# **Carnegie Mellon University** Research Showcase

**Robotics Institute** 

School of Computer Science

1-1-2005

# Analysis and Removal of Artifacts in 3-D LADAR Data

J. Tuley Carnegie Mellon University

Nicolas Vandapel Carnegie Mellon University, vandapel@ri.cmu.edu

Martial Hebert Carnegie Mellon University, hebert@ri.cmu.edu

Follow this and additional works at: http://repository.cmu.edu/robotics



Part of the Robotics Commons

# Recommended Citation

Tuley, J.; Vandapel, Nicolas; and Hebert, Martial, "Analysis and Removal of Artifacts in 3-D LADAR Data" (2005). Robotics Institute. Paper 371.

http://repository.cmu.edu/robotics/371

This Conference Proceeding is brought to you for free and open access by the School of Computer Science at Research Showcase. It has been accepted for inclusion in Robotics Institute by an authorized administrator of Research Showcase. For more information, please contact researchshowcase@andrew.cmu.edu.

# Analysis and Removal of Artifacts in 3-D LADAR Data

J. Tuley, N. Vandapel and M. Hebert Robotics Institute Carnegie Mellon University Pittsburgh, Pennsylvania 15213

Abstract—Errors in laser based range measurements can be divided into two categories: intrinsic sensor errors (range drift with temperature, systematic and random errors), and errors due to the interaction of the laser beam with the environment. The former have traditionally received attention and can be modeled. The latter in contrast have long been observed but not well characterized. We propose to do so in this paper. In addition, we present a sensor independent method to remove such artifacts. The objective is to improve the overall quality of 3-D scene reconstruction to perform terrain classification of scenes with vegetation.

#### I. STATEMENT OF PROBLEM

Our goal is to devise methods for fast and accurate three-dimensional (3-D) representation of natural scenes containing vegetation for autonomous robot operations. We consider both ground and aerial platforms using a LADAR<sup>1</sup> as the primary sensor. Tasks envisioned are obstacle detection, path planning, scene understanding, target detection and tracking in cluttered 3-D environments. We have already reported related work in several publications; see [13], [14], [7]

In such context, it is important to obtain an accurate 3-D reconstruction of the scene in order to be able to capture thin details such as branches, wires or obstacles hidden inside the grass. However, LADAR tends to produce erroneous range measurements. One source of error is due to the LADAR itself: range drift due to temperature variations, systematic errors and spurious random noise, etc. Large efforts have been devoted to the characterization [8] and the modeling of such errors [6]. Such aspect is out of the scope of this paper. The second source of errors is due to the interaction of the laser beam with the environment, for example when the beam hits simultaneously two targets at different distances. Such errors are known in the literature as mixed pixels. We choose to call them mixed measurements to avoid confusion with 2-D image processing and also to avoid the impression that we are using a sensor that produces a dense range image. If early papers recognized their existence, there has been little effort to characterize them exhaustively. In this paper we do so and present a complete taxonomy of mixed measurements. We refer to measurements which are not mixed as correct measurements to emphasize the fact that these returns correctly represent the scene. This analysis was used to develop a sensor-independent method to filter mixed measurements from data. The several variants of our approach use the specific distribution of mixed measurements in 3-D to detect them. By increasing the overall accuracy of the representation of the scene, this method allowed us to improve classification performance.

In the next section we will present the state of the art in sensing in terrain with vegetation and in detection and filtering of mixed measurements. Following that, we introduce the mixed measurement phenomena we have observed and how each was observed. Then we describe our approach to filtering mixed measurements, and finally discuss the performance of these methods. The reader interested in additional details of our work will find a longer version of this paper available as a technical report [12].

#### II. STATE OF THE ART

# A. Sensing

Depending on the application, several sensing modalities can be used in terrain containing vegetation: 1) color and multispectal cameras, for 2-D scene classification to identify vegetative cover; 2) millimeter wave (MMW) radar or synthetic aperture radar, to see through the vegetation to detect targets or obstacles even at long range but with low resolution. Global properties of the area illuminated can be extracted from the pulse returns but cannot be used for high fidelity reconstruction [10]; 3) echolocation to see through grass at short range and low resolution [11]; 4) stereo vision with specialized pre-processing image filters to reconstruct forest scenes in 3-D at short range [4], [5]; 5) LADAR sensors, which are the most commonly used.

#### B. Mixed measurement detection and filtering

The existence of mixed measurements is a long standing problem but very few papers explicitly address it.

Continuous Wave (CW) LADAR are particularly prone to such phenomena. In [9] Kweon presents a characterization of the Perceptron laser rangefinder. This work is extended to the Erim laser in [8]. In [16] mixed measurement observations from a SICK laser similar to the one we use are reported but no interpretation or analysis is provided.

In this paper, we follow a similar experimental procedure to the one introduced by Matthies in [10]. Mixed measurements for a SICK laser are reported and some

<sup>&</sup>lt;sup>1</sup>Light Amplification for Detection And Ranging

partial characterization is provided in terms of range of observability of the phenomenon. We extend this work and in addition propose a method to filter the mixed measurements.

In [1], [2] Adams explicitly models the mixed phenomenon for a scan line CW LADAR and proposes a method to detect discontinuities. In [3], Adams proposes a feature detection scheme based on a smoothness constraint for data from outdoor scenes. An Extended Kalman Filter is used to predict range measurements. Such an approach is difficult to generalize for lower density lasers and even more so when the robot navigates in terrain with non-continuous surfaces but with higher clutter due to vegetation and thin structures, such as in our case.

#### III. MIXED MEASUREMENT PHENOMENA

Our experiments reveal several classes and types of mixed measurement phenomena. Some of these are the result of using LADAR in general, while some are sensor dependent. As such, we will first describe the sensor which recorded these phenomena (III-A), the conditions which generated them (III-B), and finally describe our observations (III-C).

# A. Sensor Description

The experiments described in this paper used a SICK LMS-291-SO5, a Time of Flight (TOF) pulsed laser mounted on a pushcart. We used several other sensors, such as a Z+F laser, a Riegl LD90-3000, a Riegl LMS-Z210i and also the GDRS mobility LADAR used on the eXperimental Unmanned Vehicle (XUV), but we only report experiments using the SICK laser due to its availability and common usage in robotics sensing research.

1) Physical configuration: The SICK laser is a onedimensional scanning LADAR. To obtain a 3-D scene we mount two SICKs together and independently actuate them with Animatics SM3420D Smart Motors. The SICK itself has several scanning modes; we have chosen to scan 100 points at 0.25° separation. A typical scan consists of 800 steps of the motor. One unit is mounted in an upright position at a height of approximately 127 cm and actuated on its x axis, so that as the laser scans horizontally the mounting sweeps vertically. The second is mounted on its side at approximately 138 cm, sweeping around the z axis, so that the scan line moves horizontally. The laser height was chosen to produce a large Field of View (FOV). All of the data presented here was collected with the sidemounted laser. This configuration has a field of view which is 115° wide and 100° high, producing an  $800 \times 400$  range measurements.

2) Beam characteristics: The SICK instrument projects a beam which is rectangular in shape, with two regions of high intensity at either end of the beam, as in Fig. 1. Such beam shape is consistent with generation by a double stacked laser diode. The image was captured using an infrared camera<sup>2</sup>. The beam is swept by a rotating mirror,

set at 45° to the original beam. This has the effect of rotating the projection of the beam, so that at the edges the high-intensity regions are perpendicular to the scan line. At the center, these regions are parallel to the scan line. Fig. 1 illustrates the projection of the beam at extreme left (top image) and center through right areas (lower two images). In the configuration we use, the projected beam is 0.68° wide and scans through a 100° field of view with 1° angular resolution.

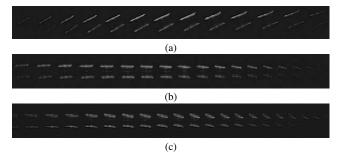


Fig. 1. **Portion of SICK scan line.** Captured with an infrared camera, this sequence show the pair of high intensity area produced by the beam and its rotation as the laser sweeps.

3) Data interpretation: The data returned by the laser is converted to rectangular coordinates with (x, y, z) representing the right, forward and upward offsets (respectively), i.e. the sensor is at the origin. All of our analysis and filtering is performed on this data, which only contains the points in 3-D space. We do not keep the connectivity (row and column identification) of the points. If the data is collected from a moving sensor platform and stored as absolute coordinates, we also keep the location of the sensor with each scan. Thus, our approach may be used with data from other 3-D TOF laser scanners by considering single scans or by accumulating data over time as the scanners move.

These points are read into a sparse voxel representation and accessed using a hash key based on the relative position of the points. This structure combines the essential features of several points into one, saving memory and reducing the computational complexity of subsequent use of the points [14].

4) Mathematical model: In order to better understand the observed mixed measurement phenomena, we employ a model of the return signal of the laser beam. By considering the portion of the beam reflected back by the target and that returned by the background (ground plane or vertical wall), it is possible to explain qualitatively the phenomena observed.

To model the energy waveform we choose a simple reflectance model:

$$E(t) = \sum_{i \in targets} \delta(t - t_i) A_i \frac{\rho_i \cos(\theta_i)}{r_i^2}$$
 (1)

with:  $\rho$ , the reflectance; r, the range;  $\theta$ , the laser incidence angle; A, the normalized area illuminated;  $\delta$ , the Dirac function. We normalize the area illuminated to conserve the flux of energy.

<sup>&</sup>lt;sup>2</sup>This camera was generously provided by the NAVLAB group

The non-zero length of the pulse and the dynamic of the receiver will smooth the energy waveform. We represent both effects as a Gaussian  $G(t_i, \sigma_i)$ .

$$E(t) = \sum_{i \in targets} \frac{1}{\sqrt{2\pi}\sigma_i} e^{-\frac{(t-t_i)^2}{2\sigma_i}} A_i \frac{\rho_i \cos(\theta_i)}{r_i^2}$$
(2)

Several techniques can be used to detect pulses [15]: threshold, center of gravity, maximum, zero crossing of second derivative and constant fraction discriminator (CFD). CFD methods are not sensitive to signal amplitude change and are robust to pulse width change; these methods look for the zero crossing of the difference between the attenuated signal waveform and a version of the original waveform delayed by a fraction of the pulse width. We use such an approach.

This model improves our understanding of several phenomena, but it is not a perfect solution. One problem is that the beam intensity is not truly homogeneously distributed. Another is the lack of complete manufacturer's technical data on essential characteristics of the laser such as beam duration and trigger condition, which increases the error of the current model. However, the model is still suitable for explaining most phenomena, which will be further discussed in subsection III-C.

#### B. Experimental Setup

To simulate vegetation-like structures in controlled conditions, small PVC pipes and dowel rods were used in a variety of scenes, both indoors and outdoors. Figs. 2, 3, 4, and 5 show sample data from each set of experiments. A similar rod set-up was used in [10] and inspired us. Table I contains a short description of the experiments performed. Additional details can be found in [12].

TABLE I
Test series description

Series	Parameters tested	Target	Clutter
1	large linear structures	board	PVC pipes
2	thin linear structures	board	dowel rods
3	sensor location	dowels	ground, background
4	background distance	dowels	background
5	natural clutter	dowels	vegetation
6	target, background distance	dowels	background

We used a  $26 \times 41$  cm Plywood board, 6 cm PVC pipes and dowel rods varying in size between 3.2 and 22.2 mm. To simulate vegetation clutter with the PVC pipes or the rods we use a Poisson distribution to determine their locations. Usage of wood targets and clutter ensures that differences in reflectivity are not the source of phenomena wellocathenomena observed can be divided into several classes: 1) mixed measurements occurring because of overlap of the laser beam onto two objects at the edge of an object; 2) mixed measurements occurring due to target structures thinner than the beam; 3) mixed measurements occurring due to particular characteristics of the SICK laser; 4) and mixed measurements from ground interaction. Fig. 6 shows the different phenomena observed. The LADAR and two rods (the thinnest and the widest)

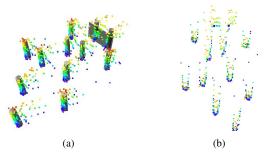


Fig. 2. Example data from test series 1. (a) Side view of the data and model, (b) Top view of data only. Color indicates height (blue is low, red high; scale is relative). Positive directions are: x right, y upward, z out of the page.

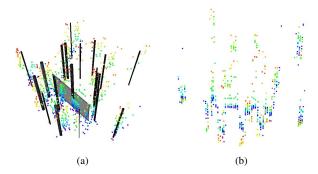


Fig. 3. Example data from test series 2. (a) Side view of the data and model, (b) Top view of data only. Color indicates height (blue is low, red high; scale is relative). Positive directions are: x right, y upward, z out of the page.

are static. The background is positioned at three different locations. The bottom part of Fig. 6 shows a side view of the set-up. The top part of the same figure shows different mixed measurement phenomena.

1) At edges: Mixed measurements appearing at the edge of objects are the most common. These occur when a beam is projected across the edge of an object, and (with a TOF laser such as the SICK) appear between the two objects. The actual position of the return is dependent on the relative strength of the return from each of the objects and the absolute distance to the target. The strength of the return is (as we model it) dependent on distance, reflectivity and orientation. If either the target or background return is sufficiently strong, no mixed measurement occurs.

This type of mixed measurements occurs at the edge of the target's range shadow. For linear structures, this is a plane (with tolerances allowing for noise), as in Fig. 7.

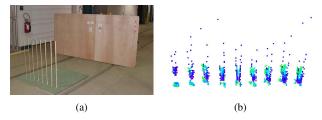


Fig. 4. **Example data from test series 4.** (a) Scene, (b) Top view of data only. Color indicates height (blue is low, red high; scale is relative). Positive directions are: x right, y upward, z out of the page.

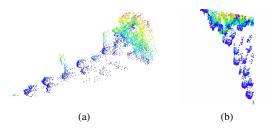
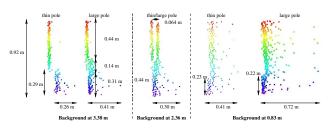


Fig. 5. Example data from test series 5. (a) Side, (b) Top views. Note dowels (green regions in left/center of (b)) and vegetation. Color indicates height (blue is low, red high; scale is relative). Positive directions are: x right, y upward, z out of the page.



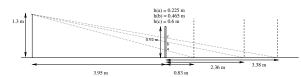


Fig. 6. Mixed measurement phenomena.

When multiple scan lines are considered, as in our case, we extend our area of interest to a thin box about this plane, due to noise and differing interactions at each beam. Because of their shape around thin structures, we refer to these phenomena as "V-shaped" mixed measurements, including when only one leg of the V is present, e.g. for a flat board.

2) With thin structures: In the cases where the target has less angular width than the width of the beam, the beam is never fully on the target (Fig. 8). In these situations, it is possible that no measurement reports the true position of the target, but instead all points may be considered mixed measurements.

As with most thin structures, the V shape often occurs, but the body of the return may be erroneous, with the edges in the typical V-shape. For example, the shape of a thin rod may be maintained, but the entire measurement is dislocated.

3) SICK-specific: As noted in Section III-A.2, the beam of the SICK laser has two high intensity edges.

For thin structures, it is possible that these regions fall to either side of the target. Then the reflected beam is affected

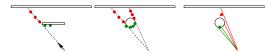


Fig. 7. **Mixed measurement taxonomy: edge-occurring.** Beams and the points which they produce are colored red for mixed or green for valid measurement.

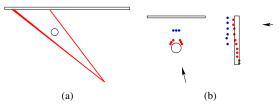


Fig. 8. **Mixed measurement taxonomy: thin structure.** (a) SICK beam (red) overlapping and bracketing a rod. (b) Points corresponding to normal mixed measurements are colored red; ghost measurements are shown in blue. On the left, a top view. On the right a side view.

more by the background. However, as the sensor sweeps, other scans will fall with one or the other high intensity edge on the target. These measurements are mixed, but contain more information from the target. Thus it can occur that the central portion of the target is measured to be farther away than the edges (we observe as much as 6-10 cm of displacement), which may be at correct (or nearly correct) positions.

Depending on the beam density of the sensor, this phenomenon can give the appearance of separated objects where only one exists. Specifically, it is possible for an "object" to appear in the range shadow of the actual target, with some returns from the correct location as well. We term this type of error a *ghost* return, and observe these only with the SICK laser and thin structures which are perpendicular to the sensor's scan line.

It is also usual for the ghost to appear as a real return, including V-shape mixed measurements behind the ghost. The edge returns also have V-shape measurements on the outside edge. With structures of width close to the gap in the beam, the center points are nearly in the correct location, but on either side ghosting occurs. This gives rise to a bottle-shaped ghost instead of one which emulates the actual target shape.

4) Ground interaction: The mixed measurements formed when a beam is split between an object and the ground (or a similar horizontal surface), have the same cause as the above classes, but exhibit different characteristics, as illustrated in Fig. 9.

As the sensor scans, the change in distance with each beam is greater for a horizontal structure than a vertical one (since beam is closer to horizontal than vertical in the space under consideration). For instance, if the target is a board, the scan projects beams which first intersect only the floor, then a combination of floor and target, and then only the target. Mixed measurements occur when the beam reflects from both. These measurements describe a slightly curved join between the floor and board. The top edge of the board is also curved, where the beams overlap between the target and the floor beyond.

The interaction between thin structures and the ground plane is also characteristic. In this case, since the laser intersects the target and the floor behinds (unlike with a board, where some forward interaction occurs), the structure appears to be disjoint, with the lower portion sloped backwards and the upper portion correctly located (assuming no other background to produce mixed measurements).

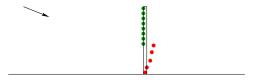


Fig. 9. Mixed measurement taxonomy: ground interaction. Valid points are shown in green; mixed points in red.

#### D. Configurations Producing Mixed Measurements

Mixed measurements do not occur in every scene; they occur only when objects are sufficiently near to each other. Our experimentation shows that the foreground and background objects which are hit by the beam must be within 2.5 m for mixed measurements to appear. However, our results do not indicate a lower bound on the occurrence of mixed measurements; on the contrary, when the foreground and background are separated by less than 1 m the mixed measurements observed spread completely between the foreground and the background in a solid group, making filtering more difficult. Fig. 10 shows where mixed measurements were observed using individual rods of several diameters.

#### IV. MIXED MEASUREMENT FILTERING

#### A. Approach

Because all types of mixed measurements observed exhibit the characteristic V in the direction of the laser ray at a point, we filter for mixed measurements only in the region expected to contain mixed measurements. Note that this technique is not capable of displacement of the body of the return (e.g. ghosts, displaced thin structures). These classes of mixed measurements, while erroneous as observed, could be produced by certain objects. Thus we do not wish to automatically remove them from a scene at this time.

Our algorithm examines each point in the scene in turn. To determine if a point is mixed, we generate a bounding box whose major axis is the ray from the sensor (origin) to the point of interest. The point of interest is at the center of the front face of the box. We compute the saliency features for the points falling in this region. In this regard, we follow the approach of [14] but with a non-isotropic interest region. The original method looks at the local point distribution in space and uses a Bayes classifier to produce the probability of belonging to three classes vegetation, solid surface and linear structure. The method takes as input a sparse set of 3-D points. At each point the scatter matrix is computed using a predefined isotropic support region. The decomposition in principal components of this matrix leads to the definition of three saliency features characterizing the 3-D points' spatial distribution as "random", "linear" and "surface". This information is used in two ways to decide whether or not a given point is a mixed measurement. In one variation, the algorithm triggers when the "surfaceness" of the region is higher than

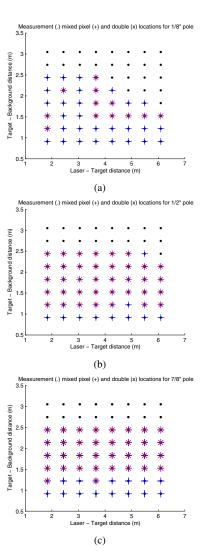


Fig. 10. Appearance of Mixed Measurements. Black dots  $(\cdot)$  indicate configurations where measurements were taken. Blue crosses (+) indicate mixed measurements and  $\times$  (red) denote ghost-type mixed measurements. (a) 3.2 mm pole. (b) 12.7 mm pole. (c) 22.2 mm pole.

the randomness and linearness. The other variation triggers when the normal of the surface described by the region is orthogonal to the laser ray. These methods are described in more detail below.

# B. Algorithm Variations

- 1) Surfaceness triggered: This group of algorithms computes the saliency features for the interest region of a point, and begins marking if the features indicate that the region is a surface. Marking techniques are outlined in subsection IV-C (all marking variations were employed in algorithms with this trigger).
- 2) Surface normal triggered: To generalize the classes of mixed measurements which can be filtered (i.e. rely less on the orientation of the interest region), this group of algorithms computes the saliency features but then compares the surface normal vector to the laser ray. If the angle between these vectors is "close" (according to some pre-determined threshold angle) then the trigger is set and

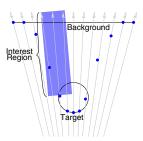


Fig. 11. **Filtering of mixed measurements.** Grey lines represent laser beams, blue points are detected points, and the blue rectangle is the interest region. (Not to scale.)

marking occurs. With this trigger the naïve, aggressive and exclusive blind window methods of marking were used. This variation was created because we realized that regions where the mixed measurements do not occur in a vertical plane would likely not represent a vertical surface. This trigger mechanism allows for surfaces of all orientations.

3) Surfaceness and surface normal triggered: Because the surfaceness and surface normal triggered algorithms differ in which types of mixed measurements they are able to filter, one variation examined combined the Aggressive algorithms from both groups (if either triggered, then the interest region was marked as mixed). This algorithm is called Surfaceness/Surface Normal – Aggressive.

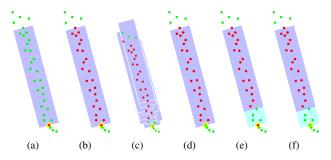


Fig. 12. **Marking variations.** In each figure, the blue area is the interest region. Green points are left as correct points; red are marked as mixed. The point labeled in yellow is the interest point. (a) Naïve. (b) Aggressive. (c) Accumulator (requiring four votes). (d) Exclusive. (e) Blind Window. (f) Exclusive Blind Window.

#### C. Marking variations

Once the trigger condition is met, one of the following variations is used to mark measurements as mixed (shown in Fig. 12):

- Naïve marks only the interest point itself.
- Aggressive marks all points in the interest region.
- Accumulator counts the number of times each measurement is reported as mixed during the traverse of all points, and then reports measurements as mixed only if some threshold of votes is reached.
- Exclusive marks all points in the region except the interest point.
- Blind window marks as aggressive, except that it leaves a "blind window" near the original point unmarked.

 Exclusive blind window is similar to blind window except that it also leaves the point of interest unmarked.

#### V. PERFORMANCE CALCULATION

The filtering algorithms and their variations described in Section IV-B exhibit varying degrees of success in detecting mixed measurements and avoiding false detection of correct measurements. This section presents the performance data for each variation.

It should be noted that the behavior of each algorithm varies significantly with its computation parameters. The variables which may be altered are the cell size of the sparse voxels, which affects the precision; the box width, height, depth, which control the region of points admitted for calculation; box roll (about the beam axis), for finding mixed measurements in non-vertical planes, blind window depth, and maximum angular deviation of surface normal from orthogonal to laser ray.

#### A. Description of Metrics

To evaluate the performance of the various filters, a 25scene corpus was first manually separated into correct and mixed measurements. The filters were tested against this data with a range of parameters to determine those which gave the best results. After filtering, the number of correct and incorrect classifications of both types was counted. For the purposes of these metrics, we define a "positive" result to be classifying a measurement as mixed, and a "negative" result to be classifying as correct. So we have true positive (tpos), false positive (fpos), true negative (tneq), and false negative (fneq) counts. The total number of positive and negative measurements are denoted by nposand nneg, respectively, and the counts for the ground truth are labeled gpos and gneg. The total number of measurements classified  $N_C = tpos + fpos + tneq + fneq$ and the total number of measurements in the original scene  $N_G = gpos + gneg$  also are used as intermediate values in computing other metrics.

The true positive rate tpr = tpos/(tpos + fneg) represents the percentage of measurements correctly identified as mixed. The false positive rate fpr = fpos/(fpos + tneg) is the portion of correct measurements which are wrongly marked as mixed. Accuracy  $acc = (tpos + tneg)/N_C$  is the percentage of correctly-marked measurements of both types. Because computing saliency requires at least three points, some interest regions may not contain enough support for the computations. In these cases, the points are discarded and are not reported as either correct or mixed, unless part of a region being aggressively marked. The loss rate,  $loss = N_C/N_G$ , is the measure of how many measurements in a scene are not reported. Generally, this is only a factor with Naïve algorithms, since very few are discarded by marking all of the points in the interest regions.

Finally, in an effort to condense analysis, a scoring system was developed by which correctly classified measurements raise the score while false reports lower the

TABLE II **Algorithm performance.** These data represent the best performance each algorithm exhibited during testing.

Algorithm	tpr	fpr	acc	Score
Surfaceness triggered				
Naïve	0.222	0.268	0.713	0.417
Aggressive	0.803	0.643	0.729	0.513
Accumulator	0.516	0.389	0.701	0.466
Blind Window	0.764	0.422	0.750	0.548
Exclusive	0.792	0.609	0.725	0.505
Exclusive Blind Window	0.668	0.147	0.812	0.657
Surface Normal triggered				
Naïve	0.188	0.217	0.722	0.410
Aggressive	0.811	0.629	0.764	0.573
Exclusive Blind Window	0.686	0.153	0.814	0.640
Combination triggered				
Aggressive	0.936	0.736	0.744	0.596

score, as do lost measurements. The number of "points" earned by each measurement is divided by the maximum possible score for the scene (i.e., for perfect classification), yielding an estimate of relative performance. By losing points when making mistakes, an algorithm's score will always be lower than its accuracy (unless both are 1). In the data shown, 5 points were added to a score for a correct classification, but 3 were removed for a false negative, 4 for a false positive, and 1 for a loss. These scoring choices are somewhat arbitrary, based on the fact that our scene classification algorithm (which uses the scene data after the mixed-measurement classifier cleans it) is more susceptible to error when mixed measurements are present than when some accurate data is lost. Thus we more heavily penalize a mixed measurement going unnoticed than a correct or non-classifiable measurement being removed.

#### B. Classification performance

Table II shows the best average (over the 25 scenes used in testing) performance for each algorithm. Note that each score in the table does not always come from the same set of parameters to the algorithm, with the exception of tpr, which is presented with its corresponding fpr (thus these are not necessarily the best false positive rates obtainable by each algorithm).

#### C. Efficiency

Written in C++, our proof-of-concept implementation is able to compute localized saliency features for approximately 6000 voxels per second. Naïvely marking when triggering on surfaceness is done in quadratic time and, with our hardware<sup>3</sup> and test data (averaging roughly 1000 points per scene, with ground removed *a priori*), can be computed for approximately 1.3 million points per second; aggressive marking can be performed in cubic time on 7500 points per second. When triggering on the surface normals, naïve marking runs on approximately 1 million points per second; running aggressively processes 7700

points per second. For comparison, the GDRS laser on the XUV records on the order of 100,000 points per second.

#### D. Performance against Point Clouds

One concern with the approach which we present here is the tendency to always find planar surfaces when considering a near-planar interest region. Our results show this is clearly the case for linear structures, but not for point clouds (e.g. vegetation); thus we tested on uniformly distributed random points. Triggering on surfaceness, no mixed measurements were reported by the Naïve algorithm with the most successful parameters for five spaces of 1000 m<sup>3</sup> with 300, 1,500, 2,000, 6,000 or 15,000 points in each region. Testing on the same data sets, surface normal-triggered Naïve reported 9 mixed measurements only on the most dense cloud.

#### E. Performance Overview

As can be seen in Table II, the optimal algorithm examined thus far is the exclusive blind window with either triggering method (surface normal-triggered is more accurate, though surfaceness-triggered shows fewer false positives). We hypothesize that this is due to the higher variation in the normals of the surfaces fitted to the points, rather than the saliency properties of the points.

With the point density given by our laser, this method is fast enough for real-time scenarios. However, given a peak 80% accuracy rate, with 15% of correct measurements misclassified most of which are at object edges and are unlikely to affect the process of finding objects, though they may affect estimates of size), the process is capable of removing mixed measurements from a scene. We can also say with reasonable certainty that the algorithm behaves as desired and only reports mixed measurements when planar structures exist, and is not so over-sensitive that it detects them in data which cannot be known to contain mixed measurements.

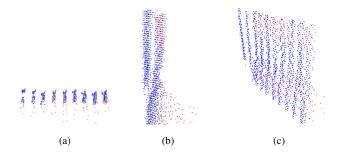


Fig. 13. **Sample filtered data.** Correct measurements denoted in blue, mixed in red. Data taken from test series 4 and filtered with the surfaceness trigger and Exclusive Blind Window marking. (a) Top, (b) side, and (c) perspective views.

### VI. CONCLUSIONS

It is possible to quickly and accurately enhance scene quality by removing mixed measurements. As seen in Fig. 13, the filtering techniques presented here are able to remove most of the mixed measurements from a scene

<sup>&</sup>lt;sup>3</sup>A dual Intel Xeon (2.8 GHz) machine with 3 GB of memory and Red Hat Linux version 9. Note: our software does not make use of both processors.

involving thin linear structures. Using existing techniques for discovery and removal of ground and vegetation, threedimensional laser imagery becomes a useful and effective tool for object recognition and feature detection.

Efficiency is as yet an issue. The algorithm is not fast enough to be run on entire scenes, so at this time we first remove the ground plane and crop the scene to our region of interest. When the entire operation of cropping the scene and filtering mixed measurements is performed, a single frame may require one to ten seconds to process. Thus its use in real-time applications is questionable.

As the project progresses, we will continue to improve algorithm efficiency and accuracy. One optimization is to filter by voxels instead of by the raw points themselves. We continue to investigate more effective filtering algorithms and marking methods and the user of other sensors, which would serve to both validate our methods and to allow us to increase the generality and sensor-independence of our algorithms.

#### ACKNOWLEDGMENTS

Prepared through collaborative participation in the Robotics Consortium sponsored by the U.S Army Research Laboratory under the Collaborative Technology Alliance Program, Cooperative Agreement DAAD19-01-209912. The authors would like to thank J. Ryan Miller, Arne Suppe and Christoph Mertz for sharing their insights about the SICK laser.

# REFERENCES

- M.D. Adams. amplitude modulated optical range data analysis in mobile robotics. In *IEEE International Conference on Robotics and Automation*, 1993.
- [2] M.D. Adams. The interpretation of phase and intensity data from a.m.c.w. light detection sensors for reliable ranging. *International Journal of Robotics Research*, 15(5), 1996.
- [3] M.D. Adams. On-line gradient based surface discontinuity detection for outdoor scanning range sensors. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2001.
- [4] A. Ansar, A. Huertas, L. Matthies, and S. Goldberg. Enhancement of stereo at range discontinuities. In CTA conference, 2004.
- [5] A. Ansar and L. Matthies. Improvement of stereao in complex scenes. In Collaborative Technology Alliance Conference, 2003.
- [6] Y.D. Chen and J. Ni. Dynamic calibration and compensation of a 3d laser radar scanning system. *IEEE Transactions on Robotics and Automation*, 9(3):318–323, 1993.
- [7] M. Hebert and N. Vandapel. Terrain classification techniques from ladar data for autonomous navigation. In *Collaborative Technology Alliances Conference*, May 2003.
- [8] Martial Hebert and Eric Krotkov. 3d measurements from imaging laser radars: how good are they? *International Journal of Image* and Vision Computing, 10(3):170–178, 1992.
- [9] In So Kweon, Regis Hoffman, and Eric Krotkov. Experimental characterization of the perceptron laser rangefinder. Technical Report CMU-RI-TR-91-01, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, January 1991.
- [10] L. Matthies, C. Bergh, A. Castano, J. Macedo, and R. Manduchi. Obstacle detection in foliage with ladar and radar. In *International Symposium on Robotic Research*, 2003.
- [11] J. Murray, R. Rosasco, and M. Holmes. Active acoustic sensing of stumps in tall grass. In *Collaborative Technology Alliance Conference*, 2003.
- [12] J. Tuley, N. Vandapel, and M. Hebert. Analysis and removal of artificats in 3-d ladar data. Technical Report TR-04-44, The Robotics Institute, Carnegie Mellon University, 2004.

- [13] N. Vandapel, R. Donamukkala, and M. Hebert. Experimental results in using aerial ladar data for mobile robot navigation. In *International Conference on Field and Service Robotics*, 2003.
- [14] N. Vandapel, D. F. Huber, A. Kapuria, and M. Hebert. Natural terrain classification using 3-d ladar data. In *IEEE International* Conference on Robotics and Automation, 2004.
- [15] W. Wagner, A. Ullrich, T. Melzer, C. Briese, and K. Kraus. From single-pulse to full-waveform airborne laser scanners: Potential and practical challenges. In *International Society for Photogrammetry* and Remote Sensing XXth Congress, 2004.
- [16] C. Ye and J. Borenstein. Characterization of a 2d laser scanner for mobile robot obstacle negotiation. In *IEEE International Conference* on Robotics and Automation, 2002.