

# Improving the Detection of Explosive Hazards with LIDAR-Based Ground Plane Estimation

A. Buck\*, J. M. Keller\*, M. Popescu\*\*

\*University of Missouri, Electrical and Computer Engineering Department, Columbia, MO

\*\*University of Missouri, Health Management and Informatics Department, Columbia, MO

## ABSTRACT

Three-dimensional point clouds generated by LIDAR offer the potential to build a more complete understanding of the environment in front of a moving vehicle. In particular, LIDAR data facilitates the development of a non-parametric ground plane model that can filter target predictions from other sensors into above-ground and below-ground sets. This allows for improved detection performance when, for example, a system designed to locate above-ground targets considers only the set of above-ground predictions. In this paper, we apply LIDAR-based ground plane filtering to a forward looking ground penetrating radar (FLGPR) sensor system and a side looking synthetic aperture acoustic (SAA) sensor system designed to detect explosive hazards along the side of a road. Additionally, we consider the value of the visual magnitude of the LIDAR return as a feature for identifying anomalies. The predictions from these sensors are evaluated independently with and without ground plane filtering and then fused to produce a combined prediction confidence. Sensor fusion is accomplished by interpolating the confidence scores of each sensor along the ground plane model to create a combined confidence vector at specified points in the environment. The methods are tested along an unpaved desert road at an arid U.S. Army test site.

**Keywords:** LIDAR, sensor fusion, ground plane estimation

## 1. INTRODUCTION

Detecting explosive hazards in uncertain environments is an ongoing research problem with a history that spans several decades. Many sensing modalities and algorithms have been proposed to tackle this challenging task. Two of the most recent approaches involve using a forward looking ground penetrating radar (FLGPR) system<sup>1</sup> and a side looking synthetic aperture acoustic (SAA) system.<sup>2</sup> Both of these approaches have shown considerable success in detecting explosive hazards along the side of a road. While the FLGPR system was originally used to detect buried landmines, it has been adapted in this work to detect above-ground targets. Conversely, the SAA system has been designed from the outset as a cost-effective solution for detecting and imaging roadside targets and foliage. Up to this point, the two sensing modalities have been examined independently, with a specialized vehicle designed for each, and tested in a desert environment at a U.S. Army test site.

In this work, we examine the potential use of a LIDAR point cloud imaging system as an additional sensing modality to improve the detection performance of the two aforementioned approaches. LIDAR uses pulses of laser light to produce an extremely sensitive detector of physical objects in an environment.<sup>3</sup> The data is returned as a three-dimensional cloud of points, each with a unidimensional intensity value indicating the magnitude of the response. Traditionally, LIDAR has been employed in a variety of applications, including surveying,<sup>4</sup> meteorology,<sup>5</sup> and mobile navigation systems.<sup>6</sup> When mounted on a moving vehicle, a LIDAR system can generate a three-dimensional model of the environment that can assist with navigation and scene understanding.

For this study, we use a LIDAR system in conjunction with the FLGPR and SAA detection systems to improve the detection rate of above-ground roadside targets on a prototypical desert lane. We begin by demonstrating how each sensing modality can be used independently to generate a set of prediction locations for a lane and how these hits are scored against known ground truth. Next, we propose a method for filtering the predicted hit locations into a reduced set by estimating the local ground plane from the LIDAR data and eliminating all locations that do not meet certain height range criteria. Finally, we demonstrate how the three sensing modalities can be fused together into a combined confidence surface to improve the overall detection rate. All of the methods presented in this paper are examined on a single dataset as an

independent case study, leaving a more thorough investigation of the application of this methodology to the general target detection problem for future work.

## 2. INDIVIDUAL SENSING MODALITIES

The motivation for this paper is to explore the potential for including LIDAR point cloud data in the traditional processing workflow for FLGPR and SAA detection systems. As such, we consider only a single dataset and the best detection results from the existing systems as the basis for the experiments in this study. The following sections describe the dataset used and the methods applied to generate detection results for each sensing modality independently.

### 2.1 Data characteristics

We consider data from a single lane of a U.S. Army test site, henceforth called the backbone lane, which contains 28 targets of various types spread over approximately 0.5 kilometers. The area considered is 2642.9 m<sup>2</sup> along the side of a dirt road, and while some targets are exposed for calibration purposes, the majority are placed in bushes, shrubs, branches, and embankment itself. The three different sensing modalities were captured at different times, but each covered the same spatial region. The sensors included a forward looking ground penetrating radar (FLGPR) system, a side looking synthetic aperture acoustic sensor (SAA), and a LIDAR system. Of particular interest for this work is the LIDAR data, which is the newest addition to the sensor suite. We first describe the standard processing methodology for the FLGPR and SAA data and then show how the LIDAR magnitude can be used as an independent sensing modality.

### 2.2 Traditional processing workflow for FLGPR and SAA

The traditional processing workflow for FLGPR and SAA data involves an initial beamforming step, in which the raw sensor data is transformed into a georeferenced intensity image. Each pixel in this image corresponds to a location in world space and the intensity represents the magnitude of response. The beamformed images of the FLGPR sensor are generated at a spatial resolution of 2.5 cm whereas those of the SAA sensor are generated at a resolution of 1 cm. As the vehicle moves down the lane, a sequence of beamformed images are created, typically with some overlapping regions. While all of the methods presented in this paper can be adapted for use in causal settings, we assume that data from the entire length of the run is available a priori.

A size-contrast filter is applied to each beamformed image using the method presented by Stone et al.<sup>7,8</sup> In this approach, a pair of concentric sliding windows passes over the image, defining an inner and outer window region for each pixel location. The difference between the distribution statistics of the inner and outer regions gives the firing strength of the filter at a particular location. In this work, we assume an inner window radius that corresponds to 0.5 m and an outer window radius of 1 m. We then compute the univariate Mahalanobis distance between the pixel intensities of the inner window region and the area between the inner and outer windows,

$$D_M(p, q) = \sqrt{\frac{(\mu_p - \mu_q)^2}{\sigma^2}}. \quad (1)$$

Here,  $\mu_p$  and  $\mu_q$  are the mean values of the inner and outer regions respectively, and  $\sigma^2$  is the variance of the outer region. A large distance between the two regions indicates the presence of an anomaly in the beamformed image. The image produced by this filter acts as a confidence map that indicates the strength of belief that there is a target at a particular location.

Next, a non-maximum filter is applied to the confidence map by dilating the image with a rectangular structuring element, corresponding to 1 m<sup>2</sup> in size and removing all points that are not equal to the local maximum. The remaining points correspond to local peaks and are saved as a set of hit locations for the beamformed image. This process is illustrated in Figure 1. The height of each peak corresponds to the firing strength of the size-contrast filter and hence, the prediction confidence. Each hit location is saved as a three-element vector corresponding to a UTM easting and northing coordinate and a confidence value used for scoring. Note that we do not typically use any elevation data from the beamformed image as it may not always be available.




Figure 1. Extracting hit locations from a beamformed image using a size-contrast filter. The original beamformed image is shown in the bottom layer. The top layer shows the confidence map produced by the size-contrast filter (opaque) and the non-maximum filter (transparent). The local peaks of the confidence map are saved as potential hit locations and are projected back into world coordinates (bottom). The size of each red circle overlaid on the beamformed image corresponds to the hit confidence.

The hits from each beamformed image in the lane are aggregated using the mean-shift algorithm<sup>9,10</sup>, which produces the final hit list for the detector. This algorithm operates by iteratively shifting points toward the average location of all points that lie within a specified radius (0.5 m for this study), and merging points as they come together. In our implementation, the confidence values of the points are summed when they are merged. The resulting hits for the FLGPR and SAA detectors on the backbone lane are shown in Figure 2. We refer to this as the set of sparse hits for each detector. The hit locations are compared to known ground truth and any hit falling within a 0.5 m radius of a true target location is marked as a confirmed hit. All other hits are false alarms. The detector seeks to assign high confidence to true hits and low confidence elsewhere so that when a threshold is applied, only the true targets remain. The receiver operating characteristic (ROC) curve for the detector indicates the probability of detecting all targets for a given false alarm rate, measured in false alarms per meter squared.

### 2.3 Target detection with LIDAR magnitude

The LIDAR sensing system used in this study returns a dense cloud of points by illuminating the environment in front of the vehicle with laser light. The angle and distance of each beam is recorded, along with an illumination factor that is dependent on the material upon which the laser is reflected. After processing to transform the data into UTM space, each point in the dataset has four attributes: easting, northing, elevation, and magnitude. While the three spatial coordinates are most useful for generating a three-dimensional model of the scene, the illumination magnitude can be used to assist in target detection. As an independent sensor, it is dependent on line of sight visibility to a target material that differs significantly from its surroundings. Our investigations have shown that the targets used on the backbone lane produce points with magnitude values significantly lower than the surrounding environment. This can be seen from the examples in Figure 3.

One observation from examining the distribution of magnitude responses throughout the LIDAR point cloud is that low magnitude responses are densely clustered around the targets and are more sparsely distributed elsewhere. To utilize this feature, we perform mean-shift clustering using the LIDAR points that have magnitude values less than a specified threshold (set to 0.2 for this study). The points are initialized with unit confidence so that the confidence values of the resulting mean-shift hit locations reflect the local density of the LIDAR points. The resulting set of sparse hit locations is shown and scored in Figure 2.

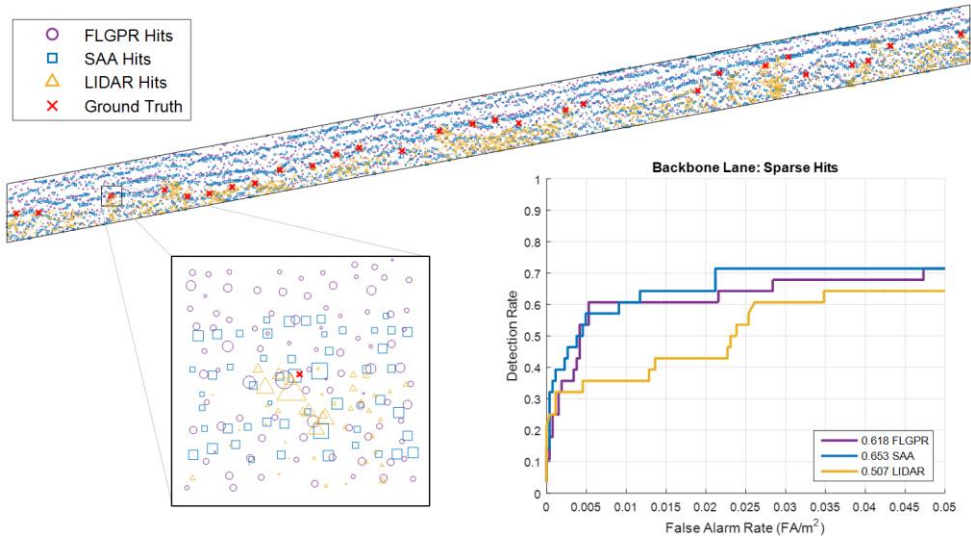


Figure 2. Results of the FLGPR and SAA detectors on the backbone lane using the traditional method and the proposed LIDAR magnitude detector. The region boundaries enclose 2642.9 m<sup>2</sup>, in which 2372 hit locations are produced using the FLGPR detector, 2234 hits are produced using the SAA detector, and 1224 hits are produced by the LIDAR magnitude detector. The full lane diagram is not drawn to scale. Plotted hit size corresponds to prediction confidence. The area under the curve (AUC) statistic for each ROC curve is given in the legend.




Figure 3. Examples of the LIDAR magnitude response at target locations. The left images show the full intensity range of the data. Notice that most points have magnitude values near the center of the spectrum, except for targets, which can be clearly identified with a much lower response. The right images highlight in red only the points with magnitude values less than 0.2.

### 3. GROUND PLANE ESTIMATION

There are many applications where an accurate model of the physical environment can provide useful information to the user. LIDAR systems have greatly advanced the capability of detecting and modeling objects in three-dimensional space. By mounting a LIDAR system on the front of a vehicle, the surrounding space can be imaged and modeled as a three-dimensional point cloud. This model can be used to enhance scene understanding and provide important elevation data where it may not otherwise exist.

As discussed previously, hits from the RX prescreener do not include elevation information. This is largely because the elevation, when available, is measured in absolute units from sea level and not relative to the local ground plane. Determining the local ground plane can be difficult in the noisy environments that make up the side of a dirt road. Bushes, ditches, and other environmental features make a flat-plane model unsuitable for this task. In the following sections, we present a method of estimating the ground plane from a LIDAR point cloud when much of what makes up the true ground plane is obscured by vegetation or other obstacles.

#### 3.1 Minimum point raster image

The ground plane of an environment is a surface that follows the local terrain, including washouts and the berm along the side of the road, but ignores vegetation and other above-ground objects. More specifically, it defines the elevation at which an object would switch between being considered a buried or above-ground target. To represent the ground plane, we define the elevation for each point in a regular grid with a fixed horizontal spacing of 10 cm throughout the environment. This is equivalent to a raster image of the world elevation on a relatively fine-grained scale. Unfortunately, many of the needed raster cells will be obscured by vegetation or be otherwise unavailable from the LIDAR data. To remedy this, we employ a weighted interpolation technique that fills in missing data with a reasonable estimate from the surrounding area.

The simplest approximation of the ground plane elevation for a particular grid cell is the minimum elevation of all LIDAR points with horizontal coordinates that lie within that cell. We refer to this as the minimum point raster image,  $P_{\min}$ . Note that  $P_{\min}$  may not have a value defined for every grid cell due to missing data.  $P_{\min}$  provides a reasonable estimate of the ground plane in open areas such as the middle of the lane or where the true ground is visible, but it may overestimate the height in areas of dense vegetation and cannot provide any estimate at all where there is no data. Nevertheless, the minimum point raster image provides the basis for our interpolation technique.

#### 3.2 Ground plane prediction classifier

Our strategy follows that of McDaniel et al.,<sup>11</sup> in which a classifier is trained to predict for each grid cell in the minimum point raster image, whether or not that point is a good estimation of the ground plane. For example, a point in the middle of the lane where the terrain is flat and all points are at roughly the same elevation would be classified as ground, whereas a point in a bush where the variation is high and the ground is not visible would be classified as not ground. In the cited work, a series of features are extracted for each grid cell  $(i, j)$ , of which we have chosen five for this study:

- $f_1$ : Number of grid cells that are occupied (no missing data) in the 3x3 neighborhood of  $(i, j)$
- $f_2$ : Minimum value in the 3x3 neighborhood of  $(i, j)$  minus the value of  $(i, j)$
- $f_3$ : Value of  $(i, j)$  minus the average value within a 10 m radius
- $f_4$ : Average value in the 3x3 neighborhood of  $(i, j)$
- $f_5$ : Elevation component of the normal vector to the best fit plane of the values in the 3x3 neighborhood of  $(i, j)$

In addition to these five features, we compute the lacunarity<sup>12,13</sup> of the minimum point raster image at four different scales. Lacunarity was originally developed for use in fractal analysis as a means of pattern quantification. It has since been extended for use with grayscale images and is defined for a set of pixels as

$$L = \frac{\sigma^2}{\mu^2}, \quad (2)$$

where  $\mu$  is the mean value of the set and  $\sigma^2$  is the variance. We use a sliding window to compute the lacunarity of  $P_{\min}$  in a local region with window sizes of 3x3, 7x7, 11x11, and 15x15 to give features  $f_6, f_7, f_8,$  and  $f_9$  respectively. These values can be computed very rapidly using integral images. Doing so requires that the input image has no missing values, so in

order to facilitate this computation, we replace missing values in  $P_{\min}$  with linearly interpolated values from the surrounding pixels. While this is sufficient to define a value at every point, it does not make an effort to remove existing vegetation from the ground plane model, which will be accomplished by the classifier and the final weighted interpolation.

In addition to the computed features defined above, we define three more features derived from the elevation statistics of the LIDAR points within each grid cell column  $(i, j)$ :

- $f_{10}$ : Standard deviation of the elevations within cell  $(i, j)$
- $f_{11}$ : Maximum elevation minus minimum elevation of cell  $(i, j)$
- $f_{12}$ : Number of LIDAR points within cell  $(i, j)$

In total, we have defined 12 feature layers for the environment, defined on a regular grid with 10 cm spacing. To train a classifier, we use a small 10m x 10m region of hand-labeled ground truth containing approximately 1.5 million LIDAR points, where each point is labeled as either ground or non-ground, shown in Figure 4. The data is rasterized into 10cm x 10cm columns and the minimum elevation point in each column is saved as the minimum point raster image. The labels of these points define the ground truth image, which is used to train the classifier. The above features are computed for the training data region and are shown in Figure 5.

We use a simple feedforward neural network with 20 neurons in the hidden layer as the classifier. Each pixel in the image is used as a training sample if there is a computed value for each feature. Note that to avoid edge effects, pixels near the image border are ignored. A random subset of 70% of the data from the hand-labeled region is used for training the network, 15% is used for validation, and the remaining 15% is used for testing. Training stops when no improvement is observed on the validation set for 6 iterations. The classifier achieved a classification rate of 96.2% on the testing set and the resulting output image is shown as the last image in Figure 5.

### 3.3 Interpolated ground plane model

After training the classifier, the minimum point raster image of the backbone lane,  $P_{\min}$ , is classified into ground and non-ground regions,  $G_{\text{out}}$ . The resulting image is thresholded at a value of 0.9, creating a binary mask that identifies high-confidence ground plane locations,  $G_{\text{mask}}$ . Although the output of the classifier has a high accuracy rate, there are still regions that are spatially inconsistent, particularly near the edges of vegetation. We apply an erosion filter to the binary mask using a circular structuring element with a 1 m radius to eliminate these boundary regions from the mask. The resulting estimated ground plane mask,  $G_{\text{est}}$ , keeps only the most certain locations, where  $P_{\min}$  is most likely to actually represent the ground plane. The final interpolated ground plane  $G_{\text{interp}}$  is initialized to the values in  $P_{\min}$  at the regions indicated by  $G_{\text{est}}$  and the remaining values are filled in with linear interpolation. A comparison of the original minimum point raster image and the final interpolated ground plane is shown in Figure 6. The interpolated ground plane captures the shape of the terrain, including the berm along the side of the road, washouts, and ditches, while eliminating most vegetation. A rendering of this example scene from the point of view of the vehicle is shown in Figure 7.




Figure 4. Hand-labeled ground truth for the ground plane classifier. (a) The 10m x 10m region used for training containing approximately 1.5 million LIDAR points, rendered here with an ambient occlusion filter to simulate light and shadow. (b) Class labels: ground (red) and non-ground (blue). (c) The training data is rasterized into 10cm x 10cm grid cells, with each column shown with a yellow outline.



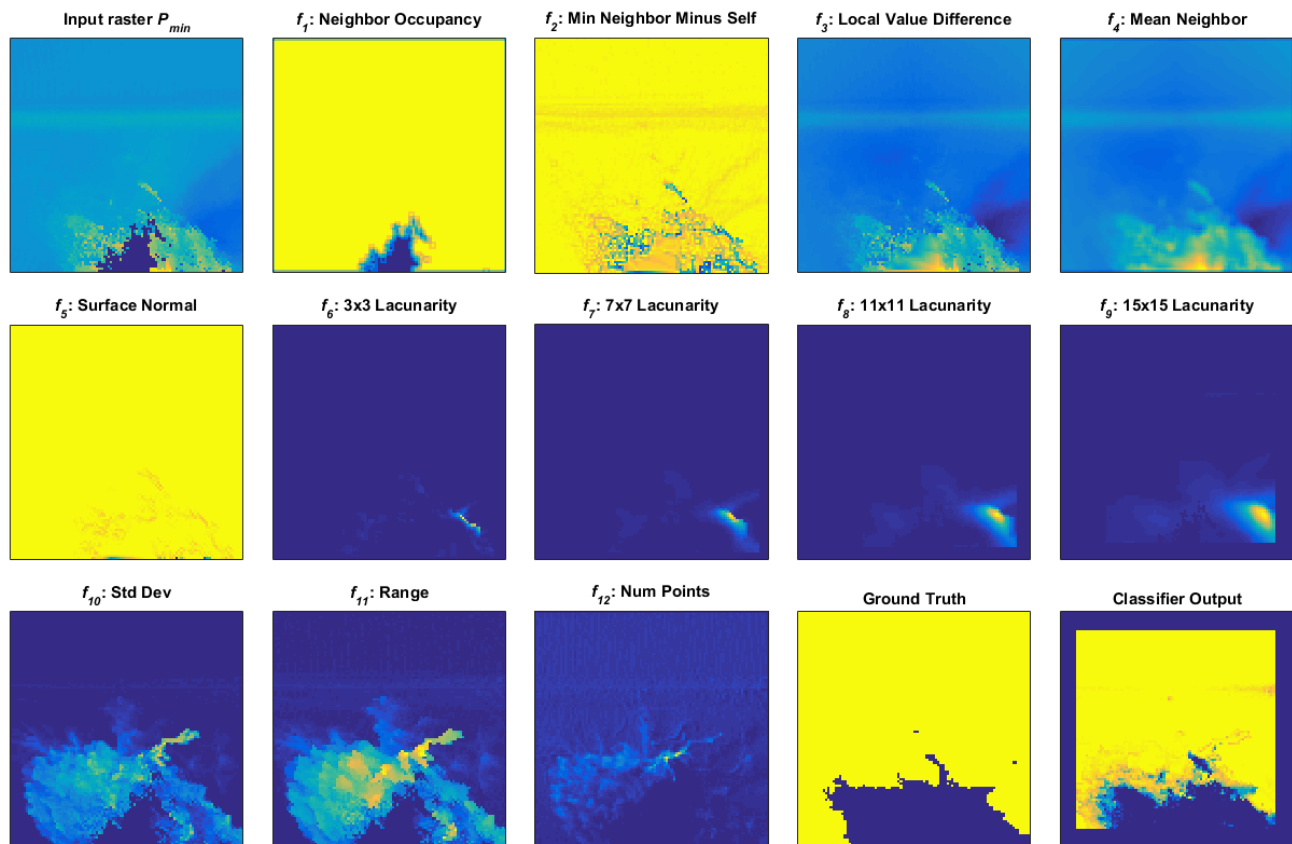


Figure 5. Features extracted from the hand-labeled training data. The top-left image shows the minimum point raster image  $P_{min}$ , followed by the extracted features  $f_1$ - $f_{12}$ . The last two images in the bottom-right show the ground truth used for training and the output of the classifier. The color scale on each image is independently normalized.

### 3.4 Ground plane filtering

The ground plane model developed here has the potential to assist in many tasks related to locating explosive hazards. In particular, it allows for a greater understanding of the three-dimensional scene. Here, we use the ground plane model as a means to filter target predictions from the hit lists of the various sensing modalities. For this study, we are interested in locating above-ground targets where there is some indication of an object above the ground plane. We use the maximum point raster image  $P_{max}$ , generated as before with  $P_{min}$ , but using the maximum instead of minimum elevation value from the LIDAR point cloud in each grid column. The difference between  $P_{max}$  and  $G_{interp}$  indicates the height range of above-ground objects. We apply a threshold  $h_t$  to the difference of these two images that indicates the minimum height range necessary for a region to be considered for scoring. This filtering mask is then dilated with a 20 cm radius circular structuring element to reduce the effects of rasterization, and all hits that fall outside of the mask are eliminated. We consider four values of  $h_t$ : 5 cm, 10 cm, 15 cm and 20 cm. A portion of the filters for the backbone lane are shown in Figure 8.

We apply ground plane filtering to the sparse hits from the previous section with various height range thresholds. The resulting ROC curves are shown in Figure 9. For small values of  $h_t$ , there is virtually no significant filtering and the filtered results are nearly identical to the unfiltered results. As  $h_t$  increases, the value of filtering becomes more apparent. Filtering removes many of the most obvious false alarms from the scoring calculation, which has the effect of shifting the ROC curves to the left. When  $h_t$  becomes too large, around 20 cm, there is too much filtering and correct hits are erroneously removed, resulting in a slight decrease of the detection rate for some of the detectors. Although filtering in this way can improve the computed performance scores, a conservative filtering approach is recommended to avoid eliminating true target hits.

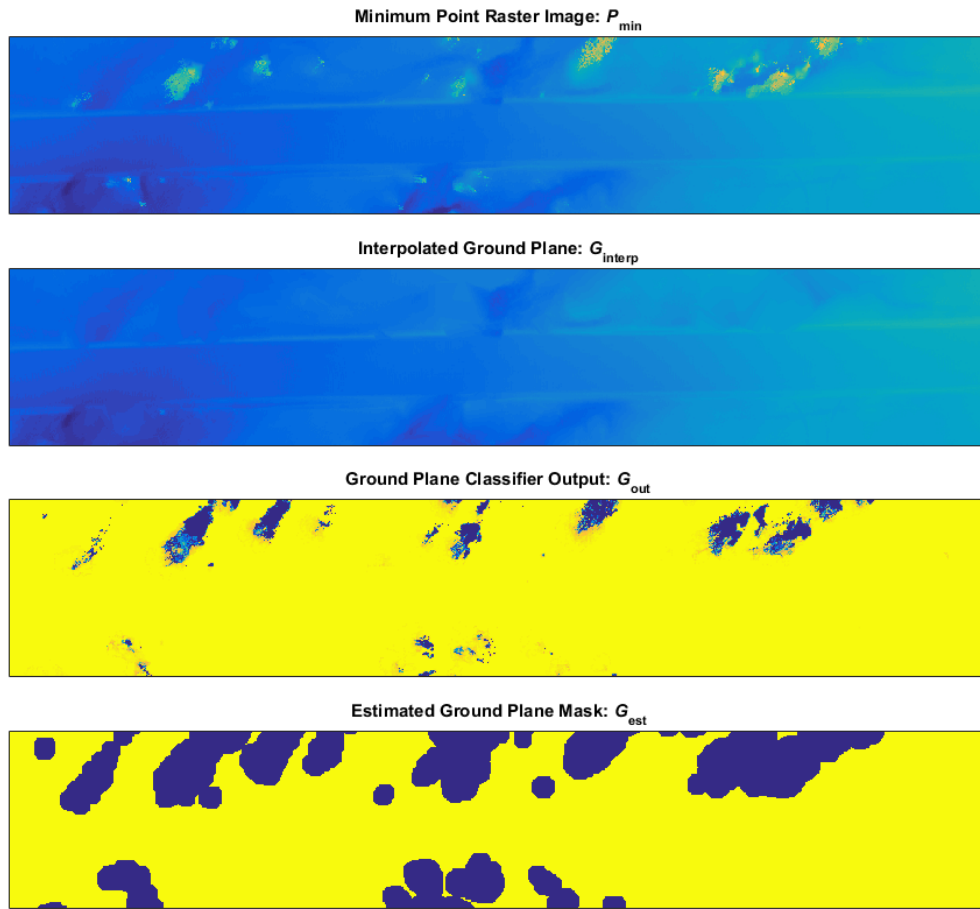


Figure 6. The interpolated ground plane,  $G_{interp}$ , is estimated from the minimum point raster image,  $P_{min}$ , which represents the minimum elevation of the LIDAR points within a grid cell. The ground plane classifier output,  $G_{out}$ , is thresholded and an erosion filter is applied to produce the estimated ground plane mask,  $G_{est}$ .  $G_{interp}$  is initialized to the values of  $P_{min}$  at the locations indicated by  $G_{est}$ , and the remaining locations are filled in with linear interpolation.




Figure 7. (a) A LIDAR rendering of the location in Figure 6 using ambient occlusion to simulate light and shadow. (b) The estimated ground plane model.



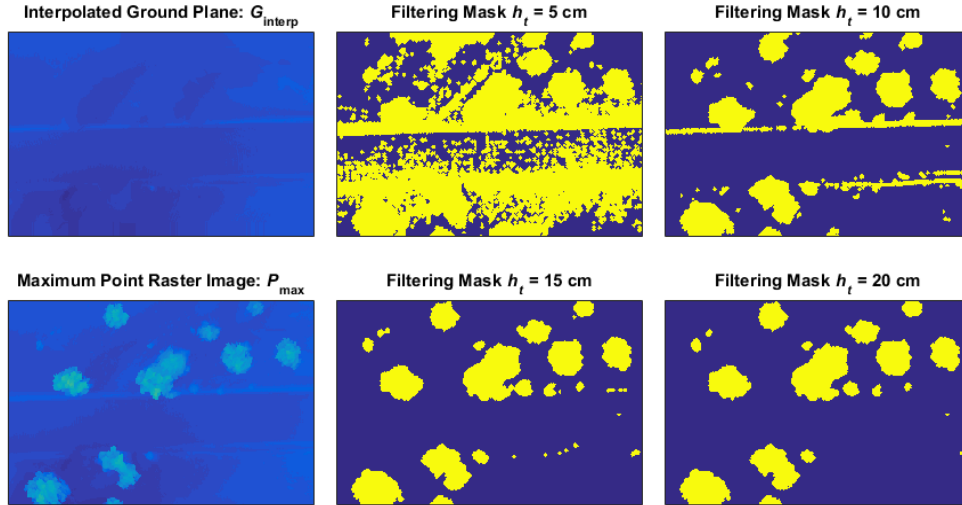


Figure 8. Ground plane filters obtained by thresholding the difference between the maximum point raster image,  $P_{\max}$ , and the interpolated ground plane  $G_{\text{interp}}$ . As the threshold  $h_t$  increases, fewer regions are considered for scoring.

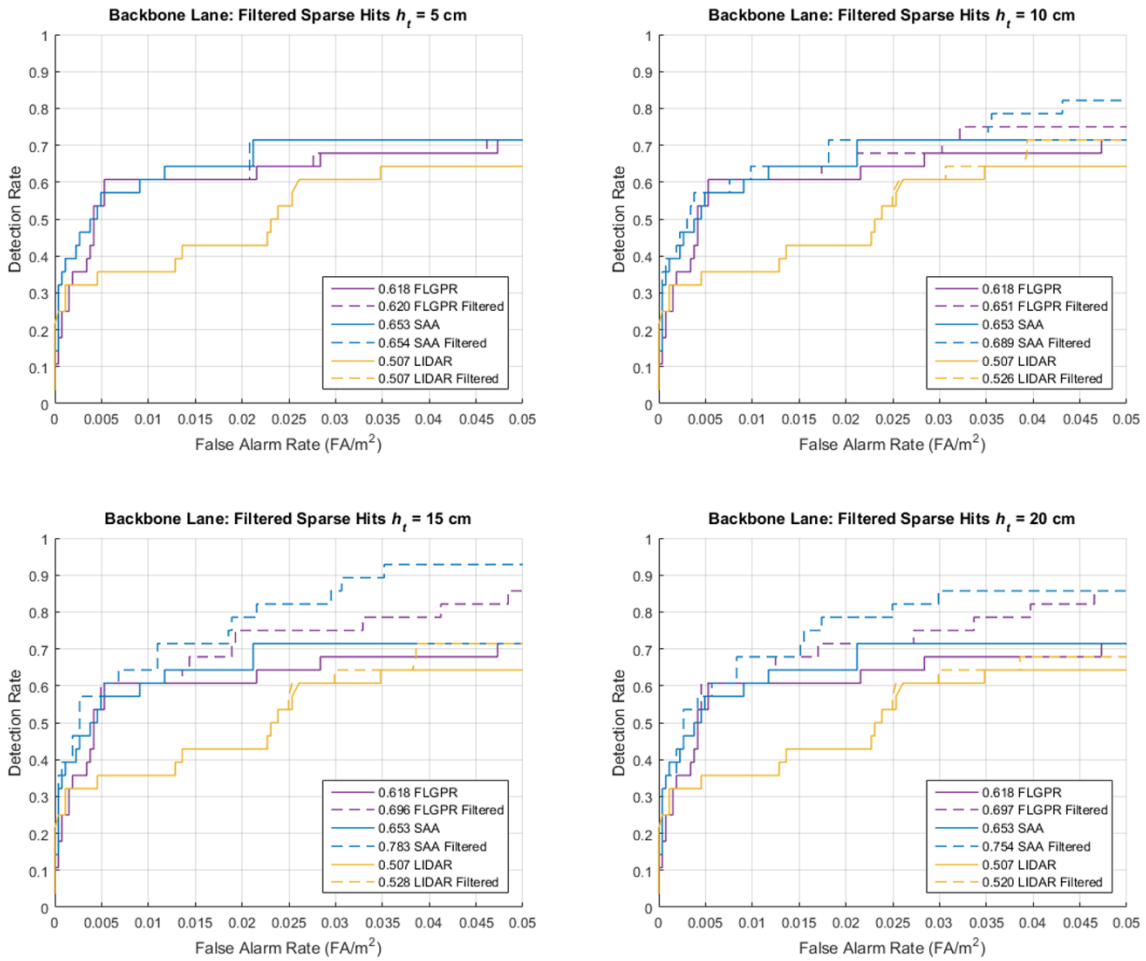


Figure 9. Results of ground plane filtering on the sparse hit lists from the backbone lane. FLGPR, SAA, and LIDAR magnitude detectors are shown for various height range thresholds.

## 4. SENSOR FUSION

In the previous sections, we have demonstrated the capability of each sensing modality independently. Now, we propose an approach to fuse the data from each sensor and generate a combined prediction. We begin by noting that each modality produces a set of hit locations, each with a confidence score. These locations and confidences are scored by labeling hits that fall within a certain radius of a ground truth target as correct hits and all others as false alarms. In this study, we use a 0.5 m radius for scoring. The hits are then ranked by confidence to produce a ROC curve describing the performance of the sensor. Our strategy in the following sections will be to define a method for generating a new set of fused hit locations from a combined confidence score of all sensing modalities.

### 4.1 Confidence surfaces

One way to combine the predictions from multiple sensors is to create a confidence surface for each sensor that spans the physical space of the environment and represents the confidence that there is a target at a particular location. With the ground plane model developed from the LIDAR, we have a foundation on which we can add additional layers of attributes. Each of the three sensing modalities considered in this work is added as an independent layer, with the values for each computed in the most appropriate way for that sensor. For the FLGPR and SAA sensors, the confidence surface can be aggregated from the RX filter responses on each beamformed image. For example, consider the beamformed image in Figure 1. The previous approach identified local maxima in the response of the RX filter and returned these as a set of sparse hit locations, transformed into world coordinates. With a minor modification, the response of each pixel in the beamformed image can be returned as a set of dense hit locations. Rather than aggregating the hits from each frame using the mean-shift algorithm, the hits can be binned into the appropriate grid cells much like points from the LIDAR point cloud. The average response of the RX filter in each grid cell produces the confidence surface for each sensor.

While the above method is appropriate for the FLGPR and SAA sensors, the confidence surface for the LIDAR magnitude is computed differently. One possible approach would be to interpolate the confidence values from the set of sparse hits for each grid cell using a method such as inverse distance weighting<sup>14,15</sup> or kernel density estimation.<sup>16</sup> This technique could be used on any of the sets of sparse hits, including those from the FLGPR and SAA sensors. However, the interpolation function may require several additional parameters to produce a good surface, such as kernel variance and weighting factors that may be difficult to learn. Instead, we use a data-driven approach for producing the LIDAR confidence surface that simply counts the number of points within each grid cell with a magnitude value less than the threshold value of 0.2. The resulting image is smoothed using a circular averaging filter with a 0.5 m radius to approximate the density of all LIDAR points with a magnitude value less than the threshold for a given grid cell.

Once a surface is defined for each sensor over the span of the lane, the surfaces are aggregated into a fused confidence surface. There are many possible approaches for aggregating these layers, and a thorough investigation of the most successful methods is the subject of ongoing research. In this study, we consider only the normalized average of the three layers, which is shown to achieve encouraging results. Each confidence surface is normalized by dividing by the mean value of the layer so that the average confidence value of each layer is 1. It is important to scale each layer appropriately, as the units of each sensor's confidence surface are not directly comparable. The average of these normalized confidence surfaces is computed as the fused confidence surface.

### 4.2 Results

We generate a set of prediction locations from each of the confidence surfaces using the same non-maximum suppression method used to generate sparse hits from the RX filter. The local maxima values of each layer within a 0.5 m radius are returned as hit locations and scored against the known ground truth locations. Additionally, we can apply the ground plane filtering technique described above to reduce the set of prediction locations. Figure 10 shows the confidence surfaces and hit locations for the three sensing modalities and the fused result for the areas around three target locations. Each modality has a unique characteristic pattern that is combined to form the fused confidence surface. The location and magnitude of each hit corresponds to the underlying confidence surface that it came from.

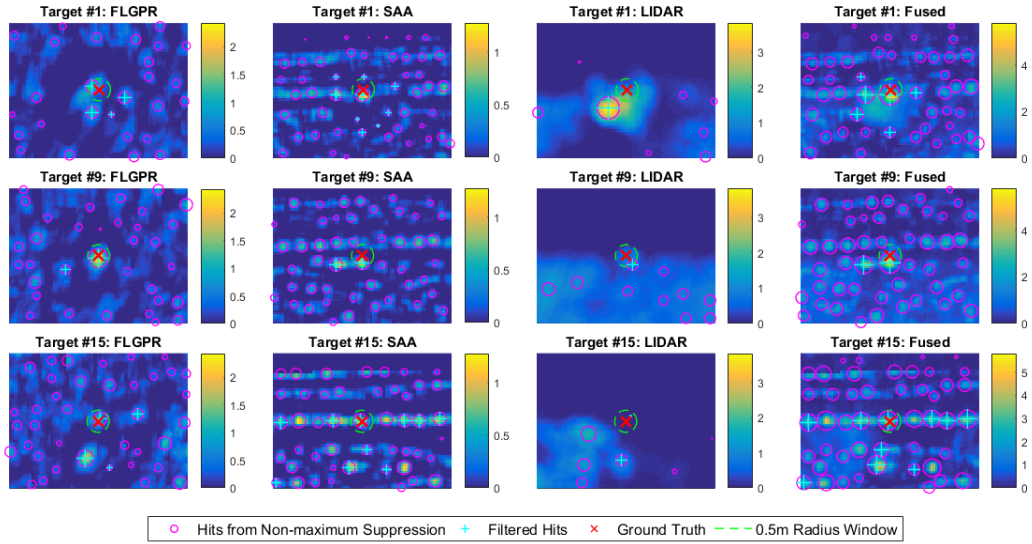


Figure 10. Confidence surfaces produced by the three sensing modalities and the fused result for the areas around three targets. Predicted hits are generated for each surface using non-maximum suppression, and are further filtered using ground plane filtering with a height range threshold of 15 cm.

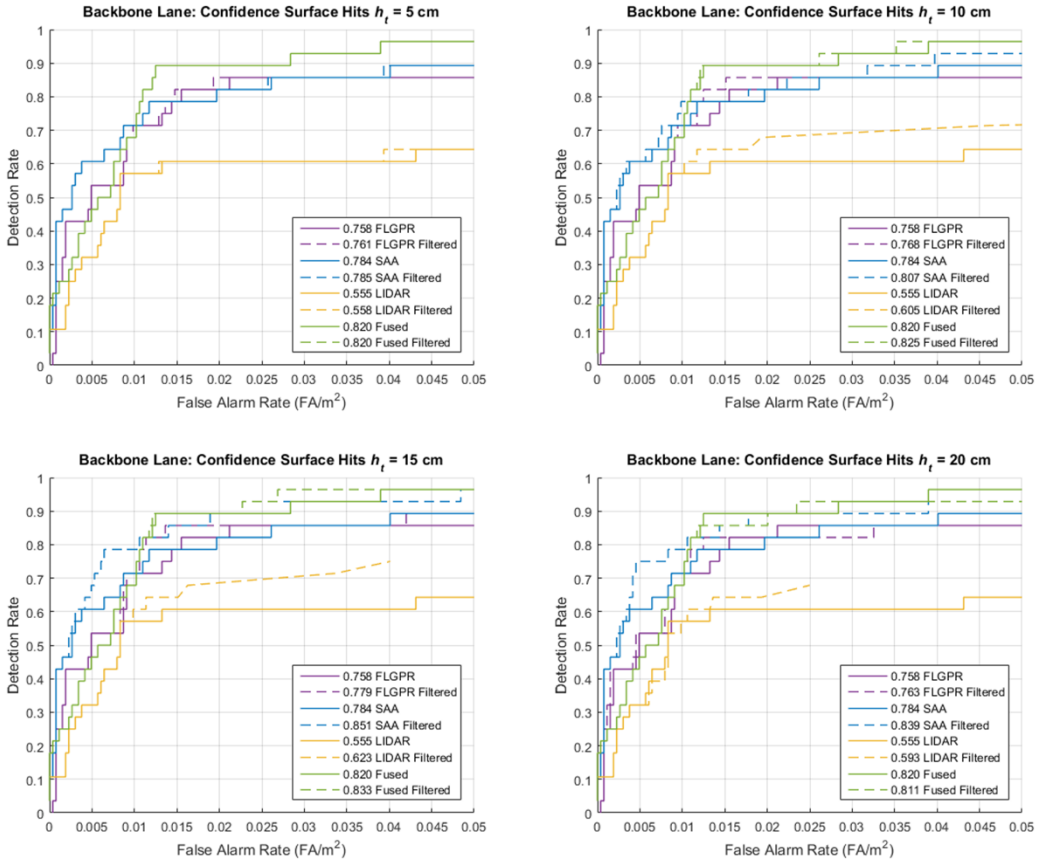


Figure 11. Results of generating hit predictions from dense confidence surfaces using non-maximum suppression on the backbone lane. Ground plane filtering is shown for various height range thresholds. FLGPR, SAA, LIDAR magnitude and fusion results are shown.

Figure 11 shows the results of scoring the predictions generated from the confidence surface layers with ground plane filtering for various height range thresholds. The fused confidence surface appears to outperform the individual sensors for false alarm rates greater than 0.01 FA/m<sup>2</sup>, and a height range threshold of 15 cm further improves the detection rate for all sensing modalities. For large height range thresholds, the ROC curve for filtered LIDAR hits ends before reaching the right limit of the graph because many of the hits are eliminated. Interestingly, all three of the individual modalities saw improved performance over the original sparse hits shown in Figure 9 by using non-maximum suppression to generate predictions from the dense confidence surface. This suggests that confidence surface processing and hit generation by non-maximum suppression may be beneficial even without sensor fusion.

## 5. CONCLUSION

This work has demonstrated an approach for integrating LIDAR point cloud data into the processing workflow for roadside explosive hazard detection in arid environments. We first showed how the LIDAR magnitude response can be used as an independent sensing modality. As an independent sensor, the LIDAR is dependent on having a clear line-of-sight view of each target, and also assumes that the target reflectance in the visual spectrum will be anomalous with respect to its surroundings.

We next presented a model for estimating the ground plane from LIDAR data and showed how it can be used to filter predicted hit locations. The ground plane model has many potential applications beyond prediction filtering that may be explored further in future work. It should be noted that filtering has the potential to remove valid predictions, implying that a conservative threshold should be used. Alternatively, it may be possible to incorporate the filtering step as an additional layer using sensor fusion.

Finally, we showed how to generate a confidence surface for each of the sensing modalities and fuse the results together, which resulted in improved detection performance. This process offers the potential for a great deal of future study to determine the optimal fusing strategy for a particular set of sensing modalities.

## ACKNOWLEDGEMENTS

This work was supported by Army Research Office grant number 57940-EV to support the U. S. Army RDECOM CERDEC NVESD.

## REFERENCES

- [1] Ton, T., Wong, D., Soumekh, M., "ALARIC forward-looking ground penetrating radar system with standoff capability," *Wirel. Inf. Technol. Syst. (ICWITS)*, 2010 IEEE Int. Conf., 1–4 (2010).
- [2] Luke, R. H., Bishop, S. S., Chan, A. M., Gugino, P. M., Donzelli, T. P., Soumekh, M., "A synthetic aperture acoustic prototype system," *Proc. SPIE, Detect. Sens. Mines, Explos. Objects, Obs. Targets XX 9454* (2015).
- [3] Cracknell, A. P., [Introduction to Remote Sensing], Taylor & Francis, London (1991).
- [4] Vosselman, G., Maas, H.-G., [Airborne and Terrestrial Laser Scanning], Whittles Publishing (2010).
- [5] Wilkerson, T. D., Schwemmer, G. K., "Meteorological Lidar developments," *Proc. SPIE, Los Alamos Conf. Opt. '81 0288* (1981).
- [6] Thrun, S., Montemerlo, M., Aron, A., "Probabilistic Terrain Analysis For High-Speed Desert Driving.," *Robot. Sci. Syst.*, 16–19 (2006).
- [7] Stone, K., Keller, J. M., Popescu, M., Spain, C. J., "Buried explosive hazard detection using forward-looking long-wave infrared imagery," *Proc. SPIE, Detect. Sens. Mines, Explos. Objects, Obs. Targets XVI 8017* (2011).
- [8] Stone, K., Keller, J. M., Anderson, D. T., Barclay, D. B., "An automatic detection system for buried explosive hazards in FL-LWIR and FL-GPR data," *Proc. SPIE, Detect. Sens. Mines, Explos. Objects, Obs. Targets XVII 8357* (2012).
- [9] Fukunaga, K., Hostetler, L., "The Estimation of the Gradient of a Density Function, with Applications in Pattern Recognition," *IEEE Trans. Inf. Theory* 21(1), 32–40 (1975).
- [10] Cheng, Y., "Mean Shift, Mode Seeking, and Clustering," *IEEE Trans. Pattern Anal. Mach. Intell.* 17(8), 790–799 (1995).

- [11] McDaniel, M. W., Nishihata, T., Brooks, C. A., Iagnemma, K., "Ground plane identification using LIDAR in forested environments," *Robot. Autom. (ICRA)*, 2010 IEEE Int. Conf., 2831-3836 (2010).
- [12] Nelson, J. D. B., Krylov, V., "Textural lacunarity for semi-supervised detection in sonar imagery," *IET Radar, Sonar Navig.* 8(6), 616–621 (2014).
- [13] Williams, D. P., "Fast Unsupervised Seafloor Characterization in Sonar Imagery Using Lacunarity," *IEEE Trans. Geoscience Remote Sens.* 53(11), 6022–6034 (2015).
- [14] Shepard, D., "A two-dimensional interpolation function for irregularly-spaced data," 23rd ACM Natl. Conf., 517–524 (1968).
- [15] Renka, R. J., "Multivariate interpolation of large sets of scattered data," *ACM Trans. Math. Softw.* 14(2), 139–148 (1988).
- [16] Scott, D. W., Sain, S. R., "Multidimensional Density Estimation," *Handb. Stat.* 24, 229–261 (2004).