

Potential Negative Obstacle Detection by Occlusion Labeling

Nicholas Heckman, Jean-François Lalonde, Nicolas Vandapel and Martial Hebert
The Robotics Institute
Carnegie Mellon University
5000 Forbes Avenue, Pittsburgh, PA 15213
{nth, jlalonde, vandapel, hebert}@cs.cmu.edu

Abstract—In this paper, we present an approach for potential negative obstacle detection based on missing data interpretation that extends traditional techniques driven by data only which capture the occupancy of the scene. The approach is decomposed into three steps: three-dimensional (3-D) data accumulation and low level classification, 3-D occluder propagation, and context-based occlusion labeling. The approach is validated using logged laser data collected in various outdoor natural terrains and also demonstrated live on-board the Demo-III eXperimental Unmanned Vehicle (XUV).

I. INTRODUCTION

Autonomous ground robot navigation in natural outdoor environment requires capabilities to sense, analyze and model the complexity of the surrounding three-dimensional terrain. One of the most popular sensing modalities is the laser range scanner (or ladar), which allows 3-D perception of a robot’s surroundings, and enables the building of high-accuracy environment maps. However, this sensor (as well as other sensors like cameras) have several limitations, such as limited field of view and range. Their biggest limitation however is that they do not “see through” objects: in a cluttered environment, it might be impossible to see the entire scene at once from any given viewpoint, thus yielding incomplete maps where a large amount of data is missing. In some cases, missing data might be explained by negative obstacles, i.e. obstacles which lie below the ground surface such as holes or ditches, for example. They are regions where the robot cannot see and might be potentially dangerous for its mobility, and thus should be treated as obstacles.

In this paper, we seek to build a richer 3-D environment map by finding plausible explanations to regions where data is missing due to occlusions or sensing limitations, and explicitly find where they might be caused by a negative obstacle. We go beyond traditional methods that rely only on existing data to interpret the scene. The proposed approach is applied to potential negative obstacle detection and demonstrated in various terrain types on-board a mobile robot equipped with a ladar (Fig. 1¹). We only label terrain regions susceptible to containing negative obstacles, so the term “potential”. We assume that, if the negative obstacle can be mapped, at close range for example, mobility assessment can be performed.

The approach leverages on existing work done on terrain classification [1], [2]. Laser data is accumulated in 3-D and



Fig. 1. Two of the terrains we experimented in with the Demo-III eXperimental Unmanned Vehicle: large ravine with a stream and terrain depression.

classified into three basic classes (linear, scatter, surface) based on the local geometric property of the point cloud. Ray tracing is performed in 3-D to determine the occluded regions in the scene. The occluder’s class is propagated through the data structure, producing a dense 3-D map which captures the occupancy, classification, visibility and nature of the occlusion for each voxel. Finally, context-based labeling of missing data is used to detect potential negative obstacle.

Our approach has several advantages over other preexisting approaches. First, we are able to work with accumulated laser data versus working with a frame by frame approach. Our approach works with complex scenes with trees, shrubs, sparse and dense vegetation, and vegetation-covered ground surface. The approach also is invariant to the slope of the ground surface. Finally, the labeling is associated with a confidence measure.

II. RELATED WORK

Since negative obstacle detection is very important for ground robotics, it has received a fair amount of attention in the literature. Our approach is original because it uses 3-D data acquired by a ladar.

Existing techniques use cameras, as in [3]. The authors rely on shadows cast by negative obstacles when illuminated by two lighting sources below and above stereo cameras. Those shadows provide cues to complement and enhance stereo vision.

Terrain depressions exhibit different thermal properties than their surroundings: they stay warmer at night. This fact was used by Matthies and Rankin [4] to detect negative obstacles. In contrast, our approach is robust to weather conditions and time of day, as we are using an active sensor, which does not depend on ambient illumination.

¹This paper is best viewed in color

Leading and trailing edges of negative obstacles produce discontinuities in images. Rosenblum et al. [5] detect such discontinuities in multi-spectral imagery as one of many cues for obstacle detection. Dima et al. [6] fuse multiple classifiers trained on texture, color, infrared and 3-D saliency features.

Negative obstacles, in some instances, are materialized by the transition sequence data / no data / data. Matthies et al. [7], [8] detect such transitions in range images. The algorithm takes into account the height of the trailing edge, the length of the data gap and the slope between the leading and trailing edges and a vehicle mobility model. In [9], the authors also take advantage of the data gap, but they do so in a 2-D occupancy map of the terrain. Our approach relies on accumulated 3-D data, takes into account a richer interpretation of the terrain (classification versus occupancy), de-couples the detection problem from the mobility assessment problem, and detects only potential negative obstacles.

High resolution overhead data (several points per square meter) was demonstrated to be suitable for negative obstacle detection [10], [11], [12]. The top-down point of view ensures that the bottom of the negative obstacle is observed, which is not the case if the obstacle is seen from a ground perspective. In contrast, our approach does not rely on the existence of such data and directly estimates the potential negative obstacles from the robot's point of view.

III. APPROACH

Our approach relies on determining the location of occlusions and identifying a subset of those where a negative obstacle might occur. Fig. 2 illustrates the process, which is decomposed into three major steps: 1) data accumulation and classification in high resolution 3-D maps (first row), 2) occluder propagation in low resolution 3-D maps (second row), and 3) context-based labeling in a low resolution 2-D map (third row).

A. Data accumulation and terrain classification

Using voxelized 3-D ladar data as input, we adopt the method introduced in [1] and perform voxel-wise classification to detect vegetation, thin structures and solid surfaces. The method relies on the use of the scatter matrix to extract features via principal component analysis (PCA). For each voxel, the approach computes the scatter matrix within a support volume and then extracts its principal components (eigenvalues). A linear combination of the components and the associated principal directions define the features. A model of the features distribution is learned off-line, prior to the mission, from manually labeled data. As the robot traverses a terrain, data is accumulated, features computed and maximum likelihood classification performed on-line.

Fast terrain classification is enabled by an efficient scrolling data structure which takes advantage of overlapping neighborhoods and the reuse of previously computed data as the algorithm scans each voxel as in [2].

B. Scene ray-tracing and occlusion labeling

We can now use the classification results to determine where occlusions occur in the scene and identify several different types of occlusion.

Classification information is first collapsed from the high-resolution 3-D map to a low-resolution 3-D map for computational tractability. Low-resolution voxels are obtained by grouping together several high-resolution voxels, and assigning the most frequent class within them. To account for noise, such as generated by sparse vegetation in the scene, we require a minimum level of occupancy of voxels in the low resolution 3-D map.

Occlusion labeling is performed by ray tracing from the current position of the laser through all the occupied voxels that exist in an area of interest in the scene. A conventional ray tracing algorithm is used [13]. The voxels that lie behind an occupied voxel (along the ray tracing direction) are considered occluded. Every occluded voxel stores the class of its occluder. Likewise, empty voxels that exist between the laser and an occupied voxel are considered visible. Figs. 2-c and 2-d illustrate this technique. Note that these figures display only occlusions by one class for clarity purposes, and occlusions by the two other classes are propagated as well. The complexity of this algorithm is $O(vk)$, where v is the number of occupied voxels and k is the distance (in voxels) from the laser position to a corner of the data structure.

C. Context-based labeling

As illustrated in Fig. 3, negative obstacles may occur only in regions where there is a transition from visible to occluded voxels. We explicitly look for those regions by searching through columns in the low resolution 3-D map.




Fig. 3. Transitions between visible unoccupied and occluded unoccupied voxels suggest the existence of a negative obstacle.

We first create a 2-D map, aligned with the low-resolution 3-D map, parallel to the ground plane. It will be used to store the location of the potential negative obstacles detected. Each 3-D column is scanned in a top-down fashion, until either: the transition between an empty voxel to an occluded voxel is found; an occupied voxel is found; or the end of the data structure is reached. If the first case (transition) is encountered, we mark the 2-D cell corresponding to the column as being a potential negative obstacle. Additionally, we also store the class of the voxel that caused the occlusion in the 2-D map. Fig. 2-f illustrates the transition regions found in that particular situation. The complexity of this step is $O(n)$, where n is the number of voxels in the 3-D low resolution map. From now on, all processing will be done in 2-D.

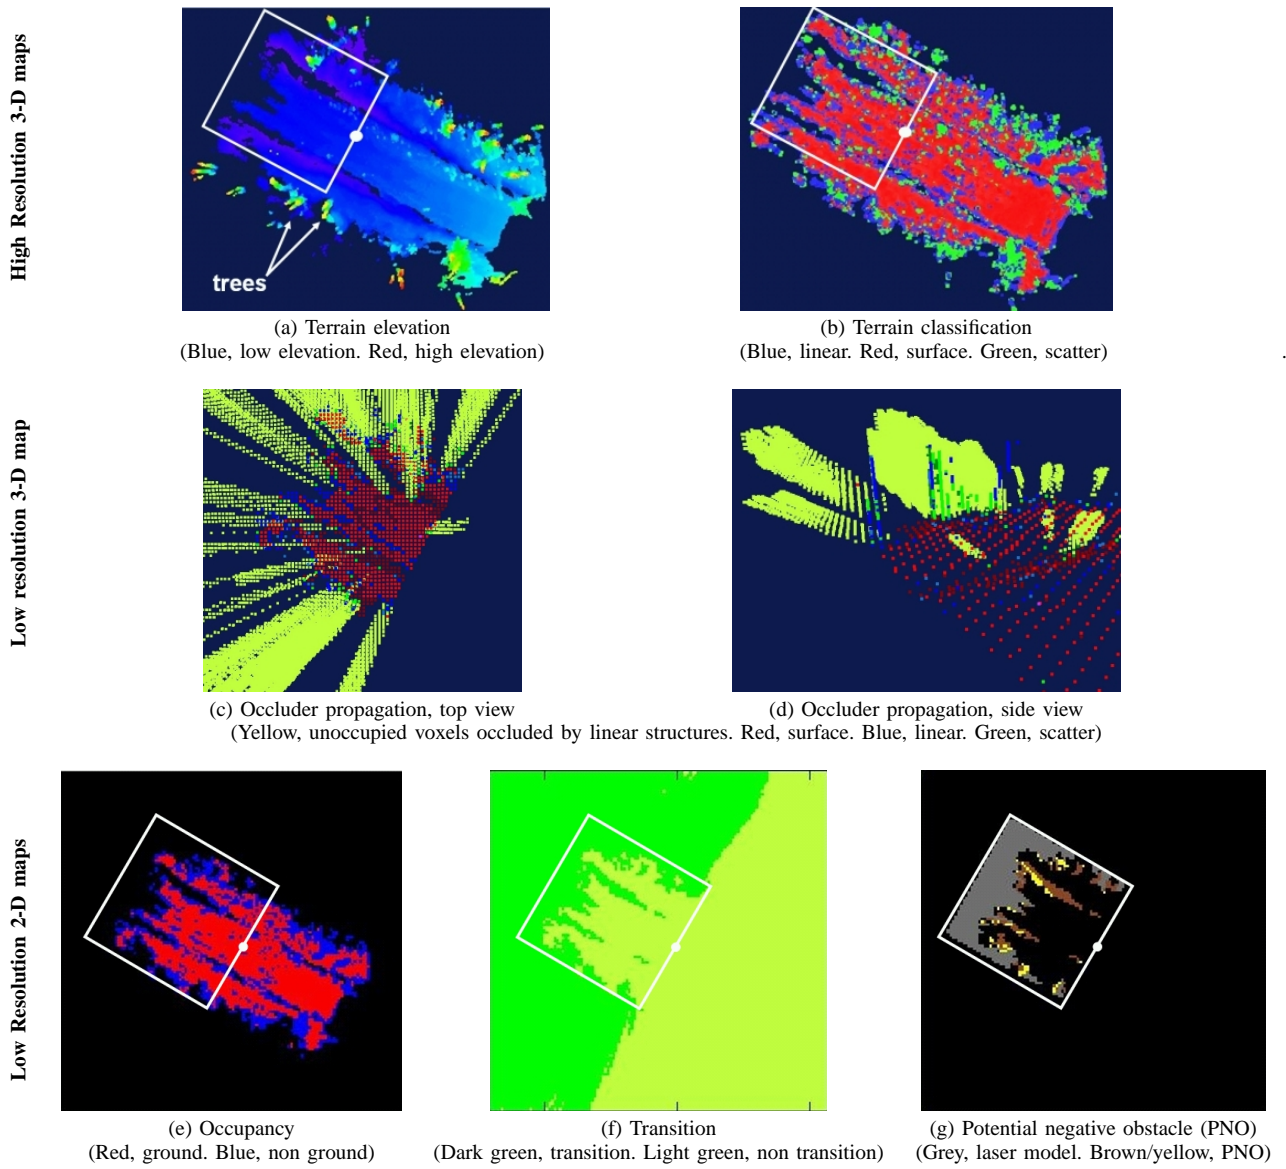


Fig. 2. Approach. The center of the vehicle is shown with a white dot, the area of interest by a white square. (g) Only labeled missing data within the area of interest is shown.

Now that we know *where* data is missing, we examine the scene context to determine *why* data is missing. Based on the type of occluder and if data can be found after each transition voxel identified in the previous step, we differentiate between four types of explanation for missing data, which are illustrated in Fig. 4.

This step is performed on each transition voxel in an interest area around the robot. First, the type of occluder is determined by looking up the class of the voxel that caused the occlusion. Next, the existence of data after the transition is determined by performing ray tracing between the projected laser position and the transition location in the 2-D map. The appropriate class can then be assigned to each transition voxel according to Fig. 4. Canonical results are shown in Fig. 2-g. Additionally, we use a model of our laser that is based on the maximum range of return on flat ground surface to determine which regions of missing data cannot

be seen. The complexity of this step is $O(v'k')$, where v' is the number of occupied voxels in the low resolution 2-D map and k' is the distance (in voxels) from the laser position to a corner of the data structure.




Fig. 4. Context-based labeling for potential negative obstacle detection

IV. EXPERIMENTS

This section presents some implementation details omitted in the previous section and experiments from off-line and on-board data processing validating the approach, including obstacle avoidance of negative obstacle.

A. Implementation, vehicle and experimental set-up

The algorithm runs on a stand-alone, off-the-shelf computer with a 3 GHz CPU and 2 GB of RAM. Our program is interfaced with the vehicle to acquire laser data and vehicle state. The voxel size of high resolution 3-D, low resolution 3-D and low resolution 2-D maps are 10 cm^3 , $40\times 40\times 10\text{ cm}^3$, and $40\times 40\text{ cm}^2$ respectively. The map dimensions are $50\times 50\times 6\text{ m}^3$, $50\times 50\times 6\text{ m}^3$ and $50\times 50\text{ m}^2$ respectively. The map is scrolling with the vehicle and maintains a constant orientation.

For live results, to interface with the XUV planner, we assign a different cost to each class of potential negative obstacle. These costs are currently engineered manually, and a probabilistic approach is the subject of future work.

Experiments were conducted in various terrains type in Central Pennsylvania during the summer and winter seasons, and with different vegetation density.

B. Off-line playback

Fig. 5 shows results produced off-line with pre-recorded data. Note that the data flow is exactly the same as the one used when processing data live on-board the vehicle except that the process is reading off a file instead of a live data stream. The vehicle is driving on a dirt road boxed by small drainage trenches, a half-meter deep and one meter wide. A bridge crosses a stream and a deep terrain depression (two meters deep and five meter wide). From left to right, the columns represent the raw 3-D data (top row) and the potential negative obstacle detection after accumulation 756, 918 and 1134 lidar frames. In the first column, the water stream depression is not yet visible but the drainage ditches are detected correctly. In the second column, the trailing edge of the depression just appeared and the corresponding negative obstacle is partially labeled. Finally, in the last column, the terrain behind the negative obstacle is clearly visible and the obstacle is correctly labeled.

In the second example, illustrated in Fig. 6, we show the robustness of the approach when the terrain is rough, includes large trees and is partially covered by tall grass. Some potential negative obstacle are detected at the edge of the perceived terrain due to a lack of data density. However, note that trees and vegetation cover are handled correctly by our approach even though they introduced range discontinuities, partial or even total occlusion of the terrain.

C. Field test results on-board the Demo-III XUV

In this section, we present results produced on-board the vehicle in real-time. First, we show in Fig. 7 results on flat ground with a change of elevation. Note the similar problem at the edge introduced by the lack of point density.

Fig. 8 shows timing results for a run performed in a forest. The graphs represent cumulative histograms of classification time, ray-tracing, context-based labeling and their total based on timing collected at each iteration of the algorithm. For example, 80% percent of the algorithm runtime takes less than 800 ms. Note that this code was run with very large data structures.

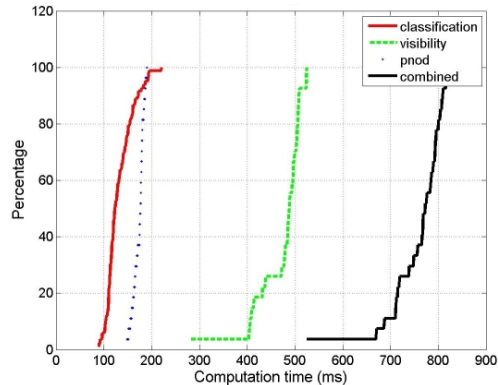


Fig. 8. Timing performance. Area of interest: 50 x 100 m. 3 GHz computer, 2 GB RAM.

Finally, we show preliminary results on how this approach can be used practically to influence the behavior of a vehicle. We exported the areas most likely to contain negative obstacles, as defined in Section III-C, into an obstacle map read by the Demo-III XUV. The vehicle was commanded to follow a path between two way points set on each side of a terrain depression. Note that this terrain feature does not challenge the mobility of the autonomous vehicle. Such obstacle was chosen for safety consideration. Note that this terrain feature fits our definition of potential negative obstacle even though it is not an obstacle. The terrain is illustrated in Figs. 1 and 9-a. As illustrated in Fig. 9-b/c, the terrain feature was characterized correctly. The vehicle stopped, tried to plan around the terrain depression based on the data available, but could not and stopped moving. The best candidate path illustrated in Fig. 9-d intersected the obstacle and was not executed. Note that these results are preliminary and need to be confirmed by further field testing.

V. CONCLUSION

Negative obstacle detection is a very challenging problem that cannot be entirely and robustly solved from a single perspective (ground versus overhead) or using a single sensing modality (radar, imagery or laser) or algorithm.

In this paper, we presented an approach for potential negative obstacle detection based on missing data interpretation that extends traditional data-driven techniques. The approach is decomposed into three steps: 1) 3-D data accumulation and low level classification, using high-resolution 3-D maps, 2) 3-D occluder propagation, using a low resolution 3-D map, and 3) context-based occlusion labeling, using low resolution 2-D maps.

Our approach has several advantages over other existing techniques. By working with accumulated data in 3-D, we

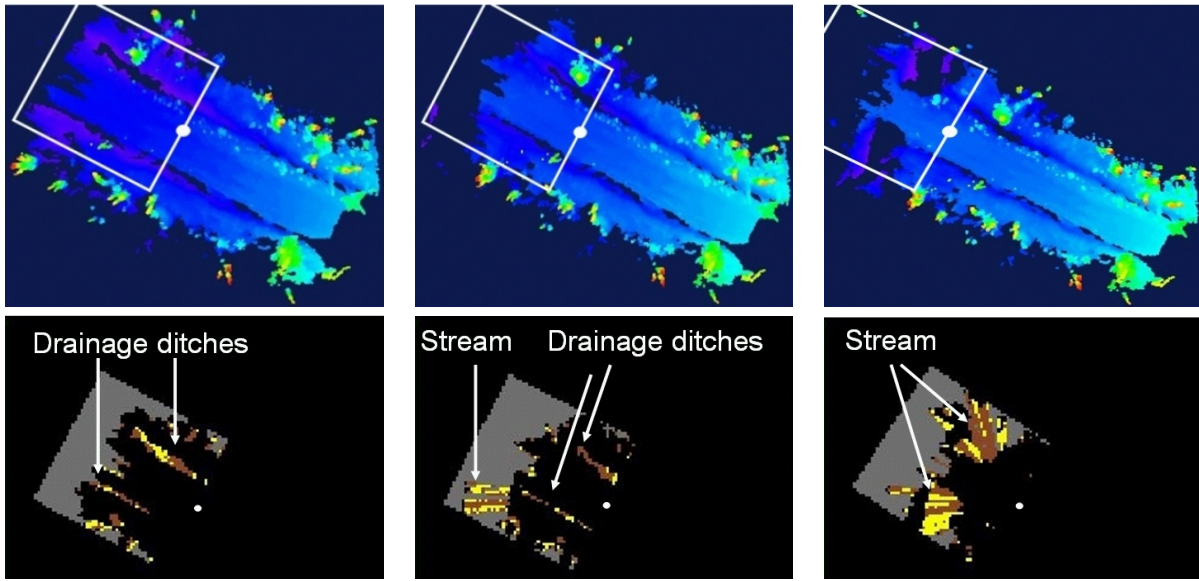


Fig. 5. Playback results. Trail with ditches and a bridge. Results after accumulation of 756, 918 and 1134 ladar frames (left, center and right column). Upper row, terrain model color coded by elevation, red for high and blue for low. The center of the vehicle is shown with a white dot, the area of interest by a white square. Bottom row, potential negative obstacle detection (grey, laser model; brown/yellow, potential negative obstacle with data after). Only labeled missing data within the area of interest is shown.




Fig. 6. Playback results. Wooded environment with a rough terrain, vegetation and large trees. Left: scene. Center: raw 3-D data color-coded by elevation from high (red) to low (blue). The center of the vehicle is represented by a white dot, the area of interest by a white square. Right: potential negative obstacle detection (grey, laser model; brown/yellow, potential negative obstacle with data after; orange/dark yellow, potential negative obstacle without data after). Only labeled missing data within the area of interest is shown.




Fig. 7. Live on-board results. Flat ground. Left: scene. Center: raw 3-D data color-coded by elevation from high (red) to low (blue). The center of the vehicle is shown with a white dot, the area of interest by a white square. Right: potential negative obstacle detection (grey, laser model; brown/yellow, potential negative obstacle with data after; orange/dark yellow, potential negative obstacle without data after). Only labeled missing data within the area of interest is shown.

are able to use a richer terrain representation. It can handle complex scenes with trees, shrubs, sparse and dense vegetation, and vegetation-covered ground surface and it is invariant to the slope of the ground surface. Finally confidence is

associated to each label.

Our approach suffers from several limitations that are currently being addressed. Water ponds are misclassified but we believe that this problem can be solved by looking at the

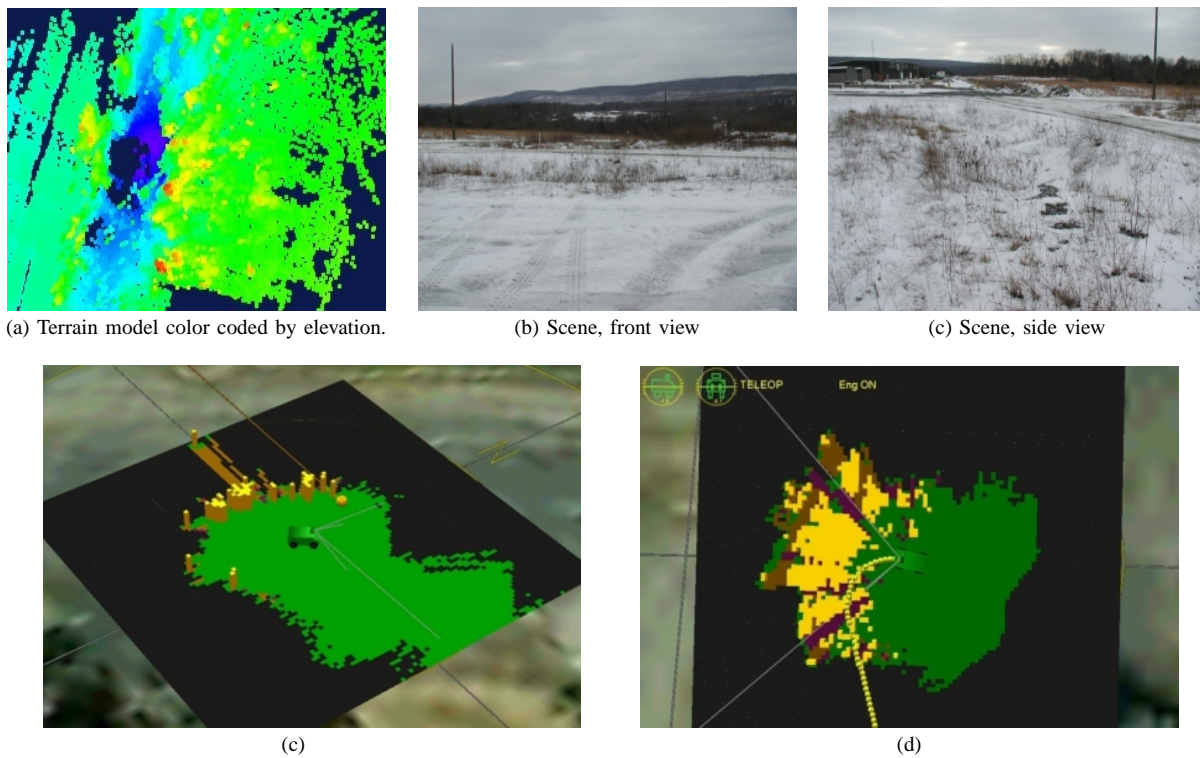


Fig. 9. On-board preliminary results. (a) The terrain model is color coded by elevation (Blue for low, red for high). (c) Terrain classification map used by the planner. In green, known traversable terrain; in yellow, obstacle with the cost coded as elevation. (d) Obstacle map and planned path for the same scene but a different run. The vehicle approached the obstacle closer than before.

elevation of the surrounding terrain because water produces horizontal obstacles. Data density will be taken into account to avoid some observed misclassification. The algorithm runs in real-time but a slow frame rate of 1 Hz, implementation improvements are under way. As mentioned before, the approach is deterministic, and a probabilistic model is currently being investigated. As with existing work, only qualitative results were presented in this paper, a quantitative evaluation is planned in order to thoroughly evaluate the reliability of such approaches.

VI. 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 General Dynamics Robotic Systems for its support.

REFERENCES

- [1] J.-F. Lalonde, N. Vandapel, M. Huber, and M. Hebert, "Natural terrain classification using three-dimensional lidar data for ground robot mobility," *Journal of Field Robotics*, 2006.
- [2] J.-F. Lalonde, N. Vandapel, and M. Hebert, "Data structure for efficient processing in 3-d," in *Robotics: Science and Systems 1*. MIT Press, June 2005.
- [3] G. Witus, R. Karlsen, D. Gorsich, and G. Gerhart, "Preliminary investigation into the use of stereo illumination to enhance mobile robot terrain perception," in *Unmanned Ground Vehicle Technology III*, 2001.
- [4] L. Matthies, "Negative obstacle detection by thermal signature," in *IEEE/RJS International Conference on Intelligent Robots and Systems*, 2003.
- [5] M. Rosenblum and B. Gothard, "A high fidelity multi-sensor scene understanding system for autonomous navigation," in *IEEE Intelligent Vehicle Symposium*, 2000.
- [6] C. Dima, N. Vandapel, and M. Hebert, "Classifier fusion for outdoor obstacle detection," in *International Conference on Robotics and Automation*, 2004.
- [7] L. Matthies, A. Kelly, T. Litwin, and G. Tharp, "Obstacle detection for unmanned ground vehicles: a progress report," in *IEEE Intelligent Vehicles Conference*, 1995.
- [8] —, "Obstacle detection for unmanned ground vehicles," in *International Symposium of Robotics Research*, 1995.
- [9] A. Rankin, A. Huertas, and L. Matthies, "Evaluation of stereo vision obstacle detection algorithms for off-road autonomous navigation," in *AUVSI*, 2005.
- [10] D. Silver, B. Sofman, N. Vandapel, J. Bagnell, and A. Stentz, "Experimental analysis of overhead data processing to support long range navigation," in *IEEE/RJS International Conference on Intelligent Robots and Systems*, 2006.
- [11] N. Vandapel, R. Donamukkala, and M. Hebert, "Unmanned ground vehicle navigation using aerial lidar data," *International Journal of Robotics Research*, vol. 25, no. 1, 2006.
- [12] A. Kelly, A. Stentz, O. Amidi, M. Bode, D. Bradley, A. Diaz-Calderon, M. Happend, H. Herman, R. Mandelbaum, T. Pilarki, P. Rander, S. Thayer, N. Vallidi, and R. Warner, "Toward reliable off road autonomous vehicles operating in challenging environments," *The International Journal of Robotics Research*, vol. 25, no. 5-6, 2006.
- [13] J. Amanatides and A. Woo, "A fast voxel traversal algorithm for ray tracing," in *Eurographics*, 1987.