Stereo-Based Tree Traversability Analysis for Autonomous Off-Road Navigation

Andres Huertas, Larry Matthies, Arturo Rankin
Jet Propulsion Laboratory
4800 Oak Grove Drive, Pasadena, California 91109
Andres.Huertas@jpl.nasa.gov

Abstract

Autonomous off-road navigation through forested areas is particularly challenging when there exists a mixture of densely distributed thin and thick trees. To make progress through a dense forest, the robot must decide which trees it can push over and which trees it must circumvent. This paper describes a stereo-based tree traversability algorithm implemented and tested on a robotic vehicle under the U.S. Department of Defense-DARPA PerceptOR program. Edge detection is applied to the left view of the stereo pair to extract long and vertical edge contours. A search step matches anti-parallel line pairs that correspond to the boundaries of individual trees. Stereo ranging is performed and the range data within trunk fragments are averaged. The diameters of each tree is then estimated, based on the average range to the tree, the focal length of the camera, and the distance in pixels between matched contour lines. We use the estimated tree diameters to construct a tree traversability image used in generating a terrain map. In stationary experiments, the average error in estimating the diameter of thirty mature tree trunks (having diameters ranging from 10-65cm and a distance from the cameras ranging from 2.5-30 meters) was less than 5 cm. Tree traversability results from the daytime for short baseline (9cm) and wide baseline (30cm) stereo are presented. Results from nighttime using wide baseline (33.5cm) thermal infrared stereo are also presented.

1. Introduction

Autonomous navigation by off-road mobile robots is an important task for many defense, in-situ planetary exploration, and rescue applications. Robotic vehicles must have the ability to operate efficiently and intelligently with minimal user interaction. Traditionally, traversability decisions are made by determining surface and geometric characteristics from data obtained passively by stereo ranging, or actively by laser ranging. Range data alone, however, is not sufficient to determine whether a detected obstacle is traversable, in spite of its size. Small bushes and thin trees may exceed the vehicle mechanical constraints size-wise, but they could be run over if they were known to be soft obstacles. We are pursuing added levels of robot intelligence by also exploiting the geometric aspects that some natural objects have, such as piecewise linear sections, direction of growth, and appearance properties. In this paper, we focus on estimating the diameters of trees in the scene, to help determine whether or not they are traversable obstacles. This process relies on edge contours extracted from the left image of the stereo pair. The edge contours are coded with their direction to limit processing to vertical (or near-vertical) object boundaries, and to match the edges on opposite boundaries of tree trunks. The matching process produces fragments that represent portions of tree trunks and other similar objects. Straight linear segments having a location, a width, a length, and an orientation in 2-D image coordinates represent these fragments. The range data provided by stereo ranging, projected onto the frame of the left camera, provide the 3-D information needed to assist the matching process and derive the range to the trees, and to estimate the diameters. An assessment of the quality of these estimates is given also, as well as examples from visible color image pairs from several diverse scenarios, and from mid-wave infrared image pairs acquired during the night.

This work is designed to estimate the diameters of trees that belong to a large class of trees having near-vertical trunks. We assume that portions of the main trunk are discernable from the background, i.e. the boundaries that delineate portions of the trunk are detectable. We call these portions, trunk *fragments*. The diameters of portions of tree trunks are assumed to vary little along the fragment. The system uses edge contrast polarity to match tree boundaries and detect fragments. Therefore, it assumes that portions of the trunk appear either brighter, or darker, than the background, and thus the boundaries have opposite contrast. The backgrounds themselves may change along the trunk of a tree, however. For example, the lower portion may be set against the ground surface while the middle section may be set against background higher vegetation, mountains, and water bodies. The upper portions may be set against tree canopies and the sky. These fluctuations do not affect fragment detection, as long as the fragments have boundaries of opposite contrast.

Different types of trees will have different appearances depending on the texture of the bark, the smoothness of the exposed core, and the density of the branches and canopy. During the night, IR emissions from tree trunks are stable and appear brighter that the surrounding material. We exploit this thermal property as well. During the day moderate to dense canopies reduce the amount of direct sun illumination, resulting in a homogeneous

appearance of the trunks. Direct illumination from certain angles in combination with the viewing angles can result in multiple fragments on a single trunk, making the diameter estimation difficult.

Our current implementation runs under real-time constraints and thus does not attempt to model the trees as complete objects. Rather, we extract fragments on each unobstructed trunk, even if they correspond to trees in the distance where range data may not be available. Of particular importance are trees that are close to the vehicle, e.g., within 15 meters, where quick reaction maneuvers are needed. In particular, estimating the diameter for the portions of the trees that are below vehicle height are crucial.

2. Related Work

Autonomous navigation has been an active area of research by space and defense agencies for several years. Ground-based and low-flying robots must have the ability to operate autonomously with minimal human assistance. The vehicle's navigation systems must be able to traverse rough, poorly modeled natural terrain, avoiding hazardous situations. Some systems employ fuzzy logic methods to mimic human behavior using computer vision-based measures of terrain roughness, slope and discontinuities [1,2]. Traversability is usually defined within the context of path planning [3]. Once the obstacles have been detected, the distance of each location in a computed terrain map to the obstacles serves as a measure of traversability for that location. The map cells typically include a measure of terrain type computed by a number of terrain classification methods [4]. Ultimately these path-planning tasks must translate into maneuvering reactions that are safe [5]. The work presented here deals with methods to derive a more accurate detection and description of the obstacles in the scene. Trees in forested scenes in particular, have geometric and appearance aspects that can be exploited to better populate the world map used to make traversability decisions. Currently, we are not aware of other systems that attempt to detect trees specifically from intensity images acquired by unmanned ground vehicles (UGV) under near real-time requirements. A real-time technique that uses active sensors, such as laser and radar, to determine traversability in terms of object density is reported in [6].

3. Approach

The tree detection module attempts to detect the location and diameter of standing trees in the scene by explicitly detecting portions of tree trunks (Fig. 1). Trunks that appear darker or brighter than their background will have boundaries of opposite contrast. Matching pairs of these edges in conjunction with stereo range data, results in a set of individual potential tree trunk fragments. Matching is performed only in the horizontal direction. With support from stereo ranging, the distance to the trees and the diameter of the trees can be estimated. Trees can then be classified as potential hazard depending on the robotic vehicle characteristics.

We have developed, integrated, and field-tested a tree hazard detection module into a real-time passive perception system that runs on a UGV. The module focuses on tree fragments that can be detected fast, rather than in computing a tree model with all its fragments integrated into a single volume. This, however, would require a grouping operation to produce an extended map or a tree survey, for example.

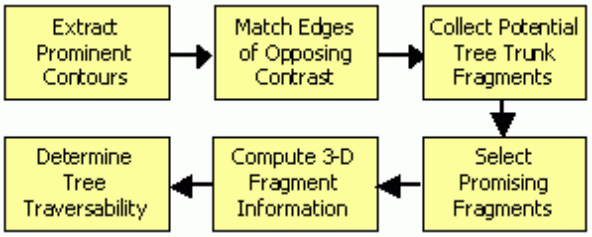

Figure 1. Steps in detection of tree trunk fragments.

Edge Detection applies a 1-D horizontal edge detector to the rectified left intensity image (Fig. 2) of a stereo pair. The intensity images are 320x240 pixels, coded to 8 bits. For expediency, we apply a 1-D edge detector followed by a 1-D non-maximal suppression step to thin the convolution output and localize the edges. The edge detector uses a first derivative Gaussian kernel having a space constant sigma=1. The edge gradient direction is used to determine the direction of the edges. We select edges that belong in contours that have vertical or near-vertical directions.

Contour extraction eliminates single pixels, bridges one-pixel gaps along the contours, and filters out weaker edges and short contours. In the example illustrated in Fig. 2, edge contours have 10 pixels or more. The color-coded contours illustrated in Fig.2 show the direction of contrast.

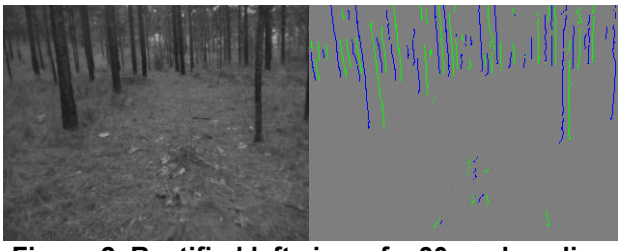

Figure 2. Rectified left view of a 30 cm baseline stereo pair and the contours extracted. Green and blue contours indicate contrast polarity.

Contour Matching looks for and matches pairs of edges of opposing contrast along the horizontal direction.

The search process starts at an edge and searches up to a maximum distance for an edge having opposite contrast. Edges having the same contrast are skipped if the underlying range data does not indicate a range discontinuity there. As a result, bark texture does not impede finding the matching tree boundary. Trees may overlap in 2-D and may be merged into a single fragment. Thin trees in the distance may be too thin to be resolved by the edge detector. These effects, however, are typically resolved as the vehicle approaches the trees to a distance where the ambiguity is reduced. The search process has been designed to incorporate the use of cues derived from other sources to help validate the matches between edges. In the example illustrated in Fig. 2, the cue comes from the intensity images themselves, by simply indicating that the tree trunks appear darker than the background. The directions associated with the edges indicate on which side of the edge the tree is. Regions having intensities darker than the mean intensity of the image, therefore, are likely to include the dark tree trunks. A similar situation occurs in nighttime infrared images, where the trees appear brighter than the background. As mentioned earlier, we also use the disparity data to verify depth discontinuities at the matched tree boundaries. Other cues can be easily incorporated, if available, such as pixel classifications of multi- or hyperspectral data. The mean intensity of the image illustrated in Fig. 2 above is 80 (out of 255), and we look for tree trunks having a projected 2-D diameter up to 30 pixels. The matching process results in the binary mask, shown in Fig.3a for our example. The same mask is overlaid in red on the input image, as shown in Fig. 3b.

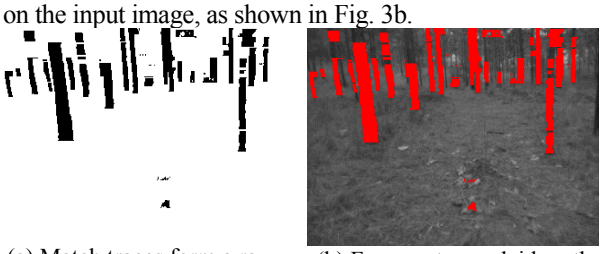

(a) Match traces form a raw fragments mask

(b) Fragments overlaid on the intensity image

Figure 3. Building the fragments mask.

Fragment selection traces the boundaries of the individual fragments and generates best-fitting ellipses, shown in red in Fig. 4a, to the individual fragments. This provides a representation that encodes the fragments shape, position, and orientation. A best -fitting ellipse [7] equates the second order central moments of the ellipse to those of the distribution of the pixels in each fragment, and thereby effectively defines both the shape and size of the ellipse. Currently, we require that these fragments represent "elongated" shapes having an aspect ratio better than 3:1, and that their orientation (angle of major axis) be consistent with a near-vertical pose (90 degrees with a 22 degree tolerance.) We allow smaller aspect ratios for large fragments that are close to the vertical direction. The ellipse parameters are used to approximate the tree fragments by parallel straight linear segments that have the same length, width and orientation than the ellipses. These are illustrated in Fig. 4b. The labels on the selected fragments correspond to the 2-D description shown in Table 1.

In general, we can expect multiple fragments on a single trunk due to branches or a non-linearly changing background. This occurs less along portions of the trees that are near the ground where branches are not usually found. The desirable result, however, is that at least one fragment per tree be found.

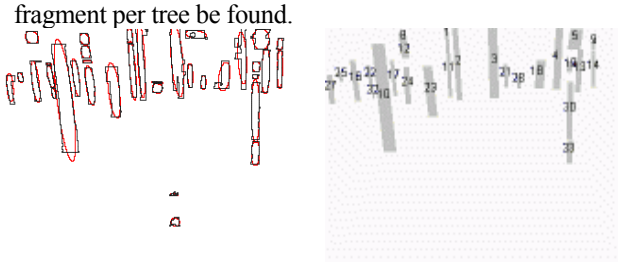

(a) Ellipse fit

(b) Tree trunk segments

Figure 4. Tree trunk fragment representation.

Frag No.	Area	Perimeter	Major Axis	Minor Axis	Angle
1	79	39.89	16.83	5.97	94.10
2	448	176.52	75.71	7.53	92.84
3	713	197.59	72.63	12.49	94.90
4	380	168.87	53.06	9.11	92.82
5	294	82.28	21.83	17.14	87.29
8	98	37.89	12.50	9.97	106.67
9	61	33.89	12.62	6.150	80.82
10	1640	263.49	125.26	16.66	97.27
11	296	121.21	60.37	6.24	94.43
12	106	41.31	13.64	9.88	104.08
13	315	101.55	47.09	8.51	87.98
14	179	84.14	41.66	5.47	88.92
16	252	109.45	43.95	7.30	97.54
17	205	84.38	39.96	6.53	98.05
18	263	90.97	29.48	11.35	79.45
19	75	33.07	11.30	8.45	80.19
21	111	53.07	25.51	5.54	91.00
22	108	53.55	24.10	5.70	101.59
23	526	129.21	56.78	11.79	96.84
24	280	89.21	40.72	8.75	97.02
25	44	26.48	10.32	5.42	103.00
27	160	75.799	29.569	6.889	93.858
28	60	33.071	14.424	5.296	86.496
30	440	134.627	64.740	8.653	89.777
31	85	35.314	11.242	9.627	28.001
32	38	25.899	11.238	4.305	101.402
33	173	62.142	25.808	8.535	87.798

Table 1. 2-D description of selected fragments.

Diameter Estimation estimates the diameter, *D*, of the tree fragments using the range data computed by stereo raging. By similar triangles: $D = \frac{dR}{f}$, where *d* is the 2-D projection of the diameter on the image, *f* is the focal length of the camera lens, and *R* is the distance, to the tree. Additional details are given below in Section 6. Fig. 5a shows the range image from stereo ranging.. Fig. 5b shows the range data overlapping the tree fragments shown above in Fig. 5b. We arbitrarily require that at least 5 % of the area of the fragment have range information available. To determine the range to the tree (*R* in the equation above), we compute the average range within the linear region representing the trunk fragment. Specifically, we do not include range data along the borders of these regions to minimize contributions near the boundaries of the trees, where the range values may incorporate background range values. These averages are illustrated pictorially in Fig. 5c, where blue colors indicate trees that are closer. Yellows and white indicate trees farthest away. Note that fragments shown in Fig. 4b that do not have sufficient range data available (Fig. 5b), are not represented in Fig. 5c. Fig. 5d shows pictorially an image of the fragments with the pixel values encoding the estimated diameters. Darker grays indicate thinner trees.

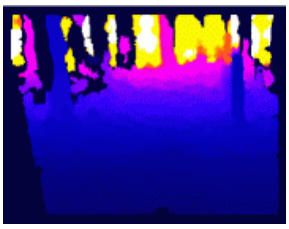

(a) Range image. Blues are closer. White is farthest.

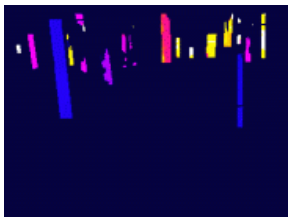

(b) Range data available for each fragment.

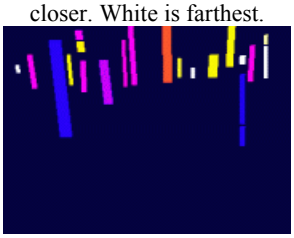

(c) Averaged range to the trees. Blue is closer, white is farthest.

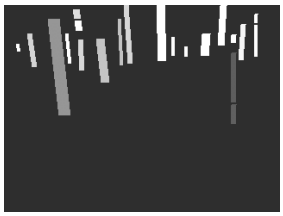

(d) Linear segments. Intensity encodes diameter. Darker values represent thinner trees.

Figure 5. Estimation of tree diameters.

4. Tree Hazard Map

Three-D estimates of tree ranges and diameters can be used to make traversability decisions as a function of the vehicle characteristics and speed of travel. Conservative assessments based on diameter alone can be made until, say, spectral information becomes available to provide information about the *type* of tree, and measures of anticipated bark and core densities can be incorporated into the traversability decisions. Fig. 6 illustrates the tree fragments as color-coded markers derived from the tree diameter estimates. Using arbitrary limits, in this example, tree fragments larger than 15 cm are coded red, or non-traversable. Fragments coded yellow are borderline fragments having a diameter between 10 and 15 cm. Fragments having diameters less than 10 cm are coded green, or traversable (see next example below.)

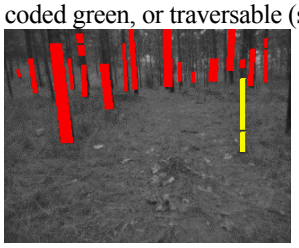

(a) Red is non-traversable. Yellow may be traversable

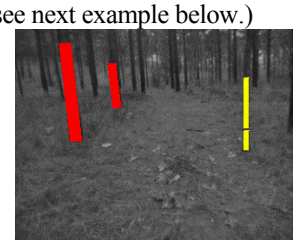

(b) Nearby trees, within 12 meters

Figure 6. Traversability markers.

The tree hazard detection module reports range to the trees and an estimate of tree diameters. This information is used to localize the trees in world coordinates, and to evaluate the hazard level as a function of vehicle characteristics. An overhead view of the north-oriented world map is shown in Fig. 7. The circles denote 10-meter intervals, and the non-traversable trees detected within 50 meters are illustrated by red marks. The gray areas indicate range data.

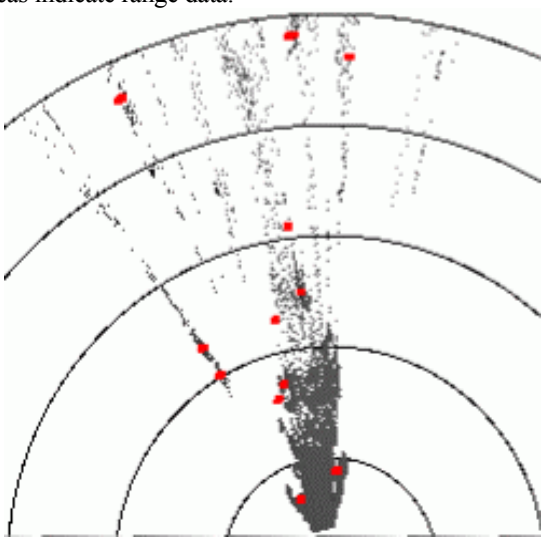

Figure 7. Overhead view of north-oriented world map with tree markers.

The example illustrated in Figures 2-7 is from a sequence of hundreds of frames from Ft. Polk, VA. Fig. 8 shows results for a few additional frames including thin and thick trees, nearby trees and trees in the distance. Note the absence of false alarms.

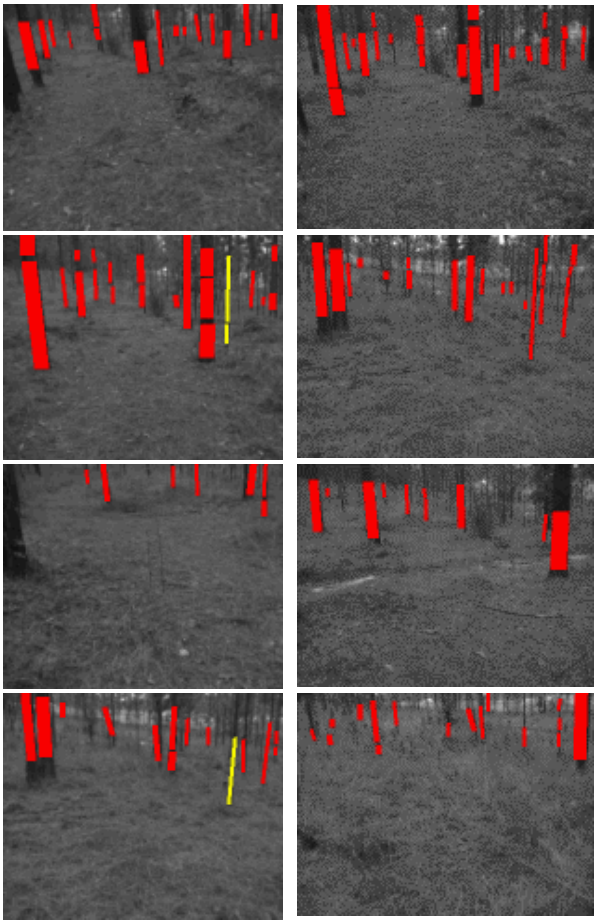

Figure 8. Result for few additional frames.

5. Nighttime Mid-Wave FLIR Example

We have tested our algorithm on mid-wave infrared wide baseline (33.5cm) stereo pairs acquired during the night. Fig. 9 shows one example. The top left portion of the image has saturation problems, but sufficient information is available to extract tree fragments and to estimate their diameters. The edges extracted from the image are illustrated in Fig. 9b. The thin stocks from the tall grass in the foreground, however, are close, and their boundaries can be resolved well. The directed linked edge contours extracted are shown in Fig. 9c. Matching the corresponding edges along these contours uses the thermal cue that trees appear bright at night. In this example, regions having intensity higher than the mean intensity help verify the presence of a tree, or other kinds of vegetation. The raw fragments are shown in Fig. 9d overlaid on the input image. The selected fragments represented by linear segments are shown in Fig. 9f. Note that the thin tall grass stock in the foreground is well represented. Fig. 9e shows the range data from stereo ranging. The selected fragments and their linear segment approximation are shown in Fig. 9f. The range data lacks density but some 3-D information is available at the tree locations (Fig. 9g.) The color-coded traversability assessment is illustrated in Fig. 9h. The diameter of the large tree on the right is overestimated because a portion of the ground vegetation appears as a smooth continuation from the bottom of the tree trunk. The current analysis does not look specifically for sudden changes in width at the bottom of the tree, as many trees may have large root systems that may be visible. The amount of overestimation, however, is not large, and tends to occur at the bottom of large trees that typically are not traversable.

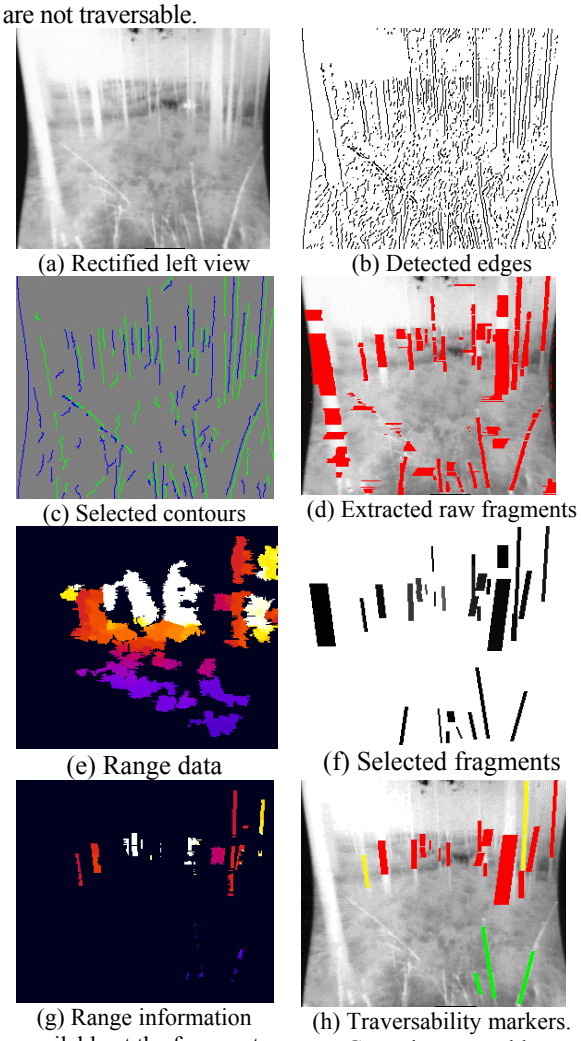

(a) Rectified left view (b) Detected edges

(c) Selected contours (d) Extracted raw fragments

(e) Range data (f) Selected fragments

(g) Range information available at the fragments (h) Traversability markers. Green is traversable

Figure 9. Nighttime FLIR example.

6. Tree Diameter Accuracy Assessment

To determine the accuracy of the tree diameter estimates, we measured the circumference of thirty trees in a forested area. Ten of the 30 measured trees are shown in Fig. 10. The range data were obtained using a narrow baseline (9.5 cm) stereo ranging.

Figure 10. Ten of the 30 trees measured.

Fig. 11 shows results for one of the scenes. The scene is cluttered and helps demonstrate the feasibility of detecting useful tree trunk fragments from this type of scene. Fig. 11b shows the range data, Fig. 11c the detected fragments, and Fig. 11d, the approximated selected fragments.

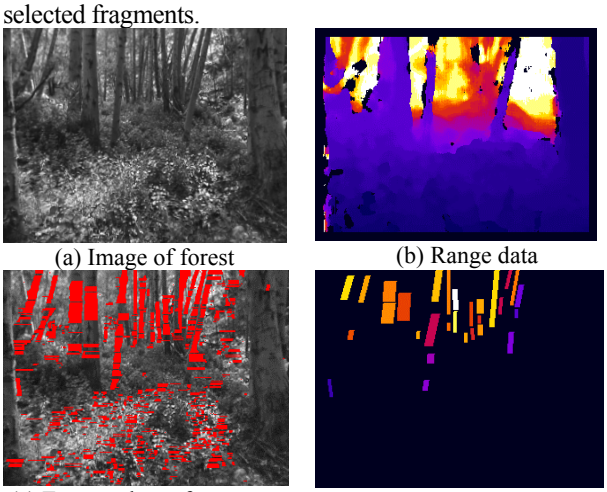

(a) Image of forest (b) Range data

(c) Extracted raw fragments (d) Tree trunk fragments

Figure 11. Example of cluttered forest scene.

The plot in Fig. 12 illustrates the measured and estimated tree diameters for all thirty samples. The horizontal axis lists the thirty trees, from left to right, sorted by range. The same plot illustrating relative measurements is shown in Fig. 13; it illustrates one outlier and a relatively small error at ranges out to 30 m.

The expected error in the tree diameter estimates is given as a function of the error in the range to the tree. This range error, is given by:

$$\sigma_{range} = \frac{R^2 kF}{B},$$ where R = range, k = *0.2* (subpixel precision), F=*0.3* mrad (IVFOV), and B=*9.5* cm (baseline.)

Fig. 14 illustrates the geometric relationships useful to derive the expected error in the diameter estimates as a function of range error. Fig. 15 illustrates a plot of the expected and actual error of tree diameter errors for the thirty trees measured.

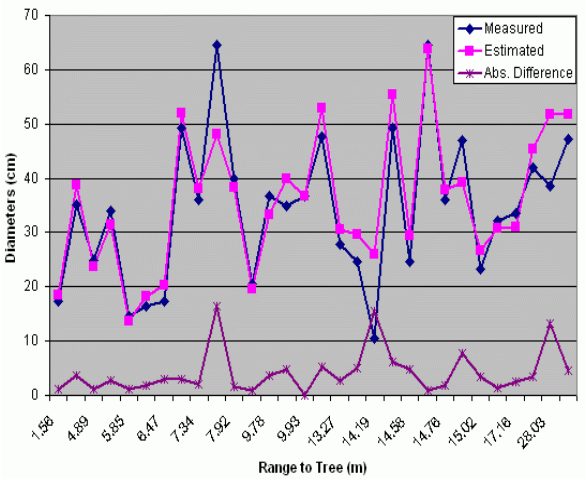

Figure 12. Absolute measurements and differences.

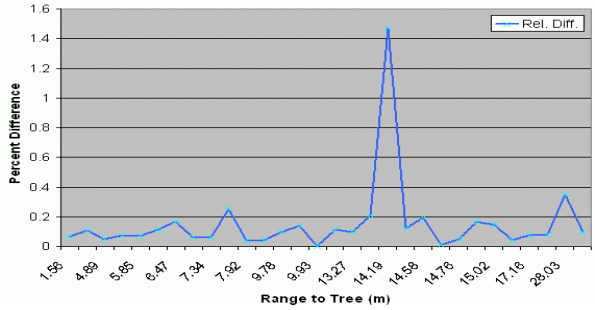

Figure 13. Relative difference as a percent of measured to estimated diameters.

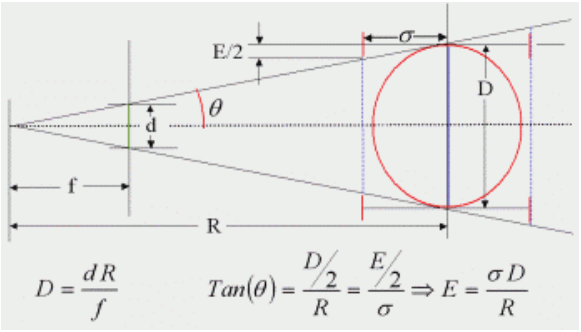

d = projection of diameter onto image
R = range to tree center
f = focal length of camera lens
D = diameter of tree

Figure 14. Expected error in diameter estimation as a function of range error.

7. More Results

At another location, the scenes consist of wooded rolling hills having many trees left over from a forest fire. The trees have more visible branches than those

illustrated earlier in Section IV. The tree model, however, remains the same. The majority of the trees appear darker than the background and reasonable edge contours, however fragmented, can be extracted from the images.

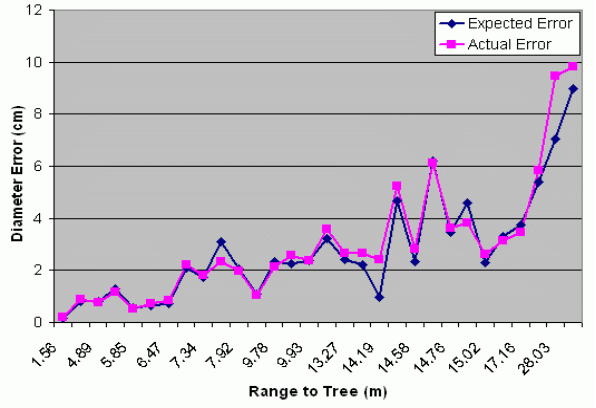

Figure 15. Expected and actual error in estimates for the thirty trees measured.

Fig. 16 illustrates a number of representative frames from a dataset collected during the fall season, and the results from the tree hazard detection module. In the left column, we show the left image from the stereo pair. The column on the right shows the detected tree trunks with the trunk fragments approximated by a color marker. The colors denote the level of hazard by comparing the computed tree diameters with an arbitrary hazard table. Red denotes a severe hazard, yellow denotes moderate hazard, green denotes a mild hazard, and cyan denotes detection where range data is not available to estimate the diameter, but the tree fragment geometry is satisfied.

Note that thin and thick trees can be detected reliably whether they are near or far away. The fragmentation of the boundaries extracted depends on the background and on the branching of the trees. A number of parameters can be adjusted to compensate for scene specific conditions. In these results, however, we used the same parameter settings used to illustrate results presented above in Section IV. Note also that this technique detects only standing trees.

A third dataset was acquired during the winter with snow on the ground and in overcast conditions. For this type of scene, the contrast of the edges along the boundaries of the trees is low and needs to be adjusted to allow significant contours to be available. Fig. 17 shows results for several frames. In the left column, there are a number of images (left view) of the scene. In the right column, we show the detected tree fragments overlaid on the images. Recall that cyan denotes a fragment detection without range data available for diameter estimation.

Note that we can expect that the lower portions of nearby trees have a high contrast against the snow on the ground, and thus improve the possibility of detection. These frames include representative conditions where some of the trees are close by and some trees are located at some distance. They also have different widths. Some frames include standing water that reflects the objects in the background. Note also that vehicle tracks on the snow are also elongated features that are thin and have boundaries of opposing contrast, thus satisfying our generic tree model. The availability of range data at these locations indicates that the stereo disparities of the features are similar to the disparities of the background, thus allowing correct disambiguation.

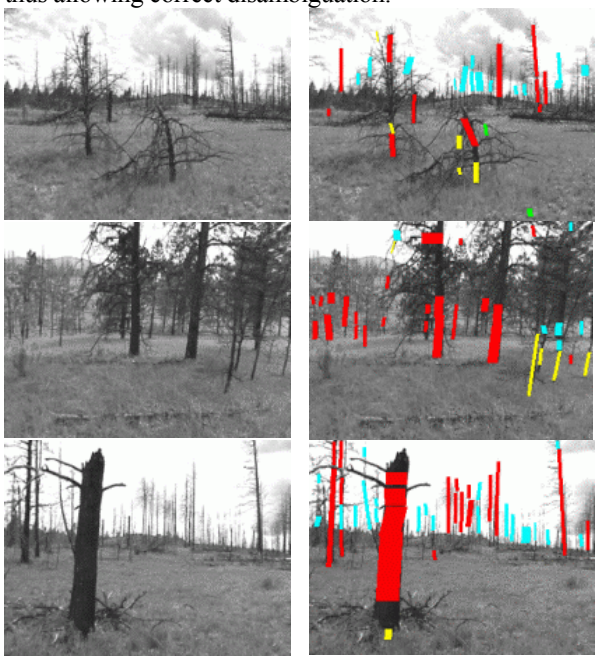

Figure 16. Tree detection and hazard determination in fall dataset.

8. Future Work

Scenes that contain some trees that appear brighter than the background and some trees that appear darker than the background require two passes and the results are simply combined. Under some illumination conditions, however, part of the same tree trunk appears bright and part appears dark. Fig. 18 shows two examples. The figure illustrates the raw "bright" and "dark" fragments that satisfy the geometric requirements for tree fragments. To verify tree trunks, we look at the range data available at the location of the fragment. The intensity discontinuities that gave raise to the trunk fragment must have corresponding depth discontinuities near by. The region "inside" the fragment must contain smooth consistent disparities, while the regions "outside" and on both sides of the fragment must have lower disparities. The fragments illustrated in Figure 18 cannot be verified by this criterion, and additional work is needed to make the verification step more general without incurring unwanted false alarms. Therefore, only tree fragments that appear wholly darker or brighter that the background can be verified by the current algorithm. A process that can detect these conditions, determine that

two passes are needed, and combine and verify these fragments is under development.

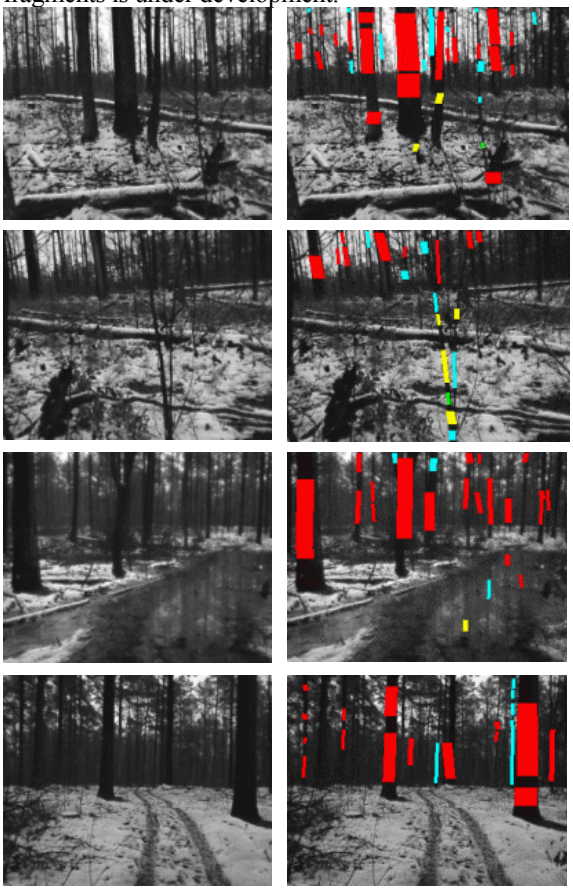

Figure 17. Tree detection and hazard determination in winter dataset.

9. Conclusion

A number of classes of obstacles in natural scenes can be described as traversable by a vehicle of certain given physical characteristics and speed of travel. This requires knowledge about the objects that can be measured by the sensors available. One such class of objects is thin trees, bushes and other low vegetation that exhibit salient geometric properties. Tree trunks in particular can be expected to stand vertically (or near vertically) and to have discernable boundaries against the background. The combination of tree types (and therefore appearance) and backgrounds, nevertheless, can make the task of detecting a full tree difficult. Portions of the trunks of trees however are in general discernable, as we have illustrated. The ability to detect and estimate the diameters of trees depends on the ability of the edge detector to resolve the tree boundaries, and the ability of the stereo algorithm to produce range information.

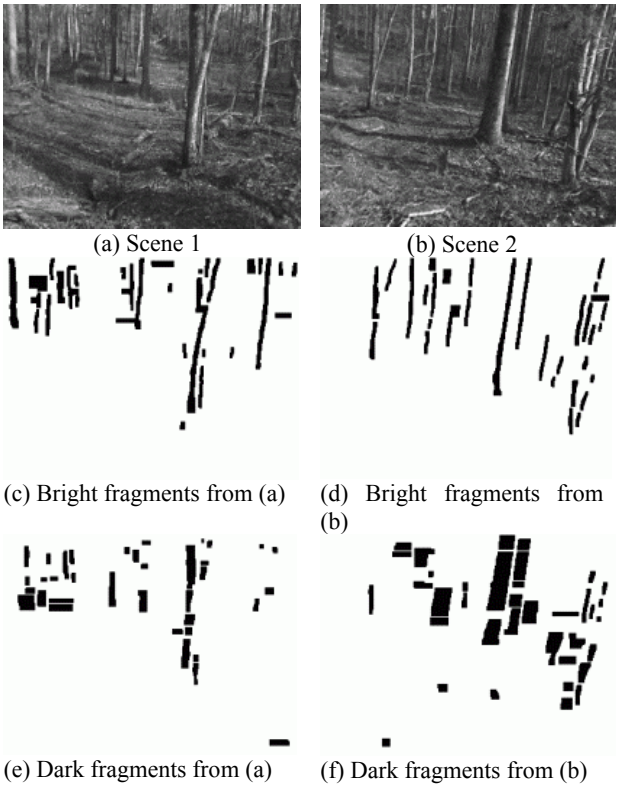

(a) Scene 1 (b) Scene 2

(c) Bright fragments from (a) (d) Bright fragments from (b)

(e) Dark fragments from (a) (f) Dark fragments from (b)

Figure 18. Mixed appearance trees.

10. References

[1] A. Howard, H. Seraji and E. Tunstel, "A Rule-Based Traversability Index for Mobile Robot Navigation," *IEEE Intl. Conf. on Robotics and Automation*, Seoul, Korea, May, 2001.

[2] A. Howard, E. Tunsel, D. Edwards and A. Carlson, "Enhancing Fuzzy Robot Navigation Systems by Mimicking Human Visual Perception of Natural Terrain Traversability," *Joint 9th IFSA World Congress and 20th NAFIPS International Conference*, Vancouver, B.C., Canada, July 2001, pp. 7-12

[3] P. Batavia and I. Nourbakhsh, "Path Planning for the Cye Personal Robot," *Proceedings of Intelligent Robots and Systems (IROS)*, Nov. 2000, pp 15-20.

[4] J. Gancet and S. Lacroix. "PG2P: A perception-guided path planning approach for long range autonomous navigation in unkown natural environments." *Proc. International Conference on Intelligent Robotics and Systems*. Las Vegas, NV, 2003.

[5] 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 8th International Symposium on Experimental Robotics (ISER)*, Italy, 2002.

[6] M. Ollis and T. Jochem, "Structural Method for Obstacle Detection and Terrain Classification," *Proceedings of the SPIE*, Vol 5083, Unmanned Vehicle Technology, Orlando, FL, Sept. 2003, pp 1-12.

[7] H. Cramer, *Mathematical Methods of Statistics*, Princeton Landmarks in Mathematics and Physics, Princeton University Press. Princeton, NJ, 1999.