November, 2013

Enhanced resolution edge and surface estimation from ladar point clouds containing multiple return data

Kevin D Neilsen
Scott E Budge

Optical Engineering

SPIEDigitalLibrary.org/oe

Enhanced resolution edge and surface estimation from ladar point clouds containing multiple return data

Kevin D. Neilsen
Scott E. Budge

SPIE

Enhanced resolution edge and surface estimation from ladar point clouds containing multiple return data

Kevin D. Neilsen
Scott E. Budge
Utah State University
Center for Advanced Imaging Ladar
Department of Electrical and Computer Engineering
4120 Old Main Hill
Logan, Utah 84322-4120
E-mail: scott.budge@ece.usu.edu.

Abstract. Signal processing enables the detection of more returns in a digital ladar waveform by computing the surface response. Prior work has shown that obtaining the surface response can improve the range resolution by a factor of 2. However, this advantage presents a problem when forming a range image—each ladar shot crossing an edge contains multiple values. To exploit this information, the location of each return inside the spatial beam footprint is estimated by dividing the footprint into sections that correspond to each return and assigning the coordinates of the return to the centroid of the region. Increased resolution results on the edges of targets where multiple returns occur. Experiments focus on angled and slotted surfaces for both simulated and real data. Results show that the angle of incidence on a 75-deg surface is computed only using a single waveform with an error of 1.4 deg and that the width of a 19-cm-wide by 16-cm-deep slot is estimated with an error of 3.4 cm using real data. Point clouds show that the edges of the slotted surface are sharpened. These results can be used to improve features extracted from objects for applications such as automatic target recognition. [DOI: 10.1117/1.OE.52.11.113103]

Subject terms: ladar; lidar; waveform processing; super-resolution; multiple return; footprint; point cloud; surface response.

Paper 130867 received Jun. 12, 2013; revised manuscript received Sep. 20, 2013; accepted for publication Oct. 17, 2013; published online Nov. 8, 2013.

1 Introduction

Digitized ladar waveforms provide an increased ability to detect multiple returns or echos. This is not only accomplished by the ability to record and store the waveform but also by means of postprocessing to obtain the surface response. The surface response is obtained by deconvolving the transmitted pulse out of the received waveform and results in an improvement in range resolution.[1,2] It was shown that for the ladar used in this work, obtaining the surface response increases the range resolution from 30 to 14 cm.[2]

When the ladar beam falls on two or more surfaces, more than one returning pulse can be detected. It is common to assume 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. 1 shows that this assumption can misplace the point on the first surface. In applications such as automatic target recognition (ATR), good results rely on precise measurements of target details and features. This shot misplacement results 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 to create a three-dimensional (3-D) surface. In applications such as 3-D modeling, displaying points might not be as useful as showing the object with either a smooth or mesh surface. 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 ladar systems, it is also common to choose the first or last return.

Courtesy of NAVAIR in China Lake, California, the vehicle integrated sensor suite for targeting applications eye-safe ladar testbed (VISSTA ELT) was used for collecting data.[2–4] The ELT is a ladar capable of digitizing the return pulse waveform at 2 GHz from a 500-MHz bandwidth receiver for offline data analysis. Mounted in the VISSTA van, it has a single sensor flying spot scanner capable of custom scanning patterns. The profile of the ELT footprint is approximated to be an isotropic Gaussian.

Although obtaining the surface response is not necessary for the processing method presented, all testing in this article assumes that the waveform data were processed using the surface response. The reason for this is to increase the range resolution, i.e., increase the number of detected multiple returns. Due to previous results for the ELT, the nonnegative least squares (NNLS) deconvolution method was chosen for this work.[2]

In this article, we propose novel theory and methods for estimating surface parameters, namely a slotted surface and an angled surface. We then apply these methods to process point clouds to more accurately place points from multiple

0091-3286/2013/$25.00 © 2013 SPIE

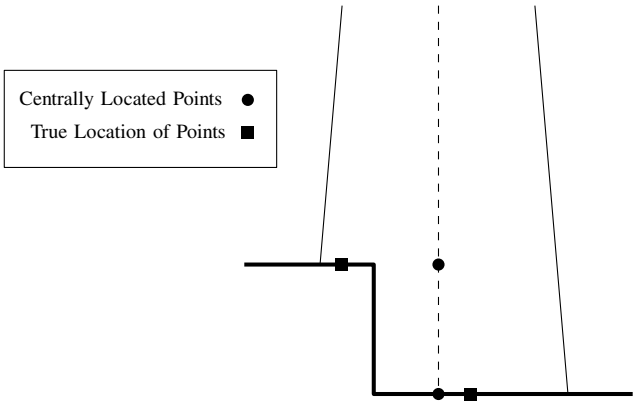

Fig. 1 Assuming that points are centered in the footprint results in incorrect point estimation.

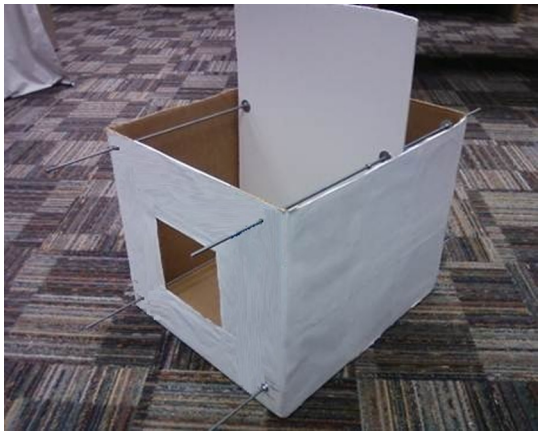

Fig. 2 The target used for testing in these experiments.

returns on the surfaces they represent. In Sec. 2, we present a method for estimating two surface parameters of interest: the width of a slotted surface and the angle of incidence of a planar surface. Section 3 builds upon the method from Sec. 2 to develop a method that solves for the locations of points within the beam footprints. Conclusions are given in Sec. 4.

2 Surface Parameter Estimation

2.1 *Slot Width Estimation*

Let the isotropic Gaussian spatial beam footprint be given by

$$f(x,y) = \frac{1}{2\pi\sigma_{\mathrm{f}}^2}\exp\left[-\frac{(x-x_{\mathrm{c}})^2+(y-y_{\mathrm{c}})^2}{2\sigma_{\mathrm{f}}^2}\right], \tag{1}$$

where $(x_{\mathrm{c}}, y_{\mathrm{c}})$ is the center of the beam and σ_{f} is the standard deviation of the footprint, which can be computed by $\sigma_{\mathrm{f}} = (\beta/4)r$ 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 $\beta/2$.

Assume that two distinct surfaces were hit by the ladar pulse and the edge is normal to the x-axis. The location of the edge can be determined by

$$x_{\mathrm{edge}} = \sigma_{\mathrm{f}} Q^{-1}\left(1-\frac{E_1}{E_1+E_2}\right)+x_{\mathrm{c}}, \tag{2}$$

where $Q(x)$ is the normal distribution Q function, E_1 and E_2 are the energies on the first and second surfaces, respectively, x_{edge} is the x-coordinate of the of the surface edge, and x_{c} is the x-coordinate of the center of the beam footprint. This equation assumes that the reflectivity of the two surfaces is the same. If this is not true, E_2 must be multiplied by the reflectivity ratio of the first surface to the second surface.

To test the ability to locate edges with this method, a target was placed 493 m from the ELT. At this distance, σ_{f} for the ELT is estimated to be 4 cm. Figure 2 shows the target with a 19×19 cm^2 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 the first surface. A second surface was placed 16 cm behind the first surface.

A point cloud of the box containing position, range, and intensity for each point was formed. When obtaining the data, the scanner was set to slowly traverse the slot in the box in the horizontal direction only. This provided many points that hit both surfaces. The edges on both sides of the slot were located according to Eq. (2). An equivalent target and scenario were created in a ladar simulator (LadarSIM) developed to model the behavior of ladar on complex target surfaces.[5] Table 1 shows the results of processing. The mean was calculated by the difference between the mean of edge 2 and the mean of edge 1.

Discrepancies between results from LadarSIM and the ELT are mostly ascribed to error in the pointing control. Because this article 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.

2.2 *Angle Estimation*

In Sec. 2.1, it was assumed that the surface hit by the laser was two discontinuous surfaces offset in range. It should be noted that the surface response of that surface consists of two offset delta functions. In this section, it is assumed that the surface hit by the laser is an inclined plane. It has been shown that for a beam with a isotropic Gaussian footprint that hits a planar surface, the surface response has a Gaussian profile.[6] The standard deviation of this surface response is given by

$$\sigma_{\mathrm{s}} = \frac{2\tan(\phi)\sigma_{\mathrm{f}}}{c}, \tag{3}$$

where c is the speed of light, and ϕ is the angle of incidence where $0 \le \phi < (\pi/2)$ rad.[6] As $\phi \to 0$, $\sigma_{\mathrm{s}} \to 0$. This indicates that the surface response of a flat surface perpendicular to the sensor's line of sight is a delta function. As $\phi \to (\pi/2)$ rad, $\sigma_{\mathrm{s}} \to \infty$, and the amplitude of the pulse goes to zero. This

Table 1 Estimate of slot width.

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

indicates that the surface response of a planar surface for large ϕ is a single, wide pulse. In practice, if the deconvolution method is under-regularized, the resulting estimate for the surface response of an angled plane may not be a single, wide pulse.

Solutions to the surface response using the NNLS method are unregularized.[7] With this method, noise can cause the surface response estimate, $\hat{s}(n)$, 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.[8] In order to preserve a non-negative surface response, the low-pass filter chosen for this method is a Gaussian filter with band edge B_0.[8] The Gaussian filter for regularization is given by

$$W(f) = \exp\left[-\frac{(2\pi f)^2}{2B_0^2}\right], \tag{4}$$

where B_0 is a parameter controlling the amount of regularization. It is determined so that $W(f)$ filters those frequencies outside of the support of the system transfer function.

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 the surface response to pass, it was observed that the surface response estimate of two offset surfaces still represents the true surface response accurately.[6] However, the surface response of angled surfaces is not well represented. A similar result was noticed in other works.[9] If the angle of incidence is small (<60 deg), the surface response appears identical to the surface response from hitting a single flat surface at 0 deg. The reason for this can be understood from Eq. (3). 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 deg, a change in angle causes a large change in the standard deviation of the surface response. For example, for a value of $\sigma_f = 4$ cm, the difference between the width of the theoretical surface response at $\phi = 25$ deg and $\phi = 30$ deg is 0.03 ns. The difference between the width of the surface response at $\phi = 75$ deg and $\phi = 80$ deg is 0.52 ns. The ELT receiver has a sample period of 0.5 ns. Thus, the received waveform is notably different only if the angle of incidence is large.

If the angle of incidence is large (>60 deg), a problem occurs in which unwanted high-frequency noise corrupts the surface response. The surface response of a planar angled surface for the VISSTA ELT waveform has a Gaussian profile. However, Figure 3 shows an example of 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.

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, decreasing B_0 to eliminate noise also decreases the range resolution. It has been shown that two Gaussian pulses cannot be separated if their means are closer than $2B_0$.[6] For the experiments

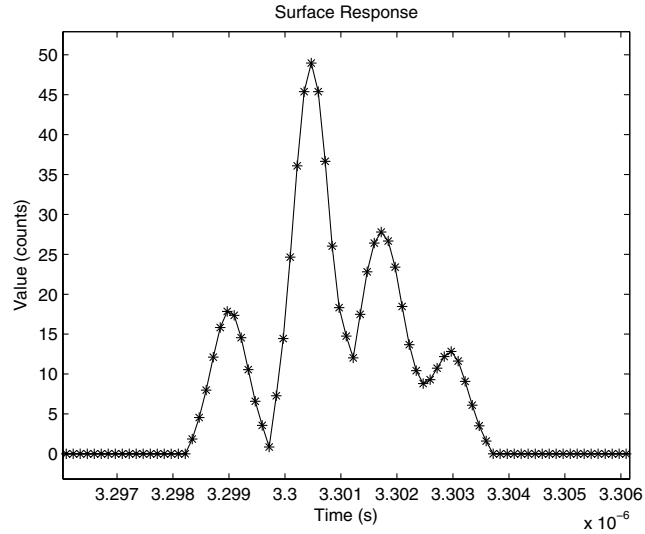

Fig. 3 An example of an under-regularized surface response resulting from hitting a planar surface at 75 deg shows artificial peaks instead of one wide peak.

reported in this article, empirical results showed that a range of values of B_0 from about $1/4$ to $3/8$ of the sampling rate could be used.

Although the surface response in Fig. 3 does not represent the true surface response, the low-frequency envelope shows what one would expect in the surface response for an angled surface, i.e., the lower frequencies of the surface response estimate that lie within the bandwidth of the received signal match the true surface response. It is only the high-frequency content that is incorrect.

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. It is possible that the ranges of the peaks in the noisy surface response correspond to some location that is truly on the surface.

Under the assumption that a planar surface was hit, a method to treat this surface response to estimate the angle of incidence is presented. 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 of even reflectivity, 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.

For a zero mean Gaussian with a standard deviation of 1, the boundaries separating energy from each peak in the footprint, as shown in Fig. 4, are computed by

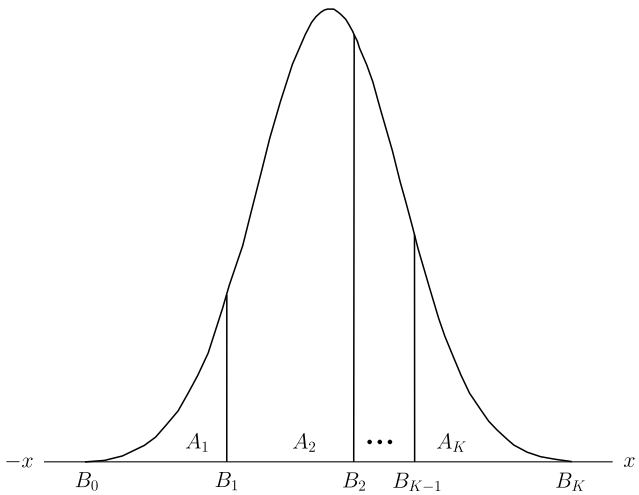

Fig. 4 Side profile of spatial beam footprint showing areas and boundaries of peaks.

$$B_k = Q^{-1}\left(1 - \sum_{j=1}^{k} E_j\right)\sigma_{\mathrm{f}}, \tag{5}$$

where 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

$$\begin{aligned} l_k &= \left[\frac{\int_{B_{k-1}}^{B_k} x \exp\left(\frac{-x^2}{2}\right)\mathrm{d}x}{\int_{B_{k-1}}^{B_k} \exp\left(\frac{-x^2}{2}\right)\mathrm{d}x}\right], \\ &= \left[\frac{\exp\left(\frac{-B_{k-1}}{2}\right) - \exp\left(\frac{-B_k}{2}\right)}{\sqrt{2\pi}E_k}\right]. \end{aligned} \tag{6}$$

The angle of incidence of the planar surface is given by

$$\phi = \tan^{-1}\left[\frac{d(r_i, r_j)}{d(l_i, l_j)}\right], \tag{7}$$

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. The angle of incidence can be obtained using any i and j. To obtain the most accurate estimate, the angles between consecutive points are computed and averaged.

To investigate the method, a test board was constructed to provide a large flat surface. The test board had an estimated reflectivity of 0.7 and was placed at 493 m. The angle, ϕ, was measured using a compass with accuracy of ~1 or 2 deg. The board was scanned at 45, 60, and 75 deg, and the angle was estimated using the method presented in this section. A setup was used in LadarSIM to match this scenario. More than 40 points were randomly selected from the middle of the target for both the ELT data and the simulated data.

The number of peaks, K, varies depending on ϕ. In the experiments, K was observed to be at most 4 on a plane where $\phi = 75$ deg. A 45-deg 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 deg. The results in Table 2 show that both the simulated and real data provide an accurate estimation of ϕ.

Table 2 Angle estimate.

Data	Truth (deg)	Estimate (deg)	Standard deviation (deg)
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

3 Point Location Estimation

Section 2 assumed that the surface being targeted was known to be either two surfaces offset in range or an angled, planar surface. However, to form a point cloud of an unknown image, the prior information about the surface being planar, slotted, or highly complex is not available. Therefore, it is not possible to distinguish between the surface response of a slotted surface and the surface response of an angled surface using the processing method used in Sec. 2.

For the method presented in this section, 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 Sec. 2.1 for the slotted 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. 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(a) shows that a surface cannot be interpolated because there are two points at one location. Figure 5(b) shows the interpolated surface if the first return is chosen. This scheme overextends the interpolated surface. Figure 5(c) shows the interpolated surface if the last return is chosen. This shows the interpolated surface cutting into the actual surface. Figure 5(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 Sec. 2.2. 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 does not place the first return above the surface.

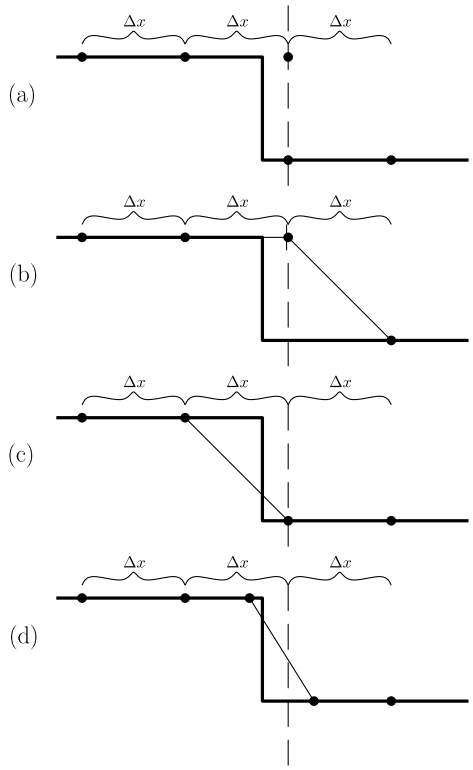

Fig. 5 The side view depiction of an interpolated discontinuous surface that contains a multiple return shows that using the planar approximation described in this section produces a sharper edge more representative of a discontinuous surface. (a) All returns. (b) First return. (c) Last return. (d) Planar approximation.

The planar assumption holds only for surfaces that do not contain very high spatial frequencies. 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 middle of the footprint would be inside the slot. Only two returns can be detected in this case. Using the method of Sec. 2.2, 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 Sec. 2.2 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. The proposed method for determining the angle of orientation is to calculate the best fit plane to the neighboring points with the constraint that the angle of incidence is known. With the best fit plane defined, the shots can be placed to lie on that plane.

The method presented 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 bidirectional reflectance distribution function.

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.

$$a_1(x - x_0) + a_2(y - y_0) = (z - z_0), \tag{8}$$

where (x, y, z) is any point in the plane, (x_0, y_0, z_0) is the position of some known point on the plane, and a_1 and a_2 are the parameters that describe the plane. This can be defined in a spherical coordinate system using ϕ from Sec. 2.2 and θ where $0 \leq \theta < 2\pi$. The parameters of the plane are given by

$$a_1 = \cos(\theta)\tan(\phi), \quad a_2 = \sin(\theta)\tan(\phi). \tag{9}$$

To solve for a_1 and a_2, the only free variable is θ because ϕ is determined from the method in Sec. 2.2. In order to obtain θ, neighboring points must be known. In a flash ladar, neighboring points would be given by using adjacent elements in a focal plane array. However, for a flying spot scanning ladar such as the ELT, this is not immediately known. Furthermore, noise on the azimuth and elevation measurements makes determining neighboring points even more difficult. A nearest neighbor search must be performed. In this work, the nearest neighbor search is accomplished using the K-D tree search method.[10] 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.

The points on the grid are interpolated using inverse distance squared weighting. Other methods could be used for this. The range at location (x_i, y_i) is given by[11]

$$r(x_i, y_i) = \frac{\sum_{m=1}^{M} r(x_m, y_m)[(x_m - x_i)^2 + (y_m - y_i)^2]^{-1}}{\sum_{m=1}^{M} [(x_m - x_i)^2 + (y_m - y_i)^2]^{-1}}, \tag{10}$$

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, the value of M must be high. However, if the shots are adjacent, a smaller value can be used. For the experiments in this work, $M = 64$ 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 eight points, then the error is given by

$$e = |X\mathbf{a} - \mathbf{r}_p|_2^2, \quad = \mathbf{a}^T X^T X\mathbf{a} - 2\mathbf{a}^T X^T \mathbf{r_p} + \mathbf{r_p}^T \mathbf{r_p}, \tag{11}$$

where $X = [\,\mathbf{az}_p \quad \mathbf{el}_p\,]$ and $\mathbf{a} = [\,a_1 \quad a_2\,]^T$. The value of $\mathbf{a}$ is constrained by Eq. (9), 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 $\mathrm{d}e/\mathrm{d}\theta = 0$ is computed. This is derived as follows:

$$\frac{\mathrm{d}e}{\mathrm{d}\theta} = \frac{\mathrm{d}\mathbf{a}^T}{\mathrm{d}\theta}\frac{\mathrm{d}e}{\mathrm{d}\mathbf{a}}, \tag{12}$$

$$\frac{\mathrm{d}\mathbf{a}^T}{\mathrm{d}\theta} = [\,-\sin(\theta) \quad \cos(\theta)\,]\tan(\phi), \tag{13}$$

$$\frac{\mathrm{d}e}{\mathrm{d}\mathbf{a}} = 2X^T X\mathbf{a} - 2X^T \mathbf{r_p}, \tag{14}$$

$$\begin{aligned}\frac{\mathrm{d}e}{\mathrm{d}\theta} &= 2[\,-\sin(\theta) \quad \cos(\theta)\,]\tan(\phi)X^T X\begin{bmatrix}\cos(\theta)\\ \sin(\theta)\end{bmatrix}\tan(\phi)\\ &\quad - 2[\,-\sin(\theta) \quad \cos(\theta)\,]\tan(\phi)X^T\mathbf{r_p} = 0.\end{aligned} \tag{15}$$

This leads to solving for θ from

$$\begin{aligned}[\,-\sin(\theta) \quad \cos(\theta)\,]X^T\mathbf{r_p} &= [\,-\sin(\theta) \quad \cos(\theta)\,]\\ &\quad \times \tan(\phi)X^T X\begin{bmatrix}\cos(\theta)\\ \sin(\theta)\end{bmatrix}.\end{aligned} \tag{16}$$

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

$$\begin{aligned}[\,-\sin(\theta) \quad \cos(\theta)\,]X^T\mathbf{r_p} &= \mathrm{tr}\left\{[\,-\sin(\theta) \quad \cos(\theta)\,]\right.\\ &\quad \left.\times \tan(\phi)X^T X\begin{bmatrix}\cos(\theta)\\ \sin(\theta)\end{bmatrix}\right\},\\ &= \mathrm{tr}\left\{\begin{bmatrix}\cos(\theta)\\ \sin(\theta)\end{bmatrix}[\,-\sin(\theta) \quad \cos(\theta)\,]X^T X\right\}\tan(\phi),\\ &= \mathrm{tr}\left\{\begin{bmatrix}-\cos(\theta)\sin(\theta) & \cos^2(\theta)\\ \sin^2(\theta) & \cos(\theta)\sin(\theta)\end{bmatrix}X^T X\right\}\tan(\phi).\end{aligned} \tag{17}$$

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 2×2 identity matrix. This would not be possible if the noisy data were used directly. This simplification allows Eq. (17) to be written as

$$\begin{aligned}&[\,-\sin(\theta) \quad \cos(\theta)\,]X^T\mathbf{r_p}\\ &\quad = \mathrm{tr}\left\{\begin{bmatrix}-\cos(\theta)\sin(\theta) & \cos^2(\theta)\\ \sin^2(\theta) & \cos(\theta)\sin(\theta)\end{bmatrix}\right\}24\beta^2\tan(\phi), = 0,\end{aligned} \tag{18}$$

where θ is solved from Eq. (18) to obtain

$$\theta = \tan^{-1}\left(\frac{b_2}{b_1}\right), \tag{19}$$

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 down the center of the footprint as shown in Fig. 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 6(a) shows the side view of an angled surface to demonstrate the ranges of the points. Figure 6(b) shows the top view to demonstrate the spatial locations of the points. Then, the locations of the points are adjusted via the method from Sec. 2.2, as shown in Fig. 7. The angle ϕ is now set, but θ may be anywhere along the dotted circle shown in Fig. 7(b). Next, the final locations of the points are determined by rotating the points about the line of sight to θ, as shown in Fig. 8.

3.1 *Slotted Surface Results*

A point cloud using the ELT data of the slotted surface from Sec. 2.1 is shown in Fig. 9. This point cloud shows the same data that were used to obtain the statistics for Sec. 2.1. 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 9(a) shows the point cloud when the data are not processed to obtain the surface response. Because the surface response is 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 9(b) shows the point cloud when the surface response is computed and the

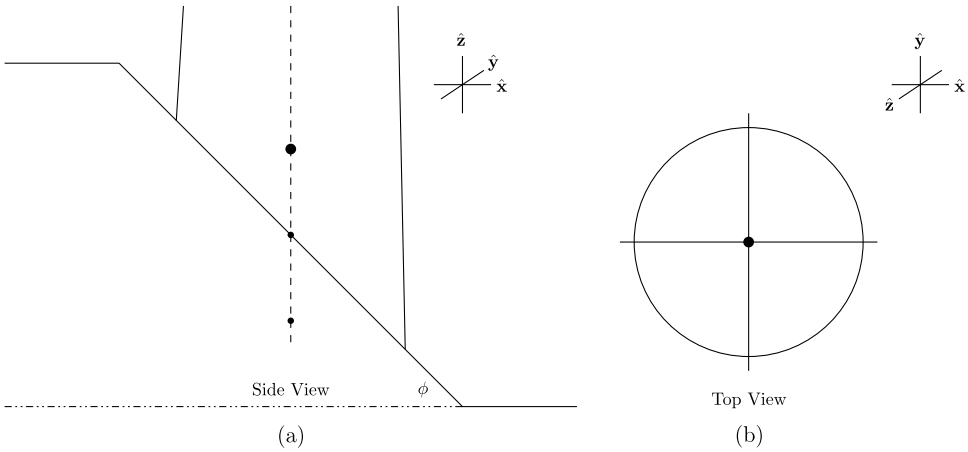

Fig. 6 The shot on a planar surface before adjusting the locations of the points incorrectly represents the surface. (a) Side view of sloped surface showing point locations, and (b) top view.

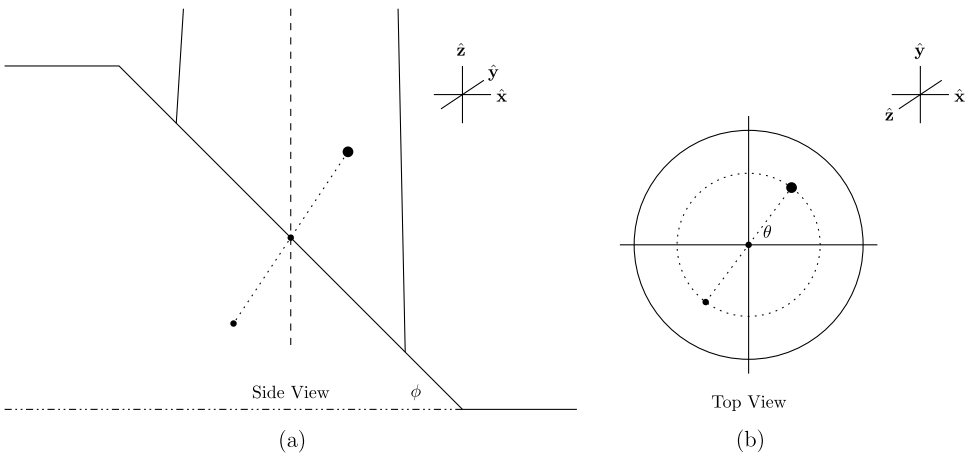

Fig. 7 The shot on planar surface after solving for ϕ represents an angled surface but not properly oriented. (a) Side view of sloped surface showing point locations, and (b) top view.

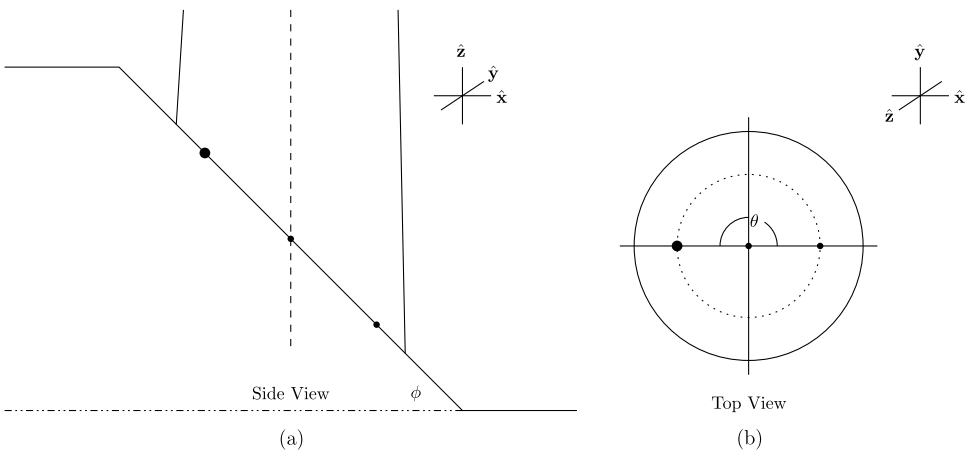

Fig. 8 The shot on planar surface after rotating by θ correctly represents the angled surface. (a) Side view of sloped surface showing point locations, and (b) top view.

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 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. 9(a) and removes the effect seen in the point cloud of Fig. 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

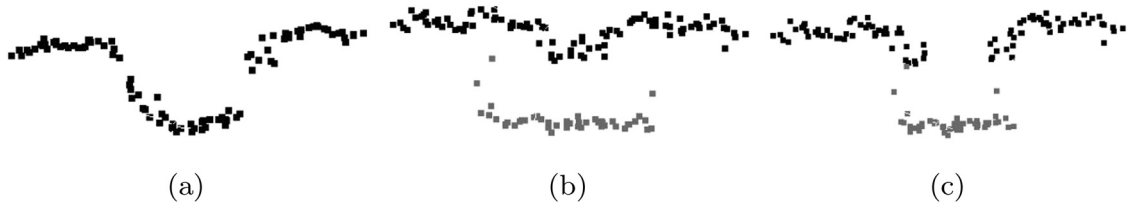

Fig. 9 Side view of eye-safe ladar testbed (ELT) point cloud of slotted surface. (a) Processing without obtaining the surface response results in a curved surface. (b) Processing using the surface response and assuming the points to be located in the center of the footprint results in two surfaces. (c) Processing using the surface response and solving for the locations of the points within the beam footprint results in two flat surfaces.

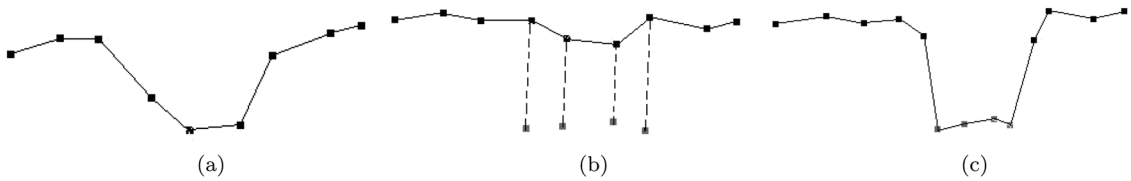

Fig. 10 Side view of ELT point cloud of slotted surface with a practical shot spacing. (a) Processing without obtaining the surface response results in a curved surface. (b) Processing using the surface response and assuming the points to be located in the center of the footprint results in two surfaces. (c) Processing using the surface response and solving for the locations of the points within the beam footprint results in two flat surfaces.

surface but does not represent a practical scan. Figure 10 shows the scan of the slotted surface with shots spaced so that the beam footprints are barely overlapping. The points have been connected so that the side view of the interpolated surface can be visualized. Figure 10(b) also shows dotted lines that connect the multiple returns. These points are those that cause problems when interpolating a surface because they have the same values for the azimuth and elevation.

3.2 *Angled Surface Results*

A point cloud using the ELT data of the angled surface from Sec. 2.2 is shown in Fig. 12. This is a top view of the angled surface. Two axes were drawn with a 75-deg line passing through the point cloud. The arrow indicates the line of sight. The target for these experiments is shown in Fig. 11. Figure 12(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 12(b) shows the point cloud when the surface response is computed and the coordinates of the returns are aligned down

Fig. 11 Target used for data collection.

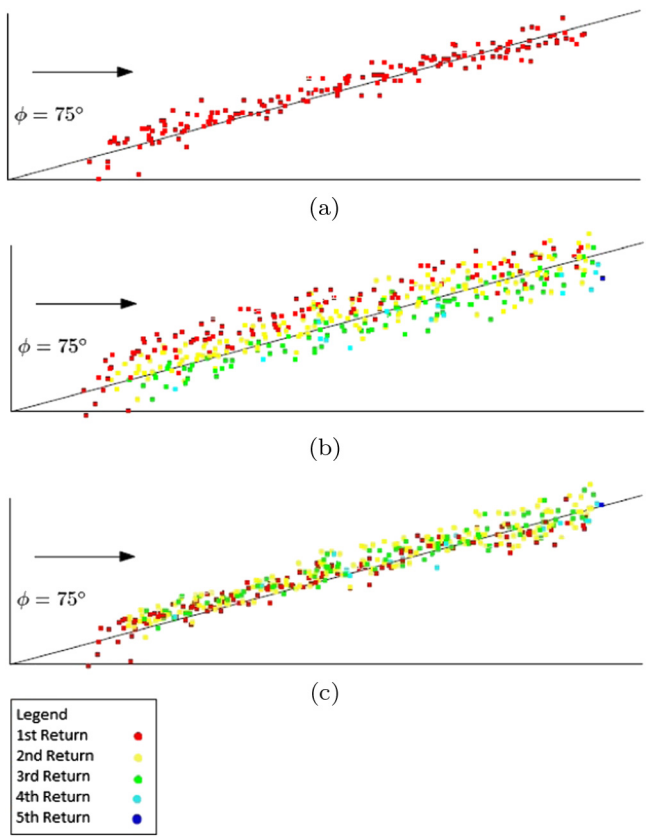

Fig. 12 Top view of ELT point cloud of angled surface. (a) Processing without obtaining the surface response results in a surface with low variance from the plane. (b) Processing using the surface response and assuming the points to be located in the center of the footprint results in a surface with high variance from the plane. (c) Processing using the surface response and solving for the locations of the points within the beam footprint results in a surface with low variance from the plane and with an increased number of points on the plane.

Table 3 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

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 12(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. 12(a) but less than the point cloud of Fig. 12(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. 12. The standard deviation of the points from the best fit plane was computed, and results are shown in Table 3.

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 angled surfaces.

An estimate of the angle ϕ was given in Table 2, but an estimate of θ was not given. The same data used in Sec. 2.2 were used to estimate θ. To provide more accurate test conditions that reflect a shot spacing with adjacent, nonoverlapping points, points were removed from the point cloud. Every fifth row and column were used to create an approximate shot spacing of 80 μrad. At the time of data collection, the angle θ was not considered. All tests were conducted with the true value of $\theta \approx 0$ deg. To show that the method can estimate θ over the span of [0 deg, 360 deg], 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 $\phi = 60$ deg in Fig. 13 and with $\phi = 75$ deg in Fig. 14. The mean estimate is plotted with error bars showing the standard deviation of the estimate.

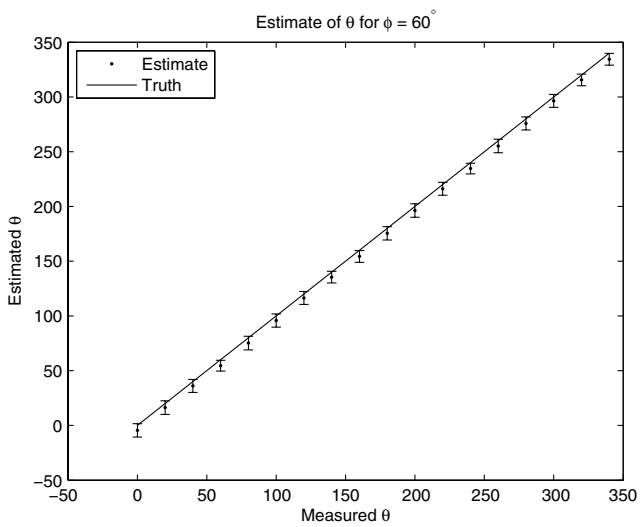

Fig. 13 The estimate of θ for $\phi = 60$ deg matches truth within the uncertainty.

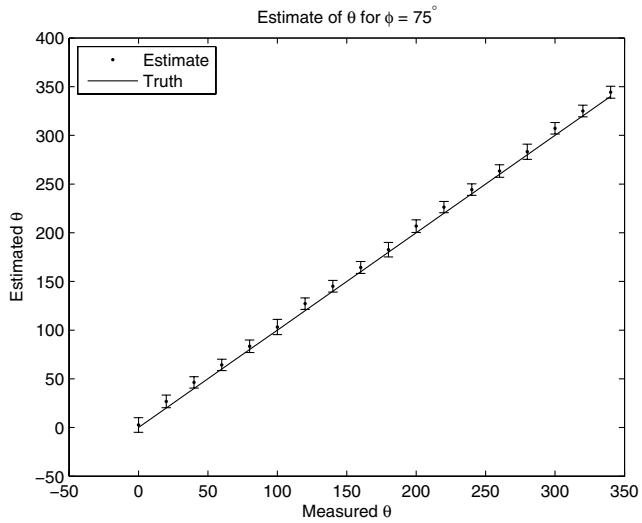

Fig. 14 The estimate of θ for $\phi = 75$ deg matches truth within the uncertainty.

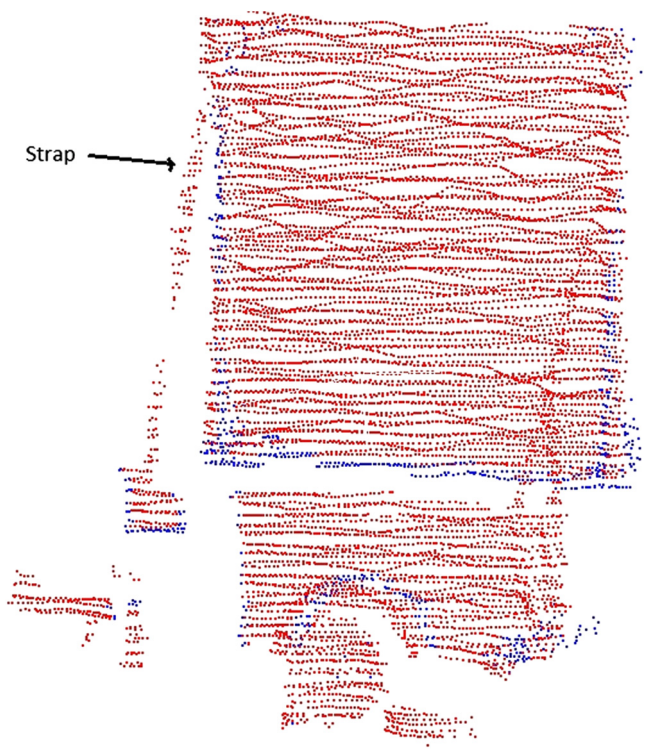

Fig. 15 Test board scanned by the ELT.

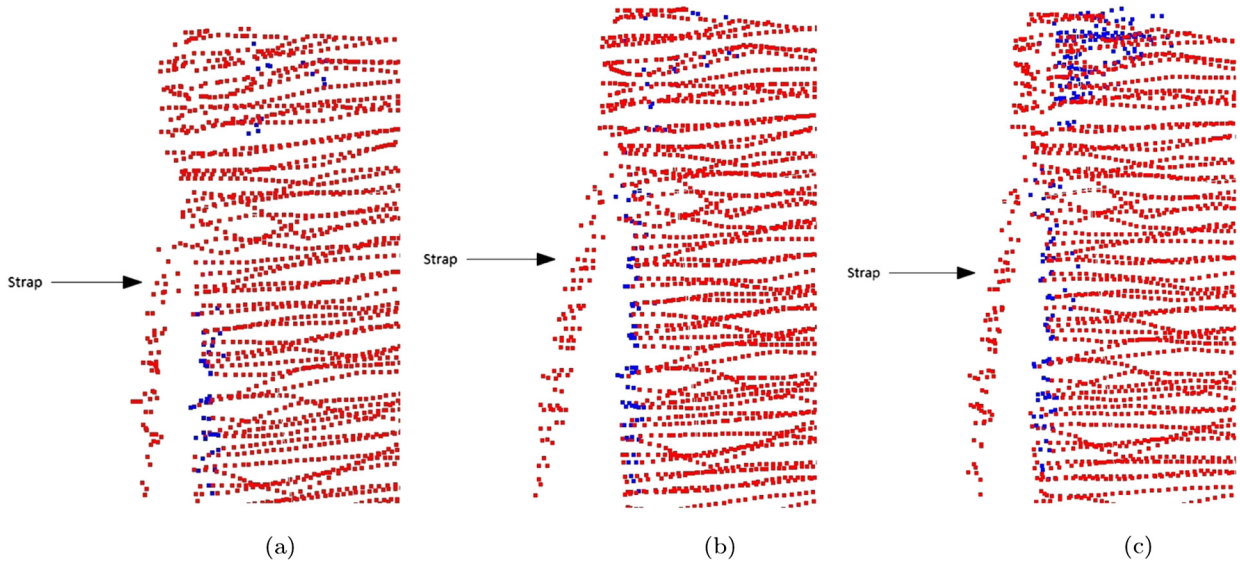

Fig. 16 Strap on the test board. (a) When processed without obtaining the surface response, the area where the strap is close to the test board is poorly represented. (b) When processed using the surface response and assuming the points to be located in the center of the footprint, the points represent the strap as being wider than it is in reality. (c) When processed using the surface response and solving for the locations of the points within the beam footprint, the points on the strap are narrower and closer to the truth.

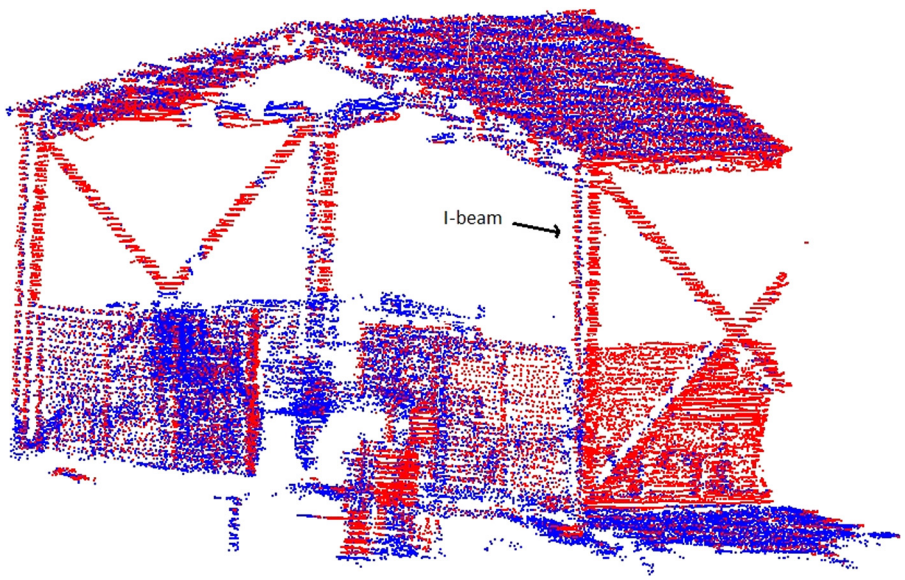

Fig. 17 Scan of shed with an I-beam.

From Figs. 13 and 14, the error on the estimate of θ is on the order of 5 to 8 deg. 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 few errors in the pointing control were used, the amount of error on θ may be less than the error in this system.

3.3 *Complex Objects Results*

Sections 3.1 and 3.2 showed point clouds of the two main surfaces used in this work. This section shows the effect of the method on point clouds of other objects that may appear in a scene.

Figure 15 shows a scan of the test board used in experiments from Sec. 3.2. 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 16 shows a closer view of the test board. Figure 16(a) shows the point cloud when the data are not processed to obtain the surface response. Figure 16(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. Figure 16(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 toward the top of the strap. Another detail to notice is that the strap is slightly thicker in Fig. 16(b) than it is in the other point clouds. This result supports the drawing in Fig. 1.

A building that contains an I-beam was scanned and shown in Fig. 17. A close-up view is shown in Fig. 18. The arrow indicates the location of the I-beam. The color

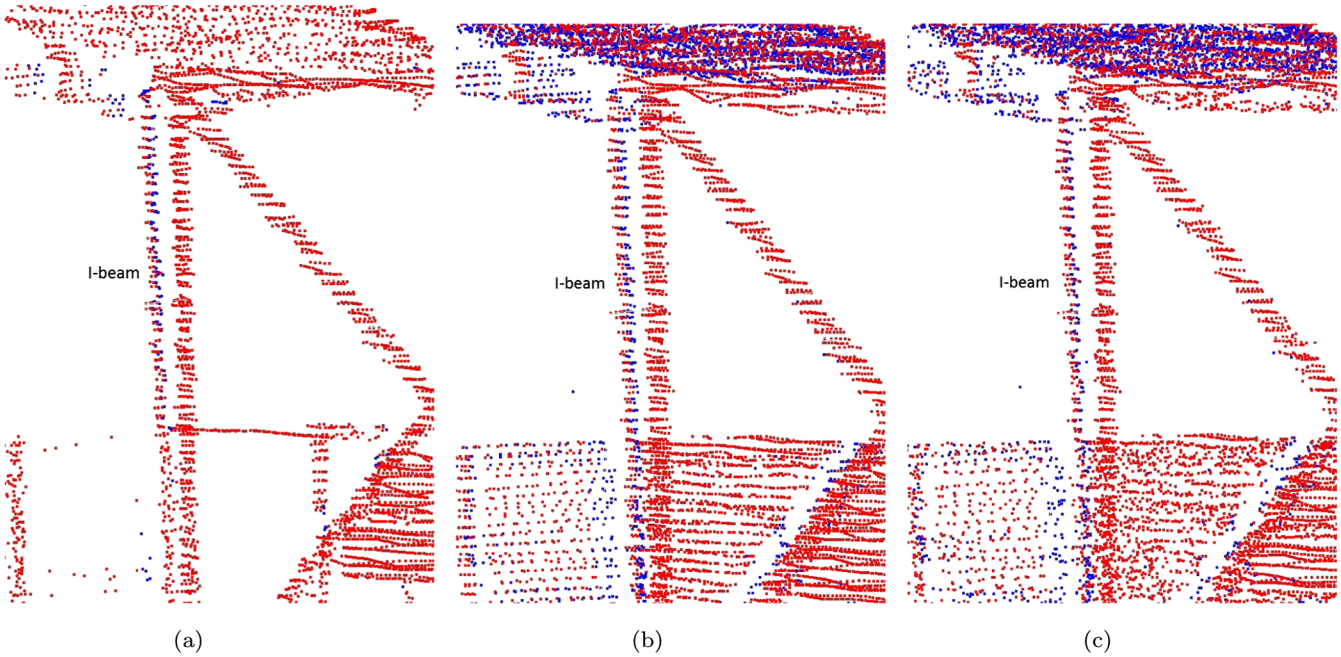

Fig. 18 I-beam scanned by the ELT. (a) When processed without obtaining the surface response, detects fail to show both surfaces of the I-beam at the bottom where the two surfaces are closer together. (b) When processed using the surface response and assuming the points to be located in the center of the footprint, both surfaces of the I-beam are visible. (c) When processed using the surface response and solving for the locations of the points within the beam footprint, the edge of the I-beam is narrower and more representative of the true I-beam.

indicates return number. Red shows the first return, and blue shows other returns. Figure 18(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 18(b) shows that multiple returns are detected all along the I-beam when deconvolution is applied. Figure 18(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. 18(b). This is due to solving for the locations of the returns within the beam footprint. Again, this result supports the drawing in Fig. 1.

Another detail that can be observed in this point cloud is the fence along the bottom. The fence is barely detected in Fig. 18(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. 18(a) used a higher threshold setting is because of ringing in the waveform. Ringing in the waveform can cause 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. 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.

4 Conclusion

A technique that exploited the locations of neighboring points in the point cloud to enhance the resolution on the edges of surfaces was developed. The problem of representing a highly sloped surface with a single point was solved by forming an estimate for the angle of incidence. In situations where the underlying surface is to be interpolated, the 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 by computing the surface response and accurately place the points within the footprint via the method presented in this article. This has direct application to object recognition tasks such as ATR by increasing the number of points on the target, thus providing better understanding of the target features.

The processing method presented provides a means of deferring regularization from deconvolution to a stage in the processing where more information is available to determine the shape of the surface. This enables more returns to be detected on the edges of surfaces where the increase in resolution is needed. The method provides a means of finding a better estimate of those locations than simply using the center of the footprint. In areas that are flat and θ is small, only one return was detected in the testing of Sec. 2.2. On surfaces that contain low spatial frequencies such as these, resolution is not an issue.

The proposed method increases the resolution with data spatially sampled around Nyquist. Applications that can utilize multiple images may work better in a traditional super-resolution framework. A method was given for determining the width of a slot in a surface given that it is assumed to be a slot. In ATR, if a target in question is known to have a slot over a specific area, this method may enable using the width of the slot as a feature to confirm the target. Also, the method was shown to be a means of estimating the angle of incidence of a planar surface using a single waveform.

References

1. B. Jutzi and U. Stilla, "Range determination with waveform recording laser systems using a Wiener filter," *ISPRS J. Photogramm. Remote Sens.* **61**(2), 95–107 (2006).
2. K. D. Neilsen, S. E. Budge, and R. T. Pack, "Signal processing on waveform data from the eyesafe ladar testbed (ELT)," *Proc. SPIE* **7684**, 76840R (2010).
3. K. D. Neilsen et al., "Design and validation of the eyesafe LADAR testbed (ELT) using the LadarSIM system simulator," *Proc. SPIE* **7323**, 73230B (2009).
4. R. T. Pack et al., "Eyesafe LADAR test-bed with coaxial color imager," *Proc. SPIE* **7323**, 732303 (2009).
5. R. T. Pack et al., *USU LadarSIM Release 3.0 User's Manual*, Utah State University/Space Dynamics Laboratory, Logan, Utah (2007).
6. K. Neilsen, "Signal processing on digitized LADAR waveforms for enhanced resolution on surface edges," Master's Thesis, Utah State University, Logan, Utah (2011).
7. C. Lawson and R. Hanson, *Solving Least Squares Problems*, Chapter 23, pp. 158–165, Prentice Hall, Englewood Cliffs, New Jersey (1974).
8. M. Bertero and P. Boccacci, *Introduction to Inverse Problems in Imaging*, Institute of Physics Publishing, South Independence Mall West, Philadelphia (1998).
9. S. Harsdorf and R. Reuter, "Stable deconvolution of noisy lidar signals," in *Proc. EARSeL-SIG-Workshop LIDAR*, Dresden, Germany (2000).
10. J. H. Friedman, J. L. Bentley, and R. A. Finkel, "An algorithm for finding best matches in logarithmic expected time," *ACM Trans. Math. Softw.* **3**(3), 209–226 (1977).
11. D. Shepard, "A two-dimensional interpolation function for irregularly-spaced data," in *ACM '68: Proc. 1968 23rd ACM National Conf.*, pp. 517–524, ACM, New York (1968).

Kevin Neilsen worked as a research assistant at the Center for Advanced Imaging Ladar at Utah State University. He received a bachelor's and master's degree in electrical engineering at Utah State University in December 2010. His areas of interest are signal processing, image processing, and digital communications.

Scott E. Budge received his BS, MS, and PhD degrees in electrical engineering from Brigham Young University in 1984, 1985, and 1990, respectively. Scott joined the faculty at Utah State University in January of 1989, where he has been involved in research into image data compression algorithms for transmission systems and space-based observation platforms. His current work involves research on methods for exploitation of full-waveform LADAR and EO images for 3-D image applications. He is also involved in the development of high-performance image processing algorithms designed for implementation in hardware and VLSI systems.