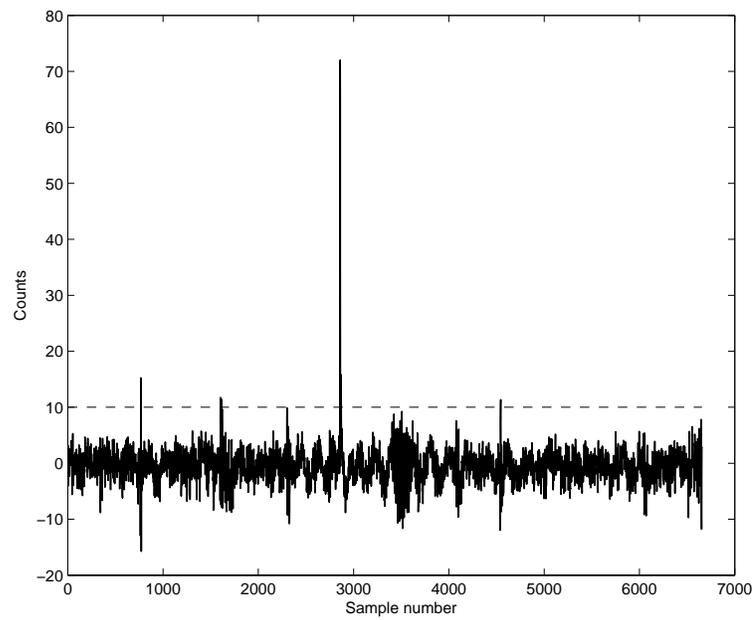
Since the pulse spectrum only extends to half of the fold over frequency, a significant amount of the white noise signal $N(f)$ is in the stop band of the matched filter and is therefore reduced by the filter. Figure 3.4 shows the magnitude response of $L(f)$ and the power spectrum of $R(f)$.

3.3 Interpolation

The ELT digitizes data at 2 GHz which corresponds to 7.5 cm/sample in range. Relative to the desired measurement accuracy, 7.5 cm/sample is very coarse and requires interpolation to a higher sample rate. The two stages in the processing that require interpolation are deconvolution and signal discrimination. Deconvolution or deblurring amplifies high frequency content in the waveform which is limited to Nyquist. Thus, increasing the sample rate by means of interpolation, allows a better representation of a delta function. If it were not for interpolation, range estimates would be binned at 7.5 cm. Sub-centimeter range error is only obtainable if the signal is interpolated. An example of a pulse interpolated by a factor of $N = 4$ (which is the interpolation factor used for the experiments in this thesis) is given in Fig. 3.5. This interpolation factor is adequate for processing through the deconvolution stage. However, for detection, it still falls short of sub-centimeter range



(a) The waveform before the matched filter.



(b) The waveform after the matched filter.

Fig. 3.3: Waveform of a shot at a sufficient range to show high noise levels.

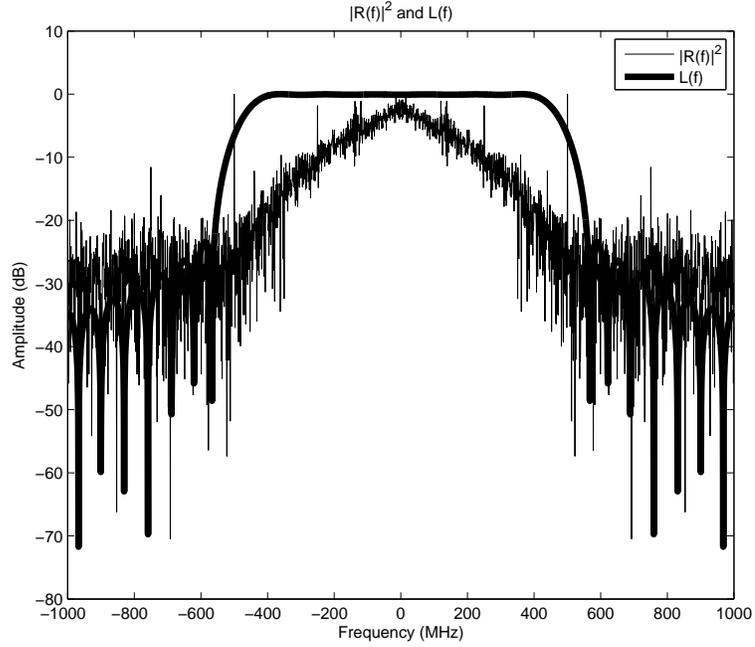
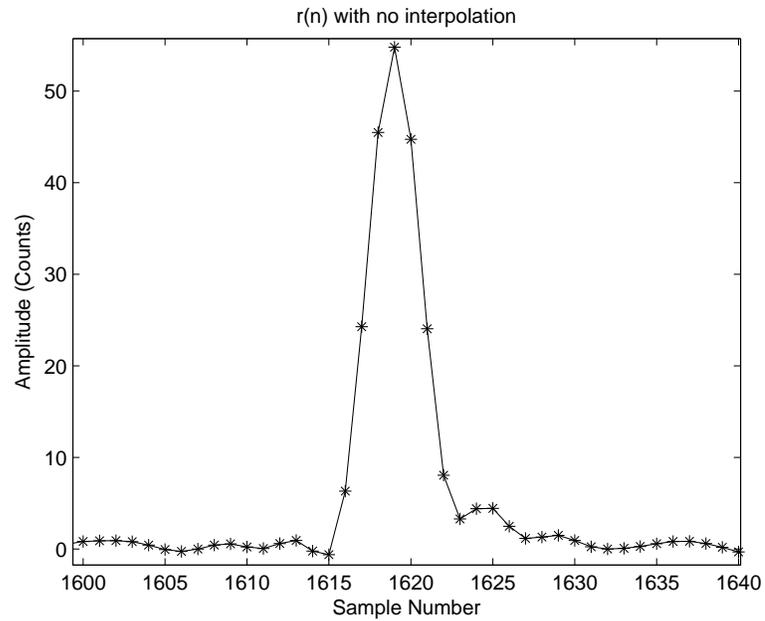


Fig. 3.4: Plot of the power spectrum of the received signal and the low-pass filter.

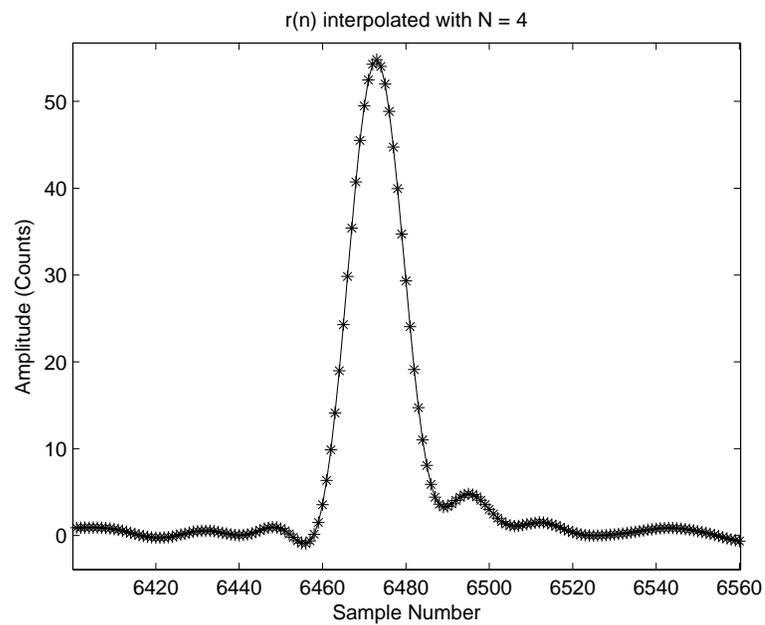
error; thus, linear interpolation is applied in the discrimination stage to reduce range error at very little computational cost.

Figure 3.6 shows the process of interpolation. Interpolation by a factor of N is performed by first upsampling the signal, i.e. inserting $N-1$ zeros between every sample in the signal. It can be shown that upsampling a signal causes the signal's spectrum to be compacted by a factor of N such that a frequency f_0 maps to $\frac{f_0}{N}$ as shown in Fig. 3.6(d) [23]. An artifact of the frequency compaction is that spectral copies of the original pulse are centered at $f = \frac{1}{N}, \frac{2}{N}, \frac{3}{N} \dots$ in normalized frequency. Next, the signal of interest is obtained by filtering the signal with a linear phase low-pass filter, $I(z)$, that has a bandwidth of $\frac{1}{2N}$ (normalized frequency). This eliminates the spectral copies and leaves only the interpolated signal with the power spectrum shown in Fig. 3.6(f).

It may be argued that the low-pass filter $L(z)$ is unnecessary due to the low-pass filter $I(z)$ used for interpolation. However, it is advantageous to use both filters to reduce computation time. Applying a filter $L(z)$ of order k to the uninterpolated signal is equivalent

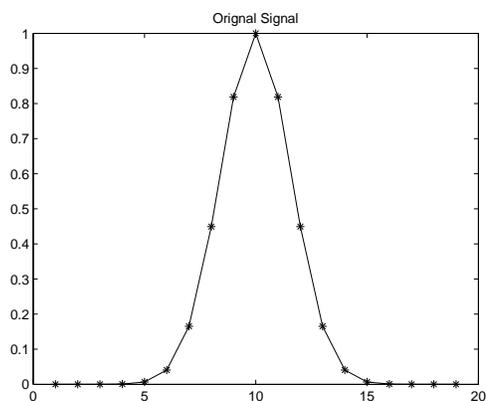


(a) No interpolation.

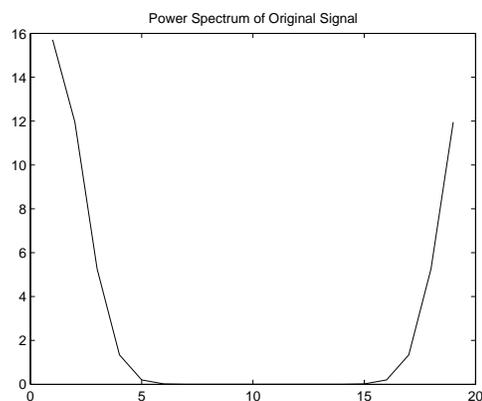


(b) Interpolation by a factor of 4.

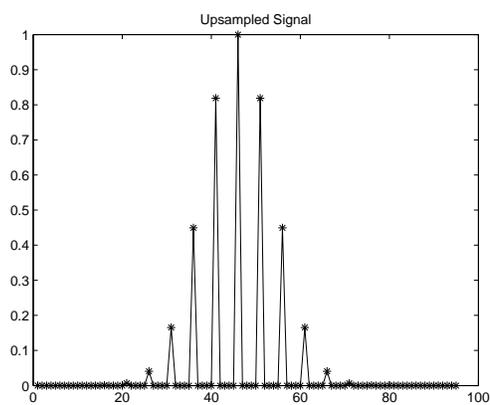
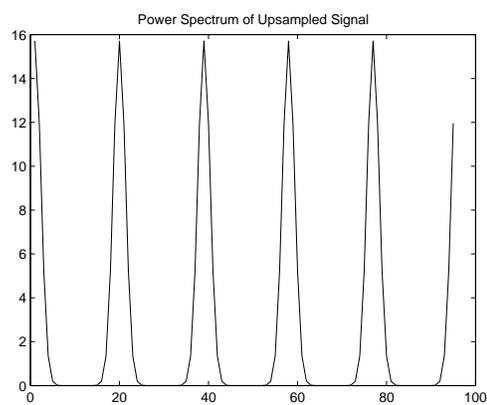
Fig. 3.5: Waveform of interpolated pulse.



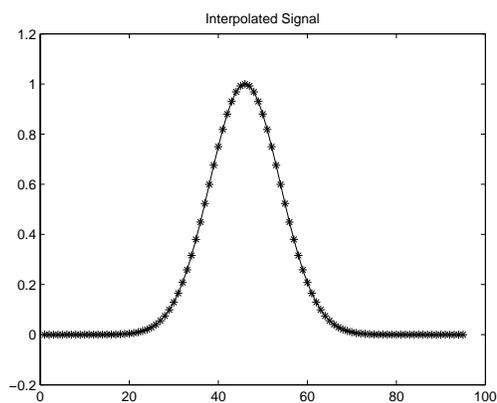
(a) Original signal to be interpolated.



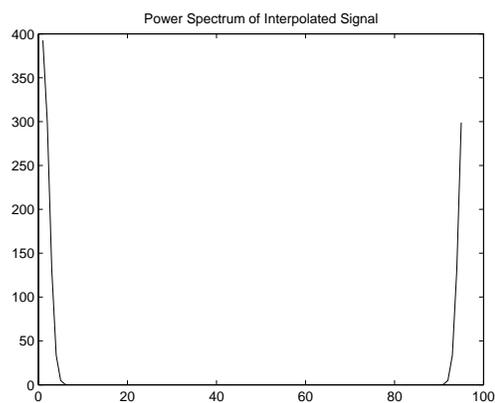
(b) Power spectrum of original signal.

(c) Signal after inserting $N-1$ zeros for upsampling.

(d) Power spectrum of upsampled signal.



(e) Upsampled signal after low-pass filtering.



(f) Final power spectrum of interpolated signal.

Fig. 3.6: The process of interpolation.

to applying a filter $L(z)$ of order kN to the interpolated signal. This is because interpolation causes the frequencies in the spectrum $R(f)$ to be compacted. Therefore, the filter is required to have a steeper bandedge. By filtering the noise on an uninterpolated signal, the filter $I(z)$ can have a more gradual transition band without passing high levels of noise through the system. Hence, the more gradual transition band allows the filter $I(z)$ to be a lower ordered filter [26].

3.4 Deconvolution

As the laser pulse travels to and from the scene, the return waveform is modified as it potentially hits multiple surfaces at various depths and angles within the beam footprint [11]. Deconvolution is a tool used to estimate the surface response of the target. Because deconvolution is a deblurring operation, the surface response is used to obtain higher range resolution on targets such as those with multiple surfaces within the beam footprint. Applying the low-pass filters $L(z)$ and $I(z)$ to (3.1) leads to

$$R(z)L(z)I(z) = (P(z)S(z)H(z) + N(z)H(z) + N_q(z))L(z)I(z), \quad (3.4)$$

which can be rearranged to obtain the target surface response,

$$S(z) = \frac{\tilde{R}(z)}{\tilde{H}(z)} + \tilde{N}(z), \quad (3.5)$$

where

$$\tilde{R}(z) = R(z)L(z)I(z), \quad (3.6)$$

$$\tilde{H}(z) = P(z)H(z)L(z)I(z), \quad (3.7)$$

and

$$\tilde{N}(z) = \frac{N(z)}{P(z)} + \frac{N_q(z)}{P(z)H(z)}. \quad (3.8)$$

Deconvolution is possible only if $\tilde{R}(z)$ and $\tilde{H}(z)$ are known and effects from $\tilde{N}(z)$ can be minimized. $\tilde{r}(n)$ is the measured waveform after filtering and interpolation. To reduce computation, only segments of $\tilde{r}(n)$ containing returns are processed. The middle of a segment is taken to be the range of the return and is denoted by r_m . This is found using a first derivative crossover discriminator on $r(n)$ [19, 25]. The beginning of the segment is $r_m - 30N$ and the end of the segment is $r_m + 30N$ where N is the sampling factor. If two segments overlap, they are merged into one segment that starts at the beginning of the first segment and stops at the end of the second segment. Each segment is then windowed with a Tukey window before being processed individually.

An estimate of $\tilde{H}(z)$ (the blur function) is determined using data obtained by holding the ELT scanner stationary while aimed at a flat target surface oriented normal to the ELT. This scenario allows the transmitted pulse to return to the sensor unaltered by the surface. This indicates that the surface response is a delta function for a flat surface perpendicular to the sensor. Several pulses are recorded, and each waveform is registered with respect to range. The waveforms are then averaged which substantially reduces noise and provides an accurate estimate of $\tilde{H}(z)$ as shown in Fig. 3.7. Once $\tilde{H}(z)$ is obtained, it may be used for any ELT experiment that shares the same receiver electronics characteristics. Alternatively, the transmitted pulse for each shot could have been used to estimate the blur function [11], but noise on each transmitted pulse waveform, combined with differences in the electronics used to capture the transmitted and received pulse, produce a less accurate estimate of the blur in the received waveforms than the averaging method used.

Once $\tilde{R}(z)$ and $\tilde{H}(z)$ are obtained, the surface response is computed using one of many deconvolution algorithms. Although many deconvolution techniques exist, three were tested with the ELT waveform. These are the Wiener filter (3.9), RL deconvolution (3.11), and NNLS minimization (3.12).

The Wiener filter is given by

$$W(f) = \frac{H(f)^* P_s(f)}{|H(f)|^2 P_s(f) + P_N(f)}, \quad (3.9)$$

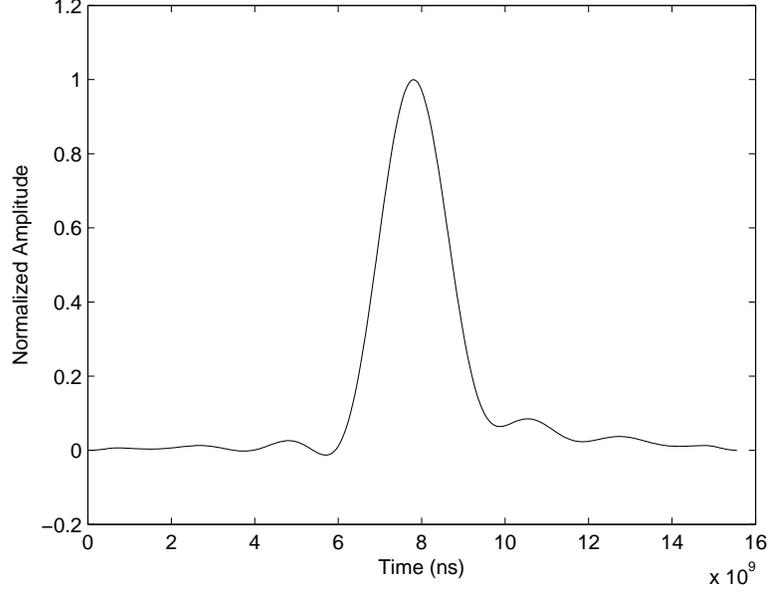


Fig. 3.7: Plot of the impulse response $\tilde{h}(t)$.

where $P_s(f)$ is the power spectral density of $s(n)$, and $P_N(f)$ is the power spectral density of the noise in the waveform. The surface response is then given by

$$\hat{S}(f) = W(f)R(f), \quad (3.10)$$

where $\hat{S}(f)$ is the fourier transform of the surface response estimate $\hat{s}(n)$.

The RL algorithm is an iterative technique where $\hat{s}(n)$ is computed by

$$\hat{s}(n)^{i+1} = \hat{s}(n)^i \sum_m \frac{r(m)h(m-n)}{\sum_n \hat{s}(n)^i h(m-n)}, \quad (3.11)$$

where $\hat{s}(n)^i$ is $\hat{s}(n)$ on the i th iteration [14].

The NNLS algorithm minimizes

$$\|H\mathbf{s} - \mathbf{r}\|_2^2 \quad s.t. \quad \mathbf{s} \geq 0, \quad (3.12)$$

where \mathbf{r} is a column vector created with the elements of $r(n)$, \mathbf{s} is a column vector created with the elements of $s(n)$, and H is the convolution matrix representation of $h(n)$ [27]. H

has a Toeplitz structure given by

$$H = \begin{bmatrix} h(0) & 0 & \cdots & 0 \\ h(1) & h(0) & \cdots & 0 \\ \vdots & \vdots & & \vdots \\ h(N) & h(N-1) & \cdots & h(0) \\ \vdots & \vdots & & \vdots \\ 0 & 0 & \cdots & h(N) \end{bmatrix}. \quad (3.13)$$

To comment on the computational cost, deconvolution is the process that is most time consuming for creating the point cloud. A high density point cloud of over 200,000 shots that required 358 s for processing without deconvolution took 756 s with the Wiener filter, 493 s with the NNLS method, and 644 s with the RL method. The majority of the computation in the Wiener filter case comes from gathering the noise statistics on every pulse. Much of the computational cost could be eliminated by assuming the same statistics on every pulse. Solving for the locations of returns in the footprint was not benchmarked because it was observed to never take more than 5 s to process the entire point cloud. The machine used to process the data has a 2.66 GHz Intel Core 2 Duo processor with 8 Gb of RAM. The code was written in Matlab utilizing C-mex functions for the more demanding processes.

3.5 Range Discrimination

Several range discrimination methods were tested on the surface response. These include the crossover, constant fraction, and Gaussian decomposition methods. Based on experimentation, the discrimination method chosen for the work in this thesis is the first derivative crossover [19]. This method provided the most accurate results for our experiments. Results of this testing go beyond the scope and purpose of this thesis and are therefore not presented.

To provide more precision in the range estimate than is available with the current sample rate, linear interpolation is used at this stage. A similarity exists between the approach used in section 3.3 and simple linear interpolation. Linear interpolation is equivalent to

using $L(f) = \text{sinc}^2(f)$. The passband of $\text{sinc}^2(f)$ rolls off quickly, thus altering the high frequencies in the signal's spectrum. As was seen in Fig. 3.5, a filter with better passband characteristics than $L(f) = \text{sinc}^2(f)$ was used at the interpolation stage. Since the signal has already been interpolated, it now contains lower frequency content in the sampled frequency domain and linear interpolation at this stage provides desirable results.

Chapter 4

Surface Response Estimation

4.1 Surface Response Derivation

In this section, the surface response is derived to understand what it is and what it looks like for different surfaces. The surface response is the impulse response of the target's surface. For example, if a delta function were transmitted from the laser and the receiver had infinite bandwidth with no noise, the profile of $r(n)$ would be the surface response. It is assumed that the beam divergence is small so that the laser beam is nearly parallel. With this assumption, the range from the sensor to the target is the same over the entire footprint when a flat, perpendicular surface is hit. The profile of the spatial beam footprint of the laser is denoted by $f(x, y)$. The profile of the ELT footprint is approximated to be a bivariate Gaussian distribution

$$f(x, y) = \frac{1}{2\pi\sigma_f^2} \exp\left(\frac{-((x - x_c)^2 + (y - y_c)^2)}{2\sigma_f^2}\right), \quad (4.1)$$

where (x_c, y_c) is the center of the beam and σ_f is the standard deviation of the circularly Gaussian footprint which can be computed by

$$\sigma_f = \frac{\beta}{4}r, \quad (4.2)$$

where r is the range from the sensor to the surface and β is the divergence of the laser defined to give a value of e^{-2} at $\beta/2$. The surface is represented with the function $g(x, y)$ that returns the true distance from the sensor to the surface at location (x, y) . Applying (2.1) to $g(x, y)$ results in the time of flight from the sensor to the surface. If a beam with no spatial width were to be transmitted at (x, y) , the surface response would result as a

delta function offset by $\frac{2}{C}g(x, y)$. Let the footprint be divided into adjacent beamlets so that the footprint is the union of all beamlets. If the number of beamlets increases so that the width of each beamlet goes to zero, the surface response is given by

$$s(t) = \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \delta \left(t - \frac{2}{C}g(x - x_c, y - y_c) \right) f(x, y) dx dy. \quad (4.3)$$

For the ELT pulse,

$$s(t) = \frac{1}{2\pi\sigma_f^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \delta \left(t - \frac{2}{C}g(x, y) \right) \exp \left(\frac{-(x^2 + y^2)}{2\sigma_f^2} \right) dx dy, \quad (4.4)$$

is the theoretical surface response.

4.1.1 Surface Response of Two Offset Surfaces

One surface of interest is a surface that is discontinuous inside the footprint as shown in Fig. 4.1. Let the edge run parallel to the y -axis so that

$$g(x, y) = \begin{cases} r_1, & x \leq x_{edge} \\ r_2, & x > x_{edge} \end{cases}, \quad (4.5)$$

where r_1 is the range of the first surface, and r_2 is the range of the second surface. Figure 4.1(a) shows the side view of the Gaussian footprint, which has a 1-D Gaussian profile.

The surface response of this surface with the VISSTA ELT footprint is given by

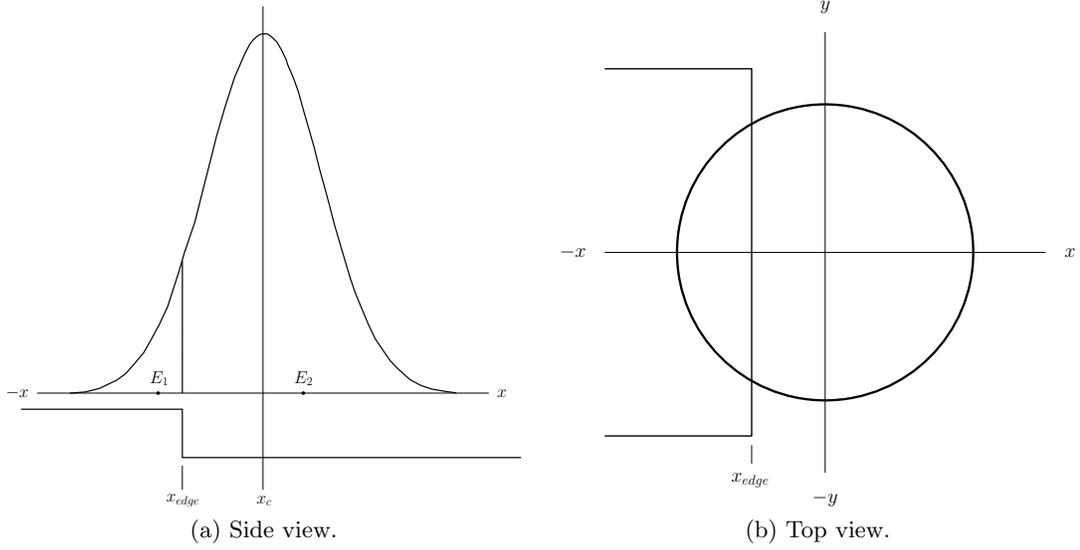


Fig. 4.1: The footprint covering a discontinuous surface that gives two returns.

$$\begin{aligned}
 s(t) &= \frac{1}{2\pi\sigma_f^2} \delta\left(t - \frac{2r_1}{C}\right) \int_{-\infty}^{\infty} \int_{-\infty}^{x_{edge}} \exp\left(\frac{-(x^2 + y^2)}{2\sigma_f^2}\right) dx dy, \\
 &+ \frac{1}{2\pi\sigma_f^2} \delta\left(t - \frac{2r_2}{C}\right) \int_{-\infty}^{\infty} \int_{-x_{edge}}^{\infty} \exp\left(\frac{-(x^2 + y^2)}{2\sigma_f^2}\right) dx dy, \\
 &= \delta\left(t - \frac{2r_1}{C}\right) \frac{1}{\sqrt{2\pi}\sigma_f} \int_{-\infty}^{x_{edge}} \exp\left(\frac{-x^2}{2\sigma_f^2}\right) dx dy, \\
 &+ \delta\left(t - \frac{2r_2}{C}\right) \frac{1}{\sqrt{2\pi}\sigma_f} \int_{-x_{edge}}^{\infty} \exp\left(\frac{-x^2}{2\sigma_f^2}\right) dx dy, \\
 &= E_1 \delta\left(t - \frac{2r_1}{C}\right) + E_2 \delta\left(t - \frac{2r_2}{C}\right), \tag{4.6}
 \end{aligned}$$

where E_1 is the energy landing on the first surface, and E_2 is the energy landing on the second surface. This shows that the surface response of two offset surfaces contains two delta functions. Furthermore, the range from the sensor to the target is found by the locations of the delta functions. It should be noted that if $r_1 = r_2$, a single delta function results. This represents a single flat surface perpendicular to the sensor.

Another result from this derivation is that given infinite bandwidth and no noise, there

is no limit to the range resolution. This is not the case with real data, so it is assumed that the delta functions take on a Gaussian pulse with a standard deviation, σ . The value $\sigma = 0$ represents the ideal case which results in a delta function. It is desired to know the limit of detecting two Gaussian pulses in a waveform. To obtain a best case scenario, assume that two Gaussian pulses have the same amplitude and standard deviation. The pulses are separated by 2τ as shown in Fig. 4.2.

Let the summation of the two pulses be given by

$$f(t) = \frac{A}{\sqrt{2\pi}\sigma} \left(\exp\left(\frac{-(t-\tau)^2}{2\sigma^2}\right) + \exp\left(\frac{-(t+\tau)^2}{2\sigma^2}\right) \right), \quad (4.7)$$

where A is the amplitude and σ is the standard deviation of the pulses. When τ is small, $f(t)$ appears as one pulse with a peak at $t = 0$. As τ increases, $f(t)$ separates into two distinct pulses. If $\frac{d^2f}{dt^2}|_{t=0} > 0$, then two peaks are detectable with a first order differentiator. The resolving limit is where $\frac{d^2f}{dt^2}|_{t=0} = 0$. This can be solved by

$$\frac{d^2f}{dt^2}|_{t=0} = \frac{2(\tau^2 - \sigma^2) \exp(\frac{-\tau^2}{2\sigma^2})}{\sigma^4} = 0 \rightarrow \tau = \sigma. \quad (4.8)$$

Therefore, if the separation between two Gaussian pulses is less than 2σ , they will appear as one pulse to a first order differentiator. This result was derived for a noiseless, continuous pulse. In the presence of noise or sampling, the separation must be larger in order to detect the two pulses.

4.1.2 Surface Response of a Planar Surface

The second surface to consider is a planar surface that is not necessarily perpendicular to the sensor. Let the angle of this surface be a rotation about x -axis as shown in Fig. 4.3. ϕ is the angle of incidence, $0 \leq \phi < \frac{\pi}{2}$, and r_c is the center of the footprint.

For this surface, $g(x, y)$ is given by

$$g(x, y) = r_c - y \tan(\phi), \quad (4.9)$$

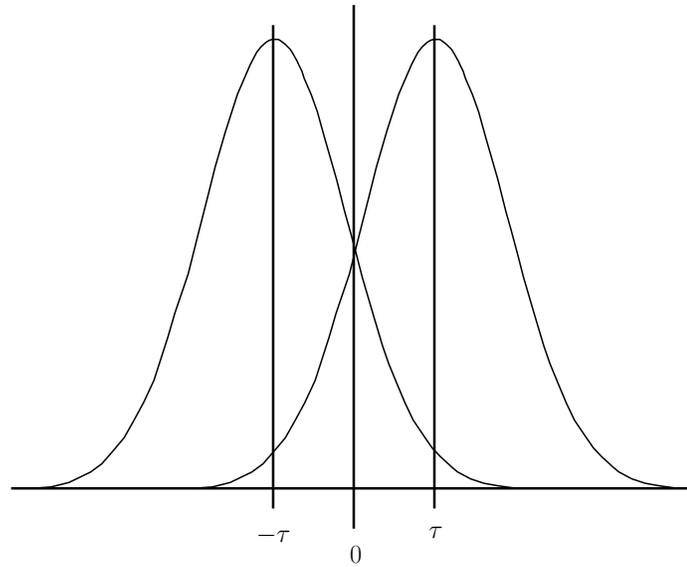
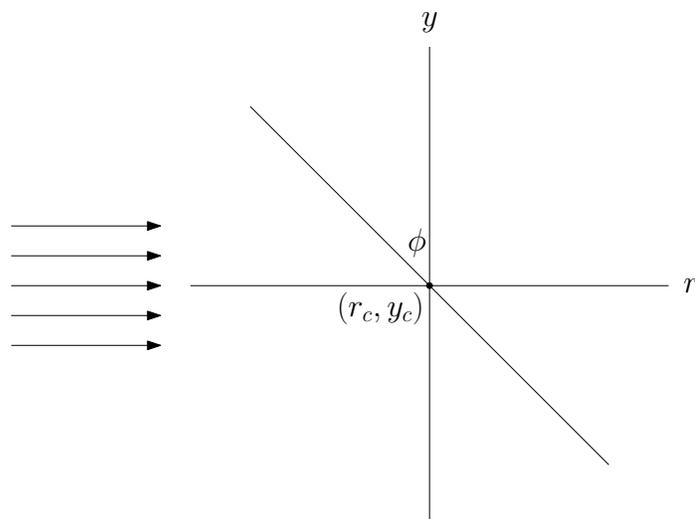


Fig. 4.2: Separation of two Gaussian pulses.

Fig. 4.3: Planar surface at an angle of ϕ .

which leads to the surface response for the VISSTA ELT

$$\begin{aligned}
s(t) &= \frac{1}{2\pi\sigma_f^2} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \delta\left(t - \frac{2}{C}(r_c - y \tan(\phi))\right) \exp\left(\frac{-(x^2 + y^2)}{2\sigma_f^2}\right) dx dy, \\
&= \frac{1}{\sqrt{2\pi}\sigma_f} \int_{-\infty}^{\infty} \delta\left(t - \frac{2}{C}(r_c - y \tan(\phi))\right) \exp\left(\frac{-y^2}{2\sigma_f^2}\right) dy, \\
&= \frac{C}{2\sqrt{2\pi} \tan(\phi)\sigma_f} \exp\left(\frac{-\left(\frac{C(t - \frac{2}{C}r_c)}{2 \tan(\phi)}\right)^2}{2\sigma_f^2}\right), \\
&= \frac{1}{\sqrt{2\pi} \left(\frac{2 \tan(\phi)\sigma_f}{C}\right)} \exp\left(\frac{-(t - \frac{2}{C}r_c)^2}{2 \left(\frac{2 \tan(\phi)\sigma_f}{C}\right)^2}\right). \tag{4.10}
\end{aligned}$$

This result shows that the resulting surface response has a Gaussian profile. The standard deviation of the surface response is given by

$$\sigma_s = \frac{2 \tan(\phi)\sigma_f}{C}. \tag{4.11}$$

As $\phi \rightarrow 0$, $\sigma_s \rightarrow 0$. This indicates that the surface response of a flat surface perpendicular to the sensor's line of sight is a delta function. When $\phi = 0$, this result agrees with the result from section 4.1.1 where $r_1 = r_2$ by returning a single delta function. As $\phi \rightarrow \frac{\pi}{2}$ radians, $\sigma_s \rightarrow \infty$, and the amplitude of the pulse goes to zero. This indicates that the surface response of a planar surface for large ϕ is a single, wide pulse. Equation (4.11) also agrees with intuition. As the beam footprint hits a surface at an angle, it is spread or stretched across the surface. A measure of the pulse stretch in the waveform is provided by σ_s .

4.2 Regularization

Inverse problems such as deconvolution are ill-posed and lead to undesirable solutions [28]. For the NNLS method, the convolution matrix, \mathbf{H} is ill-conditioned, therefore, perturbations due to noise are severely amplified in the solution. This means that in prac-

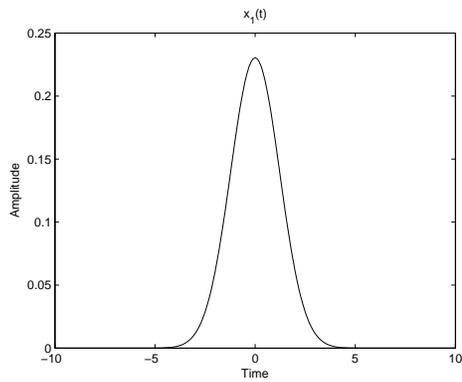
tice, obtaining the true surface response is not always possible. It has been shown that only information inside the system bandwidth can be recovered unambiguously [28].

The ambiguity problem can be seen through an example of two signals passing through a system. The first signal, $x_1(t)$, is a Gaussian pulse and the second signal, $x_2(t)$, is an upsampled Gaussian pulse. $x_2(t)$ was constructed so that the non-zero values lie on $x_1(t)$. The system, $h(t)$, is a sinc^2 pulse which transforms to a triangle function in the frequency domain. Figure 4.4 shows these signals in the time domain. Figure 4.4 also shows the convolution of the signals with the system. The result of $x_1(t) * h(t)$ appears to be equal to the result of $x_2(t) * h(t)$. Figure 4.5 shows that taking the difference, $x_1(t) * h(t) - x_2(t) * h(t)$ results in a very small residual relative to the original amplitudes. Even the quantization noise for the ELT receiver is larger than this residual. If the signals of interest were surface responses, then this example shows that two distinct surface responses can result in the same received waveform when passed through the system.

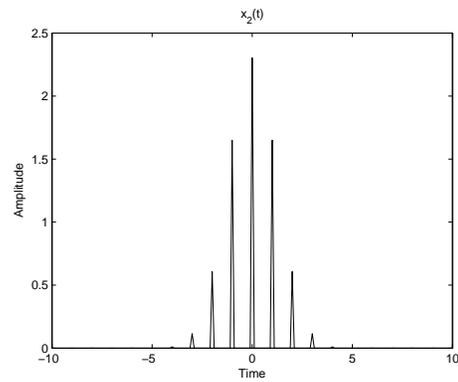
The problem can be explained by observing the Fourier transforms of the signals in the example. Figure 4.6 shows the Fourier transforms of the signals from Fig. 4.4. When $X_2(f)$ is multiplied by $H(f)$, the high frequency content in the signal is lost due to $H(f) = 0$ for high frequencies. Within the band limits of $H(f)$, $X_1(f) = X_2(f)$. Therefore the convolution of the system and each signal provides two indistinguishable results.

Regularization is a method that incorporates prior information in order to select the most feasible solution. A common approach when working with low-pass systems is to add a constraint to the problem that chooses the solution with minimal energy [28]. Such a constraint provides smooth results, thus high frequency noise is eliminated. However, high frequency content in the signal is removed as well.

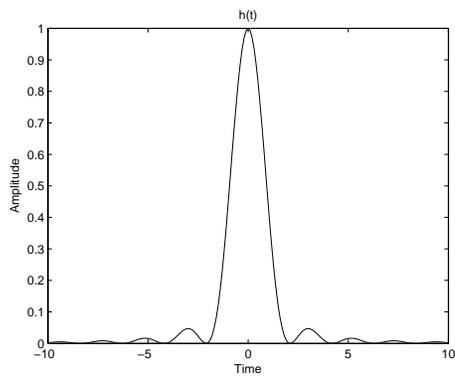
Regularization requires a parameter that determines the amount of regularization [29]. More regularization leads to higher error in the solution, but also increases its feasibility [29]. The regularization parameter may appear as a band edge in a filter, a parameter included in the cost function, or even as the number of iterations in an iterative technique [28]. The amount of regularization for iterative methods decreases with more iterations.



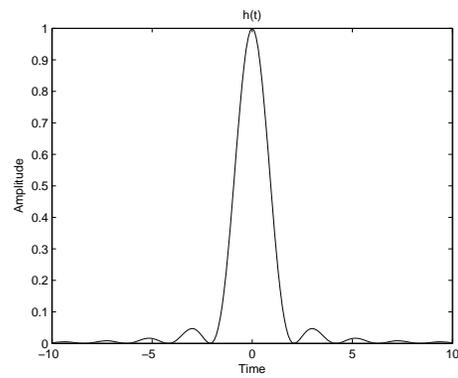
(a) Gaussian pulse.



(b) Upsampled Gaussian pulse.



(c) System impulse response.



(d) System impulse response.

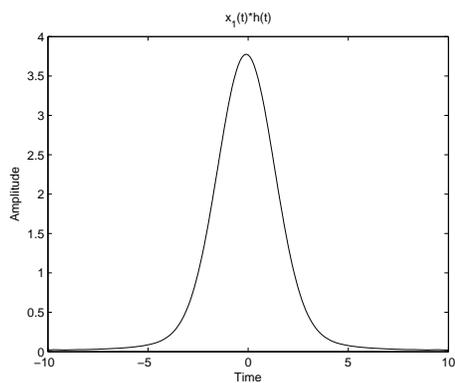
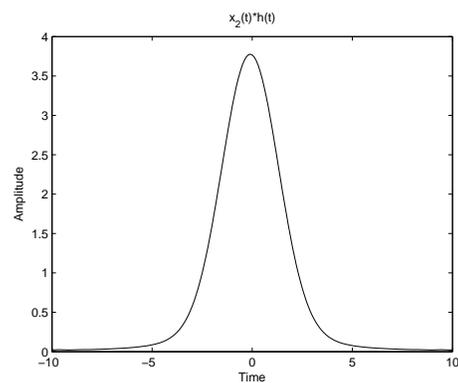
(e) Signal resulting from $x_1(t) * h(t)$.(f) Signal resulting from $x_2(t) * h(t)$.

Fig. 4.4: Convolution of two signals with a system in the time domain.

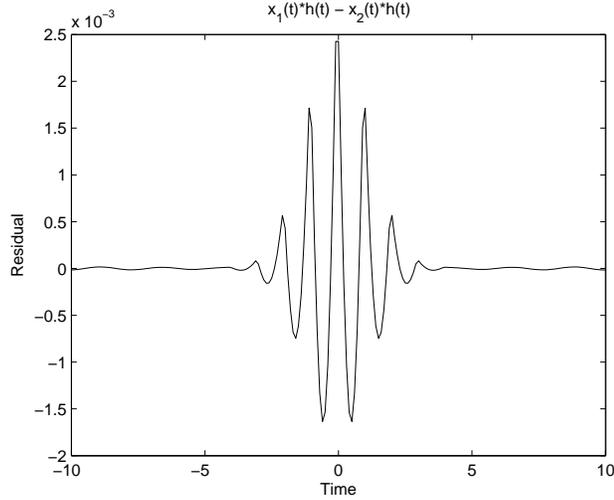


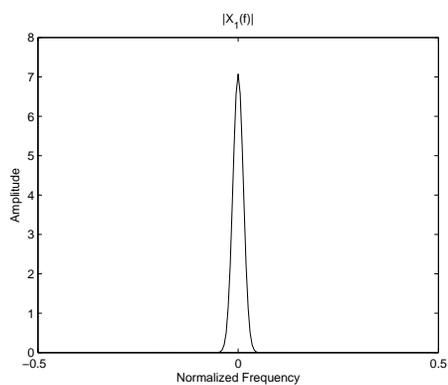
Fig. 4.5: Difference between the convolution between the system and each signal.

Deconvolution using the Wiener filter regularizes the signal naturally [28]. The amount of filtering is determined by the SNR; so for any f where $\tilde{H}(f) = 0$, $\hat{S}(f) = 0$ also. The regularization parameter is $P_{\tilde{N}}(f)$ from (3.9). Although the Wiener filter is regularized, it does not incorporate the positivity constraint like the other two methods. For the experiments in this thesis, positivity was applied to $\hat{s}(n)$ after filtering by setting all negative values to 0.

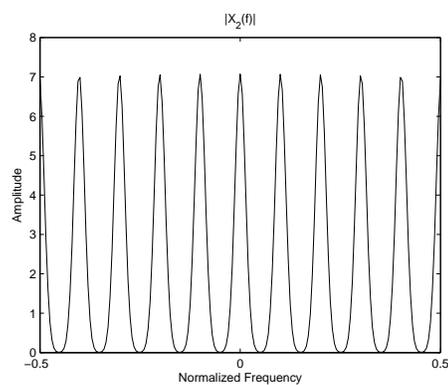
The RL algorithm is also naturally regularized, where the regularization parameter is determined by i from (3.11). As i increases, $\hat{s}(n)$ approaches the unregularized solution. It was determined that $i = 25$ provided the best results for the VISSTA ELT waveform.

The NNLS method is unregularized. With this method, noise can cause $\hat{s}(n)$ to appear as two close returns even when only one flat surface was hit by the laser. This problem can be addressed by low-pass filtering $\hat{s}(n)$ for regularization after it has been estimated. In order to preserve a non-negative surface response, the low-pass filter chosen for this method is a Gaussian filter where the band edge, B_0 , acts as the regularization parameter [28]. The Gaussian filter for regularization is given by

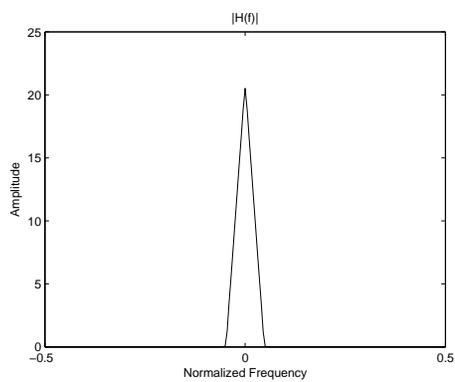
$$Y(f) = \exp\left(-\frac{(2\pi f)^2}{2B_0^2}\right). \quad (4.12)$$



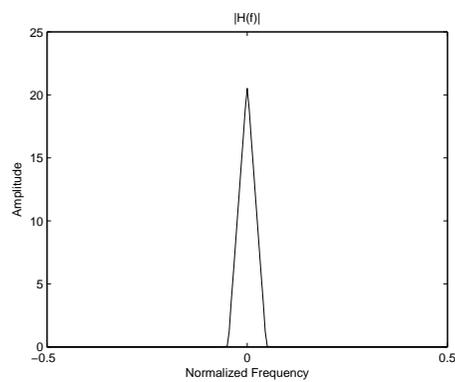
(a) Magnitude response of Gaussian pulse.



(b) Magnitude response of upsampled Gaussian pulse.



(c) Magnitude response of system.



(d) Magnitude response of system.

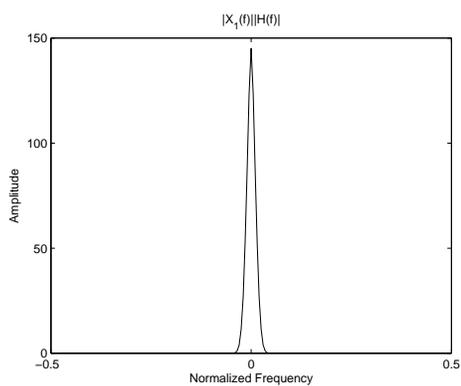
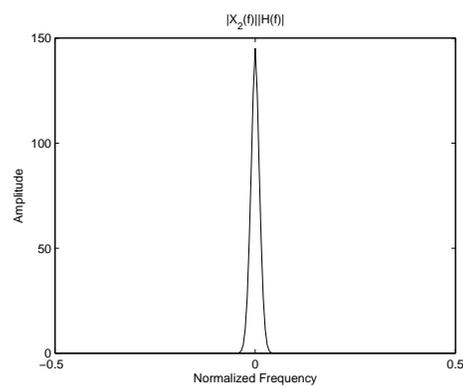
(e) Magnitude response of $|X_1(f)||H(f)|$.(f) Magnitude response of $|X_2(f)||H(f)|$.

Fig. 4.6: Convolution of two signals with a system in the frequency domain.

B_0 is determined so that $Y(f)$ filters those frequencies outside of the support of $\tilde{H}(f)$.

Increasing B_0 gives less error to the solution, but allows high frequency noise to increase in the solution. Upon increasing B_0 to allow high frequencies in $\hat{S}(f)$ to pass, it was observed that the surface response estimate of two offset surfaces still represents the true surface response accurately. However, the surface response of angled surfaces is not well represented. A similar result was noticed in other works [13]. If the angle of incidence is small ($< 60^\circ$), the surface response appears identical to the surface response from hitting a single flat surface at 0° . The reason for this can be understood from (4.11). The standard deviation of the Gaussian surface response is proportional to the tangent of the angle. At small angles, a change in angle causes a small change in the standard deviation of the surface response. As the angle gets closer to 90° , a change in angle causes a big change in the standard deviation of the surface response. For example, the value of σ_f for the ELT pulse at a range of 500 m is estimated to be 4 cm. Using that value in the equation, the difference between the width of the surface response at $\phi = 25^\circ$ and $\phi = 30^\circ$ is .03 ns. The difference between the width of the surface response at $\phi = 75^\circ$ and $\phi = 80^\circ$ is .52 ns. The ELT receiver has a sample period of .5 ns. Thus, the received waveform is notably different only if the angle of incidence is large.

If the angle of incidence is large ($> 60^\circ$), a problem occurs in the surface response. Unwanted high frequency noise corrupts the surface response. It was shown in section 4.1.2 that the surface response of a planar angled surface for the VISSTA ELT waveform has a Gaussian profile. However, Fig. 4.7 shows the surface response estimate when B_0 is set too high. The high frequency noise causes ripples in the surface response that appear as multiple peaks. This surface response could be mistaken for the surface response of four distinct surfaces at different ranges. This is an example of two surface responses that cannot be recovered uniquely without introducing prior information about the surface.

The typical solution would be to decrease B_0 or equivalently to increase the amount of regularization, thus trading error in the solution for feasibility. However, increasing B_0 to eliminate noise also decreases the range resolution. Section 4.1.1 shows that two Gaussian

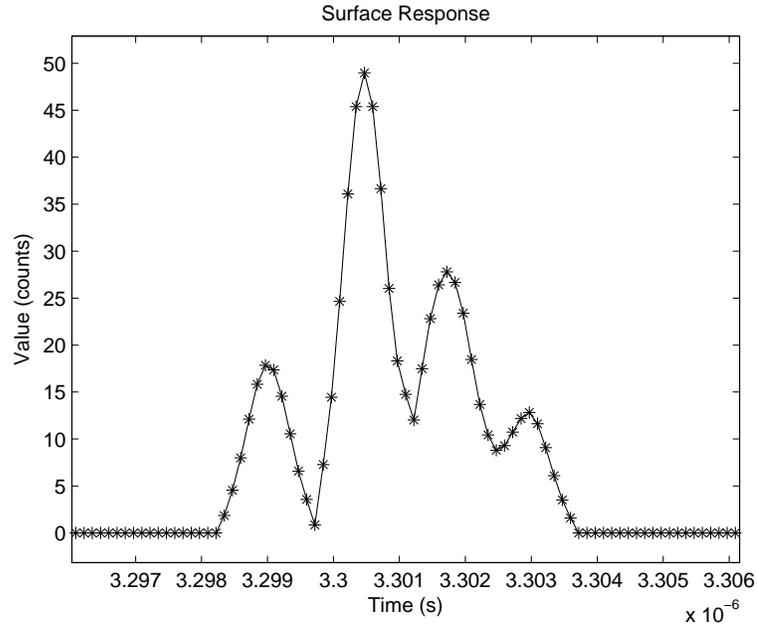


Fig. 4.7: Under-regularized surface response resulting from hitting a planar surface at 75° .

pulses cannot be separated if their means are closer than $2B_0$. For the experiments in this thesis, B_0 was set to .32 ns or equivalently 4.8 cm when (2.1) is applied. This means that the absolute best range resolution in the absence of noise would be 9.6 cm.

Although the surface response in Fig. 4.7 does not represent the surface response derived in section 4.1.2, the low frequency envelope shows what one would expect in the surface response for an angled surface. It is only the high frequency content that is incorrect. Under the assumption that a planar surface was hit, section 4.3.4 shows a method to treat this surface response to estimate the angle of incidence.

4.3 Surface Response Results

4.3.1 Range Error on a Single Surface

Range error is an important measure of the performance of a lidar system. This error can be estimated statistically by [30]

$$\sigma_r = \frac{c}{2} \cdot \frac{\sigma_d(t_p)}{\left| \frac{d}{dt} v(t) \right|_{t=t_p}}, \quad (4.13)$$

where c is the speed of light, $\sigma_d(t_p)$ is the detection noise at the timing point t_p , and $v(t)$ is the pulse waveform. This equation indicates that the range error depends on the slope and noise level of the waveform at the point of detection. The surface response provides a signal with a higher slope, but amplified noise due to the high-pass nature of deconvolution. For this reason, three deconvolution techniques were tested to discover the trade off between range resolution and range error. The methods tested are the Wiener filter, RL deconvolution, and NNLS minimization.

To test the range error, a flat, white Lambertian target with estimated reflectivity of 0.7 was constructed and placed at a range of 493 m. The beam was held stationary while sending multiple pulses. The data were processed using the three deconvolution methods as well as using the waveform without deconvolution. The standard deviations of the range for each test were computed from the data, and are found in Table 4.1.

These results show that for the Wiener filter and NNLS algorithm, range error was not an issue, but for the RL algorithm, range resolution came at the cost of range error. It was observed that as the number of iterations for the RL method increased, the range error increased to 4.5 mm and remained there even at 100 iterations. Prior work showed less difference between the methods. However, it used only 10 iterations for the RL method as well as a different discrimination method [31].

Table 4.1: Range error at 493 meters.

Method	Range error (mm)
No Deconvolution	1.8
Wiener filter	1.8
RL	4.5
NNLS	1.8

4.3.2 Resolution of Two Surfaces

Multiple returns can be detected easily when the surfaces that were hit are farther apart in range than the width of the pulse, but it is a more difficult task for close surfaces. A purpose of deconvolution is to increase the ability to detect and estimate the surfaces that are closer than the pulse width.

To explore the ability to resolve two surfaces, simulated data from LadarSIM were collected. Simulated data were used to isolate the effects of noise on the processing methods. This eliminated other sources of error such as non-uniformity in the surface and blur ($\tilde{H}(z)$) estimation errors. Models were created to match the target and the performance of the ELT. After creating a point cloud, a shot whose energy was evenly split across the two surfaces was selected as a representative waveform. Thirty-five waveforms that differed only in noise were generated from the representative waveform and used for the statistics. The distance between the two surfaces was calculated. Figure 4.8 shows the results of the experiment as the distance between the two surfaces was varied. The mean estimate of each method is plotted with error bars that show the standard deviation of the measurement. Not all methods were capable of resolving two surfaces at all the different surface separations tested. If a point is not shown on the plot, it indicates that the method could not resolve the surfaces at that separation.

The pulse was resolved with no deconvolution at 28 cm. The error at this separation was higher than any measurement from the other methods. The Wiener filter provided 4 cm improvement in resolution, but could not resolve the surfaces past 24 cm. The NNLS method estimated the separation of the surfaces the best with regards to the mean, standard deviation in measurement, and ability to resolve the two surfaces at the 14 cm separation.

Data were also collected from the ELT to show the performance of each method on a real instrument. To measure the ability to resolve two surfaces on real data, the same target as described in section 4.3.1 was again placed 493 meters from the ELT. Figure 4.9 shows the target with a 19 cm X 19 cm square hole cut out from the middle of the front side. Surrounding the hole, sufficient area remained on the border of the target to act as

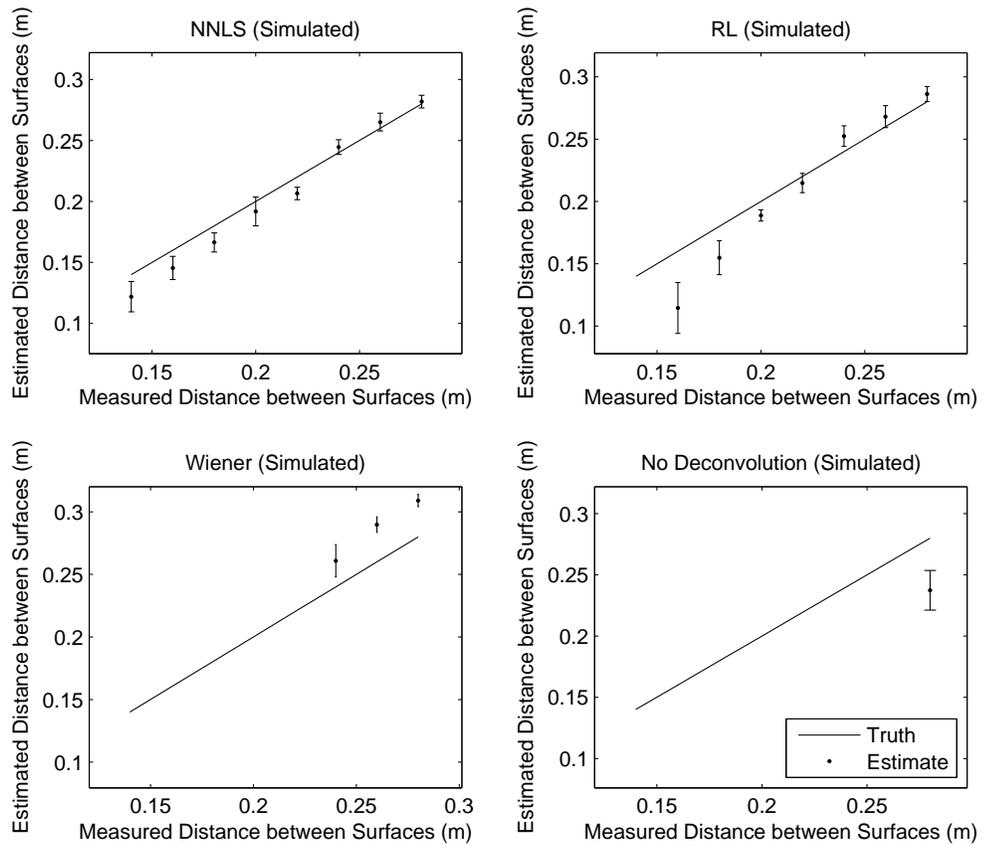


Fig. 4.8: Simulation measurements from returns hitting two surfaces as the distance of separation between the surfaces varied.

the first surface. Directly behind the hole, a second surface was placed capable of varying in range. This provided two surfaces that have a variable distance between them.

To obtain shots that hit both surfaces, the ELT scanner was set to carefully scan across the hole in the target. The data were processed to create four point clouds—one for each processing method. For all shots that detected both surfaces at the range of the target, the distance between the surfaces was recorded. To ensure that the same points were used for statistics, a shot was used only if all methods tested at that separation detected two returns in the waveform. Thus, points that were barely detectable by only one method were discarded. The ideal data points for the purposes of testing would be those that lie centered on the edge of the discontinuous surface. Although this method did not obtain only those points, it avoided points where the footprint’s energy rested heavily on only one surface. At least 400 points were used at every separation except for the 14 cm test which only used 68 points. The distance between the two surfaces was calculated, and the results are shown in Fig. 4.10.

Plots of an example waveform from the ELT after deconvolution are shown in Fig. 4.11. This provides an indication of the resolving power of each method.

For ELT data, the pulse was resolved at 30 cm without deconvolution. The Wiener filter again could not resolve the two surfaces past 24 cm. The NNLS method estimated the mean separation of the surfaces the best, and the RL method had the lowest standard deviation in measurement. Both the NNLS and RL algorithms resolved the 14 cm surface separation which was found to be the limit.

The Wiener filter did not perform as well as the other methods. A reason for this is the lack of the positivity constraint on the filter. This constraint was found to be necessary to provide a feasible result for the least squares solution. In comparison to previous work, the Wiener filter did not deconvolve the signal to the same extent with respect to the original pulse width [11]. A reason for this is the difference in the bandwidth of the receiver electronics [28]. The 500 MHz receiver bandwidth of the ELT is half of the 1 GHz receiver used by Jutzi and Stilla [11].

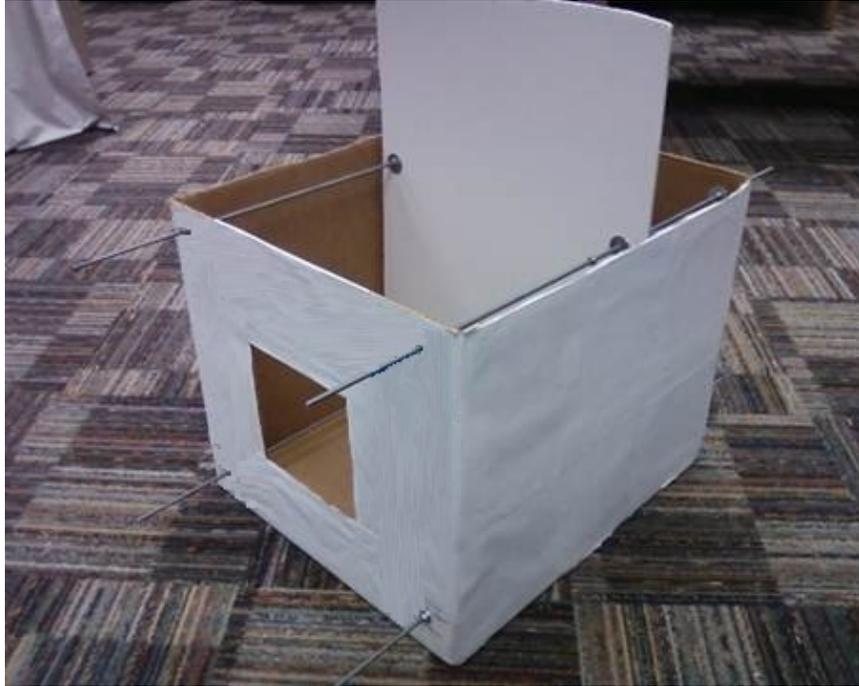


Fig. 4.9: The target used for testing in these experiments.

4.3.3 Effect of SNR on Range Resolution

Since deconvolution amplifies noise, the effect of SNR on the surface range estimation was tested. SNR for these tests is defined as

$$SNR = \frac{r_{max}(n)}{\sigma_{\tilde{n}}}, \quad (4.14)$$

where $\sigma_{\tilde{n}}$ is the standard deviation of the noise and $r_{max}(n)$ is the maximum of the received waveform.

For simulation, noise levels were adjusted in the electronics and APD models to allow the SNR to vary. The same test as described in section 4.3.2 was then performed at a surface separation of 20 cm for the RL and NNLS methods. The Wiener Filter was not tested due to the inability to resolve two surfaces at 20 cm. Figure 4.12 shows the results for the simulated data.

To experimentally test SNR effects on real data, the surface separation on the target was set to 20 cm, and the amount of light returning to the detector was manipulated by

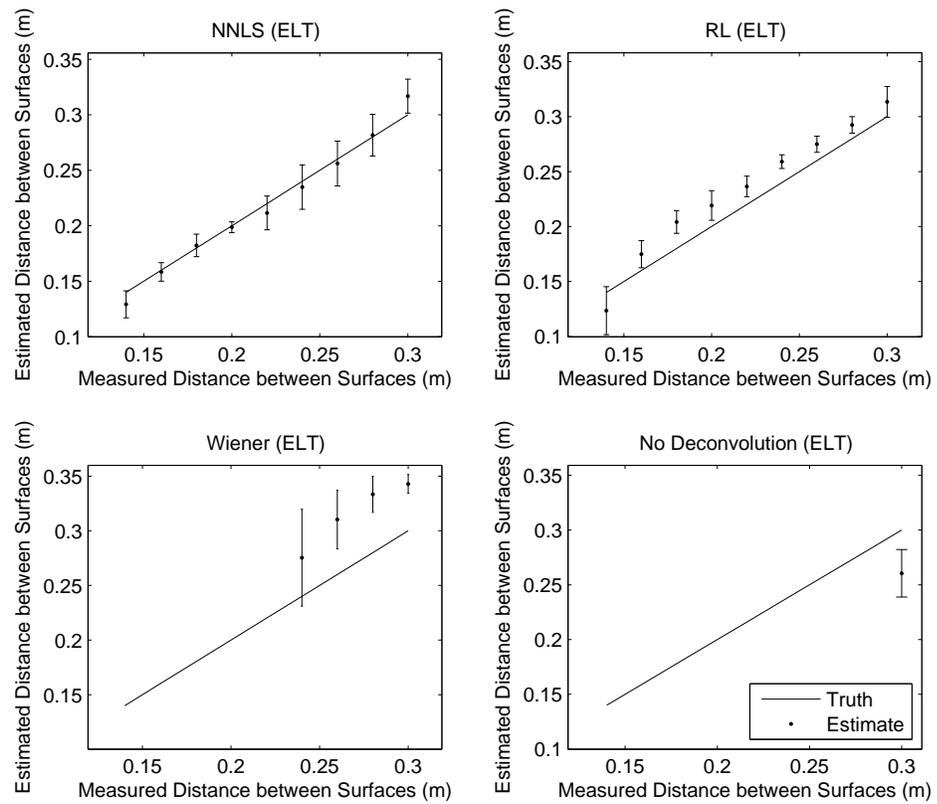


Fig. 4.10: ELT measurements from returns hitting two surfaces as the distance of separation between the surfaces varied.

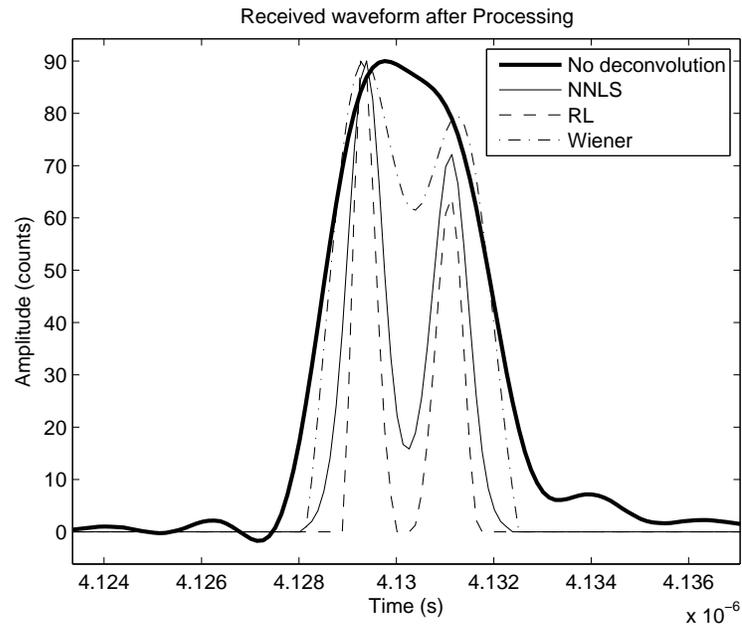


Fig. 4.11: The ELT received waveform processed with the different deconvolution methods.

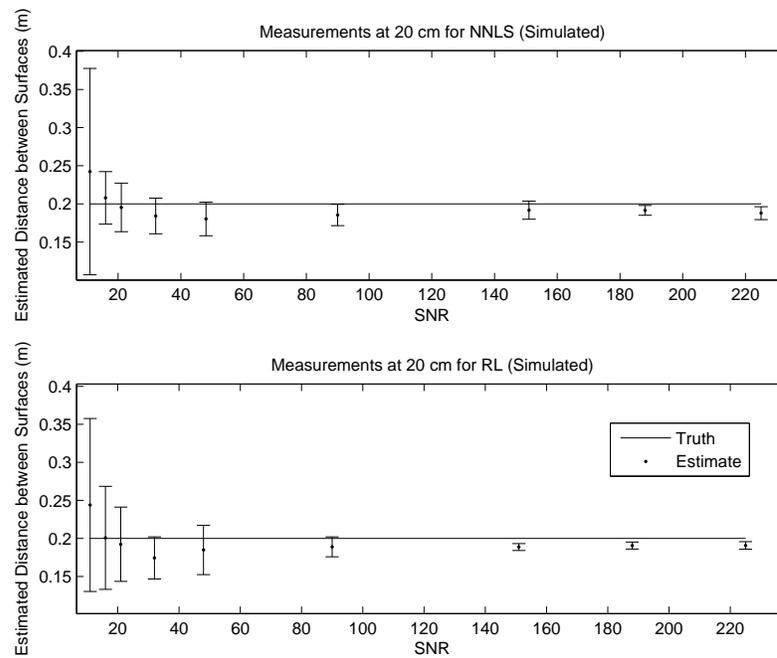


Fig. 4.12: Simulation measurements from returns hitting two surfaces at 20 cm as the SNR varied.

controlling an alignment mirror in the optical path. Gain settings in the electronics were kept constant. This provided data with varying SNR without affecting the receiver transfer function. The same processing procedure as in section 4.3.2 was carried out for the RL and NNLS methods. The Wiener filter was not tested due to the inability to resolve two surfaces at 20 cm. Figure 4.13 shows the results at the varying SNR levels.

Obtaining the surface response increased the range resolution for the VISSTA ELT even at the lower SNR's tested. Both the NNLS and RL methods detected the two surfaces at an SNR of 33, which was the lowest SNR tested with real data. At an SNR of 33, the standard deviation increased only slightly, and the mean of the estimate did not decrease in fidelity. The simulated data were processed at even lower SNR levels than the real data. At an SNR of 20 and below, the tested methods began to have difficulty deconvolving the signal. Not only did the error increase, but also the ability to detect two returns decreased. Not all shots at the lower SNR's of the simulated data were detected to have two returns; therefore, they were not used in the statistics. A conclusion to be drawn is that a high sample rate and SNR may be helpful, but not necessary provided that the SNR is adequate. For this experiment, the necessary SNR for the ELT data is approximately 30.

4.3.4 Angle Estimation

In current lidar systems, one range is reported for a shot that hits at an angle. However it should be noted that within the beam footprint there is not just one range that corresponds to the surface. Photons from one side of the footprint return before photons from the other side of the footprint. This indicates that there is a time window over which the photons return to the receiver.

Figure 4.7 shows the surface response of an angled surface computed from the NNLS deconvolution method. It is possible that the ranges of the peaks in the noisy surface response correspond to some location that is truly on the surface. Furthermore, because $\hat{s}(n)$ is not corrupted at low frequencies, the low frequency envelope of $\hat{s}(n)$ appears to be valid, e.g. the middle peak in Fig. 4.7 has more energy than the outer peaks. It was observed that the surface response estimates of angled surfaces had the highest peak in the

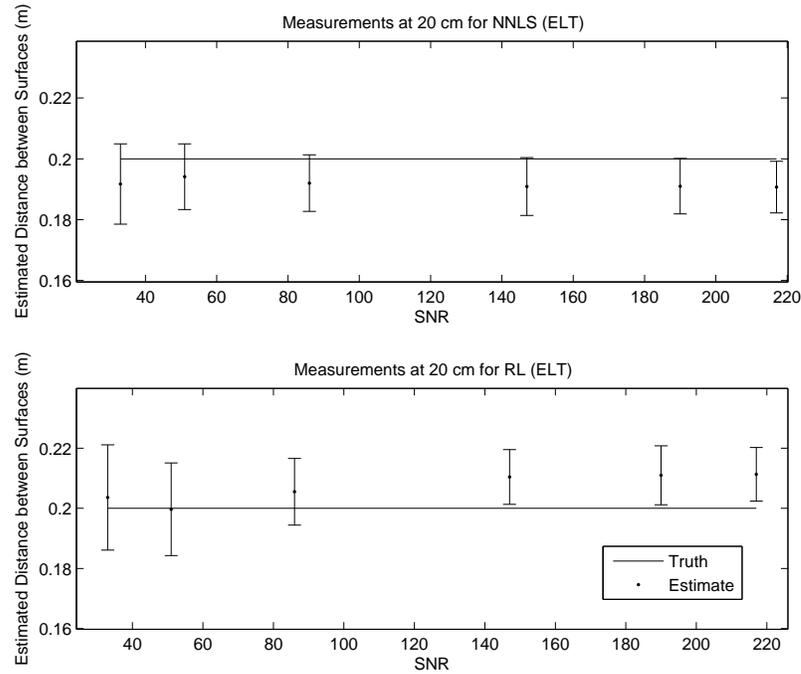


Fig. 4.13: ELT measurements from returns hitting two surfaces at 20 cm as the SNR varied.

center and the amplitudes of other peaks decreased as they moved away from the center of the surface response.

The former two paragraphs can be summarized to help form a method to estimate the angle of a surface given that the surface is planar. First, somewhere within the beam footprint there is a location that corresponds to the range of each peak in the surface response. The requirement of the angle estimation method is to decide where on the surface each peak corresponds to. Second, due to the low frequency envelope in the surface response, the energy in each peak can be estimated and used as if the energy of the beam footprint is concentrated at the locations of the peaks.

Assume that a shot hits a planar surface and the surface response results in K peaks. By assuming a planar surface, it is known that energy from the first peak comes from one extreme of the footprint, energy from the last peak comes from the other extreme, and energy from intermediate peaks comes from the middle of the footprint in order as shown in Fig. 4.14.

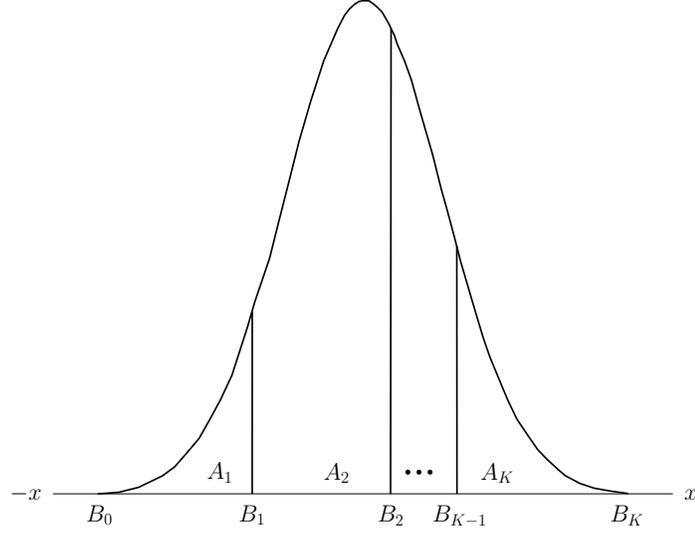


Fig. 4.14: Side profile of spatial beam footprint showing areas and boundaries of peaks.

For a zero mean Gaussian with a standard deviation of one, the boundaries separating energy from each peak in the footprint, as shown in Fig. 4.14, are computed by

$$B_k = Q^{-1} \left(1 - \sum_{j=1}^k E_j \right) \sigma_f, \quad (4.15)$$

where

$$Q(x) = \frac{1}{\sqrt{2\pi}} \int_x^{\infty} \exp\left(-\frac{t^2}{2}\right) dt, \quad (4.16)$$

and E_k is the normalized energy contained in each peak or the area of each region A_k . With boundaries defined, the mean value or expected location of peak k is given by

$$\begin{aligned} l_k &= \left(\frac{\int_{B_{k-1}}^{B_k} x \exp\left(\frac{-x^2}{2}\right) dx}{\int_{B_{k-1}}^{B_k} \exp\left(\frac{-x^2}{2}\right) dx} \right), \\ &= \left(\frac{\exp\left(\frac{-B_{k-1}^2}{2}\right) - \exp\left(\frac{-B_k^2}{2}\right)}{\sqrt{2\pi} E_k} \right). \end{aligned} \quad (4.17)$$

The angle of incidence of the planar surface can be deduced from the geometry in Fig. 4.15 and is given by

$$\phi = \tan^{-1} \left(\frac{d(r_i, r_j)}{d(l_i, l_j)} \right), \quad (4.18)$$

where $d(r_i, r_j)$ returns the distance between any two peaks in range and $d(l_i, l_j)$ is the distance between the corresponding peaks in the footprint. Figure 4.15 shows that ϕ can be computed using any two peaks. To obtain the most accurate estimate, the angles between consecutive points are computed and averaged.

To test the method, a test board shown in Fig. 4.16 was constructed to provide a large flat surface. The test board had an estimated reflectivity of 0.7 and was placed at 493 meters. The angle, ϕ , was measured using a compass with accuracy of approximately one to two degrees. The board was scanned at 45° , 60° , and 75° , and the angle was estimated using the method presented in this section. A setup was used in LadarSIM to match this scenario. Over 40 points were randomly selected from the middle of the target for the both the ELT data and simulated data.

The number of peaks, K , varies depending on ϕ . In the experiments, K was observed to be at most four on a plane where $\phi = 75^\circ$. A 45° angled surface was tested, but it was found that at this angle, only one peak was detected for both the VISSTA ELT and LadarSIM. After deconvolving the signal, it was indistinguishable from a return that hit at 0° . The results in Table 4.2 show that both the simulated and real data provide an accurate estimation of ϕ .

Table 4.2: Angle estimate.

Data	Truth (degrees)	Estimate (degrees)	Standard Deviation (degrees)
VISSTA ELT	60	64.4	4.9
VISSTA ELT	75	76.0	1.4
LadarSIM	60	59.6	4.1
LadarSIM	75	72.9	1.8

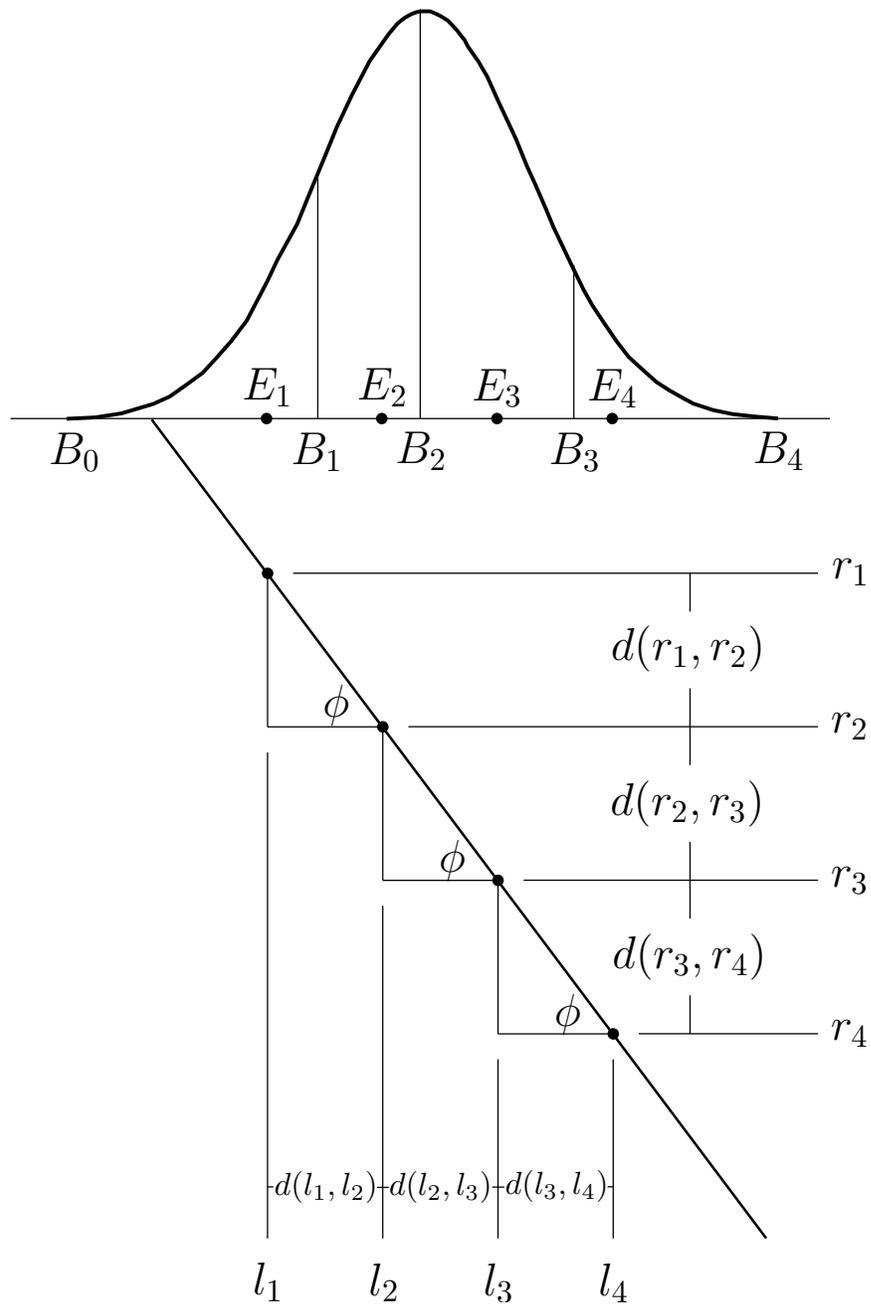


Fig. 4.15: Relation of peaks to the spatial locations within the footprint.



Fig. 4.16: Test board used in experiments.

4.3.5 Slot Width Estimation

Under the assumption that a target with two discontinuous surfaces to form a slotted surface is being targeted, a method similar to the angle estimation method can be used to estimate the width of the slot from section 4.3.2. However, for slot width estimation it is the boundaries, B_k , that are desired. The width of the slot is wider than the beam footprint, so each shot may hit only one of the edges.

Let the 2-D Gaussian spatial beam footprint be the footprint given in (4.1). Assume that two distinct surfaces were hit by the ladar pulse, and the edge is perpendicular to the x-axis as shown in Fig. 4.1. The location of the edge is determined by

$$x_{edge} = \sigma_f Q^{-1}(1 - E_1) + x_c. \quad (4.19)$$

To test the ability to locate edges with this method, data from section 4.3.2 were used for a slot depth of 16 cm. The measured width of the slot was 19 cm. An equivalent target and scenario was created in LadarSIM. Table 4.3 shows the results of processing. The mean was calculated by the difference between the mean of edge 2 and the mean of edge

1. Because the two edges are independent events, the standard deviation was computed by adding the variance of the measurement on edge 1 with the variance of the measurement on edge 2 and taking the square root.

Discrepancies between results from LadarSIM and the ELT are mostly ascribed to error in the pointing control. Because this paper is concerned with improving the waveform analysis, it was assumed that there was no pointing error in the LadarSIM model. Therefore, the LadarSIM results indicate the reliability of this method when considering only errors coming from noise in the waveform. Scintillation is another possible source of error for this test, but is likely to have had minimal effect due to data collection at temperatures below 0°C.

Table 4.3: Estimate of slot width.

Data	Truth (cm)	Mean (cm)	Standard Deviation (cm)
LadarSIM	19	19.20	1.13
VISSTA	19	18.58	3.44

Chapter 5

Estimation of Spatial Location Within the Footprint

Work in previous chapters dealt with a single shot at a time. All processing was performed on waveforms independent of other waveforms. In this chapter, it is assumed that the surface response has been computed for each shot. This chapter describes improvements that can be made to the point cloud using the surface response of the shots in the scan.

It has heretofore been assumed that each multiple return is spatially located in the middle of the beam footprint. When a point cloud is created, all returns from one waveform share the same values for azimuth and elevation. Although aligning all returns down the center of the beam footprint is reasonable in some applications, Fig. 5.1 shows that this assumption can misplace the point on the first surface in applications such as ATR. This situation is from the shot hitting two surfaces when the center of the shot does not line up with the edge. In this case, the first shot appears to be floating above the surface. Thus, a false surface is reported from that shot.

Another problem that stems from the presence of multiple returns occurs when the point cloud is interpolated. Interpolating a point cloud may be desired in various applications. For a 3-D model of an object, displaying points might not be as useful as showing the object with either a smooth or mesh surface. To create a surface from points, the points need to be connected so that the space between the points is filled in via interpolation. For interpolation, the data can be viewed as a range image where the azimuth and elevation are coordinates, and the range is the value at those coordinates. For a range image, each azimuth and elevation pair can only have one range value; therefore a problem occurs with multiple returns. There is not a one-to-one correspondence of spatial values to range values. To overcome this problem, one may choose the average range or the range corresponding to the return of highest intensity to be the range represented in the range image. In analog

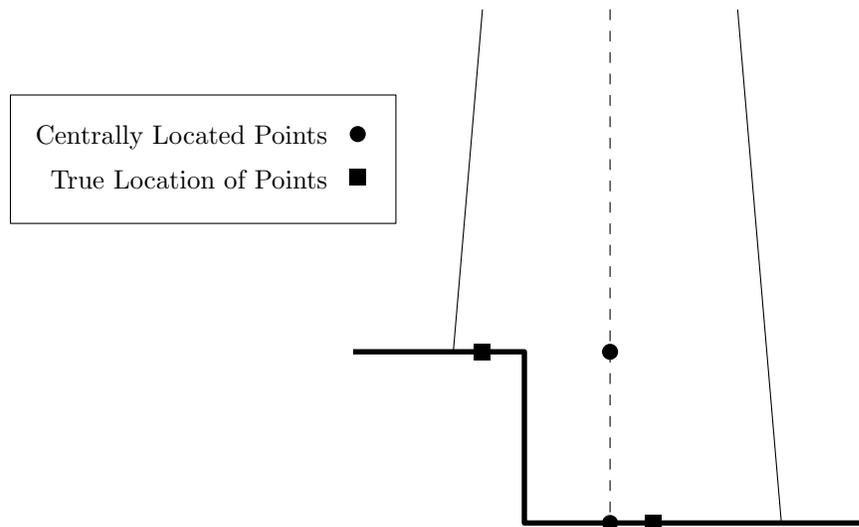


Fig. 5.1: Incorrect point estimation when the center of the footprint is assumed as the location of points.

ladar systems, it is also common to choose the first or last return.

The method used for angle estimation in section 4.3.4 can be built upon to obtain a solution to both of these problems. The angle estimation method assumes that the ladar shot hit a planar surface at an angle. Under this assumption, the locations of the peaks are solved for. Thus, each peak has a different value for azimuth, elevation, and range that can be used in the range image. This avoids the problem of having multiple range values for each coordinate in a range image, because the azimuth and elevation estimates of the peaks are not simply aligned down the center of the beam footprint. Therefore, the range image can be interpolated to obtain the underlying structure of the surface. It also makes it possible to correct the bad point estimation shown in Fig. 5.1.

Work in Chapter 4 assumed that the surface being targeted was known to be either two surfaces offset in range or an angled, planar surface. This may be useful in some applications. For example, in an ATR setting, a target of interest that is known to have an indentation at a particular location might be interrogated. It could be assumed that there is a slot at the particular location, and the slot could be measured for its width. If indeed there is a slot, the width could be used as a feature in the recognition. However, to form a point cloud of an unknown image, the prior information about the surface being

planar, discontinuous, or highly complex is not available. Therefore, it is not possible to distinguish between the surface response of a discontinuous surface and the surface response of an angled surface using the processing method used in prior chapters of this thesis.

For the method presented in this chapter, it is assumed that all surfaces can be approximated to be planar surfaces. Therefore, surfaces with multiple returns are treated as though the multiple returns are peaks in the surface response of a planar surface. Under this assumption, the precise location of the edge cannot be determined as was done in section 4.3.5 for the discontinuous surface. However, it removes the need to know the type of surface being targeted and provides a general method for processing the entire point cloud.

Sharp edges on a surface contain high spatial frequencies. The effect of assuming a shot on the edge of the target to be a shot on an angled surface is essentially a low-pass filtering operation. It smooths the edge of the surface. At some point in the processing, a low-pass filter must be applied if the surface is to be interpolated. Applying a low-pass filtering operation by the planar surface assumption preserves the edge better than postponing it until the after the point cloud is formed as depicted in Fig. 5.2. This figure shows an example of an interpolated surface under different methods for determining which point to use for the interpolation. The shot spacing is denoted by Δx . Figure 5.2 (a) shows that a surface cannot be interpolated because there are two points at one location. Figure 5.2 (b) shows the interpolated surface if the first return is chosen. This scheme over extends the interpolated surface. Figure 5.2 (c) shows the interpolated surface if the last return is chosen. This shows the interpolated surface cutting into the actual surface. Figure 5.2 (d) shows the interpolated surface treating the shot on the edge as if it were a shot on an angled surface processed by the method from section 4.3.4. This method has an advantage that both the first and last points can be used. This causes the edge to be steeper than any of the other methods and doesn't place the first return above the surface.

The planar assumption holds only for surfaces with spatial frequencies lower than the width of the footprint. It can misplace points otherwise. For example, a surface with a very narrow slot might contain energy from the laser on both sides of the slot. Energy from the

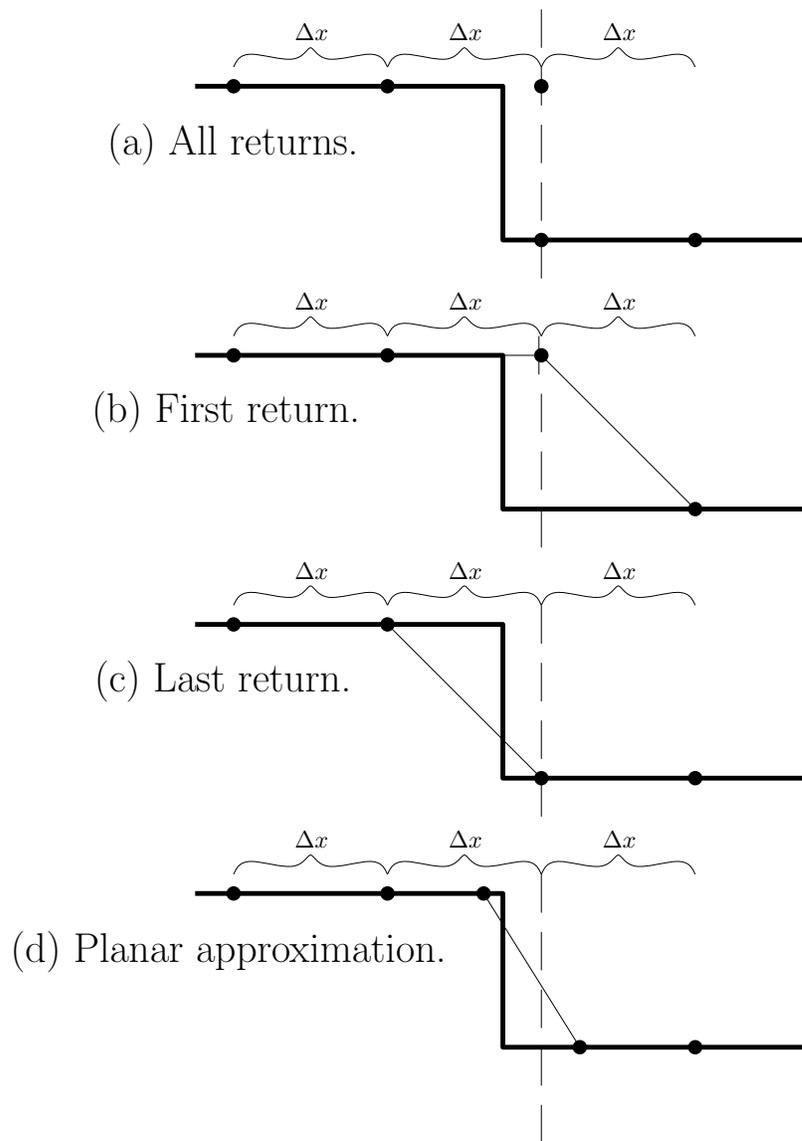


Fig. 5.2: Side view depiction of an interpolated discontinuous surface that contains a multiple return.

middle of the footprint would be inside the slot. Only two returns can be detected in this case. Using the method of section 4.3.4, spatial locations of those returns would be solved assuming only one boundary between them. In this situation, doing nothing to solve for the location of the points within the beam footprint would also misplace the points.

Although section 4.3.4 presented a method to determine the angle of incidence, there is no way to determine the orientation of the angled surface from a single ladar shot. There is still a degree of freedom that has not been estimated. Thus, the coordinates of neighboring points must be incorporated to determine the orientation of the surface. If two discontinuous surfaces are hit, it is likely that some neighboring points lie on both of the surfaces. The orientation of the plane can be determined by choosing the orientation that places the points from the current shot closest to the neighboring points while preserving the estimated angle of incidence. This is equivalent to choosing the orientation that results in the smoothest interpolated surface. For an example, Fig. 5.3 shows an edge of a surface. Neighboring points are shown on the surface as smaller points. The shot on the edge is shown to have two returns with a line indicating the angle of incidence. The first figure shows a good choice of surface orientation based on the locations of the neighboring points. The second figure shows a bad choice of surface orientation. Shown on the figures, the distance between the points in the first figure is less than the distance between the points in second figure.

5.1 Method for Locating Returns Within the Footprint

The method presented in this chapter relies on the following assumptions:

1. The range, intensity, and pointing direction of each return is known;
2. The range of the returns in the waveform does not affect the normalized intensity;
3. The coordinates of neighboring points are known;
4. The surface can be represented as if entire footprint lies on a single planar surface;
5. The surface has constant reflectivity and a Lambertian bi-directional reflectance distribution function (BRDF).

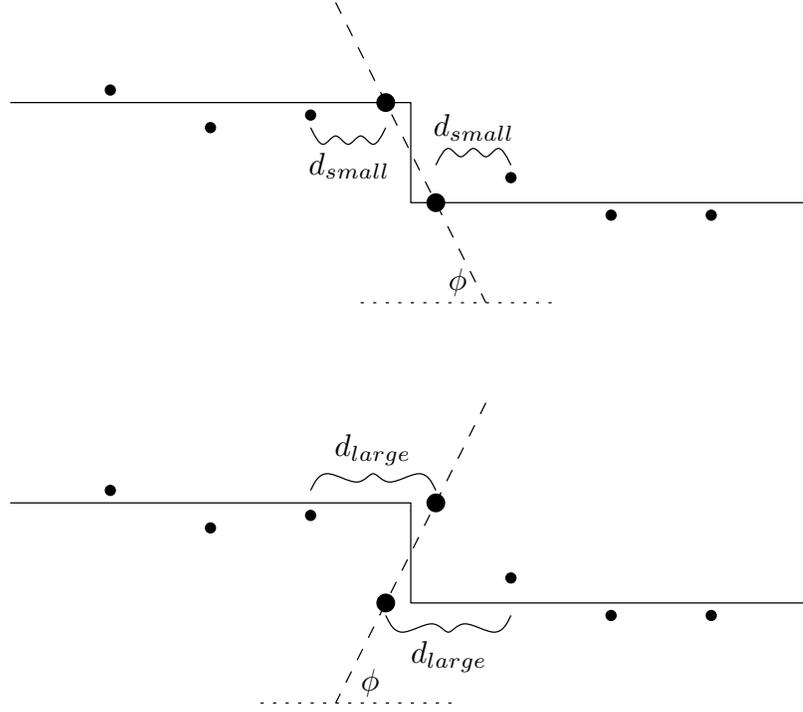


Fig. 5.3: Choosing the orientation of the surface to promote smooth surfaces.

Since the location of the estimates depend on the amount of energy in each return, assumption 2 states that the attenuation in the waveform due to range is the same for all returns. The intensity could be calibrated to eliminate this assumption by taking into account the attenuation of the pulse due to the change in range.

The first step in the proposed method is to compute the best fit plane of the neighboring points. A plane is defined by the dot product of a normal vector and points on the plane.

$$\mathbf{n} \cdot (\mathbf{v} - \mathbf{v}_0) = 0, \quad (5.1)$$

where $\mathbf{n} = n_x \hat{\mathbf{x}} + n_y \hat{\mathbf{y}} + n_z \hat{\mathbf{z}}$ is the normal vector, $\mathbf{v} = x \hat{\mathbf{x}} + y \hat{\mathbf{y}} + z \hat{\mathbf{z}}$ is any point in the plane, and $\mathbf{v}_0 = x_0 \hat{\mathbf{x}} + y_0 \hat{\mathbf{y}} + z_0 \hat{\mathbf{z}}$ is the position of some known point on the plane. The equation can be rewritten as

$$a_1(x - x_0) + a_2(y - y_0) = (z - z_0), \quad (5.2)$$

where $a_1 = -\frac{n_x}{n_z}$ and $a_2 = -\frac{n_y}{n_z}$.

This can be defined in a spherical coordinate system using ϕ from section 4.1.2 and θ where $0 \leq \theta < 2\pi$. Let $\hat{\mathbf{x}}$ be a unit vector in the azimuth direction, $\hat{\mathbf{y}}$ be a unit vector in the elevation direction, and $\hat{\mathbf{z}}$ be a unit vector in the direction of the range. Let ϕ be the rotation about $\hat{\mathbf{y}}$ and θ be the rotation about $\hat{\mathbf{z}}$. Then x_0 is the azimuth, y_0 is the elevation, and z_0 is the mean of the range weighted by the intensity for the current shot. Applying a rotation by ϕ followed by a rotation by θ to $\mathbf{n} = \hat{\mathbf{z}}$ gives

$$\begin{aligned} \mathbf{n} &= \begin{bmatrix} \cos(\theta) & -\sin(\theta) & 0 \\ \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \cos(\phi) & 0 & -\sin(\phi) \\ 0 & 1 & 0 \\ \sin(\phi) & 0 & \cos(\phi) \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}, \\ &= \begin{bmatrix} -\cos(\theta) \sin(\phi) \\ -\sin(\theta) \sin(\phi) \\ \cos(\phi) \end{bmatrix}, \end{aligned} \tag{5.3}$$

which leads to

$$\begin{aligned} a_1 &= \cos(\theta) \tan(\phi), \\ a_2 &= \sin(\theta) \tan(\phi). \end{aligned} \tag{5.4}$$

To solve for a_1 and a_2 , the only free variable is θ , because ϕ is determined from the method in section 4.3.4. In order to obtain θ , neighboring points must be known. Ideally, neighboring points would be given by using adjacent elements in the FPA. However, for a flying spot scanning ladar such as the ELT, this is not given. Furthermore, noise on the spatial measurements makes determining neighboring points even more difficult. A nearest neighbor search must be performed. In this thesis, the nearest neighbor search is accomplished using the K-D tree search method [32]. Figure 5.4 shows an example of the nearest neighbor selection. It is desirable to have neighboring points on all sides of the

current point for which θ is being solved. It is undesirable to use neighboring points that are far from the current shot, because the correlation of points decreases with distance. Another difficulty is the presence of multiple returns in the neighboring points. In the case of multiple returns in the neighboring points, the range is computed as the average range weighted by the intensity.

An algorithm that selects the appropriate neighbors to use in any situation for this application would be difficult to develop. It can also be said that the correct number of neighbors, M , to use may vary. For example, noise may cause the M neighbors to be concentrated on one side of the current shot. In this case, more neighbors would be needed to ensure that there are points on all sides of the current point. One solution to this problem is to use interpolation to obtain regularly spaced data. This can be done by interpolating eight points on a grid that surrounds the current shot. This is not interpolation for estimating the surface, but is simply for obtaining regularly spaced data so that the best fit plane can be easily obtained. The distance between points on the grid for the experiments was set to 2β as shown in Fig. 5.5.

The points on the grid are interpolated using inverse distance squared weighting [33]. The range at location (x_i, y_i) is given by

$$r(x_i, y_i) = \frac{\sum_{m=1}^M r(x_m, y_m) ((x_m - x_i)^2 + (y_m - y_i)^2)^{-1}}{\sum_{m=1}^M ((x_m - x_i)^2 + (y_m - y_i)^2)^{-1}}, \quad (5.5)$$

where (x_i, y_i) are the coordinates of the point to be interpolated, (x_m, y_m) are the coordinates of the neighboring points, and $r(x, y)$ is range at (x, y) .

Interpolating points on a grid provides easy data to work with. It also makes the choice of M simpler and more robust. If the value of M is set too high, the points that are far from the current shot will not be weighted as much as the shots that are closer. This is because of the inverse distance squared weighting. Thus M can be set to ensure that there are sufficient neighbors to be less susceptible to problems from having all M neighbors on one side of the point. The value of M depends on the shot spacing and should be selected to ensure that neighbors are selected on all sides of the current shot. If the shots overlap,

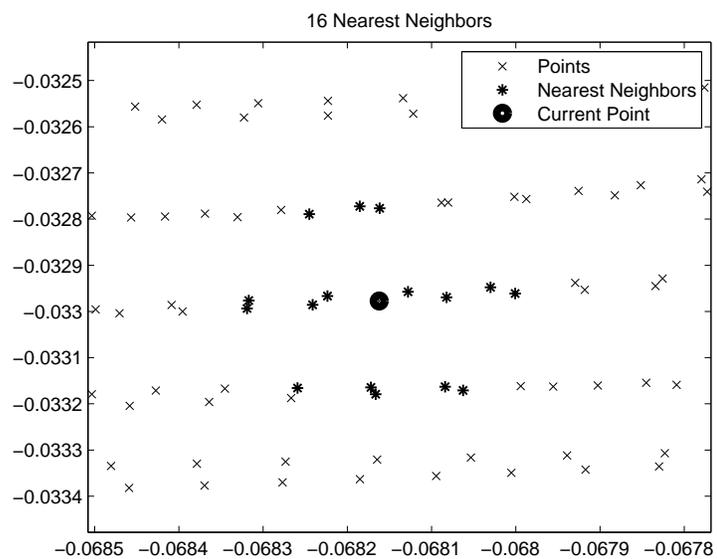


Fig. 5.4: Nearest neighbor selection.

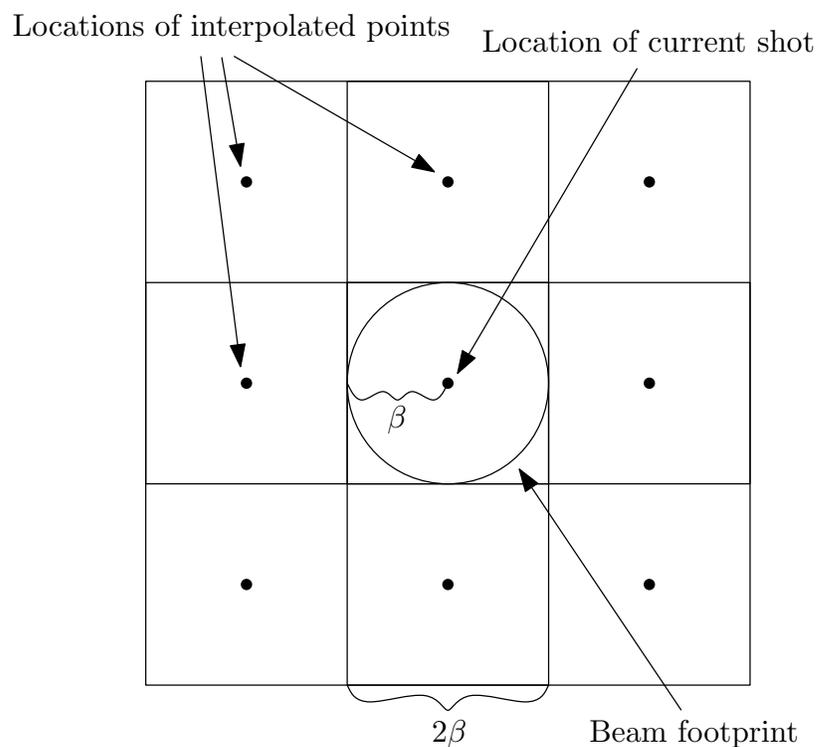


Fig. 5.5: Pixel spacing of interpolated points.

the value of M must be high. However, if the shots are adjacent, a smaller value can be used. For the experiments in this thesis, $M = 64$ was found to work well on dense point clouds, and $M = 16$ was found to work well on point clouds with a shot spacing set so that shots are roughly adjacent. This is more consistent with a shot spacing that would be used in an application such as ATR.

The points on the grid are used to derive a best fit plane to the data by minimizing the squared error. If \mathbf{az}_p , \mathbf{el}_p , and \mathbf{r}_p are vectors containing the azimuth, elevation, and range of the 8 points, then the error is given by

$$\begin{aligned} e &= \|X\mathbf{a} - \mathbf{r}_p\|_2^2, \\ &= \mathbf{a}^T X^T X \mathbf{a} - 2\mathbf{a}^T X^T \mathbf{r}_p + \mathbf{r}_p^T \mathbf{r}_p, \end{aligned} \quad (5.6)$$

where $X = [\mathbf{az}_p \ \mathbf{el}_p]$ and $\mathbf{a} = [a_1 \ a_2]^T$. The value of \mathbf{a} is constrained by (5.4) where ϕ is known. This constraint gives rise to a solution different than the traditional best fit plane solution. To minimize the error, the value of θ for which $\frac{de}{d\theta} = 0$ is computed. This is derived as follows.

$$\frac{de}{d\theta} = \frac{d\mathbf{a}^T}{d\theta} \frac{de}{d\mathbf{a}}, \quad (5.7)$$

$$\frac{d\mathbf{a}^T}{d\theta} = [-\sin(\theta) \ \cos(\theta)] \tan(\phi), \quad (5.8)$$

$$\frac{de}{d\mathbf{a}} = 2X^T X \mathbf{a} - 2X^T \mathbf{r}_p, \quad (5.9)$$

$$\begin{aligned} \frac{de}{d\theta} &= 2[-\sin(\theta) \ \cos(\theta)] \tan(\phi) X^T X \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix} \tan(\phi) \\ &\quad - 2[-\sin(\theta) \ \cos(\theta)] \tan(\phi) X^T \mathbf{r}_p = 0. \end{aligned} \quad (5.10)$$

This leads to solving for θ from

$$[-\sin(\theta) \quad \cos(\theta)]X^T \mathbf{r}_p = [-\sin(\theta) \quad \cos(\theta)] \tan(\phi) X^T X \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix}. \quad (5.11)$$

Because the terms are scalars, the equation can be rewritten using the trace operator as

$$\begin{aligned} [-\sin(\theta) \quad \cos(\theta)]X^T \mathbf{r}_p &= \text{tr} \left([-\sin(\theta) \quad \cos(\theta)] \tan(\phi) X^T X \begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix} \right), \\ &= \text{tr} \left(\begin{bmatrix} \cos(\theta) \\ \sin(\theta) \end{bmatrix} [-\sin(\theta) \quad \cos(\theta)] X^T X \right) \tan(\phi), \quad (5.12) \\ &= \text{tr} \left(\begin{bmatrix} -\cos(\theta) \sin(\theta) & \cos^2(\theta) \\ \sin^2(\theta) & \cos(\theta) \sin(\theta) \end{bmatrix} X^T X \right) \tan(\phi). \end{aligned}$$

This equation yields no simple analytical solution. However, the data to form X comes from the coordinates of the interpolated grid. Because of this, X is an orthogonal matrix, so $X^T X = 24\beta^2 I$, where I is a 2x2 identity matrix. This would not be possible if the noisy data were used directly. This simplification allows (5.12) to be written as

$$\begin{aligned} [-\sin(\theta) \quad \cos(\theta)]X^T \mathbf{r}_p &= \text{tr} \left(\begin{bmatrix} -\cos(\theta) \sin(\theta) & \cos^2(\theta) \\ \sin^2(\theta) & \cos(\theta) \sin(\theta) \end{bmatrix} \right) 24\beta^2 \tan(\phi), \\ &= 0. \quad (5.13) \end{aligned}$$

θ is solved from (5.13) to obtain

$$\theta = \tan^{-1} \left(\frac{b_2}{b_1} \right), \quad (5.14)$$

where $[b_1 \ b_2]^T = X^T \mathbf{r}_p$.

At this point, the best fit plane is known, and the points can be adjusted so that they lie on the best fit plane. Originally, the points all lie along $\hat{\mathbf{z}}$ as shown in Fig. 5.6. This illustrates a shot with three peaks in the surface response. One shot appears above the surface, one appears on the surface, and the other appears below the surface. Figure 5.6 (a) shows the side view of an angled surface to demonstrate the ranges of the points. Figure 5.6 (b) shows the top view to demonstrate the spatial locations of the points. All shots are originally located in the center of the beam footprint. Then, the locations of the points are adjusted via the method from section 4.3.4 as shown in Fig. 5.7. This result is equivalent to a rotation by ϕ about $\hat{\mathbf{y}}$. The angle ϕ is now set, but θ may be anywhere along the dotted circle shown in Fig. 5.7 (b). Next, the final locations of the points are determined by rotating the points to θ about $\hat{\mathbf{z}}$ as shown in Fig. 5.8.

5.2 Results of Locating Returns Within the Footprint

Results from section 4.3 quantified the ability to measure characteristics of the surface if the surface shape was known. However, it is more difficult to quantify the ability of the method from section 5.1, because the true surface is not known for the targets tested. For this reason, many of the results in this section are visual only.

LadarSIM is capable of generating waveforms of single selected points from a point cloud. However, it currently cannot produce waveform data for an entire point cloud. Since the method presented in section 5.1 requires the use of neighboring points, simulated data were not generated for the experiments in this section.

5.2.1 Slotted Surface

A point cloud using ELT data of the slotted surface from section 4.3.5 are shown in Fig. 5.9. This point cloud shows the same data that were used to obtain the statistics for section 4.3. The width of the slot of this target was constructed to be just larger than the width of the footprint. The points display the return number. Black represents a first return, and gray represents a second return. Figure 5.9 (a) shows the point cloud when the data are not processed to obtain the surface response. Because the surface response is

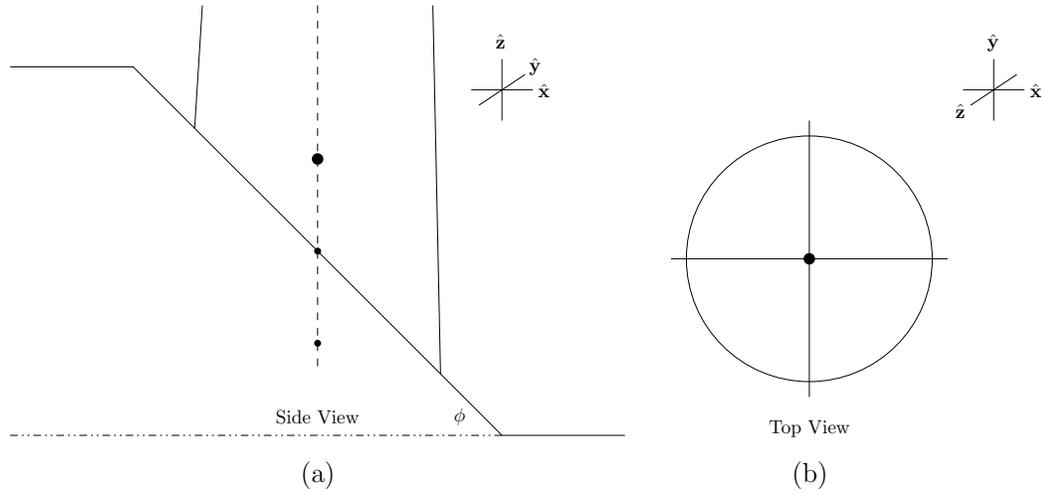


Fig. 5.6: Shot on planar surface before adjusting the locations of the points.

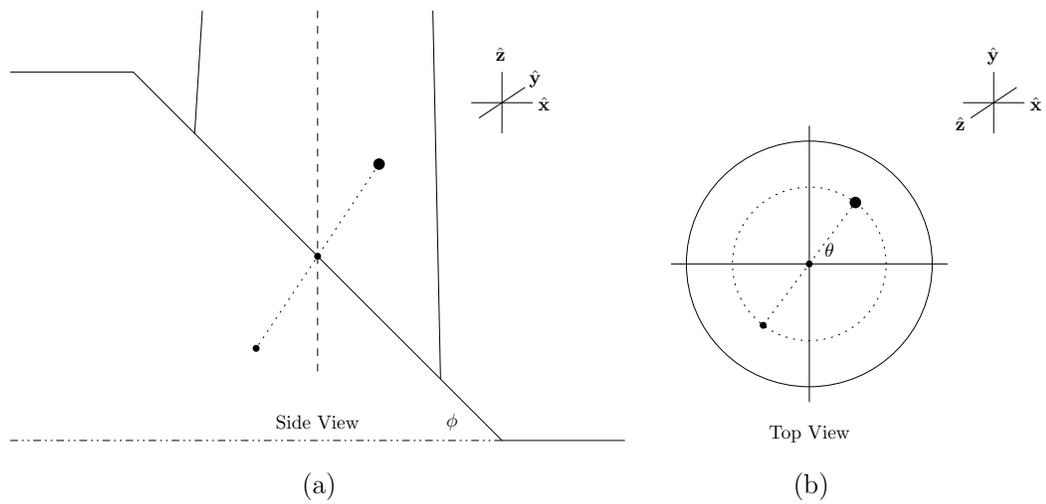


Fig. 5.7: Shot on planar surface after solving for ϕ .

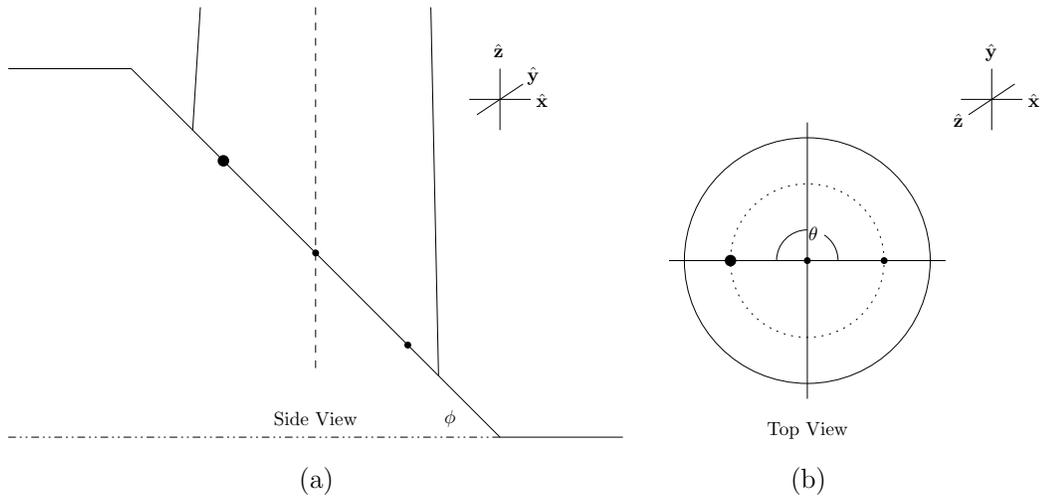


Fig. 5.8: Shot on planar surface after rotating by θ .

not computed, the waveform contains only one resolvable pulse. This causes the points to gradually transition to the second surface. Thus, the point cloud reports a sloped surface leading into the slot. Figure 5.9 (b) shows the point cloud when the surface response is computed and the coordinates of the returns are aligned down the center of the beam footprint. This point cloud appears to show two surfaces, one behind the other. Figure 5.9 (c) shows the point cloud when the surface response is computed and the new method is applied. This point cloud shows a sharper edge than the point cloud of Fig. 5.9 (a) and removes the effect seen in the point cloud of Fig. 5.9 (b).

In order to obtain a sufficient number of samples, the point clouds were obtained with a high point density. This provides insight to the quality of representing the slotted surface, but does not represent a practical scan. Figure 5.10 shows the scan of the slotted surface with points removed so that shots are barely overlapping. The points have been connected so that the side view of the interpolated surface can be visualized. Also an approximation to the true surface geometry has been drawn with dotted lines. Since the multiple returns in Fig. 5.10 (b) share the same azimuth and elevation values, the last return was chosen as the point to use in the interpolation.



(a) Processed without obtaining the surface response.

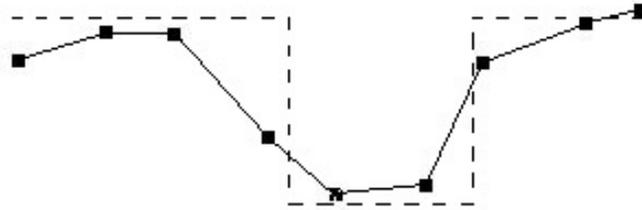


(b) Processed using the surface response and assuming the points to be located in the center of the footprint.

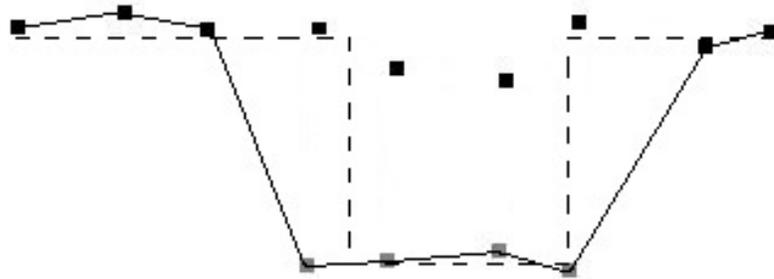


(c) Processed using the surface response and solving for the locations of the points within the beam footprint.

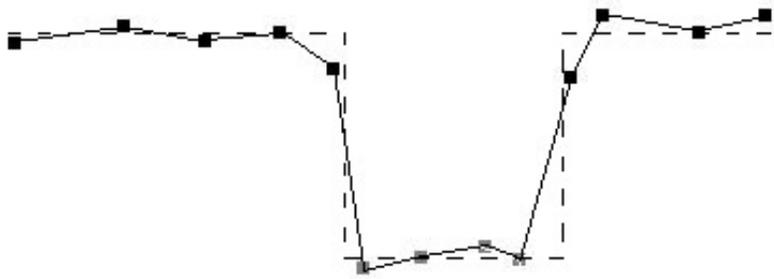
Fig. 5.9: Side view of ELT point cloud of slotted surface.



(a) Processed without obtaining the surface response.



(b) Processed using the surface response and assuming the points to be located in the center of the footprint. Interpolation is done on the last return.



(c) Processed using the surface response and solving for the locations of the points within the beam footprint.

Fig. 5.10: Side view of ELT point cloud of slotted surface with a practical shot spacing. The dotted line shows the true surface. The solid line shows the surface estimate by interpolation.

5.2.2 Angled Surface

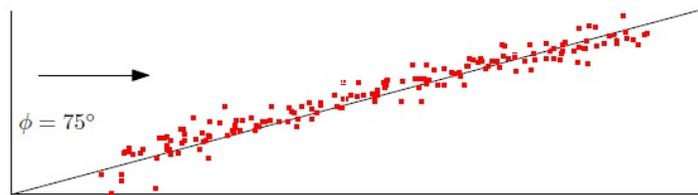
A point cloud using ELT data of the angled surface from section 4.3.4 are shown in Fig. 5.11. This is a top view of the angled surface. Two axes were drawn with a 75° line passing through the point cloud. The arrow indicates the line of sight. The target for these experiments is shown in Fig. 4.16. This point cloud shows the same data that were used to obtain the statistics for section 4.3. Figure 5.11 (a) shows the point cloud when the data are not processed to obtain the surface response. This point cloud shows that the angled surface is well represented, because the deviation of points from the plane is small. Figure 5.11 (b) shows the point cloud when the surface response is computed and the coordinates of the returns are aligned down the center of the beam footprint. Distinct return number levels are seen in this point cloud. This comes from the multiple peaks that exist in the surface response of the angled surface. Figure 5.11 (c) shows the point cloud when the surface response is computed and the new method is applied. The deviation of the points in this point cloud is greater than the point cloud of Fig. 5.11 (a), but less than the point cloud of Fig. 5.11 (b). This shows that adjusting the locations of the peaks within the beam footprint places the points closer to the surface than simply assuming the locations to be down the center of the beam footprint.

A best fit plane was computed for the points on the angled surface. This was taken directly from the point clouds in Fig. 5.11. The standard deviation of the points from the best fit plane was computed, and results are shown in Table 5.1.

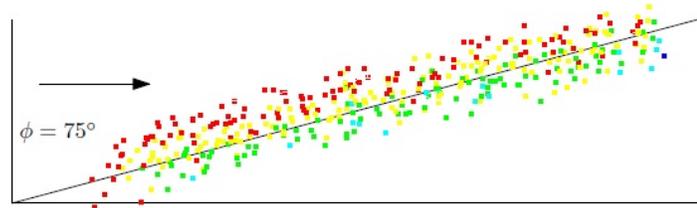
Results show that when the under-regularized surface response is computed on an angled surface, the error on the surface increases. However, by solving for the locations of the returns inside the footprint, this error is reduced almost to the original value. This indicates that the resolution can be increased with a small trade-off in the fidelity of representing

Table 5.1: Error of points from angled surface fitted to a plane.

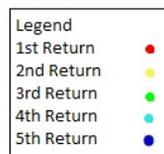
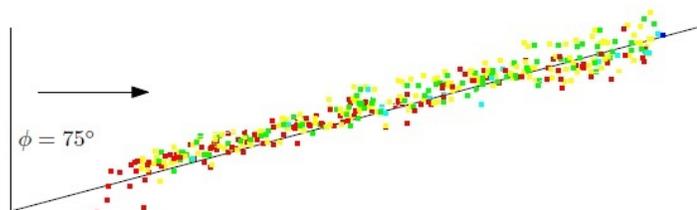
Data	Standard Deviation (cm)
VISSTA ELT (no surface response)	2.37
VISSTA ELT (points located in center of footprint)	4.02
VISSTA ELT (points located using new method)	2.39



(a) Processed without obtaining the surface response.



(b) Processed using the surface response and assuming the points to be located in the center of the footprint.



(c) Processed using the surface response and solving for the locations of the points within the beam footprint.

Fig. 5.11: Top view of ELT point cloud of angled surface.

angled surfaces.

An estimate of the angle ϕ was given in Table 4.2, but an estimate of θ was not given. The same data used in section 4.3.4 were used to estimate θ . To provide more accurate test conditions that reflect a shot spacing with adjacent, non-overlapping points, points were removed from the pointcloud. Every 5th row and column were used to create an approximate shot spacing of 80 microradians. At the time of data collection, the angle θ was not considered. All tests were conducted with the true value of $\theta \approx 0^\circ$. To show that the method can estimate θ over the span of $[0^\circ, 360^\circ)$, the data of the flat surface were rotated by various angles unknown to the algorithm. Variation in the results comes from interpolating the data onto the grid for obtaining the best fit plane. Points were selected from the center region of the test board and θ was estimated. The mean estimate and standard deviation were computed. Results are shown with $\phi = 60^\circ$ in Fig. 5.12 and with $\phi = 75^\circ$ in Fig. 5.13. The mean estimate is plotted with error bars showing the standard deviation of the estimate.

From Fig. 5.12 and Fig. 5.13 the error on the estimate of θ is on the order of 5° to 8° . Errors in the estimation of θ may come from noise in the scanner. This noise can cause variation in the position of neighboring points and affect the best fit plane from which θ is derived. If a system with little error in the pointing control were used, the amount of error on θ may be less than the error in this system.

5.2.3 Complex Objects

Sections 5.2.1 and 5.2.2 show point clouds of the two surfaces used throughout this thesis. This section shows the effect of the method on point clouds of other objects that may appear in a scene.

Figure 5.14 shows a scan of the test board used in experiments from section 5.2.2. The red points show first returns and the blue points show other returns. The strap on the left attaches to the test board at the top. Figure 5.15 shows a closer view of the test board. Figure 5.15 (a) shows the point cloud when the data are not processed to obtain the surface response. Figure 5.15 (b) shows the point cloud when the surface response is computed and

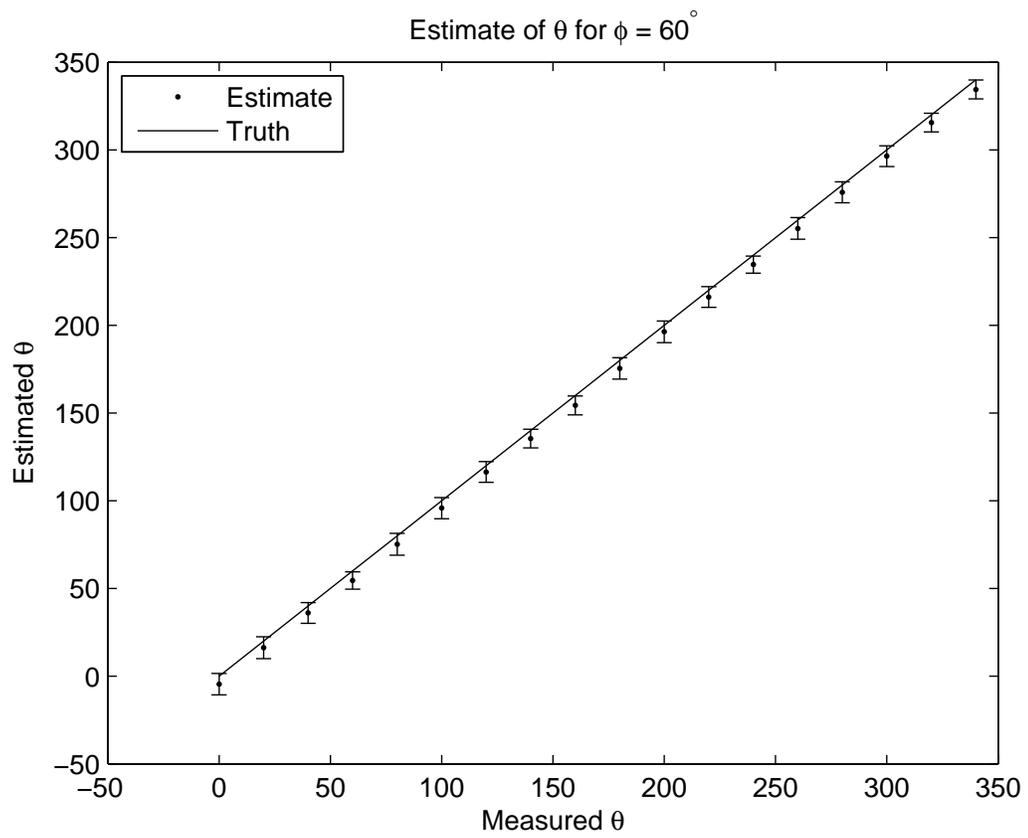


Fig. 5.12: Estimate of θ for $\phi = 60^\circ$.

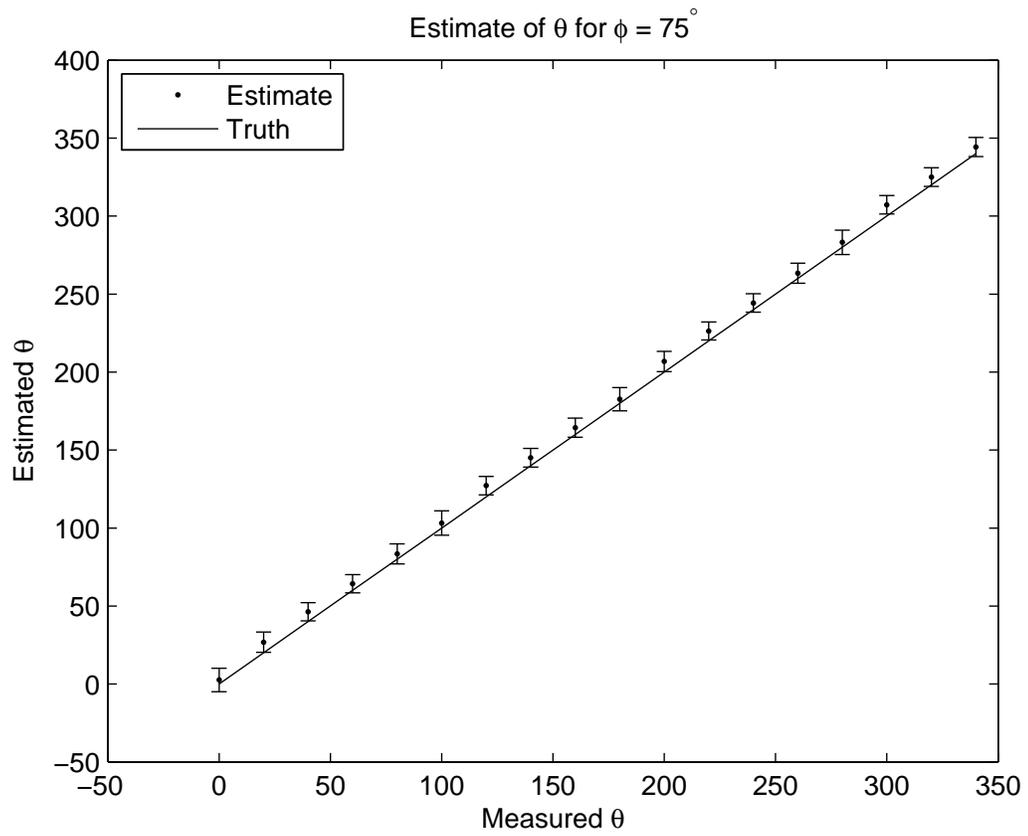


Fig. 5.13: Estimate of θ for $\phi = 75^\circ$.

the coordinates of the returns are aligned down the center of the beam footprint. Figure 5.15 (c) shows the point cloud when the surface response is computed and the new method is applied. When the surface response is not computed, there are second returns that are missed towards the top of the strap. Another detail to notice is that the strap is slightly thicker in Fig. 5.15 (b) than it is in the other point clouds. This result supports the drawing in Fig. 5.1.

Figure 5.16 shows a large cube that was an object in a scan. Two smaller cubes rest on opposite corners on top of the large cube. The angle of incidence on one side of the cube was high ($\approx 85^\circ$). Figure 5.17 shows the top view of the cube. The arrow in the figures indicates the line of sight of the ELT. Figure 5.17 (a) shows the least amount of variation on the surface. Figure 5.17 (b) shows the most variation on the surface and Figure 5.17 (c) shows an improvement upon Figure 5.17 (b).

A building that contains an I-beam was scanned and shown in Fig. 5.18. A close-up view is shown in Fig. 5.19. The arrow indicates the location of the I-beam. Color indicates return number. Red shows the first return, and blue shows other returns. Figure 5.19 (a) shows multiple returns at the top of the I-beam. However, the second surface of the I-beam is undetected at the bottom. This is because the distance between the two surfaces of the I-beam decreases. Figure 5.19 (b) shows that multiple returns are detected all along the I-beam when deconvolution is applied. Figure 5.19 (c) shows that when the new method is applied, the edge of the I-beam is narrower than the edge of the I-beam in Fig. 5.19 (b). This is due to solving for the locations of the returns within the beam footprint. Again, this result supports the drawing in Fig. 5.1.

Another detail that can be observed in this point cloud is the fence along the bottom. The fence is barely detected in Fig. 5.19 (a). This is due to the threshold setting for the discrimination method. The point clouds were created according to a threshold level that provided the best results. The point clouds that were created from the surface response used a lower threshold setting. The reason that the point cloud in Fig. 5.19 (a) used a higher threshold setting can be seen in Fig. 5.20. Ringing in the waveform can cause

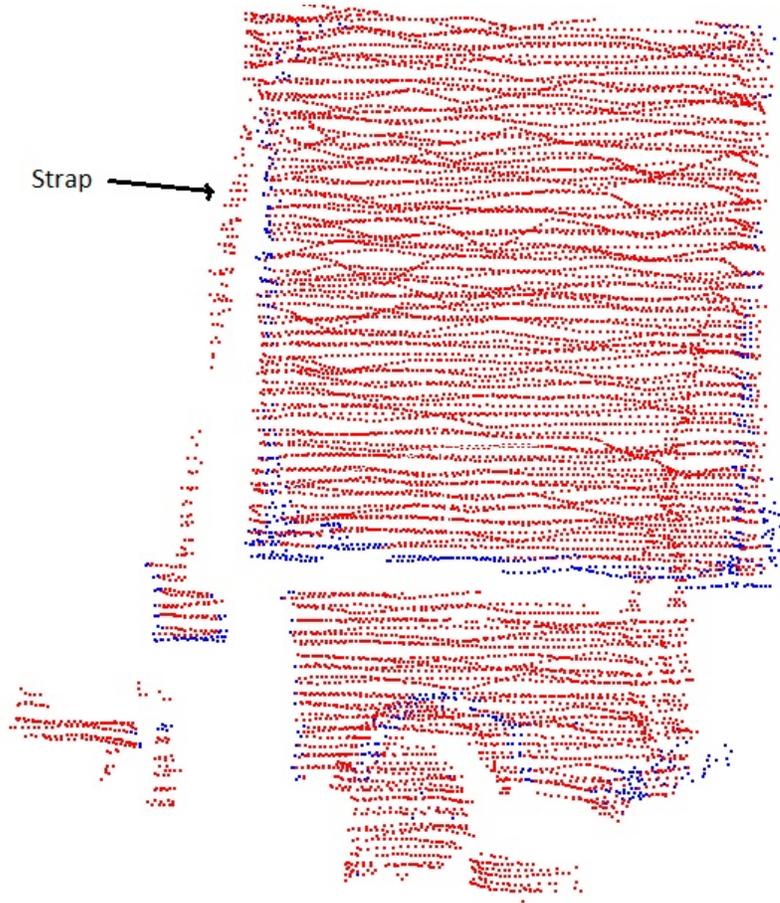
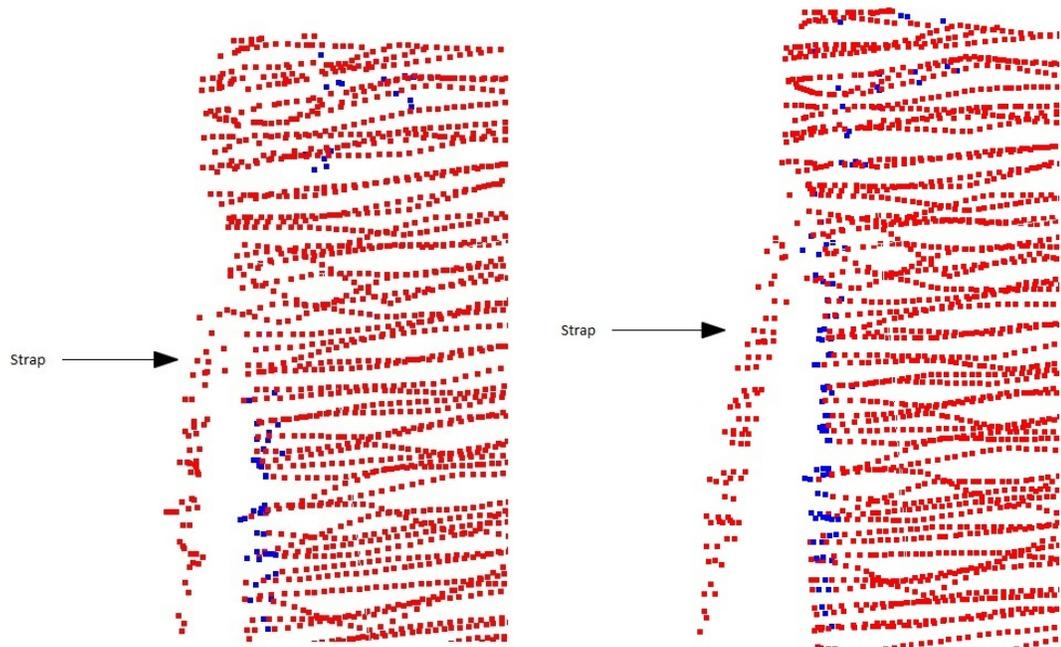
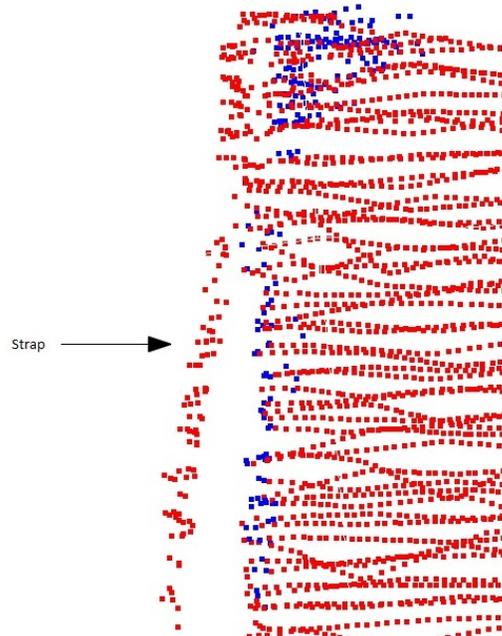


Fig. 5.14: Testboard scanned by the ELT.



(a) Processed without obtaining the surface response.

(b) Processed using the surface response and assuming the points to be located in the center of the footprint.



(c) Processed using the surface response and solving for the locations of the points within the beam footprint.

Fig. 5.15: Strap on the testboard. First returns are shown in red. Other returns are shown in blue.

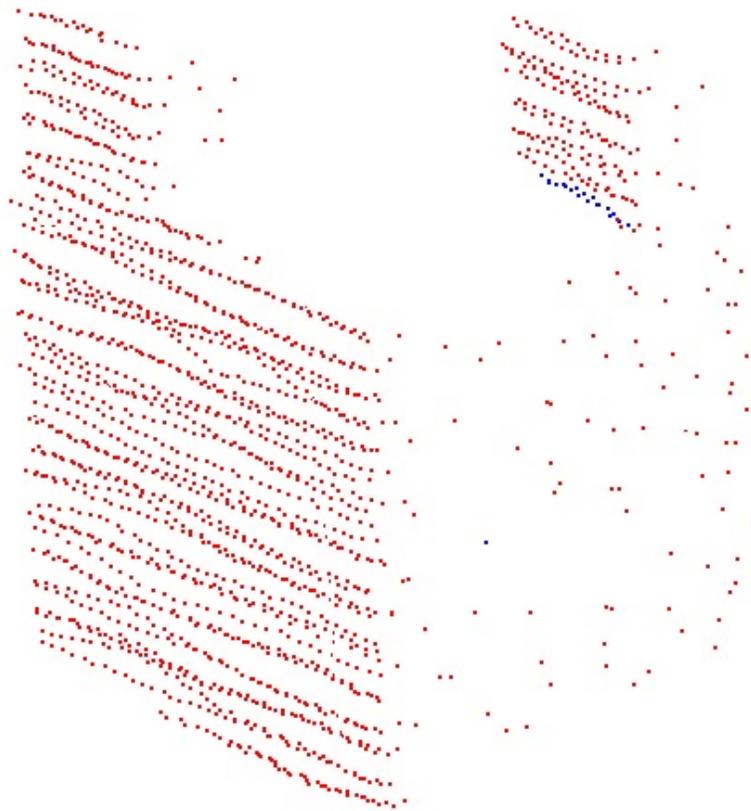
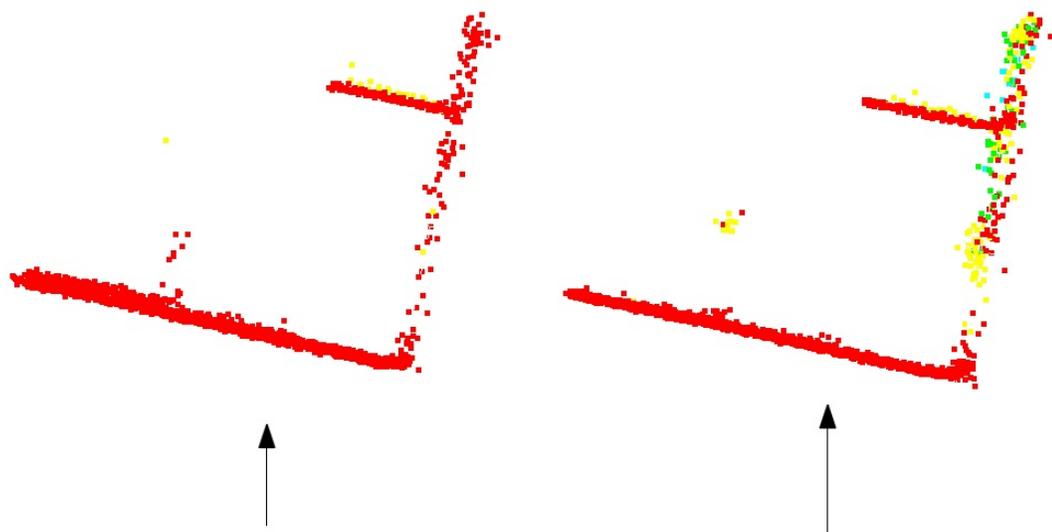
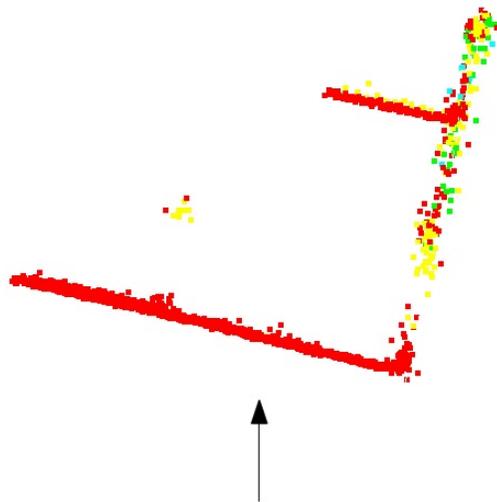


Fig. 5.16: Cube scanned by the ELT.



(a) Processed without obtaining the surface response.

(b) Processed using the surface response and assuming the points to be located in the center of the footprint.



(c) Processed using the surface response and solving for the locations of the points within the beam footprint.

Fig. 5.17: Top view of cube scanned by the ELT.

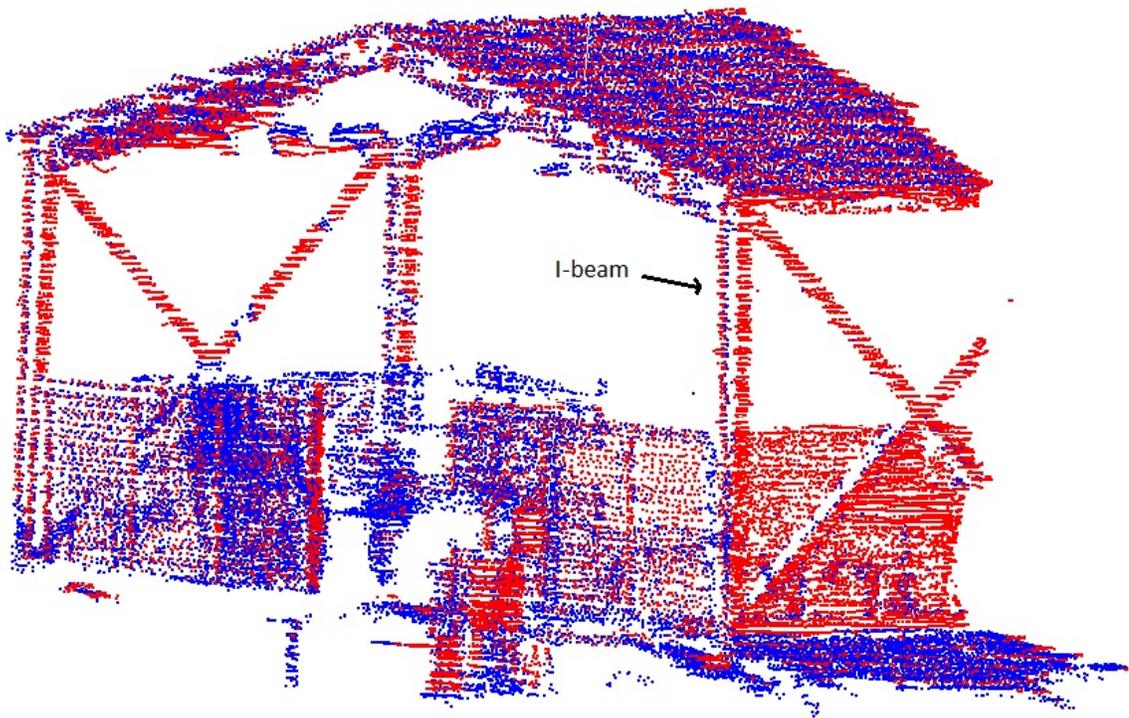
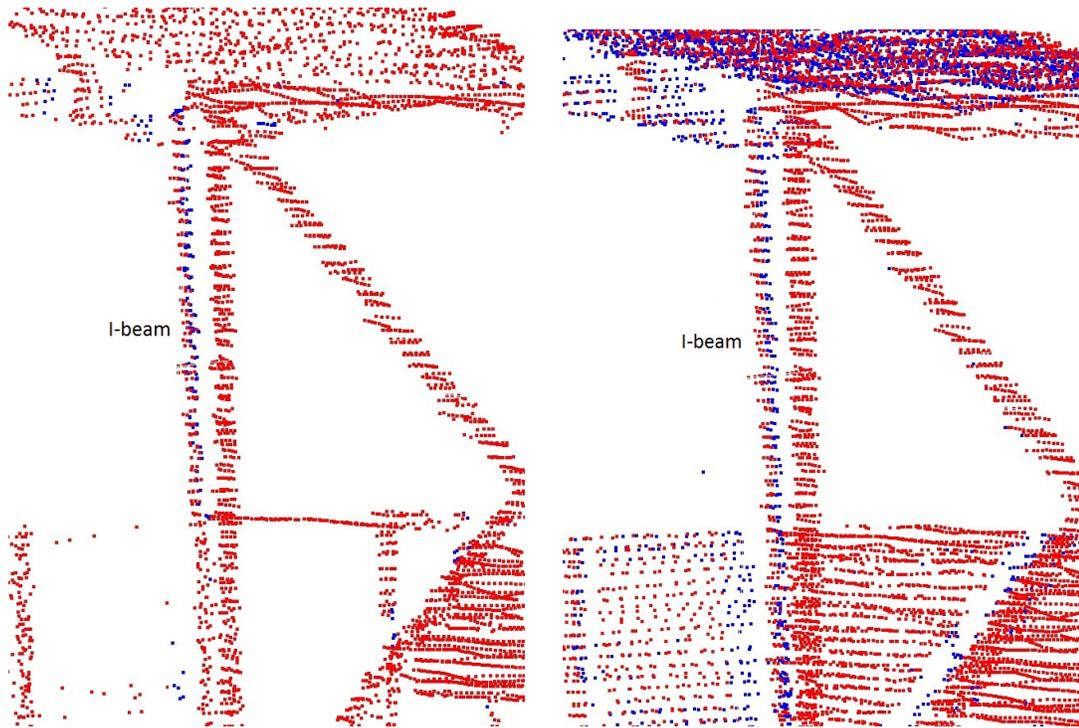


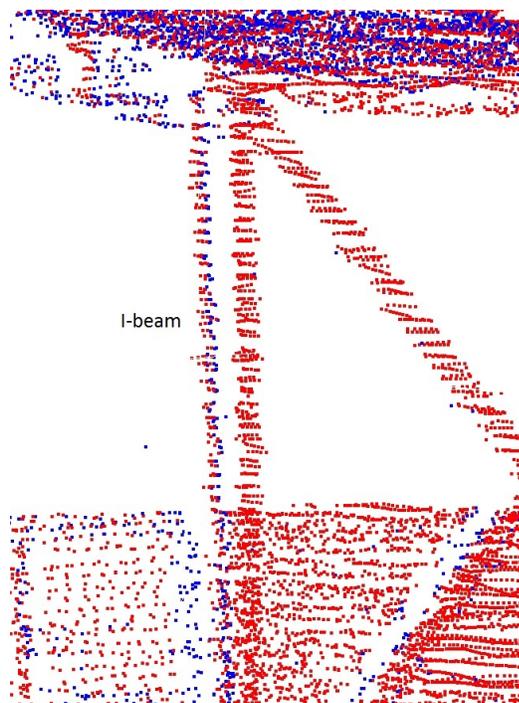
Fig. 5.18: Scan of shed with an I-beam.

multiple detections if the threshold is set too low. Detections from a ringing pulse are not desirable because they create double images in a point cloud. Figure 5.21 shows the surface response of this pulse. The ringing was removed when the waveform was deconvolved to obtain the surface response. In this point cloud, obtaining the surface response enabled a lower threshold setting to be used. This increased the ability to detect weak returns, thus providing more detail in the point cloud.



(a) Processed without obtaining the surface response.

(b) Processed using the surface response and assuming the points to be located in the center of the footprint.



(c) Processed using the surface response and solving for the locations of the points within the beam footprint.

Fig. 5.19: I-beam scanned by the ELT.

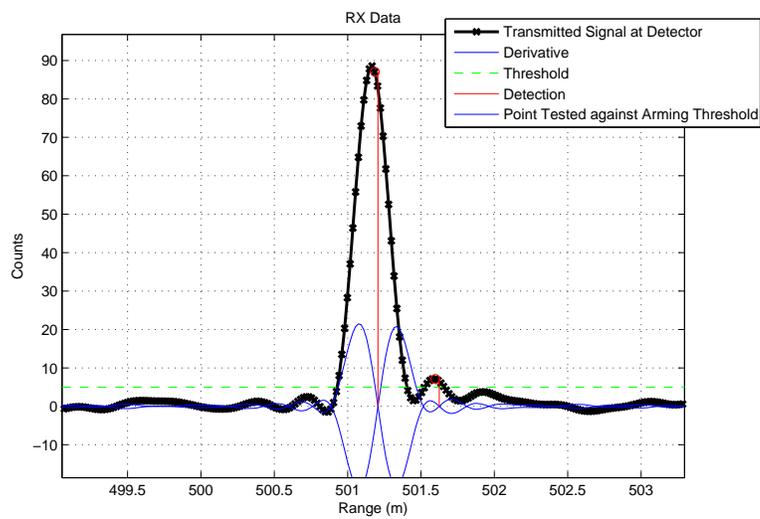


Fig. 5.20: ELT waveform with a low detection threshold.

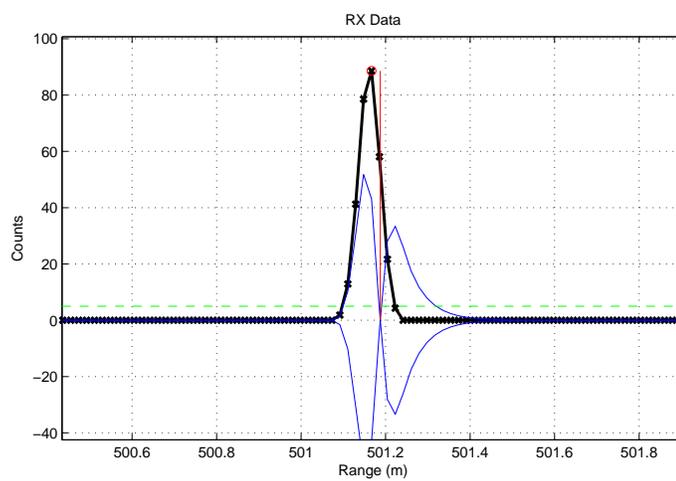


Fig. 5.21: Surface response of ELT waveform with a low detection threshold.

Chapter 6

Conclusion

6.1 Summary

A technique that exploited the digitized waveform and neighboring points in the point cloud to enhance the resolution on the edges of surfaces was developed in this thesis. The problem of representing a highly sloped surface with a single point was solved by forming an estimate for the angle of incidence. Treatment of multiple returns has been simplified by solving for the locations of the returns within the beam footprint. The advantage of the method is the ability to increase the number of points on the edges of targets and accurately place them within the footprint. This has direct application to ATR scenarios by increasing the number of points on the target.

Chapter 3 defined the signal processing model and strategy to obtain the surface response. Obtaining the surface response is not new to lidar processing; however, the VISSTA ELT is a practical lidar with a waveform sample rate low enough to sustain the high data rates required for forming a 3-D point cloud.

Chapter 4 derived the surface response in general as well as the surface response of an angled surface and two offset surfaces. These surface responses were used to compare the surface response estimate from the ELT and LadarSIM data using three methods of deconvolution. The preferred method based on range error and resolution was the NNLS method. By under-regularizing the surface response at this stage in the processing, the range resolution increased, and a method was formed that allows the angle of incidence to be estimated.

Described in Chapter 5, the needed regularization occurs after all points have been processed. At this stage in the processing, neighboring points are used to obtain an estimate

of the locations of each point within the footprint. Because the information is gleaned from neighboring points, this is a low-pass filtering or regularizing operation. To obtain the surface from a point cloud, points must be connected. Assuming all surfaces to be planar allows sharper edges due to the increase of points as shown in Fig. 5.2 and the results from section 5.2. The increase in points only occurred at edges of targets and high-angled slopes. It is at these areas where increased resolution is needed. Other locations can be well represented with low resolution point clouds. Using the method to obtain better estimates of the locations of the returns within the beam footprints allows the user of the data to obtain a better interpolation of the surface using the interpolation method of their choice.

6.2 Further Research

This thesis provided a small step toward a better representation of edges on complex targets. Assumptions from section 5.1 made this problem tractable. The first step would be to eliminate assumption 2 by calibrating the intensity in each return based on the range of the return. Second, assumption 5 should be removed. This assumption may be difficult to handle, but is very important for ATR applications. If a camouflaged target were hit, the reflectivity of the surface under different areas of the footprint would change significantly. The developed method relies heavily on knowing the amount of energy contained in each portion of the footprint. Camouflaged surfaces would be poorly represented with the current method. It is possible that neighboring points could be used to obtain a reflectivity estimate for different areas under the footprint.

Studying more types of surfaces would strengthen the processing method presented in this thesis for both the surface response estimation and solving for the locations of returns in the beam footprint. Results in this thesis include only angled and discontinuous surfaces; however, these surfaces do cover the extremes in surface complexities. The planar surface is represented with low spatial frequencies, and the discontinuous surface is composed of high spatial frequencies. Although the discontinuous surface was shown to be well represented by a plane at a large angle, a better approach might be possible by studying more surfaces such as cylindrical surfaces, round surfaces, and surfaces on the corners of objects.

It was shown in this thesis that the angle of incidence can be estimated and two close surfaces can be resolved if it is known *a priori* which surface is hit. Section 4.2 discussed that only information within the system bandwidth can be recovered unambiguously. Studying range resolution with bandwidth and noise power as variables would be beneficial. This could enable range resolution to be used as a design specification. It could also provide those who use ladar for applications such as ATR to know the resolving capabilities of a ladar.

Chapter 5 presented a method to process the point cloud without the knowledge of the surface hit by the beam footprint. This was done by assuming that any surface response with multiple returns came from a planar surface. This was used to sharpen the edges of the underlying interpolated surface; however, a measurement of the location of an edge cannot be taken due to the planar assumption. Neighboring points could be used to determine the type of surface being hit. For example, if the ranges of proximate shots were binned at two locations, it is highly probable that a shot in the center of the neighborhood with two returns does not represent an angled surface, but rather a discontinuous surface. The study of more surfaces would allow further research to process a point cloud in this manner.

References

- [1] F. Bretar, A. Chauve, C. Mallet, and B. Jutzi, "Managing full waveform lidar data: A challenging task for the forthcoming years," *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 37, pp. 415–420, 2008.
- [2] J. Jung and M. M. Crawford, "A two-stage approach for decomposition of ICESat waveforms," in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, vol. 3, pp. 680–683, July 2008.
- [3] M. A. Hofton, J. B. Minister, and J. B. Blair, "Decomposition of laser altimeter waveforms," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 38, pp. 1989–1996, 2000.
- [4] W. Wagner, A. Ullrich, V. Ducic, T. Melzer, and N. Studnicka, "Gaussian decomposition and calibration of a novel small-footprint full-waveform digitising airborne laser scanner," *ISPRS Journal of Photogrammetry & Remote Sensing*, vol. 60, pp. 100–112, 2006.
- [5] A. Neuenschwander, L. Magruder, and R. Gutierrez, "Signal processing techniques for feature extraction and classification using small-footprint full-waveform airborne LIDAR," in *Proceedings of IEEE International Geoscience and Remote Sensing Symposium*, vol. 3, pp. 676–679. IEEE, July 2008.
- [6] A. Chauve, C. Mallet, F. Bretar, S. Durrieu, M. Deseilligny, and W. Puech, "Processing full-waveform lidar data: Modelling raw signals," *ISPRS Workshop on Laser Scanning*, 2007.
- [7] C. Mallet and F. Bretar, "Full-waveform topographic lidar: State-of-the-art," *ISPRS Journal of Photogrammetry & Remote Sensing*, vol. 64, no. 1, pp. 1 – 16, Jan. 2009.
- [8] N. Vlassis and A. Likas, "A greedy EM algorithm for Gaussian mixture learning," *Neural Processing Letters*, vol. 15, pp. 77–87, 2002.
- [9] S. Hernandez-Marin, A. M. Wallace, and G. J. Gibson, "Bayesian analysis of lidar signals with multiple returns," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 29, no. 12, pp. 2170–2180, Dec. 2007.
- [10] C. Mallet, F. Lafarge, F. Bretar, U. Soergel, and C. Heipke, "Lidar waveform modeling using a marked point process," in *ICIP'09: Proceedings of the 16th IEEE international conference on Image processing*, pp. 1693–1696. Piscataway, NJ, USA: IEEE Press, 2009.
- [11] B. Jutzi and U. Stilla, "Range determination with waveform recording laser systems using a Wiener Filter," *ISPRS Journal of Photogrammetry & Remote Sensing*, vol. 61, no. 2, pp. 95 – 107, 2006 [Online]. Available: <http://www.sciencedirect.com/science/article/B6VF4-4M3RP2G-1/2/43f1af56ebe69683f1cf814354e61239>.

- [12] T. K. Moon and W. C. Stirling, *Mathematical Methods and Algorithms for Signal Processing*, ch. 17, pp. 717–743. Upper Saddle River, NJ 07458: Prentice Hall, 2000.
- [13] S. Harsdorf and R. Reuter, “Stable deconvolution of noisy lidar signals,” in *Proceedings of EARSeL-SIG-Workshop LIDAR*, June 2000.
- [14] W. H. Richardson, “Bayesian-based iterative method of image restoration,” *Journal of the Optical Society of America*, vol. 62, no. 1, pp. 55–59, Jan. 1972.
- [15] L. A. Magruder and A. L. Neuenschwander, “Lidar full-waveform data analysis for detection of faint returns through obscurants,” M. D. Turner and G. W. Kamerman, Eds., vol. 7323, no. 1, p. 73230L. SPIE, 2009 [Online]. Available: <http://link.aip.org/link/?PSI/7323/73230L/1>.
- [16] U. Stilla, W. Yao, and B. Jutzi, “Detection of weak laser pulses by full waveform stacking,” *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 36, no. 3, pp. 25–30, 2007.
- [17] M. Kirchhof, B. Jutzi, and U. Stilla, “Iterative processing of laser scanning data by full waveform analysis,” *ISPRS Journal of Photogrammetry & Remote Sensing*, vol. 63, pp. 99–114, 2008.
- [18] B. Jutzi, J. Neulist, and U. Stilla, “Sub-pixel edge localization based on laser waveform analysis,” in *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 36(Part /W8), pp. 109–114. JSPRS, 2005.
- [19] A. Kilpela, *Pulsed Time-of-Flight Laser Range Finder Techniques for Fast, High Precision Measurement Applications*. Ph.D. dissertation, University of Oulu, Finland, 2004.
- [20] B. Jutzi and U. Stilla, “Laser pulse analysis for reconstruction and classification of urban objects,” *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 34, pp. 151–156, 2003.
- [21] K. D. Neilsen, S. E. Budge, R. T. Pack, R. R. Fullmer, and T. D. Cook, “Design and validation of the eyesafe LADAR test-bed (ELT) using the LadarSIM system simulator,” in *Laser Radar Technology and Applications XIV*, M. D. Turner and G. W. Kamerman, Eds., vol. 7323, p. 73230B. SPIE, May 2009 [Online]. Available: <http://link.aip.org/link/?PSI/7323/73230B/1>.
- [22] R. T. Pack, P. Israelsen, and K. Sealy, “A co-boresighted synchronized ladar/EO imager for creating 3D images of dynamic scenes,” in *Laser Radar Technology and Applications X*, G. W. Kamerman, Ed., vol. 5791, pp. 42–50. SPIE, May 2005.
- [23] J. G. Proakis and D. G. Monolakis, *Digital Signal Processing*, ch. 11, pp. 760–762. Upper Saddle River, NJ 07458: Prentice Hall, 2007.
- [24] S. E. Budge, B. C. Leishman, and R. T. Pack, “Simulation and modeling of return waveforms from a ladar beam footprint in USU LadarSIM,” in *Laser Radar Technology and Applications XI*, G. W. Kamerman and M. D. Turner, Eds., vol. 6214, p. 62140N. Orlando (Kissimmee), FL, USA: SPIE, Apr. 2006 [Online]. Available: <http://link.aip.org/link/?PSI/6214/62140N/1>.

- [25] R. T. Pack, R. R. Fullmer, S. E. Budge, *et al.*, *USU LadarSIM Release 3.0 User's Manual*, Utah State University/Space Dynamics Laboratory, Logan, UT, Mar. 2007.
- [26] P. P. Vaidyanathan, *Multirate systems and filter banks*. Upper Saddle River, NJ, USA: Prentice-Hall, Inc., 1993.
- [27] C. Lawson and R. Hanson, *Solving Least Squares Problems*, ch. 23, pp. 158–165. Englewood Cliffs, NJ: Prentice Hall, 1974.
- [28] M. Bertero and P. Boccacci, *Introduction to Inverse Problems in Imaging*. South Independence Mall West, Philadelphia, PA 19106: Institute of Physics Publishing, 1998.
- [29] A. Neumaier, “Solving ill-conditioned and singular linear systems: A tutorial on regularization,” 1998 [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.53.96>.
- [30] T. Ruotsalainen, P. Palojarvi, and J. Kostamovaara, “A wide dynamic range receiver channel for a pulsed time-of-flight laser radar,” *IEEE Journal of Solid-State Circuits*, vol. 36, no. 8, pp. 1228–1238, Aug. 2001.
- [31] K. D. Neilsen, S. E. Budge, and R. T. Pack, “Signal processing on waveform data from the eyesafe ladar testbed (ELT),” in *Laser Radar Technology and Applications XV*, M. D. Turner and G. W. Kamerman, Eds., vol. 7684, p. 76840R. Orlando, Florida, USA: SPIE, Apr. 2010 [Online]. Available: <http://link.aip.org/link/?PSI/7684/76840R/1>.
- [32] J. H. Friedman, J. L. Bentley, and R. A. Finkel, “An algorithm for finding best matches in logarithmic expected time,” *ACM Trans. Math. Softw.*, vol. 3, no. 3, pp. 209–226, 1977.
- [33] D. Shepard, “A two-dimensional interpolation function for irregularly-spaced data,” in *ACM '68: Proceedings of the 1968 23rd ACM national conference*, pp. 517–524. New York, NY, USA: ACM, 1968.