Nighttime negative obstacle detection for off-road autonomous navigation

Arturo L. Rankin*, Andres Huertas, Larry H. Matthies
Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive, Pasadena, CA, USA 91109

ABSTRACT

Detecting negative obstacles (ditches, holes, wadis, and other depressions) is one of the most difficult problems in perception for unmanned ground vehicle (UGV) off-road autonomous navigation. One reason for this is that the width of the visible portion of a negative obstacle may only span a few pixels at the stopping distance for vehicle speeds UGV programs aspire to operate at (up to 50kph). The problem can be further compounded when negative obstacles are obscured by vegetation or when negative obstacles are embedded in undulating terrain. Because of the variety of appearances of negative obstacles, a multi-cue detection approach is desired. In previous nighttime negative obstacle detection work, we have described combining geometry based cues from stereo range data and a thermal signature based cue from thermal infrared imagery. Thermal signature is a powerful cue during the night since the interiors of negative obstacles generally remain warmer than surrounding terrain throughout the night. In this paper, we further couple the thermal signature based cue and geometry based cues from stereo range data for nighttime negative obstacle detection. Edge detection is used to generate closed contour candidate negative obstacle regions that are geometrically filtered to determine if they lie within the ground plane. Cues for negative obstacles from thermal signature, geometry-based analysis of range images, and geometry-based analysis of terrain maps are fused. The focus of this work is to increase the range at which UGVs can reliably detect negative obstacles on cross-country terrain, thereby increasing the speed at which UGVs can safely operate.

Keywords: Negative obstacle detection, stereo vision, perception, autonomous navigation, thermal infrared

1. INTRODUCTION

There are two primary classes of unmanned ground vehicle (UGV) negative obstacles: depressions and drop-offs. Depressions are areas that are sunken below the surrounding terrain. Some of the terms used to describe depression type negative obstacles include ditch, trench, hole, and wadi. Some of the terms used to describe drop-off type negative obstacles include ledge, cliff, and steep decline. In terrain maps built from range data, large areas of no data typically exists where there are drop-offs. As large no data regions should be avoided, drop-off type negative obstacles don't explicitly need to be detected as a negative obstacle. In this paper, we focus solely on depression type negative obstacles. Detecting negative obstacles from an unmanned ground vehicle (UGV) in time to stop at vehicle at speeds UGV programs aspire to operate at (up to 50kph) remains a difficult problem. For a UGV traveling at a speed of 24kph, the stopping distance is approximately 11.9 meters. This analysis assumes a 30% down grade, a coefficient of friction of 0.65 for dry soil, 0.5 seconds to detect the trench and fully apply the brake, and a 1.8 meter safety buffer between the stopped vehicle and the trench [1]. Other references suggest a more conservative stopping distance of 17.8 meters to account for differences in terrain and soil conditions and reaction time [2]. For all-terrain sized wheeled UGV's, a negative obstacle slightly wider than the wheel diameter (~0.5 meters) is large enough to cause vehicle damage at moderate speeds. At higher speeds, however, a UGV may be able to traverse narrow negative obstacles by jumping over them [3]. Determining if a narrow negative obstacle can be jumped over (rather than stopping or turning to avoid it) is beyond the scope of this paper.

In [4], linear negative obstacles are detected by combining geometric analysis of range data and an intensity image horizontal band detector. Most prior negative obstacle detection work, however, has focused on geometric analysis of range data since the geometry of a negative obstacle is its dominant characteristic [5-6]. In [7], occlusions are located in

* Arturo.L.Rankin@jpl.nasa.gov; phone (818) 354-9269; fax (818) 393-5007

accumulated ladar range data and missing data consistent with a depression are labeled a "potential" negative obstacle. Other techniques, which do not look for negative obstacle specific features, include assessing increasing higher cost to data gaps in a terrain map as the gaps approach a UGV [8] and training multiple classifiers using multiple sensors (color, infrared, ladar) and fusing the classification results [9].

Negative obstacles are difficult to detect beyond short distances with ranging sensors for several related reasons. First, the visible portion of a negative obstacle rapidly shrinks with range. While the angle subtended by a positive obstacle decreases by the inverse of range, the angle subtended by a negative obstacle decreases by the inverse of range squared [10]. Secondly, unless the sensors are mounted on a mast [8], it is hard to see into a depression far enough to distinguish it from terrain self occlusions that naturally occur on cross country terrain. In Demo III field experiments at Aberdeen Proving Ground, MD, a man-made 0.6 meter trench on flat terrain was detected using a Dornier EBK ladar range imaging camera (8.2mrad vertical pixel resolution) at a maximum range of 9.5 meters for a vehicle speed of 24kph [1]. Similarly, the 0.6 meter trench was detected in FLIR stereo range data (2.3mrad vertical pixel resolution) at a range of 7.3 meters and in narrow field-of-view (FOV) CCD stereo range data (1.0mrad vertical pixel resolution) at a range of 11.0 meters [11], all short of the required detection range for stopping prior to the ditch for a vehicle speed of 24kph.

In addition, it is more difficult to detect naturally occurring narrow negative obstacles than man-made narrow negative obstacles. Man-made negative obstacles are generally dug on flat terrain with a backhoe so that the ditch sides are vertical and well defined. Naturally occurring negative obstacles, however, tend to exist on undulated terrain and have eroded sides. Figure 1 illustrates the challenge of detecting naturally occurring narrow negative obstacles. The left image shows a side view of a negative obstacle from Ft. Knox that has a width of approximately the UGV wheel diameter. The middle image is from one of the UGV's forward looking color cameras during a normal approach from the right side of the trench. Here, the trailing edge of the trench is clearly visible, which aids in negative obstacle detection. The right image is from a normal approach from the left side of the trench. During the approach, the trailing edge of the trench has a lower elevation than the leading edge. Looking at the trench from a distance, one can not tell if there is a slight depression that is traversable, or a deep trench that is non-traversable. An additional cue is needed to disambiguate terrain self occlusions from narrow, but deep, negative obstacles. During the Perception for Off-Road Robots (PerceptOR) program, a thermal signature cue for negative obstacles was developed to perform such a function during the nighttime.

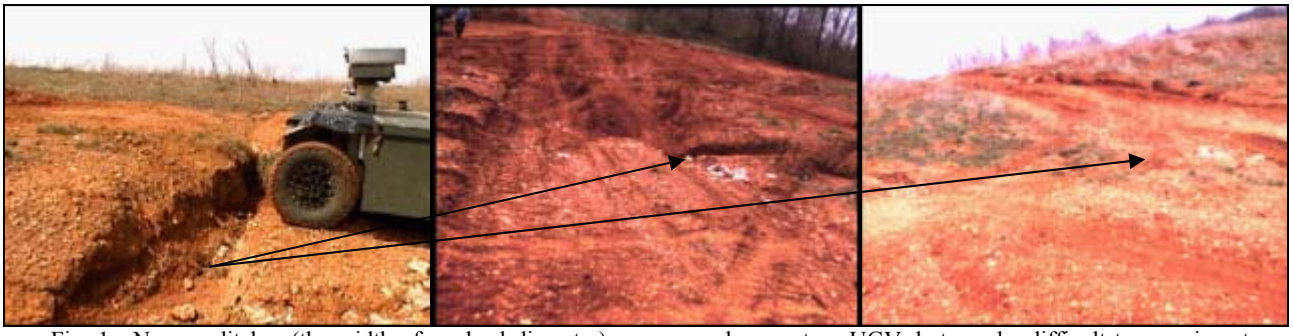

Fig. 1. Narrow ditches (the width of a wheel diameter) can cause damage to a UGV, but can be difficult to perceive at ranges required to stop (left image). The upslope on the trailing edge of a trench aids detection (middle image). When there is an absence of a visible upslope on the trailing edge, however, it is difficult to distinguish a slight depression from a deep trench (right image).

PerceptOR was one of six key supporting technology programs of the Defense Advanced Research Projects Agency (DARPA)/Army Future Combat System (FCS) program which ran from March 2001 until March 2004. It was designed to develop prototype perception approaches to advance the state-of-the-art for off-road obstacle detection, and to enable higher levels of autonomous mobility needed for FCS operations. The Science Applications International Corporation (SAIC) led team, called Team Raptor, was composed of SAIC, the Machine Vision Group (MVG) at the Jet Propulsion Laboratory (JPL), Carnegie Mellon University (CMU), Applied Perception Incorporated (API), and Visteon Corporation. Team Raptor had a thrust element to increase the range at which negative obstacles can be reliably detected with traditional geometry based negative obstacle detectors by exploiting the heat transfer characteristics of negative obstacles via analysis of thermal infrared intensity images.

A simple model of heat flux for radiative, convective, and conductive heat transfer was generated that, at night, predicts the interior of a negative obstacle should cool more slowly than surrounding terrain. Then, several twenty-four hour data collections were performed on trenches that verified that negative obstacles tend to be warmer than their surrounding terrain for most of the night. A nighttime negative obstacle detection algorithm was implemented on the Team Raptor UGV that combined cues from stereo range data and thermal signature [10, 12]. In our algorithm reported in [10], the thermal signature of negative obstacles was located by independently analyzing each column in a rectified intensity image (registered with stereo range data) for large intensity gradients. Consistency checks on the relative geometry of the start pixel, end pixel, and any range data in between were used to support or reject a negative obstacle hypothesis. This algorithm, however, was subject to false alarms on horizontal structure warmer than the ground but above the ground plane. We subsequently improved the algorithm by analyzing the thermal signature of closed contours in rectified intensity images, strengthened geometry filtering to verify that candidate negative obstacles are in the ground plane, and merged negative obstacle detection results into a world traversability cost map that can be used to plan safe paths for the UGV. This paper details the improvements to the algorithm. In section 2, we detail the new combined geometry/thermal signature approach using twenty thermal infrared stereo image pairs from a sequence acquired during the nighttime. In section 3, we discuss inserting detected negative obstacles into instantaneous and world maps. Section 4 contains some results, and section 5 contains conclusions and suggested follow up work.

2. GEMOMETRY/THERMAL CUES FOR NIGHTTIME DETECTION

During the PerceptOR program, a negative obstacle detector was developed and implemented on the Team Raptor vehicle that combines the results of traditional geometry-based negative obstacle detection (that detects a gap followed by an upslope in the range data) and a new approach that classifies negative obstacles at nighttime based on thermal signature. The new heat signature based negative obstacle detector was evaluated using the narrow trench shown in Figure 2. It uses edge detection applied to the intensity image to generate closed contour regions that are analyzed for cues for a negative obstacle. A Gaussian operator and Laplacian operator are applied to the rectified infrared intensity image in a single step (LoG convolution). Gaussian smoothing effectively suppresses the influence of the pixels that are up to a distance sigma from the current pixel. We have used a sigma of 1.75 for the results produced in this paper. The Laplacian operator approximates the second derivative and is non-directional. While the first derivative of an image function has an extreme at the position corresponding to an edge in the image, the second derivative is zero at the same position. Thus, zero crossings in the second derivative image indicate the location of region boundaries.

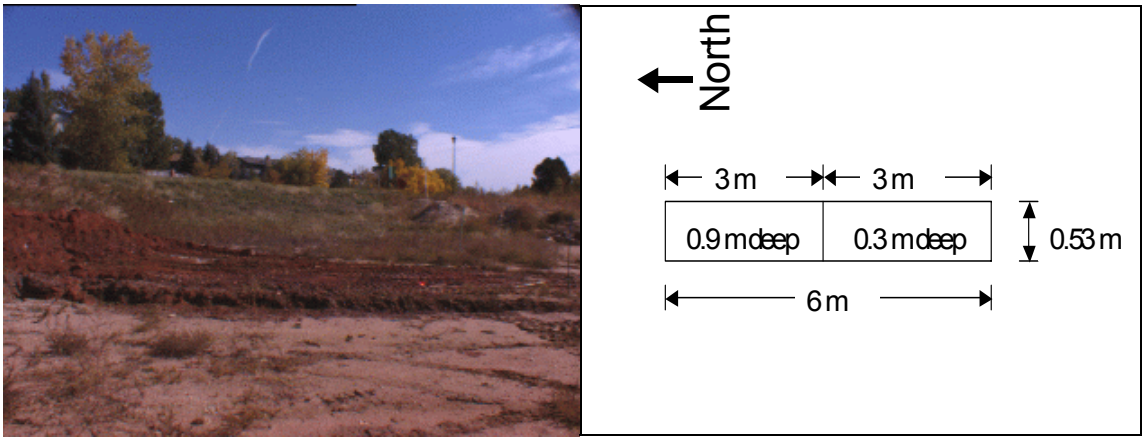

Fig. 2. A 0.53m wide trench was dug on the property East of the SAIC building in Littleton, CO, in early October 2002. Oriented north-south, the trench dimensions are shown to the right of the picture.

The output of the LoG convolution is used to produce a closed contour image. Each closed contour region is analyzed for thermal signature and geometry cues (rather than analyzing each column in the intensity image independently as was done in [10]). The range data within each closed contour is also analyzed for geometry cues. Following intensity based thresholding and geometry-based thresholding, a binary negative obstacle image is generated. The algorithm has been tested on twenty frames from a nighttime data collection sequence which was acquired on the Raptor vehicle during a

diagonal approach of the trench at 1 meter/s at 7:23pm, December 5, 2002. The trench was dug months earlier (prior to October 11) to exclude thermal artifacts related to a freshly dug trench.

Thermal infrared sensing was performed using a stereo pair of Cincinnati Electronic's NC256 mid-wave infrared (MWIR) cameras having a FOV of 67.5°x62.4°. The MWIR cameras output analog images at a resolution of 256x256 pixels. The framegrabber, however, sampled the images at a resolution of 640x480 pixels. In software, the intensity images were reduced to a resolution of 320x240 pixels [12]. This is the resolution that Team Raptor stereo and negative obstacle classification was performed at. Each pixel in the intensity image had a value that ranged from 0 to 255.

A sample rectified infrared intensity image of the trench in Figure 2 is shown in the left image of Figure 3. The images in Figures 3-8 all illustrate the steps involved in producing a negative obstacle image. For each rectified infrared intensity image pair processed, a LoG operator is applied (using a sigma=1.75), which produces an image with the negative pixels denoting bright intensity regions. The zero crossings correspond to the region boundaries in the rectified intensity image. The right image in Figure 3 is the LoG image of the left image in Figure 3. The brighter regions in the rectified intensity image have negative values (darker intensities) in the LoG image.

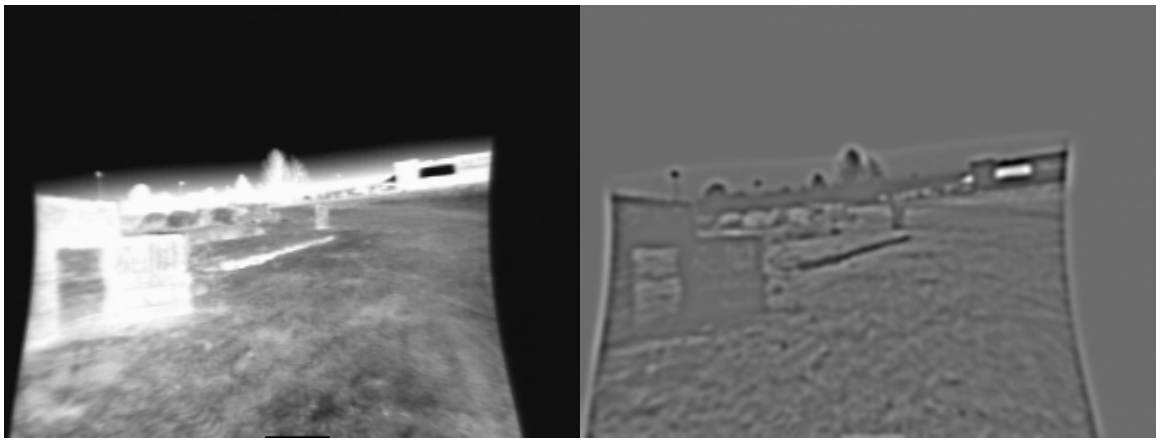

Fig. 3. Sample rectified infrared intensity image from a diagonal approach to the trench in Figure 2 (left) and its LoG image (right). Bright pixels in the intensity image have negative pixel values in the LoG image. Region boundaries in the intensity image have pixel values of zero in the LoG image. Hay bales are to the left of the trench and just beyond the right side of the trench.

In the left image in Figure 4, the pixels in the LoG image having negative values are shown in red. In the right image in Figure 4, only the pixels with negative LoG values that are not close to zero (< -1.8) are shown in red. To further reduce the number of candidate negative obstacle regions, small regions (less than 50 total pixels) were eliminated from consideration. The left image in Figure 5 illustrates this. This region image is a binary image where pixels inside each region are assigned 255. A closed contour image is generated from the region image by extracting only the outer perimeter of each region. The right image in Figure 5 illustrates the closed contour image. The left image in Figure 6 shows the closed contour image overlaid on the rectified intensity image.

To perform analysis on the candidate negative obstacle regions, each region needs to have a unique label. A connected component algorithm fills each pixel in a region with a unique value. If each pixel in a connected components image is represented by 8 bits, there can be up to 255 unique connected components in the image. The right image in Figure 6 contains the connected components image for the region image shown in Figure 5. Here, the pixel values are normalized (or stretched) to span 0-255 for display purposes. Each connected component in the image has a unique shade of gray.

The connected components image and the rectified intensity image are registered. For each connected component, an average interior intensity is calculated using the rectified infrared intensity image. Similarly, an average intensity is calculated for the boundary of each connected component. Here, a two pixel wide boundary around each connected component is used. The difference between the average intensity within each connected component and the average intensity around the border of each connected component is assigned to each connected component in an intensity

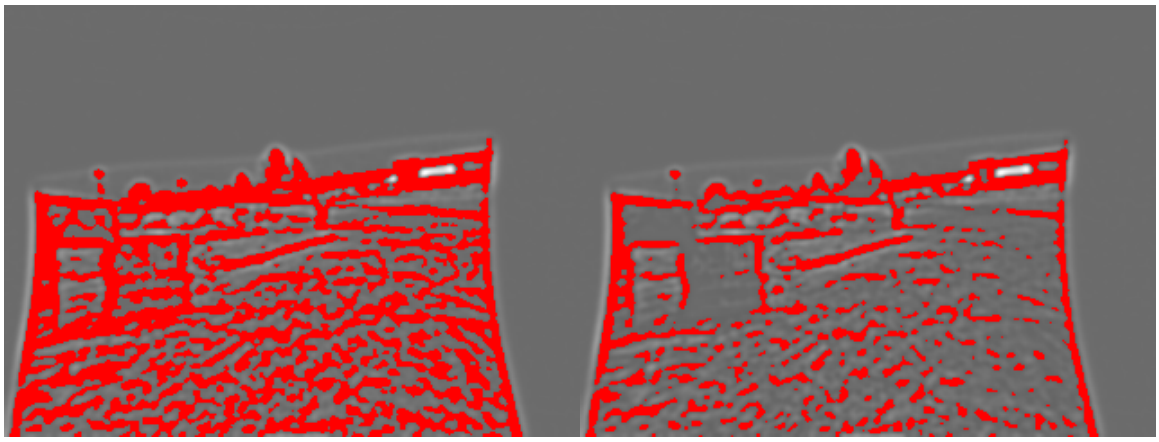

Fig. 4. In the left image, the pixels in the LoG image having negative values are colored red. In the right image, the pixels in the LoG image having negative values not close to zero (< -1.8) are colored red.

Fig. 5. The left image contains potential negative obstacle regions following the filtering in Figure 4 and removal of all regions smaller than 50 pixels. The right image is a closed contour image of the candidate negative obstacle regions.

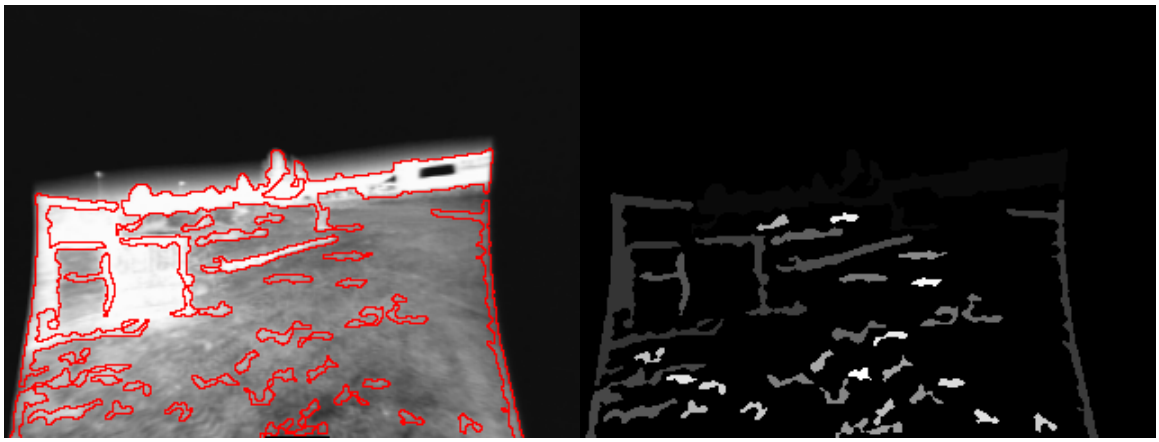

Fig. 6. In the left image, the closed contour image is overlaid on the rectified infrared intensity image. The right image is a normalized connected components image. Here, the intensity values have no meaning beyond illustrating that each connected component has a distinct label.

difference image. The left image in Figure 7 illustrates the intensity difference image. (Here, the average intensities are normalized between 0-255 to enhance the image for display purposes.) The intensity difference image is then thresholded and a binary negative obstacle image is created. The right image in Figure 7 illustrates thresholding the intensity difference image with a value of 40. Note that most of the candidate negative obstacle regions are removed with this step.

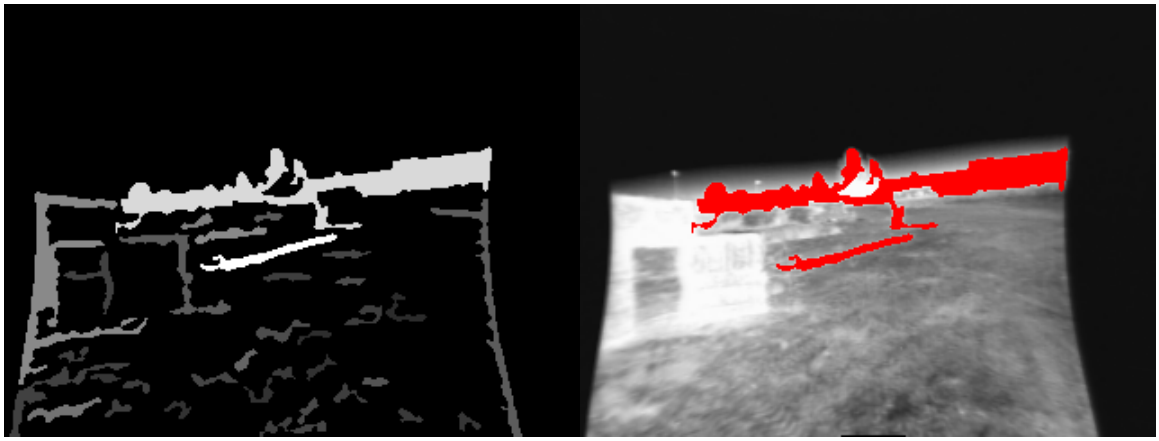

Fig. 7. Normalized intensity difference image (left) and the negative obstacle image after thresholding the intensity difference image with a value of 40 (right).

The negative obstacle image generated by thresholding the intensity difference image is then passed through a set of geometry filters. The left image in Figure 8 shows the closed contour image overlaid onto the corresponding stereo range image. Note that there are connected components for which there is no range data. There are several geometric criteria that cause a candidate negative obstacle in the negative obstacle image to be removed. First, the connected components for which less than 20% of the pixels contain range data are removed from the negative obstacle image. Second, the connected components that have a length of less than 0.67 meters, or more than 80 meters, are removed from the negative obstacle image. The length of a component is estimated by finding the bounding rectangle in the horizontal plane and calculating the distance along the diagonal.

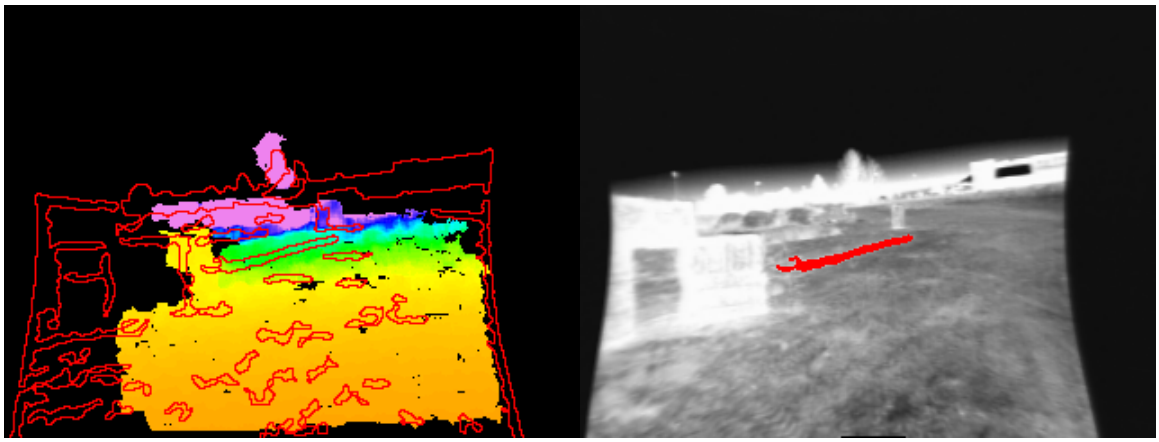

Fig. 8. Edge contour image overlaid on the range image (left) and the negative obstacle image following intensity difference thresholding and geometry thresholding (right).

Third, the connected components which have an average range of 30 meters or greater, are removed from the negative obstacle image. The size of the stereo-based local terrain map used during PerceptOR was 25 meters. Any negative obstacles detected beyond this range were not placed within the local terrain map that populates the larger world map. Fourth, the connected components that have an average width of 0.4 meters or less and a maximum width of 0.45 meters

or less are removed from the negative obstacle image. The width of a connected component is determined by differencing the ground plane range to the near edge and the far edge on a column-by-column basis.

While the connected component width filtering does not produce an accurate measure of trench width for trenches that are oriented parallel to the vehicle, it does rule out classifying as negative obstacles normal and diagonal oriented trenches that have a width less than the diameter of the vehicle tires. Columns for which there is no range data on the leading or trailing edge of a connected component are not used in calculating average or maximum connected component width. Fifth, the connected components that have an average height of 0.4 meters or more are removed from the negative obstacle image. The average height of a connected component is determined by differencing the far edge elevation and the near edge elevation on a column-by-column basis. This allows for the far edge of the connected component to be a mild embankment. This criterion removes positive obstacles, such as walls and buildings, from the negative obstacle image. The right image in Figure 8 shows that after intensity difference thresholding and geometry thresholding, only the detection on the trench remains in the negative obstacle image.

3. MAPPING DETECTED NEGATIVE OBSTACLES

When obstacles are detected in image space, they are generally transferred to a terrain map that can be used by a vehicle level path planner. For each stereo pair of infrared images, we generate a north-oriented instantaneous terrain map. (An instantaneous terrain map is a snapshot of the terrain, generated from a single stereo pair of images.) Instantaneous terrain maps are fused into a single vehicle-centered world map over time, which can be used to plan safe UGV paths. Figure 9 illustrates terrain mapping for the scene in Figure 3. The upper left image in Figure 9 shows positive obstacle detection results (in green) and negative detection results (in yellow) overlaid on the left rectified infrared intensity image. Positive obstacle detection was performed in image space using the algorithm described in [13]. Negative obstacle detection is performed 1) by the new combined thermal/geometry based approach, 2) in the stereo range image by the traditional geometry-based column detector, and 3) in the instantaneous terrain map by an omni-directional detector. Only the omni-directional negative obstacle regions that intersect the negative obstacle regions from the column detector are placed within the instantaneous terrain map [13]. In this particular example, only the new combined thermal/geometry approach detected the trench. The lower left image in Figure 9 is a 320x240 stereo range image for the given scene. The lower right image in Figure 9 is a 25 meter, 20cm resolution 3D instantaneous terrain map from the camera's perspective that contains elevation and obstacle data. The brown cells are free of obstacles. And, the upper right image in Figure 9 is an overhead view of a 50 meter vehicle-centered, north-oriented world map that merges instantaneous terrain maps over time. The gray regions indicate where range data has been placed into the world map and the red rectangle represents the vehicle's location and heading.

In Figure 9, discrete positive and negative obstacles detected in image space were transferred to instantaneous and world maps. These maps label the terrain that contains obstacles as unsafe and the UGV is prohibited from traversing there. Although the maps contain elevation data, they do not directly indicate how "good" traversable terrain is. For example, when a smooth road is right beside rough terrain that does not contain obstacles, clearly, it would be preferable to drive on the smooth road. To capture terrain roughness information, cells in a terrain map can encode a traversability cost. For the scene in the left image of Figure 10, the right image shows a 60 meter overhead view of a world traversability cost map. A traversability cost for each cell was generated by using stereo range data to determine the local slope and height of the terrain. Green terrain has a low traversability cost, red terrain is non-traversable, and intermediate colors between green and red have an intermediate traversability cost. The negative obstacle is assigned the highest traversability cost, i.e., it is not traversable. Based on their height, the map cells that contained hay bales are non traversable. Note, however, although the hay bales are no longer in the sensor's FOV, they persist in the world traversability cost map.

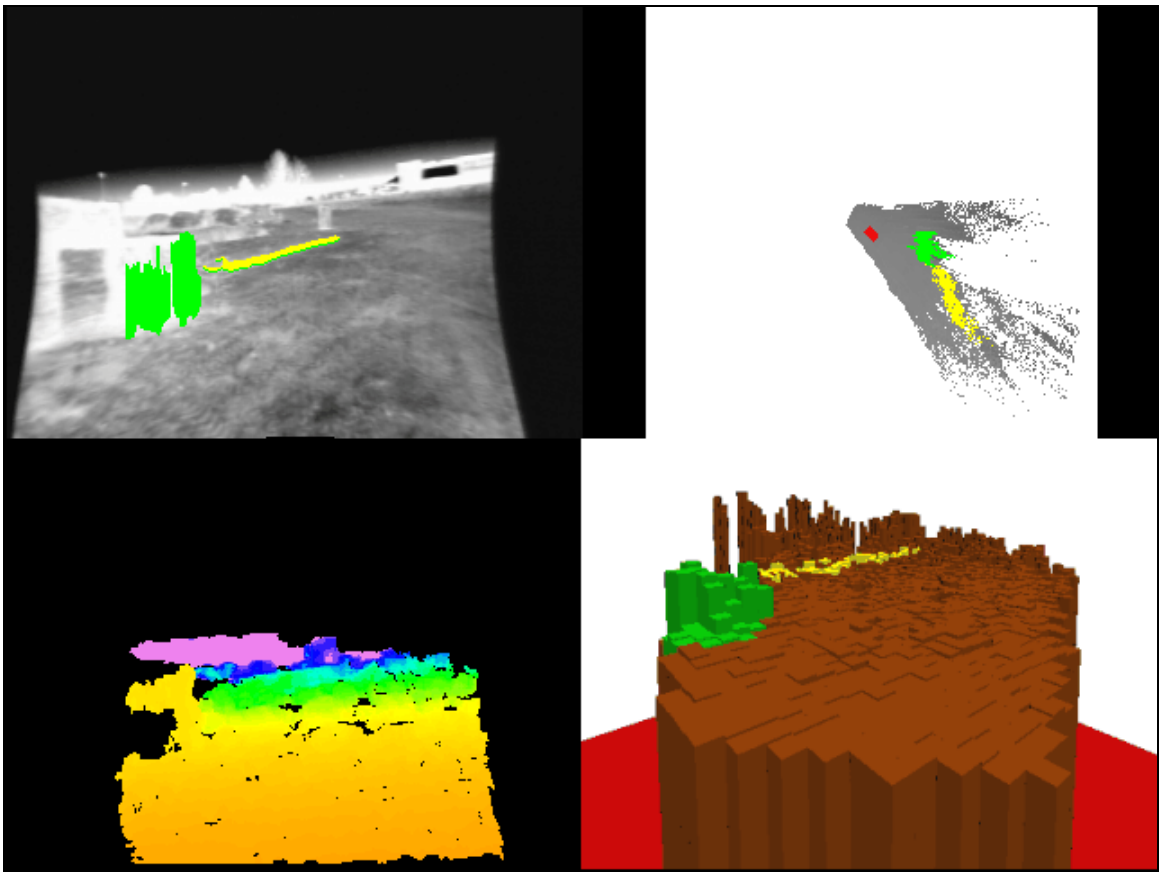

Fig. 9. The upper left image shows positive obstacle detection results (in green) and negative detection results (in yellow) overlaid on the left rectified infrared intensity image. The lower left image is a 320x240 range image for the given scene. The lower right image is a 3D instantaneous terrain map from the camera's perspective that contains elevation and obstacle data. The upper right image is an overhead view of a vehicle-centered, north-oriented world map that merges instantaneous maps over time. The red rectangle indicates the vehicle's current position and orientation.

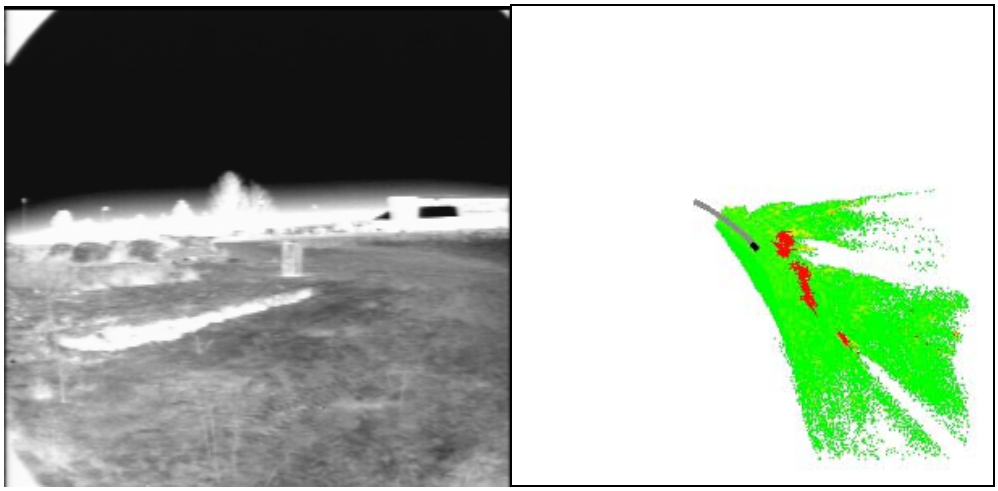

Fig. 10. Overhead view of a 60 meter, vehicle-centered, north-oriented, 20cm resolution world traversability cost map for the final image in the sequence. Here, all cells containing negative obstacle pixels have been labeled as non-traversable (red). Based on their height, the hay bails have also been detected as obstacles. Although the first set of hay bails are no longer in the FOV, that information persists in the world map.

Unless there is false-positive negative obstacle detection, the certainty of negative obstacle detection increases as a UGV nears it and the sensors begin to see inside it. At far ranges, false-positive negative obstacle detection is inevitable. To reduce the adverse effects of false-positive negative obstacle detection on local path planning, a detected negative obstacle could be assigned a traversability cost that is a function of the distance of the perceived negative obstacle from the UGV. Portions of the perceived negative obstacle that is within the stopping distance of the UGV (for its current speed) would receive labeled as non traversable. Figure 11 illustrates this using a 3D rendering of the world traversability cost map. In the left image, all of the negative obstacle pixels have been labeled as non-traversable (red). In the right image, the traverability cost of a cell containing a negative obstacle is a function of the distance from the front of the UGV to the cell.

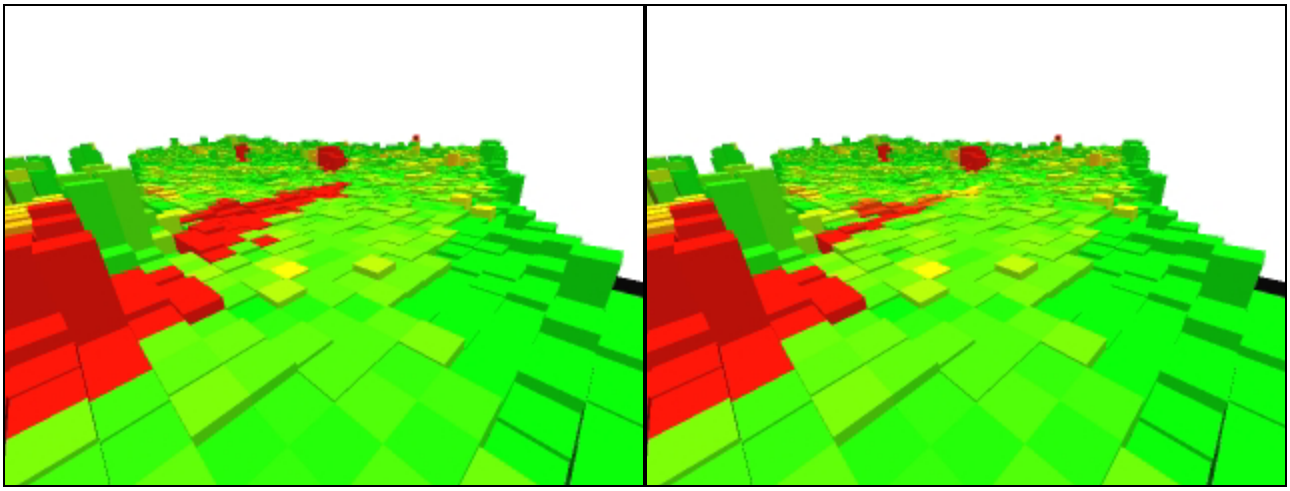

Fig. 11. 3D rendering of the world traversability cost map. In the left image, all of the negative obstacle pixels have been labeled as non-traversable (red). In the right image, the traverability cost of a cell containing a negative obstacle is a function of the distance from the front of the UGV to the cell.

4. RESULTS

During the PerceptOR program, a nighttime negative obstacle detection algorithm was implemented on the Team Raptor UGV that combined cues from stereo range data and thermal signature. The algorithm exploits the knowledge that depression-type negative obstacles tend to be warmer than their surrounding terrain during most of the night. In our original algorithm [10], the thermal signature of negative obstacles was located by independently analyzing each column in a rectified intensity image for large intensity gradients. That algorithm was susceptible to false alarms on horizontal structure warmer than the ground but above the ground surface. The algorithm was subsequently improved by analyzing the thermal signature of closed contours in rectified intensity images, strengthening geometry filtering to verify that candidate negative obstacles are in the ground plane, and by merging negative obstacle detection results into a world traversability cost map that can be used to plan safe paths for the UGV.

The new thermal signature algorithm was tested on a nighttime MWIR stereo data collection sequence of twenty frames which was acquired on the Raptor vehicle during a diagonal approach to a 0.53 meter wide trench at a speed of 1 meter/s. The negative obstacle detection results following intensity and geometry based thresholding are shown in Figure 12. The trench was reliably detected in all the frames and there were no false detections in any of the frames. Half of the trench was three times deeper than the other half. The trench was first detected at a range of 16.8 meters from the forward looking sensors. At this range, although the trench only spans 3 rows in the infrared intensity image, we expect that a UGV traveling up to 24kph could stop prior to the trench (for the nominal terrain conditions specified in the introduction). During the daytime, the same trench was detected at a range of 10.6 meters using a stereo pair of color cameras. The MWIR imagery was processed at a vertical pixel resolution of 4.5 mrad and the color imagery was processed at a vertical pixel resolution of 3.1 mrad.

Low intensity variance has been used as a cue for detecting water bodies [14]. One might expect the well-defined trench used in this study to have a uniform temperature. But in plotting the intensity variance for each closed contour regions that lay on the ground for several frames, there was no clear indication that low intensity variance could be used as a cue for detecting negative obstacles.

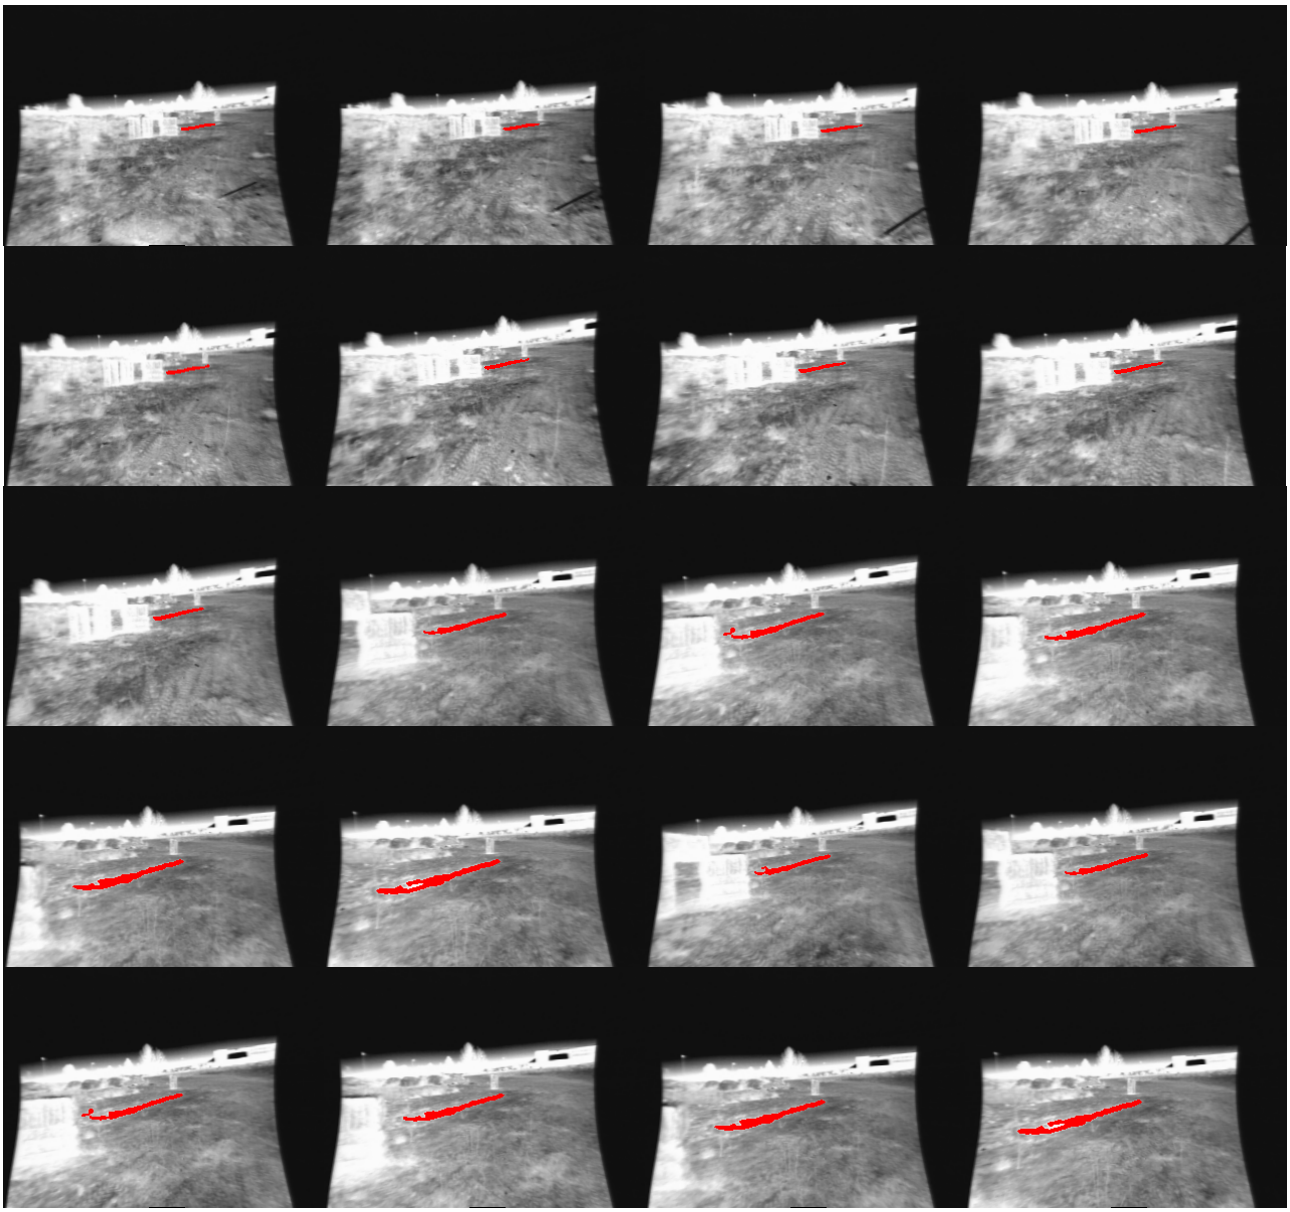

Fig. 12. Negative obstacle detection images overlaid on rectified infrared intensity images following the intensity difference and geometry based thresholdings.

5. CONCLUSIONS

Detecting depression-type negative obstacles remains one of the most difficult problems in perception for unmanned ground vehicle off-road autonomous navigation. A trench slightly wider than the wheel diameter of a UGV can cause vehicle damage at moderate speeds. The width of the visible portion of a narrow trench, however, may only span a few pixels at the stopping distance for vehicle speeds UGV programs aspire to operate at (up to 50kph). One way to improve

negative obstacle detection capability is to use sensors with a small vertical angular pixel resolution. This can be accomplished by using high resolution stereo cameras and lenses that yield a narrow FOV. A narrow FOV, however, could severely limit the swath of terrain that can be verified for higher speed traverse.

Another way to improve the range at which depression-type negative obstacles can be detected is to perform both geometric and thermal analysis. This is particularly valuable during the nighttime since negative obstacles tend to be warmer than their surrounding terrain during most of the night. In the MWIR stereo image sequence used for this study, the 0.53 meter wide trench was reliably detected in all the frames and there were no false detections in any of the frames. The nighttime combined geometry/thermal negative obstacle detector extended the maximum range of detection from 10.6 to 16.8 meters. At the maximum detection range, the trench only spanned 3 rows in the infrared intensity image. Typically, there needs to be at least 7 pixels (the height of the stereo correlation window) on a negative obstacle to be geometrically detected in stereo range data. Further gains in detection range could be realized by using a thermal infrared sensor with a smaller vertical angular pixel resolution. Uncooled thermal infrared cameras are now available having one fourth the vertical angular pixel resolution (1mrad, 30° vertical FOV) as the cooled MWIR sensors used for this study.

By detecting the thermal signature of a negative obstacle at night, the speed at which a UGV can safely operate on cross-country terrain can be increased. But there are limitations to detecting a thermal signature from negative obstacles. Most of our nighttime data collections to date have been performed after a relatively clear day where a trench was heated by the sun for a significant portion of the day. It is unclear if a nighttime negative obstacle thermal signature exists after an overcast or rainy day, when a portion of a negative obstacle is filled with rain or snow, or when the negative obstacle exists on a paved road. More data is needed under these conditions.

Following the development of this algorithm, several areas of future work were identified. Close regions can sometimes be connected in the region image. This can occur when a portion of a negative obstacle in the background is occluded by a positive obstacle in the foreground having similar intensity characteristics. When this happens to a true negative obstacle region, both the connected regions may be filtered from the negative obstacle image. Conversely, both the regions may be labeled as negative obstacles, over estimating the size of the true negative obstacle. Work needs to be done to ensure that distinct regions are separated prior to the intensity difference thresholding stage. In addition, the value used to threshold a set of intensity difference images is arbitrarily chosen. More work is needed to determine if parameters such as this can be auto selected. Finally, additional testing needs to be performed on naturally occurring negative obstacles.

ACKNOWLEDGEMENTS

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by SAIC under the DARPA PerceptOR program, through an agreement with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government or the Jet Propulsion Laboratory, California Institute of Technology.

REFERENCES

1. *"Demo III Experimental Unmanned Vehicle (XUV) Program: Autonomous Mobility Requirements Analysis"*, Technical Report, General Dynamics Robotic Systems, April 1998.
2. T. Chang, T. Hong, S. Legowik, and M. Abrams, "Concealment and obstacle detection for autonomous driving", *Proceedings of the IASTED Conference on Robotics and Applications*, Santa Barbara, CA, October 1999.
3. K. Iagnemma, D. Golda, M. Spenko, and S. Dubowsky, "Experimental study of high-speed rough-terrain mobile robot models for reactive behaviors", *Proceedings of the Eighth International Symposium on Experimental Robotics,* Sant'Angelo d'Ischia, Italy, July 2002, pp. 628-637.

4. M. Rosenblum and B. Gothard, “A high fidelity multi-sensor scene understanding system for autonomous navigation”, *Proceedings of the IEEE Intelligent Vehicle’s Symposium*, Dearborn, MI, October 2000, pp. 637-643.
5. P. Bellutta, R. Manduchi, L. Matthies, K. Owens, and A. Rankin, “Terrain perception for demo III”, *Proceedings of the IEEE Intelligent Vehicle’s Symposium*, Dearborn, MI, October 2000, pp. 326-331.
6. L. Matthies, A. Kelly, T. Litwin, and G. Tharp, “Obstacle detection for unmanned ground vehicles: A progress report”, *7th International Symposium of Robotics Research*, Munich, Germany, 1995.
7. N. Heckman, J. Lalonde, N. Vandapel, and M. Hebert, “Potential negative obstacle detection by occlusion labeling”, *Technical Report CMU-RI-TR-07-07*, Carnegie Mellon University, 2007
8. A. Kelly, et. al., "Toward reliable off road autonomous vehicles operating in challenging environments," The International Journal of Robotics Research, Vol. 25, No. 5-6, 2006.
9. C. Dima, N. Vandapel, and M. Hebert, "Classifier fusion for outdoor obstacle detection", IEEE International Conference on Robotics and Automation, Vol. 1, April 2004, pp. 665 - 671.
10. L. Matthies and A. Rankin, "Negative obstacle detection by thermal signature", *Proceedings of the IEEE Conference on Intelligent Robots and Systems (IROS)*, Las Vegas, October 2003.
11. L. Matthies, T. Litwin, K. Owens, A. Rankin, K. Murphy, et. al., “Performance evaluation of UGV obstacle detection with CCD/FLIR stereo vision and LADAR”, *IEEE Workshop on Perception for Mobile Agents,* Santa Barbara, CA, June 1998.
12. A. Rankin, C. Bergh, S. Goldberg, P. Bellutta, A. Huertas, and L. Matthies, “Passive perception system for day/night autonomous off-road navigation”, *SPIE Defense and Security Symposium: Unmanned Ground Vehicle Technology VI Conference*, Orlando, April 2005, pp. 343-358.
13. A. Rankin, A. Huertas, and L. Matthies, “Evaluation of stereo vision obstacle detection algorithms for off-road autonomous navigation”, *Proceedings of the 32nd AUVSI Symposium on Unmanned Systems*, Baltimore, June 2005.
14. A. Rankin and L. Matthies, "Daytime water detection and localization for unmanned ground vehicle autonomous navigation", *Proceedings of the 25th Army Science Conference*, Orlando, November 2006.