



**US Army Corps
of Engineers®**
Engineer Research and
Development Center




Mobility in Complex Environments

Obstacle Detection and Quantification for Vehicle Mobility in Winter Conditions

Sergey N. Vecherin, Jacob M. Shaker, and Michael W. Parker

June 2020



The U.S. Army Engineer Research and Development Center (ERDC) solves the nation's toughest engineering and environmental challenges. ERDC develops innovative solutions in civil and military engineering, geospatial sciences, water resources, and environmental sciences for the Army, the Department of Defense, civilian agencies, and our nation's public good. Find out more at www.erdclibrary.on.worldcat.org/discovery.

To search for other technical reports published by ERDC, visit the ERDC online library at www.erdclibrary.on.worldcat.org/discovery.

Obstacle Detection and Quantification for Vehicle Mobility in Winter Conditions

Sergey N. Vecherin, Jacob M. Shaker, and Michael W. Parker

*U.S. Army Engineer Research and Development Center (ERDC)
Cold Regions Research and Engineering Laboratory (CRREL)
72 Lyme Road
Hanover, NH 03755-1290*

Final Report

Approved for public release; distribution is unlimited.

Prepared for Headquarters, U.S. Army Corps of Engineers
Washington, DC 20314-1000

Under Mobility in Complex Environments
PE 622145 / Project BG2 / Task SBG202, "Cold Regions Effects in
Urban Environments"

Abstract

Recently, the focus for military operations has shifted from the desert to cold climates, causing a corresponding shift in the military's need to better understand the mobility of our current vehicle fleet in these areas. Therefore, this work investigated the effects of winter conditions on military vehicle mobility. The main objective was to detect obstacles on the scene.

This study developed and tested a method for automatic obstacle detection in the digital elevation model of a scene. The method detects statistical anomalies relative to an estimated background image that contains no obstacles. The sensitivity of the detection can be adjusted by a specified probability of false alarms, and the obstacle detection confidence is characterized by a probability of detection. The visible height of obstacles above the snow is related to the actual height of the obstacles above the ground.

Compared to other detection techniques, the developed method is fast, calibrates itself to the cluttered images, operates with a single given image, and aligns with a detection quantification adopted in the receiver operating characteristic framework. The examples considered in this paper demonstrate high efficiency and applicability of the developed approach to the military vehicle mobility missions.

Contents

Abstract	ii
Figures and Tables	iv
Preface	v
Acronyms and Abbreviations	vi
1 Introduction	1
1.1 Background.....	1
1.2 Objectives.....	1
1.3 Approach.....	1
2 Detection Methods	5
2.1 Edge detection.....	5
2.1.1 General.....	5
2.1.2 Different edge detection methods.....	5
2.1.3 Edge detection output.....	8
2.2 Texture detection.....	11
2.3 Taking into account a desired obstacle size.....	12
3 Quantification	14
4 Application	18
5 Effects of Snow Depth on Obstacle Detection	22
6 Conclusions	25
References	26
Report Documentation Page (SF 298)	28

Figures and Tables

Figures

1	Effects of varying the averaging distance for the Sobel edge detection with a default detection threshold value: (a) no smoothing, (b) 0.1 m, (c) 0.5 m, and (d) 1 m.....	9
2	Plots of the edge detection results with default thresholds: (a) image of the scene (b) Sobel's method, (c) Prewitt's method, (d) Roberts' method, (e) Canny's method, and (f) Laplacian of Gaussian method.....	10
3	Effect of applying different detection thresholds for edge detection using the Sobel method. The <i>middle</i> plot shows the default threshold results. <i>Black</i> indicates detected edges. The <i>left</i> image applies the smaller threshold, and the <i>right</i> image applies the greater threshold	11
4	Results of the texture-analysis methods with the mask size of 1 m without image decimation: (a) range mask and (b) STD mask	12
5	Detection of obstacles accounting for their size: (a) moving average two-dimensional trend and (b) detected obstacles in <i>black</i>	13
6	Statistics for the x-derivative in the Sobel method: (a) x-derivative and (b) nonparametric (KDE) and parametric probability distribution functions for the image in a	14
7	Quantification of obstacle detection: (a) detection threshold for the chi-squared distribution with the degrees of freedom equal one and two, (b) probability of detection map for $P_{fa} = 10^{-2}$, and (c) probability of detection map for $P_{fa} = 10^{-8}$	17
8	Test obstacles	18
9	Obstacle detection at different snow-layer depths.....	20
10	Detection of a river and obstacles: (a) detection at $P_{fa} = 0.01\%$, (b) P_d for $P_{fa} = 0.01\%$, (c) detection at $P_{fa} = 1\%$, and (d) P_d map for $P_{fa} = 1\%$	21
11	A cross section of 23 cm half rounds: (a) Obstacle #2 in Fig. 8 and (b) obstacle #4 in Fig. 8. The curves represent cross sections of the obstacles with snow cover (except in November). The <i>dashed curve</i> indicates no detection.....	22
12	The relationship between actual and above-the-surface obstacle heights at different snow depths.....	23

Preface

This study was conducted for the U.S. Army Corps of Engineers under Mobility in Complex Environments PE 622145 / Project BG2 / Task SBG202, “Cold Regions Effects in Urban Environments.” The technical monitor was Mr. Josh Fairly.

The work was performed by the Signature Physics Branch (SPB) and the Force Projection and Sustainment Branch (FPSB) of the Research and Engineering Division, U.S. Army Engineer Research and Development Center, Cold Regions Research and Engineering Laboratory (ERDC-CRREL). At the time of publication, Dr. Andrew Niccolai was Chief (SPB); Mr. J. D. Horne was Acting Chief, FPSB, and Division Chief. The Deputy Director of ERDC-CRREL was Mr. David B. Ringelberg, and the Director was Dr. Joseph L. Corriveau.

COL Teresa A. Schlosser was Commander of ERDC, and Dr. David W. Pittman was the Director.

Acronyms and Abbreviations

CDF	Cumulative Distribution Functions
CRREL	U.S. Army Cold Regions Research and Engineering Laboratory
DEM	Digital Elevation Model
ERDC	Engineer Research and Development Center
FPSB	Force Projection and Sustainment Branch
KDE	Kernel Density Estimator
LoG	Laplacian of a Gaussian
PDF	Probability Density Function
ROC	Receiver Operating Characteristic
SPB	Signature Physics Branch
STD	Standard Deviation
UAS	Unmanned Aerial Systems

1 Introduction

1.1 Background

For the past 30 years, the U.S. military has been developing and using vehicles designed for desert operations. These vehicles are heavily armored and intended for specific fighting conditions. Recently, the focus for military operations has shifted from the desert to cold climates, like those in the northern boreal regions. This caused a corresponding shift in the military's need to better understand the mobility of our current vehicle fleet in these areas. This includes understanding overall sensor reliability and performance in cold climates, obstacle detection in snow-covered areas, and route planning in northern regions. Work described in this report focused on detecting snow-covered obstacles and what impact they may have on vehicle mobility in these complex environments. Output from this work will feed other work units in the Maneuver in Complex Environments work package to determine the overall time to target in all-season conditions.

1.2 Objectives

This work deals with three specific technical challenges. The first is effective acquisition of digital elevation models (DEMs) of the site of interest. When snow covers the ground, lidar may become less effective, increasing the importance of this challenge. The second challenge is the development of fast obstacle-detecting algorithms and corresponding detection confidence characterization. Ideally, the data processing should be implemented in real or near-real time. And the third one is estimating actual obstacle height from its visible fraction above the snow.

This report addresses the first challenge by using unmanned aircraft system (UAS) and photogrammetry. It addresses the second challenge by the statistical anomaly detection algorithm. And it tackles the third challenge by revealing a connection between the actual and apparent heights through empirical studies.

1.3 Approach

Photogrammetry is the process of using conventional photos taken at different view angles to generate three-dimensional DEMs of a scene (Kasser and Egels 2002). The method has rapidly improved over the last several

years, and while research has been done on its accuracy (Uysal, Toprak, and Polat 2015), little progress has been made on implementing it into real-time applications. The biggest hurdle to overcome is the large computational load required to process large areas. However, barring the computational cost, photogrammetry can be cost effective compared to other remote-sensing tools, such as lidar, for acquiring a three-dimensional site model. Photogrammetry remote sensing can be done with a consumer-grade digital camera instead of requiring an expensive laser or sensor array. In some cases, photogrammetry can be used in scenarios where lidar becomes ineffective (Deems, Painter, and Finnegan 2013). In snow, lidar requires additional calibration based on water content and density. These properties are not always known beforehand, especially in remote areas. Photogrammetry does not require calibration, although it does struggle with highly uniform and flat snow surfaces, due to insufficient information obtained from different viewing angles.

Another area of interest is how snow affects the visibility of features present in a scene. What is the baseline obstacle detectability, and how does it change with snow accumulation? This paper seeks to investigate this aspect by creating routines for obstacle detection and investigating the relationship between snow accumulation and obstacle height.

There are a few approaches to detecting obstacles in a given image. These techniques include convolutional neural networks, inverse perspective mapping, using association and symmetry considerations, or slicing of a disparity matrix (Caltagirone et al. 2019; Prakash, Akhbari, and Karam 2019; Wei et al. 2019; Zebbara et al. 2019; Zhang et al. 2019). However, these methods are not suitable for the current research objective. Some are designed for specific applications (e.g., obstacles on the road for vehicle collision avoidance), others fuse multiple sensor modalities (which requires much more data and processing time), and yet others employ neural networks trained on specific types of obstacles. For example, two widely used methods, gray-level co-occurrence matrices (Jindal, Aggarwal, and Gupta 2018) and k -mean clustering (Kanungo et al. 2002), proved to be far too computationally or logistically complex for the intended application. In addition, none of these methods quantify a degree of assurance that the detected anomaly is actually present as opposed to being a natural part of the scene.

The receiver operating characteristic (ROC) (Fawcett 2006; Powers 2011) is a common and widely accepted way to quantify binary detection systems (i.e., with only two outcomes, either detected or not). ROC naturally motivates one to formulate the problem in a statistical context, which is less computationally intensive. For these reasons, this study uses the statistical anomaly approach. In this paradigm, the detection system is characterized by a ROC curve, which shows a dependence of the probability of detection, P_d , as a function of the probability of false alarm, P_{fa} . The probability of detection characterizes the probability to detect an obstacle if it is truly present there. The probability of false alarm characterizes the probability of false positive detection, that is, a detection of an obstacle when it is not truly there. An analyst should specify an acceptable level of P_{fa} ; and at this P_{fa} , obstacles will be detected, and their P_d will also be calculated. Notably P_d and P_{fa} simultaneously reach theoretical limits of zero and one; that is, zero P_{fa} is possible (no false alarms) only if P_d is also equal to zero (no detection at all). Likewise, 100% detection ($P_d = 1$) is possible only if P_{fa} is also equal to 1 (all detections are false alarms). Good detection systems have high P_d at low P_{fa} . So, a low value of P_{fa} means fewer obstacles will be detected, increasing risk of not detecting a true obstacle. Conversely, a high value of P_{fa} means more obstacles will be detected, but the risk of false detections increases.

This study performed various signal processing methods on DEMs created with the photogrammetry software Agisoft Metashape (previously Photoscan). The images were taken with a Sony NEX-7 camera mounted on the Turbo-Ace Matrix UAS. The scene was georeferenced with several ground control points measured with the Trimble R8, Model 3, RTK GPS system. The output from Metashape is a GeoTIFF file depicting the scene. These were imported and exported through the QGIS software to make them compatible with MATLAB. The features in the DEMs were detected with MATLAB by using several different methods described in this report. Because of the high resolution of the DEMs, the results were decimated using a moving mean average. This reduces the original image resolution to a specified resolution, which smooths the data and reduces computational load. Photogrammetry has proven itself effective in outputting very accurate DEMs for a variety of purposes, such as surveying and construction, allowing accurate three-dimensional models. However, an underexplored facet of photogrammetry is the detection of obstacles covered in snow. This study flew several UAS missions during winter 2017–2018, allowing the same scene and obstacles to be observed with varying levels of

snow depth. The first mission took place in November 2017 with no snow cover. With repeated DEM results from the same area, the effect of snow coverage on obstacle detection can be analyzed and understood. This paper covers the methods tested in MATLAB to detect the obstacles under various conditions.

2 Detection Methods

This study investigated two different approaches to detect features in the terrain elevation image: edge detection and texture detection (Canny 1987; Lim 1990; Parker 2010). Edge detection is best when the features have abrupt height differences to the background (e.g., a wood table above a hard wood floor). Texture detection works optimally when the feature has a different surface texture when compared to the background (e.g., a smooth rock in a field of grass). In the case of the test scene, neither method initially appeared more suitable because snow accumulation normalizes both height and surface texture.

2.1 Edge detection

2.1.1 General

The simplest and most common approach to feature detection is edge detection. This approach finds acute changes in elevation between two adjacent spatial locations, indicating the edge of a feature. One can find these locations by comparing the magnitude of spatial two-dimensional x - and y -derivatives of the scene (the gradient) with a specified gradient magnitude detection threshold. This task used the built-in MATLAB function *edge*, found in the Image Processing Toolbox. The *edge* function has multiple settings; the four most prominent edge detection methods tested were Sobel, Prewitt, Roberts, Canny, and the Laplacian of a Gaussian (LOG).

2.1.2 Different edge detection methods

2.1.2.1 Sobel, Prewitt, and Roberts

Spatial derivatives on a numerical computational grid can be calculated using the Sobel, Prewitt, or Roberts derivative approximations, which are the three most common edge detection methods. Unlike simple finite difference derivative approximation around a given location, these techniques use a 3 by 3 pixel area around each location in the image to estimate the spatial gradient. Consequently, these techniques provide smoothed gradient estimates. The difference between the methods lies in the spatial mask used to better approximate the spatial derivatives, which results in different weights prescribed to these 3 by 3 area pixels.

Given a neighborhood of values in the image,

$$\mathbf{I} = \begin{bmatrix} z_1 & z_2 & z_3 \\ z_4 & z_5 & z_6 \\ z_7 & z_8 & z_9 \end{bmatrix}, \quad (1)$$

the Sobel method applies the following masks:

$$\mathbf{M}_{Sx} = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix}, \quad \mathbf{M}_{Sy} = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix}, \quad (2)$$

for the x - (from left to right) and y - (from up to down) derivatives, respectively. To get a smoothed Sobel approximation of the two-dimensional gradient of the image at the center point (z_5), these matrices are convolved with \mathbf{I} :

$$\mathbf{g}_x = \mathbf{M}_{Sx} * \mathbf{I}, \quad \mathbf{g}_y = \mathbf{M}_{Sy} * \mathbf{I}, \quad (3)$$

where the asterisk represents the two-dimensional convolution. The magnitude of the two-dimensional horizontal gradient is given by

$$\mathbf{g} = \sqrt{(\mathbf{g}_x^2 + \mathbf{g}_y^2)}, \quad (4)$$

where the squaring of a matrix is understood on the element-by-element basis. The direction of the gradient relative to the x -axis, θ , can be found using

$$\theta = \text{atan} \frac{\mathbf{g}_y}{\mathbf{g}_x}. \quad (5)$$

$\theta = 0$ means a zero vertical gradient, indicating a vertical edge. For the Prewitt method, the x - and y -derivative masks, respectfully, are

$$\mathbf{M}_{Px} = \begin{bmatrix} -1 & 0 & 1 \\ -1 & 0 & 1 \\ -1 & 0 & 1 \end{bmatrix}, \quad \mathbf{M}_{Py} = \begin{bmatrix} -1 & -1 & -1 \\ 0 & 0 & 0 \\ 1 & 1 & 1 \end{bmatrix}. \quad (6)$$

The only difference between the Sobel and Prewitt methods is the weighting of the nearest point to the center of the matrix.

The Roberts approximation is different in that it takes simplified 45° and 135° derivatives:

$$\mathbf{M}_{Rx} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad \mathbf{M}_{Ry} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & -1 \\ 0 & 1 & 0 \end{bmatrix}. \quad (7)$$

The edge detection based on the Roberts approximation is generally less used due to its simpler nature and limited robustness overall. The masks given in equation (7) are not symmetrical and do not effectively detect 45° angles. The Roberts method has a few benefits in specific applications, aside from its speed and generally serviceable accuracy (Parker 2010).

2.1.2.2 Laplacian of a Gaussian (LoG)

The LoG method uses the Laplacian of the two-dimensional Gaussian function,

$$G(x, y) = e^{-\frac{x^2+y^2}{2\sigma^2}}. \quad (8)$$

Its Laplacian is

$$\nabla^2 G(x, y) = \frac{\partial^2 G(x, y)}{\partial x^2} + \frac{\partial^2 G(x, y)}{\partial y^2} = \left[\frac{x^2+y^2-2\sigma^2}{\sigma^4} \right] e^{-\frac{x^2+y^2}{2\sigma^2}}. \quad (9)$$

This method convolves the original image with the expression in equation (9). This operation smooths the image and results in double-edge features (positive and negative values) where there are large gradients, indicating the presence of obstacles. The actual edges are then located by finding the zero crossings between the double edges. In this case, the detection threshold applies to the deviation of LoG-filtered image values from zero. By default the detection threshold equals the standard deviation, σ .

2.1.2.3 Canny's method

The Canny method is the most complex edge detection tool built into MATLAB's *edge* function. It integrates both the Sobel and the Gaussian methods and adds a filtering algorithm to remove weak edge points adjacent to strong edge points. The Canny method can be summarized into the following four main steps:

1. Smooth the image using a Gaussian filter with a specified standard deviation σ to reduce noise.

2. Find the local gradient magnitude, equation (4), and edge direction, equation (5). Then use Sobel's method to determine the edges and the local maximum derivative with a default detection threshold.
3. Ignore edge pixels that are not local maximums. The removed locations are "weak" edge points, and the remaining edge locations are "strong." This process is called *nonmaximal suppression*.
4. Link the edges by reintegrating the weak edge pixels that surround strong edge pixels.

The Canny method is the most powerful of these edge detection tools as it accounts for the gradient leading up to and continuing after an edge, drawing clearer edges than Sobel's or Prewitt's methods alone, while also integrating smoothing similar to LoG. The result is a clearer picture of the most intense edges, connected disparate edges, and indication of weak edges, altogether giving the image a better defined shape.

2.1.3 Edge detection output

To reduce excessive computational load and to smooth outliers in the image, DEMs were decimated to a lower resolution and smoothed using a two-dimensional, nonoverlapping moving average with a specified spatial averaging distance. For example, decimating with a 1 m averaging distance requires replacing each pixel value in a 1 m² area with the average value of the pixels contained within that area. Repeating this process over the whole area effectively decimates the image to a 1 m resolution. Figure 1 shows the effect of varying averaging distances on edge detection using the Sobel method with MATLAB's default threshold value.

Figure 1 shows that a decimation of 0.5 m is optimal due to its low number of false alarms while still effectively detecting each obstacle. The resulting size of the decimated image is also much more manageable compared to the raw DEM. The rest of the edge detection results in this paper will use a decimation of 0.5 m.

Figure 1. Effects of varying the averaging distance for the Sobel edge detection with a default detection threshold value: (a) no smoothing, (b) 0.1 m, (c) 0.5 m, and (d) 1 m.

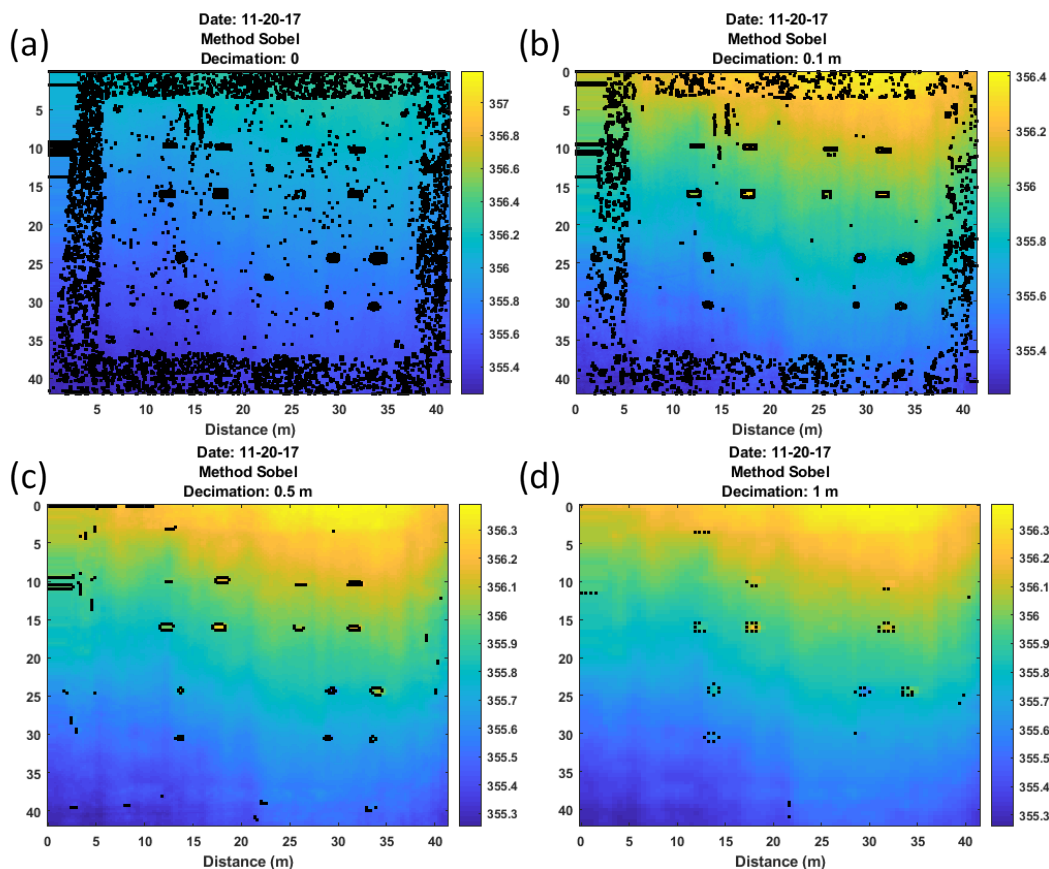


Figure 2 shows the results of each edge detection method using default thresholds, which are based on a rough estimate of the signal-to-noise ratio (Pratt 2013). Explicit threshold specification based on the accepted probability of false alarm is described in section 3.

Figure 2 shows that the Sobel, Prewitt, and Roberts methods performed quite well, identifying the obstacles with some false alarms in areas close to the boundaries. The default secondary threshold for the Canny method was too low; connecting distant noise points resulted in a cluttered image. The LoG method performed similarly to the Canny method, with many false alarms. The Sobel and Prewitt methods performed best, so this study proceeded with the Sobel method to find the baseline for other images to compare to texture detection.

Figure 2. Plots of the edge detection results with default thresholds: (a) image of the scene (b) Sobel's method, (c) Prewitt's method, (d) Roberts' method, (e) Canny's method, and (f) Laplacian of Gaussian method.

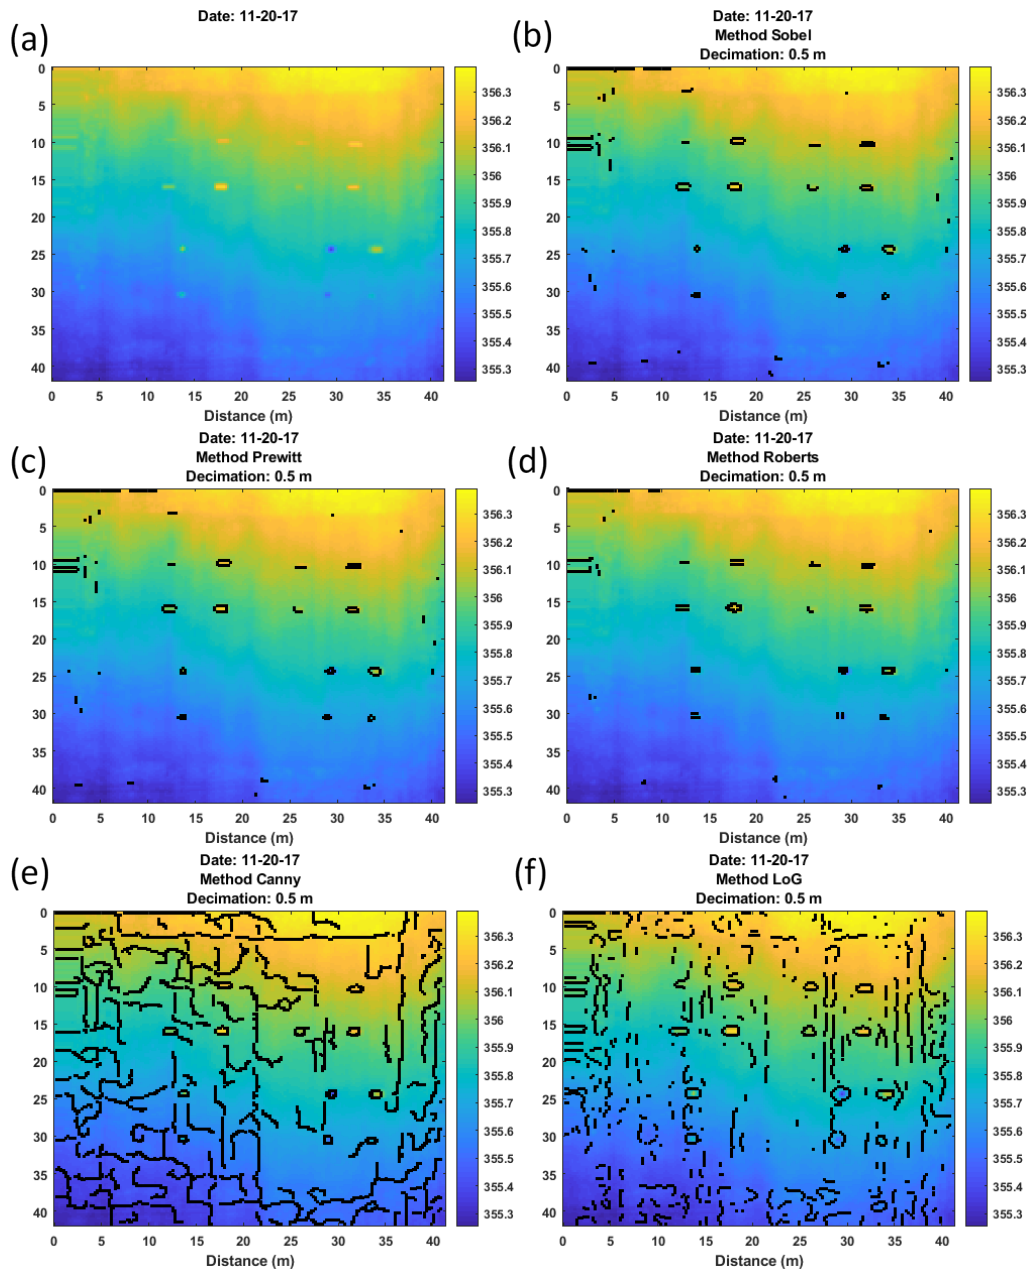
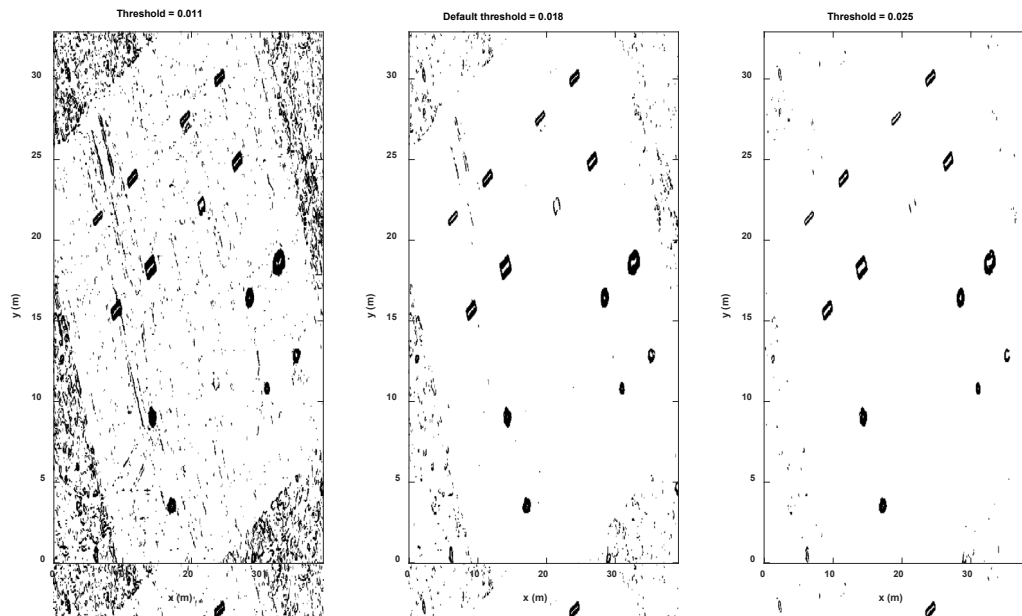


Figure 3 demonstrates the effect of using different detection thresholds for the same scene sensed on 20 November 2017.

One can see in Figure 3 that the edge detection results depend significantly on the applied detection threshold. Varying the threshold from 0.011 to 0.025 (the MATLAB default is 0.018), one can detect more or fewer edges in the same image regardless of whether they are actual obstacles or false

alarms. Similar behavior is observed for other edge detection and texture-analysis methods. This clearly indicates a necessity to develop a quantification method that allows one to choose the detection thresholds systematically, as described in section 3.

Figure 3. Effect of applying different detection thresholds for edge detection using the Sobel method. The *middle* plot shows the default threshold results. *Black* indicates detected edges. The *left* image applies the smaller threshold, and the *right* image applies the greater threshold.



2.2 Texture detection

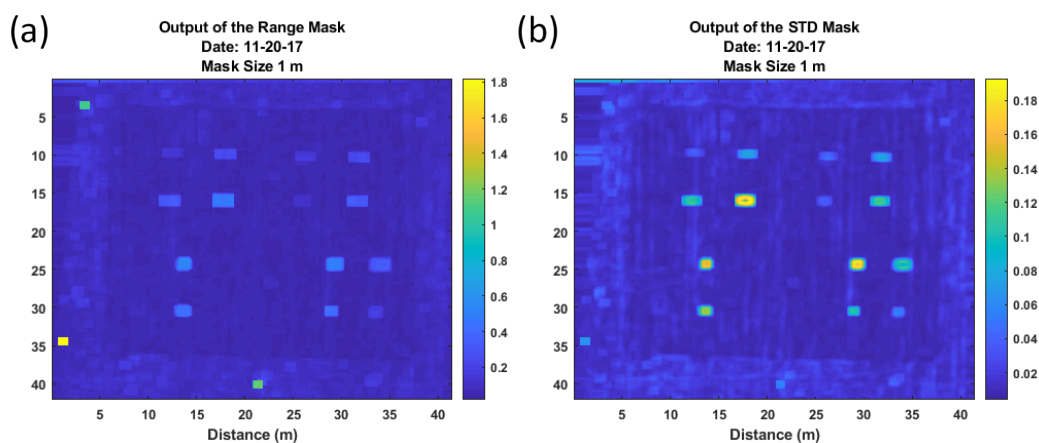
Similar to the smoothed gradient estimates of edge detection, texture detection uses spatial masks around every pixel to transform the original image to a new image. These masks highlight visible traits of the image, such as roughness or distinct patterns. Two common methods of texture detection are the range filter and the standard deviation (STD) filter. Each mask considers a $n \times m$ neighborhood around a center pixel and applies a certain operation to the neighborhood to find the new value for the center pixel. Because variation in the scene is relevant to texture detection, the decimation distance has been reduced to 0.1 m for the range and STD masks.

The range mask calculates the maximal absolute difference between the central pixel and the rest of the neighborhood, the size of which the user specifies. This method is useful when there is a pattern in the scene representing a change in texture instead of a sudden jump in intensity.

For the STD mask, one calculates the standard deviation of a $n \times m$ neighborhood of points taken around a center pixel. This mask is useful when there are statistically consistent variations in the texture; and though the variation is not enough to be qualified as an edge, the consistency and STD magnitude are enough to mark it apart from the background.

On most images analyzed in this report, the outputs of the texture masks show the obstacles much more clearly than the edge detection methods' outputs. The results of the two masks are very similar; however, the range filter is less robust. As an example, Figure 4 shows the results of texture analysis with the range and STD filters on an image with no decimation to avoid any decimation effects.

Figure 4. Results of the texture-analysis methods with the mask size of 1 m without image decimation: (a) range mask and (b) STD mask.



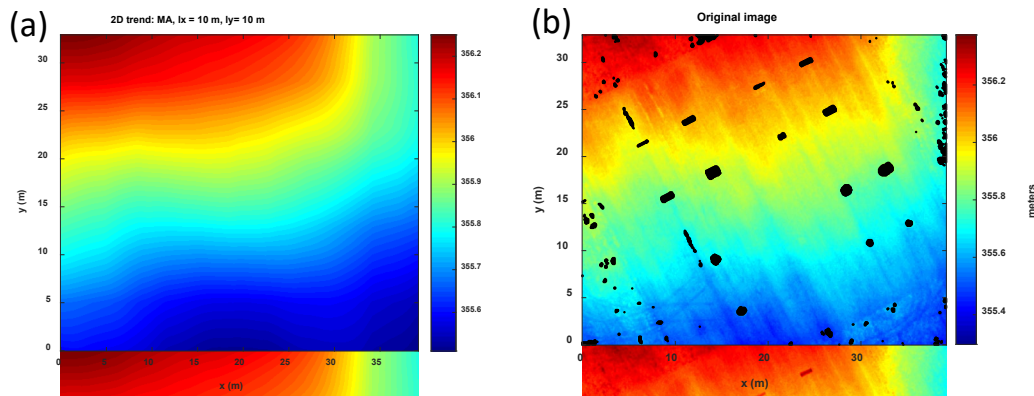
The range mask is also systemically much more sensitive to noise than the STD mask. This conclusion holds for images with decimation. For this reason, the Standard Deviation mask is considered better than the range mask, at least for the tested data sets.

2.3 Taking into account a desired obstacle size

This research effort investigated several preprocessing techniques to reduce false alarm detections. These include spatial filtration (Gaussian filter and spatial spectral filtration), image detrending (mean removal, zero-perimeter background removal, least-squared plane fit to remove possible slopes in the scene, and nonlinear background removal using a moving average), and spatial filtration to the specified obstacle size. This subsection presents results for the filtration to the specified obstacle size as being the most consistent for distinct scenes.

The spatial filtration technique considered here averages the terrain elevation values over the spatial area of the characteristic obstacle size. Indeed, for most vehicles, terrain roughness of the 5 cm horizontal scale does not pose a threat to mobility. For some other vehicles, even larger objects may not be a problem. Therefore, if one averages values over the area of this size, all small terrain elevations will be averaged out, reducing false alarms. Consequently, larger obstacles will be seen more clearly. Given the desired minimal scale of obstacle to be detected, the algorithm uses a moving average over the image and smooths out smaller height variations. Figure 5 presents the results obtained from the smoothing algorithm. Figure 5a depicts a nonlinear two-dimensional background calculated using the moving average algorithm with spatial scales of 10 m in both x - and y -directions. Figure 5b shows the detected obstacles. Note that the obstacle detection was also obtained using the moving average filter applied to the detrended scene (the original minus the trend) but using much smaller spatial scales of 50 cm. With the correctly chosen detection threshold, all obstacles were detected with very few false alarms (Figure 5). The next section describes the method for threshold determination.

Figure 5. Detection of obstacles accounting for their size: (a) moving average two-dimensional trend and (b) detected obstacles in *black*.



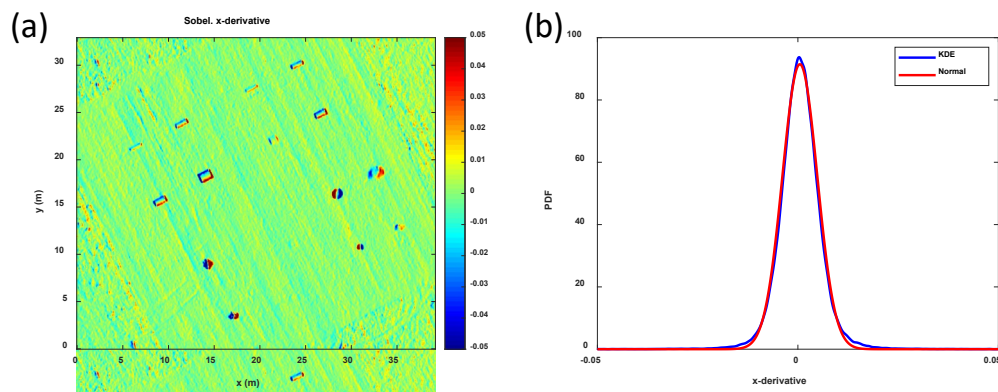
3 Quantification

The results from the previous sections clearly indicate that detection depends significantly on the applied detection threshold. As argued in the Introduction, a natural way to determine the threshold value can be found within the ROC framework. This section considers how a specified probability of false alarm can be used to determine the detection threshold and investigates how detections with this threshold can be quantified by estimating their probability of detection. This section uses the Sobel edge detection method as an example.

The first step in specifying the threshold value for the Sobel method is to estimate the probability density function (PDF) for x -derivative values for a background image (no obstacles). The problem is that there is no image available without obstacles; the one that is analyzed, most likely, contains obstacles. This problem can be resolved by using robust estimates for PDF parameters, which ignore obstacles. The procedure is described next.

Figure 6a shows the x -derivative image obtained with the Sobel method. One can see that, apart from the edges of obstacles, the derivative values are fairly homogeneous, with a slightly positive mean due to an overall scene slope. Considering these values as a set of random numbers drawn from the same probability distribution, their PDF can be estimated using the kernel density estimator (KDE), which provides a nonparametric PDF estimate, free of any assumption about the shape of the PDF (Wand and Jones 1994). The blue curve in Figure 6b shows the estimated PDF.

Figure 6. Statistics for the x -derivative in the Sobel method: (a) x -derivative and (b) nonparametric (KDE) and parametric probability distribution functions for the image in a.



The estimated PDF provides a clue to which family of analytical distributions can be used for modeling the data. Obviously, the curve is well described by the normal distribution, for which robust estimates (ignoring the outliers, which are associated with obstacles in this framework) of the mean, μ_x , and standard deviation, σ_x , can be made using sample quantiles, which are known to be robust with respect to outliers in the data (Johnson and Wichern 2002). Specifically, they are given by

$$\mu_x = q_{50}, \quad \sigma_x = (q_{75} - q_{25})/1.349, \quad (10)$$

where q_{25} , q_{50} , and q_{75} stand for 0.25, 0.5, and 0.75 quintiles, respectively. Figure 6b shows the analytical normal PDF (red) with these parameters. One can see that it matches the nonparametric KDE estimate with a high degree of accuracy. Similar estimates can be made for the y -derivative to obtain μ_y and σ_y .

Second, given the mean and standard deviations of the derivatives, their values can be standardized to obtain the following normalized sets:

$$\tilde{\mathbf{g}}_x = \frac{\mathbf{g}_x - \mu_x}{\sigma_x}, \quad \tilde{\mathbf{g}}_y = \frac{\mathbf{g}_y - \mu_y}{\sigma_y}, \quad (11)$$

where the subtraction of the means and division by the standard deviations are performed for each element of matrices \mathbf{g}_x and \mathbf{g}_y . This normalization produces $\tilde{\mathbf{g}}_x$ and $\tilde{\mathbf{g}}_y$ values distributed with zero mean and unit standard deviation.

The magnitude of the standardized spatial gradient can be characterized by

$$\tilde{\mathbf{g}}^2 = \tilde{\mathbf{g}}_x^2 + \tilde{\mathbf{g}}_y^2, \quad (12)$$

where, again, the square operation is performed for every element of the matrices.

Third, assuming that $\tilde{\mathbf{g}}_x$ and $\tilde{\mathbf{g}}_y$ are statistically independent (indeed, there are no considerations causing the spatial derivative in the x -direction to be dependent on the derivative in the y -direction) and are normally distributed (see Figure 6b), the values in $\tilde{\mathbf{g}}^2$ are distributed in accordance with the chi-squared probability distribution with two degrees of freedom, χ_2^2 .

The detection threshold, T , is found as a value of $\tilde{\mathbf{g}}^2$ such that the probability to exceed this threshold in the background image (due to random natural variations in the scene) equals the given P_{fa} . Using the properties of the cumulative distribution functions (CDF), this threshold can be found as the inverse CDF of the χ_2^2 probability distribution, taking $1 - P_{fa}$ as the argument:

$$T = \text{invCDF}_{\chi_2^2}(1 - P_{fa}). \quad (13)$$

All values of $\tilde{\mathbf{g}}^2$ above this threshold are considered to be obstacles with the probability of P_{fa} to occur naturally (to be false alarms). Note that if P_{fa} is chosen too small, the detection threshold becomes too large, which results in fewer obstacles detected. For extremely homogeneous images (like, undisturbed snow), no obstacles will be detected, which is a desirable outcome. Similar considerations apply to other detection techniques considered in this report, such as the range or STD filter and moving average intensity. In these techniques, the chi-squared distribution with a single degree of freedom should be used, χ_1^2 , because there is no second term contributing to the magnitude of the random quantity under consideration (e.g., STD), as it is the case with the two-dimensional spatial gradient, equation (12). Figure 7a shows the detection thresholds for the chi-squared distributions with one and two degrees of freedom as functions of the probability of false alarm.

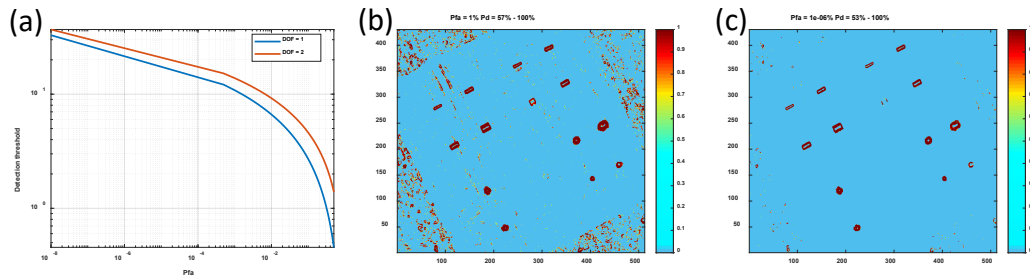
Once some pixels in the image are identified as obstacles, it is worthwhile to quantify their probability of detection, that is, the probability that the actual obstacle is detected. In this study, such an estimate assumes that the variance of each pixel in the image remains equal to the variance found from the background image. The mean of the background distribution is set to the actual values of the image at each pixel. This can be visualized as shifting the background image distribution, without changing its shape, to center it to the actual pixel value. This leads to the noncentral chi-squared probability distribution for calculating P_d .

Such a procedure has a clear practical interpretation. Namely, for values just equal to the detection threshold, the probability of detection is 0.5. That is, the detected obstacle can be either a natural feature or an actual obstacle with equal probability. This makes sense as the value lies just on the boundary of detection. The developed quantification algorithm will not detect anomalies whose probability of detection is less than 0.5. In other

words, those anomalies, which have less than 0.5 probability of being true obstacles, are not considered worthwhile for detection.

Figure 7b shows a probability of detection map, using the Sobel method, for obstacles without a snow cover for $P_{fa} = 10^{-2}$. The P_d is high for the detected obstacles since they are clearly seen relative to the background. But there are also many false alarms, notably in the corners, where P_d is significantly closer to 0.5. Figure 7c demonstrates that reducing the desired probability of false alarm to $P_{fa} = 10^{-8}$ almost completely eliminates false alarms for this image.

Figure 7. Quantification of obstacle detection: (a) detection threshold for the chi-squared distribution with the degrees of freedom equal one and two, (b) probability of detection map for $P_{fa} = 10^{-2}$, and (c) probability of detection map for $P_{fa} = 10^{-8}$.



Since the detection threshold applies to each pixel in the image individually, the image resolution may provide some insights to what P_{fa} the user should specify. Suppose, the image is 512×384 pixels. The total number of pixels is then 196,608. Therefore, $P_{fa} = 10^{-2}$ may approximately yield 1,996 false-alarm pixels. And to be extra conservative regarding false alarms in order to have just one pixel in the image, set $P_{fa} = 1/196608 \approx 5 \times 10^{-6}$. Of course, these numbers will not be satisfied exactly because one particular image may not follow exactly the estimated probabilities, but such a rule of thumb gives useful order-of-magnitude estimates.

4 Application

This section describes the results of applying the developed techniques to different scenes. A scene without a snow layer has the best detection and will serve as a reference point. Among a variety of methods, the texture detection with the STD filter proves to be a robust and fast technique for obstacle detection. Texture detection does not require image preprocessing, such as background slope removal or filtration. This study investigated preprocessing techniques for each detection method, but they were found to be unnecessary for the STD texture detection. Therefore, the results presented here are obtained with the STD filter applied directly to the DEM image. The main objective of this section is to investigate how well the method works for snow-covered scenes and how the height of detected obstacles can be characterized for a given snow depth, which is important for vehicle mobility predictions.

The test obstacles were natural and man-made objects. They consisted of eight metal half-round culverts with a varying radius from 10 to 30 cm, two holes in the ground 17 cm and 38 cm deep, two rock piles 15 cm and 14 cm high, and two dirt mounds 8 cm and 22 cm high. Figure 8 shows the obstacles and describes each.

Figure 8. Test obstacles.

Obstacle #	Height (cm)	Description
1	15	Half Round
2	23	Half Round
3	15	Half Round
4	23	Half Round
5	30	Half Round
6	41	Half Round
7	10	Half Round
8	30	Half Round
9	14	Rock Pile
10	-38	Hole
11	22	Dirt Mound
12	15	Rock Pile
13	-17	Hole
14	8	Dirt Mound

Figure 9 shows the obstacle detection at different snow depths on 27 December 2017 (32 cm of snow), 10 January 2018 (40 cm of snow), and 12 February 2018 (27 cm of snow). The first row of plots in Figure 9 depicts

the baseline reference case with no snow, taken on 20 November 2017. The pictures in the left-most column depict either visual spectrum or near-infrared spectrum images of the scene. Bühler et al. (2016, 2017) found the infrared band with wavelengths 830 nm and above to be the best visual spectrum to detect subtle variations in the snow, which appears flat in the conventional visual spectrum. The pictures in the middle depict detected obstacles at the P_{fa} specified in the picture title (from 0.01% to 1%), and the pictures in the right-most column show probability of detection maps.

As one can see, the STD texture filter reliably detects obstacles in all cases; there were only a few false-alarm detections. The increase of false-alarm detections on 10 January 2018 (the middle row in Figure 9) is due to the presence of snowdrifts. However, from the algorithms perspective, the height variations of the snowdrifts are the same as those of the true obstacles, meaning that the snow drifts might likely be considered true obstacles. A clue to the confidence of detection can be drawn from the probability of detection map (the right-most column) where snow drifts have generally smaller P_d than the actual obstacles.

A question of practical interest is whether the detection method can capture a water body, such as a river. From an operational perspective, rivers are major natural obstacles for vehicle mobility. Figure 10 depicts river and obstacle detection seen within the same scene, analyzed on 31 January 2019 at two levels of P_{fa} . As one can see, the river detection is very reliable. Nearby obstacles, however, might be missed if their height variation is too small. This is a consequence of automatic self-calibration to the average variability captured in the image. Again, higher false-alarm probability allows for the detection of more subtle height variations, as seen in Figure 10c and d.

The next section describes the connection between observable obstacle height above the snow and the true buried obstacle height at a given snow depth.

Figure 9. Obstacle detection at different snow-layer depths.

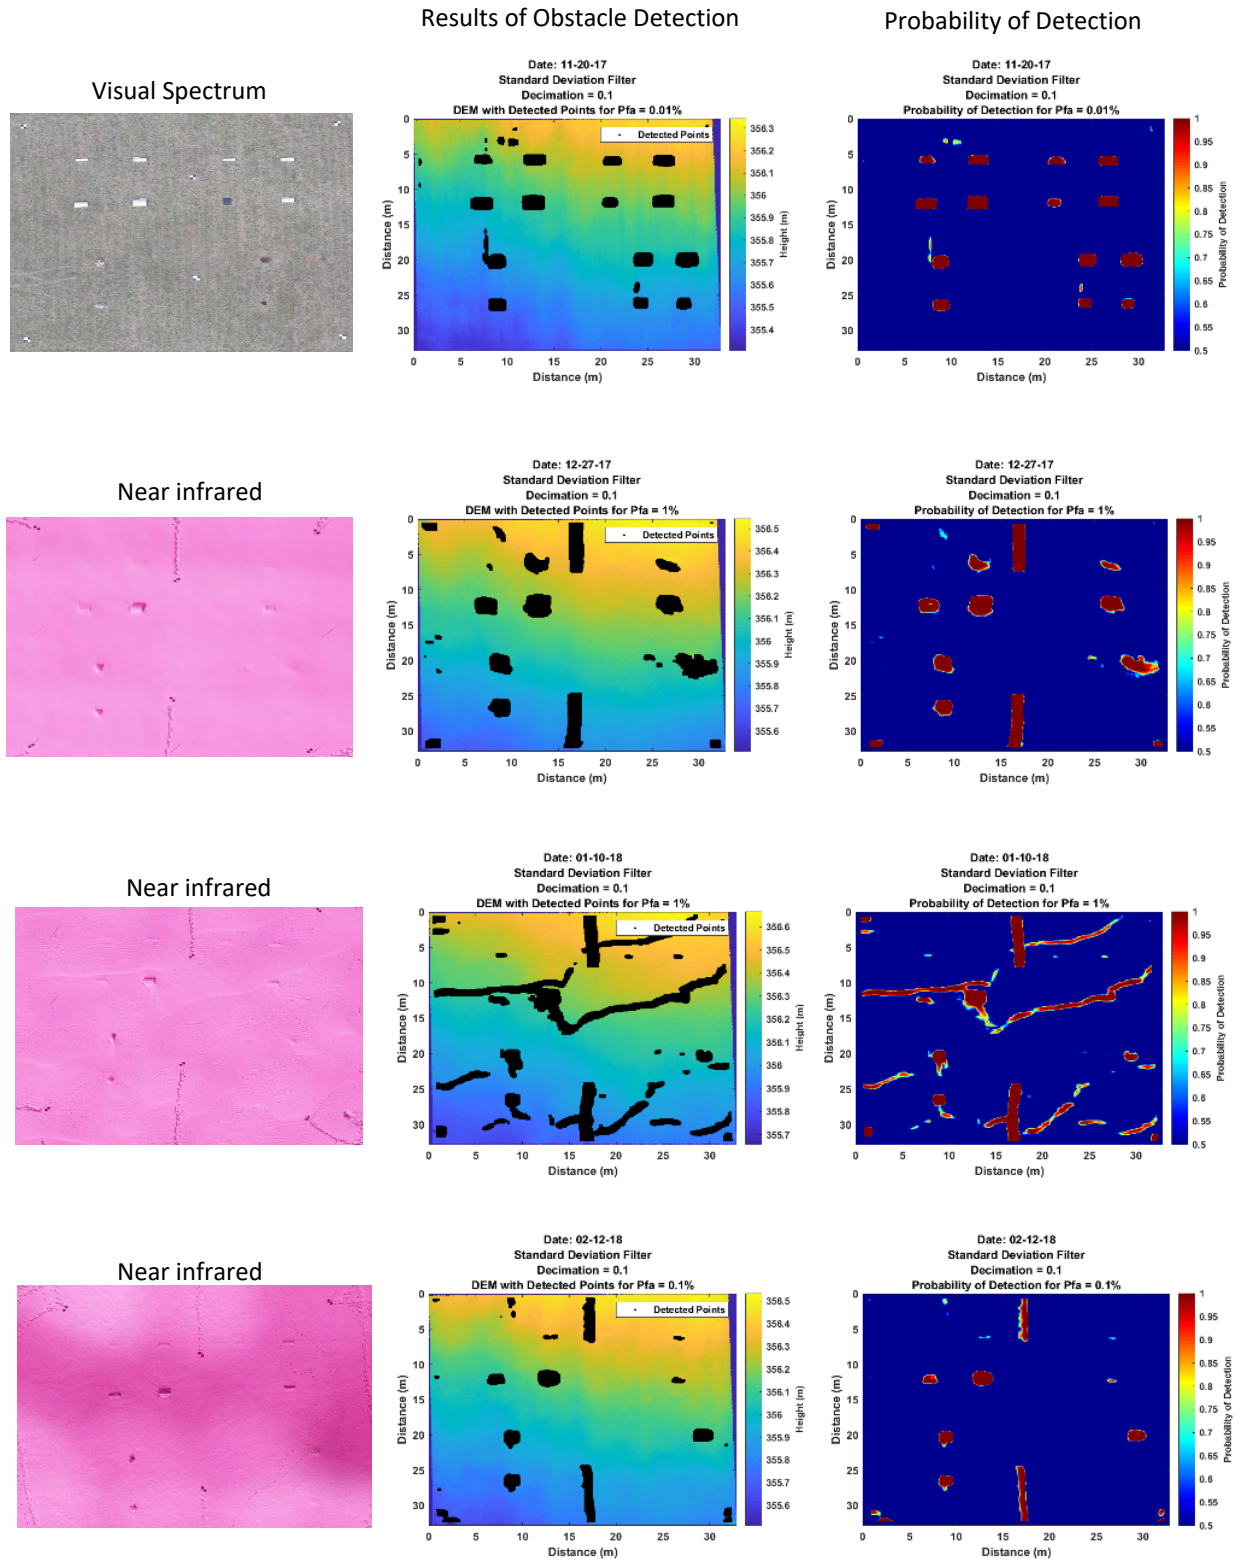
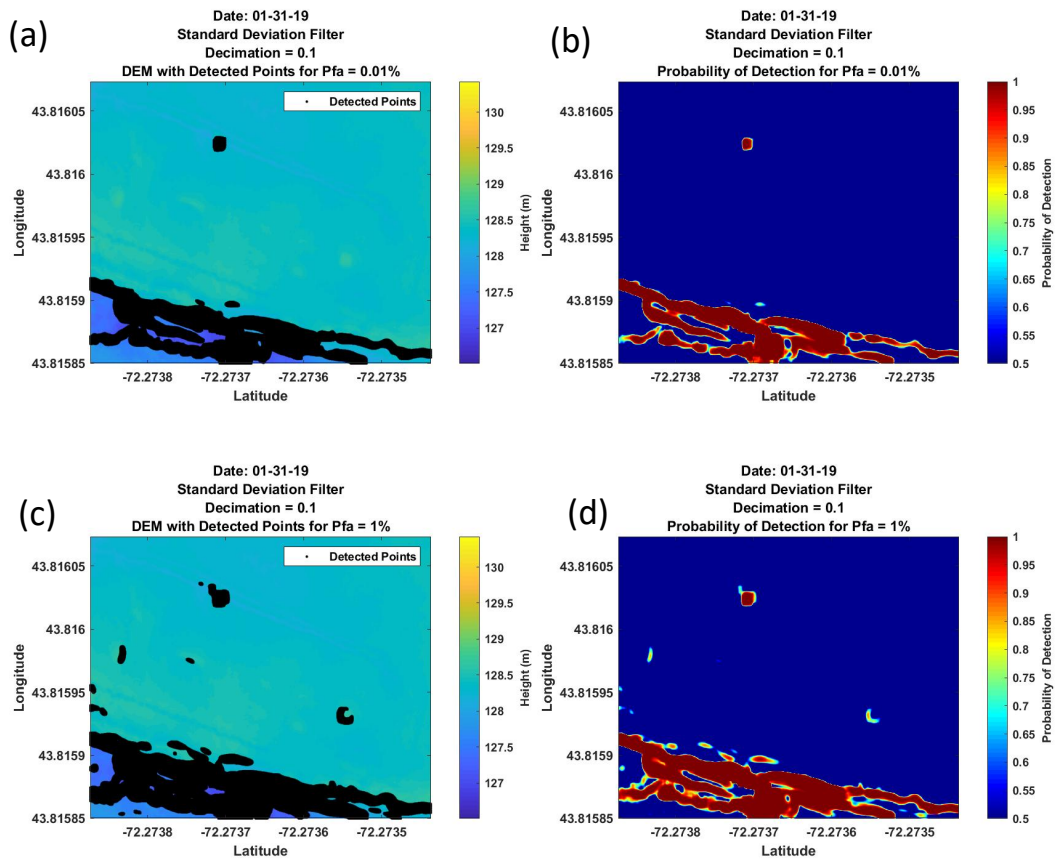


Figure 10. Detection of a river and obstacles: (a) detection at $P_{fa} = 0.01\%$, (b) P_d for $P_{fa} = 0.01\%$, (c) detection at $P_{fa} = 1\%$, and (d) P_d map for $P_{fa} = 1\%$

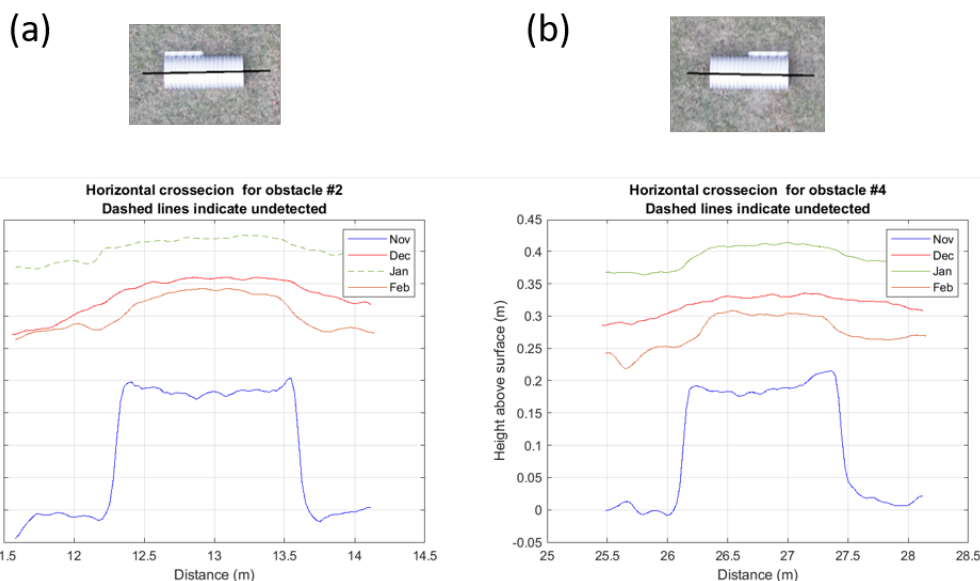


5 Effects of Snow Depth on Obstacle Detection

The experiments showed that obstacle detection depends on snow depth. As snow accumulates, it drifts and settles between obstacles, flattening the scene, eventually leading to the complete obfuscation of features. Feature geometry also affects detection and likely interacts with other snow properties not accounted for, such as moisture content, density, and accumulation rate. The final consideration is how repeated thaw and snowfall cycles may further homogenize the scene. Taking measurements of these effects was not part of this study. Given the complex and indeterminate nature of the snow properties and disregard for feature geometry effects, this study made inferences by comparing feature height and feature detectability at distinct snow depths.

Figure 11 shows the detectability of 23 cm radius culverts at varying snow depths.

Figure 11. A cross section of 23 cm half rounds: (a) Obstacle #2 in Fig. 8 and (b) obstacle #4 in Fig. 8. The curves represent cross sections of the obstacles with snow cover (except in November). The *dashed curve* indicates no detection.

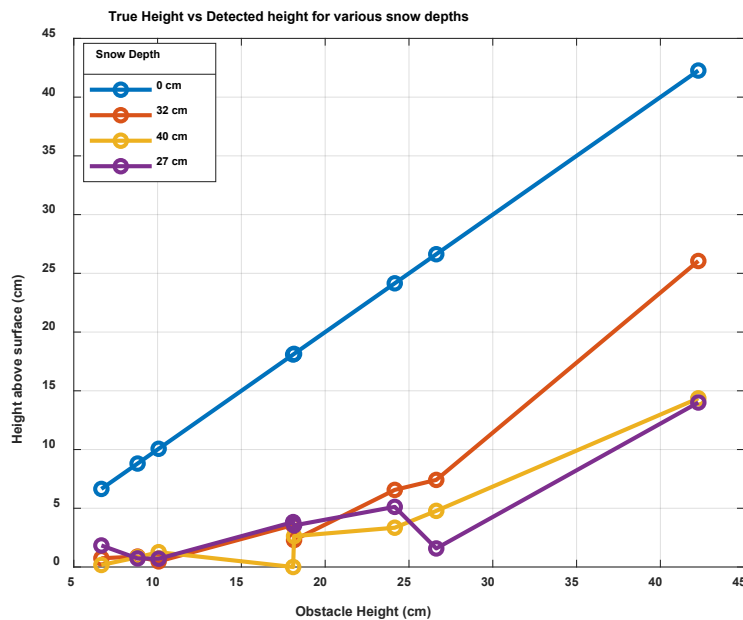


Obstacle #2, Figure 11a, was detected at all times except for the month of January, which had the highest snow depths of 40 cm. Obstacle #4, Figure 11b, was detected at all times. Figure 11 shows that obstacles of certain heights become undetectable at some snow depths. Therefore, it would be

advantageous to estimate what size obstacles will be completely covered by a certain snow depth and when obstacles are not completely covered but are detected, knowing how their apparent height (i.e., above snow) relates to their true height (i.e., above ground).

Figure 12 shows the analysis of the tested culverts. The plot relates the apparent height to the true height for varying snow-layer depths. The baseline case (no snow) will have a line drawn at 45° , meaning that the apparent height equals the true height. First, note that up to some height, obstacles are not detected at all (the detected height is zero). When they are detected, their apparent height can differ profoundly from their actual height, depending on the snow depth. More data and experiments are needed to statistically justify the results in Figure 12, but this figure serves as a guideline for the inverse problem. Namely, it addresses two practical questions: given the snow depth, (1) what is the minimal obstacle height for detection, and (2) knowing the obstacle's height above the snow, what is the actual height above the ground?

Figure 12. The relationship between actual and above-the-surface obstacle heights at different snow depths.



The experimental results demonstrate several general trends:

- Culverts are generally obscured once covered by 1.8 times their height in snow.
- The two dirt mounds were harder to detect than culverts.

- The rock piles were always detected with a maximum ratio of snow depth-to-pile height of 2.77. One would infer that they would become undetectable with several meters of snow, but the exact upper limit is still unknown.
- The two holes were never detected, even with snow-to-depth ratios less than 1. It takes less snow to obscure a hole than it does to obscure a culvert. Therefore, in general, the geometry of obstacles is an important factor.

Future efforts should focus on incorporating the results into vehicle mobility models and refining the developed tools. For example, for unknown scenes, both edge detection and texture-detection algorithms can be used because it is not known whether an obstacle will have sharp edges. Second, future studies should develop the model relating the apparent obstacle height to the actual height, either using theoretical considerations, more experiments, or both. Third, the obstacle detection can be implemented not from the point of view of statistical anomalies but from the point of view of vehicle capability. That is, if a vehicle cannot pass over an obstacle of a certain height, then all elevations above this height should be marked as obstacles.

6 Conclusions

This report characterizes vehicle mobility in the presence of snow-covered obstacles. The study developed a method for automatic obstacle detection and estimation of detection confidence that detects statistical anomalies relative to the estimated background image (i.e., the image without obstacles). The obstacle detection is based either on the anomalously large spatial gradients of the surface elevation map (edge detection) or standard deviation in the vicinity of every pixel in the image (STD filter). The probability distributions of these quantities for the background image are assumed to be chi-square distributed. The distribution parameters are estimated from the given image using robust statistics (sample quantiles, which are robust relative to the outliers). Such an approach allows one to “recreate” statistical properties of a background image using just a single given image with anomalies.

The sensitivity of detection depends on the chosen probability of false alarm, and the confidence of detection is characterized by the probability of detection. Such a quantification tightly ties to the receiver operation characteristic, which is a gold standard in the theory of detection. The advantages of the statistical anomaly approach compared to other popular methods (such as neural networks and other machine learning techniques) is its numerical efficiency, independence from the type of obstacles, automatic recalibration to capture average height variability in the image, and statistical characterization of the sensitivity and confidence of the detection. However, a drawback of the self-calibration might be missing low-height obstacles in a highly cluttered image. This is especially true when the clutter occupies most of the image. Other algorithms (e.g., fixed-height detection thresholds) should be used if detection of low-height obstacles is still important.

The cases considered in this paper highlight the efficiency and applicability of the developed method. From here, autonomous vehicle mobility programs can incorporate this automatic obstacle detection, advancing one of the most recent Army priorities. This will allow for efficient autonomous navigation and rerouting and will contribute to the Army’s next-generation military vehicles.

References

- Bühler, Y., M. S. Adams, R. Bösch, and A. Stoffel. 2016. "Mapping Snow Depth in Alpine Terrain with Unmanned Aerial Systems (UASs): Potential and Limitations." *The Cryosphere* 10 (3):1075–1088.
- Bühler, Y., M. S. Adams, A. Stoffel, and R. Boesch. 2017. "Photogrammetric Reconstruction of Homogenous Snow Surfaces in Alpine Terrain Applying Near-Infrared UAS Imagery." *International Journal of Remote Sensing* 38 (8–10): 3135–3158.
- Caltagirone, L., M. Bellone, L. Svensson, and M. Wahde. 2019. "LIDAR–Camera Fusion for Road Detection Using Fully Convolutional Neural Networks." *Robotics and Autonomous Systems* 111:125–131.
- Canny, J. 1987. "A Computational Approach to Edge Detection." In *Readings in Computer Vision*, 184–203. Elsevier.
- Deems, J. S., T. H. Painter, and D. C. Finnegan. 2013. "Lidar Measurement of Snow Depth: A Review." *Journal of Glaciology* 59 (215): 467–479.
- Fawcett, T. 2006. "An Introduction to ROC analysis." *Pattern Recognition Letters* 27:861–874.
- Jindal, A., N. Aggarwal, and S. Gupta. 2018. "An Obstacle Detection Method for Visually Impaired Persons by Ground Plane Removal Using Speeded-Up Robust Features and Gray Level Co-occurrence Matrix." *Pattern Recognition and Image Analysis* 28 (2): 288–300.
- Johnson, R. A., and D. W. Wichern. 2002. *Applied Multivariate Statistical Analysis*. Vol. 5. Upper Saddle River, NJ: Prentice Hall.
- Kanungo, T., D. M. Mount, N. S. Netanyahu, C. D. Piatko, R. Silverman, and A. Y. Wu. 2002. "An Efficient k-Means Clustering Algorithm: Analysis and Implementation." *IEEE Transactions on Pattern Analysis and Machine Intelligence* 24 (7): 881–892.
- Kasser, M., and Y. Egels. 2002. *Digital Photogrammetry*. London: Taylor & Francis Inc.
- Lim, J. S. 1990. *Two-Dimensional Signal and Image Processing*. Englewood Cliffs, NJ: Prentice Hall.
- Parker, J. R. 2010. *Algorithms for Image Processing and Computer Vision*. Indianapolis, IN: Wiley Publishing, Inc.
- Powers, D. M. W. 2011. "Evaluation: From Precision, Recall and F-Factor to ROC, Informedness, Markedness & Correlation." *Journal of Machine Learning Technologies* 2 (1): 37–63.

- Prakash, C. D., F. Akhbari, and L. J. Karam. 2019. "Robust Obstacle Detection for Advanced Driver Assistance Systems Using Distortions of Inverse Perspective Mapping of a Monocular Camera." *Robotics and Autonomous Systems* 114:172–186.
- Pratt, W. K. 2013. *Introduction to Digital Image Processing*: CRC press.
- Uysal, M., A. S. Toprak, and N. Polat. 2015. "DEM Generation with UAV Photogrammetry and Accuracy Analysis in Sahitler hill." *Measurement* 73:539–543.
- Wand, M. P., and M. C. Jones. 1994. *Kernel Smoothing*. Boca Raton, FL: CRC Press.
- Wei, Y., J. Yang, C. Gong, S. Chen, and J. Qian. 2019. "Obstacle Detection by Fusing Point Clouds and Monocular Image." *Neural Processing Letters* 49 (3): 1007–1019.
- Zebbara, K., M. El Ansari, A. Mazoul, and H. Oudani. 2019. "A Fast Road Obstacle Detection Using Association and Symmetry Recognition." In *2019 International Conference on Wireless Technologies, Embedded and Intelligent Systems (WITS)*, 3–4 April, Fez, Morocco.
- Zhang, R., Y. Zheng, F. Zhang, W. Cai, and L. Zeng. 2019. "An Outdoors Obstacle Detection Algorithm Based on Disparity Slicing." In *Proceedings, Tenth International Conference on Graphics and Image Processing (ICGIP 2018)*, 6 May, Chengdu, China.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) June 2020			2. REPORT TYPE Technical Report / Final		3. DATES COVERED (From - To) FY18-FY21	
4. TITLE AND SUBTITLE Obstacle Detection and Quantification for Vehicle Mobility in Winter Conditions					5a. CONTRACT NUMBER	
					5b. GRANT NUMBER	
					5c. PROGRAM ELEMENT 622145	
6. AUTHOR(S) Sergey N. Vecherin, Jacob M. Shaker, and Michael W. Parker					5d. PROJECT NUMBER BG2	
					5e. TASK NUMBER SBG202	
					5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) U.S. Army Engineer Research and Development Center (ERDC) Cold Regions Research and Engineering Laboratory (CRREL) 72 Lyme Road Hanover, NH 03755-1290					8. PERFORMING ORGANIZATION REPORT NUMBER ERDC/CRREL TR-20-7	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Headquarters, U.S. Army Corps of Engineers Washington, DC 20314-1000					10. SPONSOR/MONITOR'S ACRONYM(S) USACE	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.						
13. SUPPLEMENTARY NOTES Mobility in Complex Environments						
14. ABSTRACT <p>Recently, the focus for military operations has shifted from the desert to cold climates, causing a corresponding shift in the military's need to better understand the mobility of our current vehicle fleet in these areas. Therefore, this work investigated the effects of winter conditions on military vehicle mobility. The main objective was to detect obstacles on the scene.</p> <p>This study developed and tested a method for automatic obstacle detection in the digital elevation model of a scene. The method detects statistical anomalies relative to an estimated background image that contains no obstacles. The sensitivity of the detection can be adjusted by a specified probability of false alarms, and the obstacle detection confidence is characterized by a probability of detection. The visible height of obstacles above the snow is related to the actual height of the obstacles above the ground.</p> <p>Compared to other detection techniques, the developed method is fast, calibrates itself to the cluttered images, operates with a single given image, and aligns with a detection quantification adopted in the receiver operating characteristic framework. The examples considered in this paper demonstrate high efficiency and applicability of the developed approach to the military vehicle mobility missions.</p>						
15. SUBJECT TERMS Algorithms, Cold regions, Drone aircraft, Ice, Obstacles (Military science)—Detection, Photogrammetry, Snow, Trafficability						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 36	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified			19b. TELEPHONE NUMBER (include area code)	