

Obstacle Detection in Foliage with Ladar and Radar

Larry Matthies¹, Chuck Bergh¹, Andres Castano¹, Jose Macedo², and Roberto Manduchi³

¹ Jet Propulsion Laboratory, California Institute of Technology, 4800 Oak Grove Drive, Pasadena, CA, USA 91109

² Industrial and Manufacturing Engineering Department, California Polytechnic State University, San Luis Obispo, CA, USA 93407

³ Computer Engineering Department, Baskin Engineering Rm 237, University of California, 1156 High Street, Santa Cruz, CA, USA 95064

Abstract. Autonomous off-road navigation is central to several important applications of unmanned ground vehicles. This requires the ability to detect obstacles in vegetation. We examine the prospects for doing so with scanning ladar and with a linear array of 2.2 GHz micro-impulse radar transceivers. For ladar, we summarize our work to date on algorithms for detecting obstacles in tall grass with single-axis ladar, then present a simple probabilistic model of the distance into tall grass that ladar-based obstacle detection is possible. This model indicates that the ladar “penetration depth” can range from on the order of 10 cm to several meters, depending on the plant type. We also present an experimental investigation of mixed pixel phenomenology for a time-of-flight, SICK ladar and discuss briefly how this bears on the problem. For radar, we show results of applying an existing algorithm for multi-frequency diffraction tomography to a set of 45 scans taken with one sensor translating laterally 4 cm/scan to mimic a linear array of transceivers. This produces a high resolution, 2-D map of scattering surfaces in front of the array and clearly reveals a large tree trunk behind over 2.5 m of thick foliage. Both types of sensor warrant further development and exploitation for this problem.

1 Introduction

Autonomous off-road navigation is important in applications of unmanned ground vehicles in agriculture, defense, and planetary exploration. The ability to operate in all weather, all lighting conditions, and all terrain is a key goal in many of these applications. This goal requires exploiting a broader range of sensors than has been used in robotic vehicles in the past and developing a deeper understanding of the phenomenology associated with those sensors and environmental conditions. In this paper, we focus on perception for intelligent navigation in vegetated terrain, which is currently one of the most prominent problems in off-road navigation. For recent work on new sensor approaches to other aspects of complex terrain, such as water hazards and negative obstacles, see [1], [2], [3].

The problem with vegetation is two-fold: (1) vegetation can hide real obstacles and (2) soft vegetation can appear to range sensors as rough ground or as obstacles, both of which will needlessly slow the robot if it thinks it must avoid the vegetation or slow down to cross it. To solve the first problem, robots need sensors that penetrate some distance through vegetation; to solve the second, robots must recognize vegetation as such and reason about its traversability.

To our knowledge, these problems were not addressed to any significant degree in the robotics community prior to the last three years or so. Before then, there had been much work in computer vision and some work in robotics on image and terrain classification with color (RGB) imagery [4]. The remote sensing community has been using multispectral imagery extending beyond the visible spectrum for terrain classification for many years [5]; initial work on applying some of these techniques to robot navigation, using visible and near infrared imagery, were reported in [6]. The remote sensing community has also used multi-return lidar to do “vegetation removal” for terrain mapping from overhead by using the last return for each pulse to measure the ground surface, assuming that earlier returns come from tree canopy [7]. In [8], we introduced a similar concept for robot navigation with single-return lidars, whereby the spatial distribution of range data was analyzed to determine whether it was measuring a smooth surface likely to be an obstacle or a highly dispersed surface likely to be vegetation. We observed that lidar beams will “penetrate” some distance into vegetation, because the narrow beams pass between foliage elements, and that the “penetration depth”, to borrow a radar term, can be modelled probabilistically as a function of the size and frequency of the foliage elements. Similar concepts were later embodied in work at NIST with a two-axis scanning lidar and voxel-based terrain mapping algorithms in the Demo III program [9]. Several approaches to discriminating obstacles and vegetation with lidar are compared in [10], which also includes an extensive review of related literature.

Obviously, optical techniques like lidar and multispectral imaging have inherently limited range for obstacle detection in vegetation, because vegetation is opaque at these wavelengths. Radar wavelengths are more able to propagate through or diffract around vegetation. Radar has been studied extensively for detecting targets hidden in trees [11] and for remote sensing of crops and forests [12], but we have not seen published work on foliage penetrating radar for robot navigation. Thus, there is little direct experience on appropriate radar wavelengths, sensor designs, resolutions, or penetration depths for robot navigation.

There has been very little work on reasoning about the compliance of vegetation to intelligently modulate speed over it. A first attack is given in [13].

In this paper, we continue the study of lidar phenomenology for obstacle detection in vegetation by examining lidar “penetration depth” into foliage and by examining mixed pixel effects. We also present initial results on the adaptation of published radar algorithms for “multi-frequency diffraction tomography” [14] to the problem of detecting obstacles through vegetation.

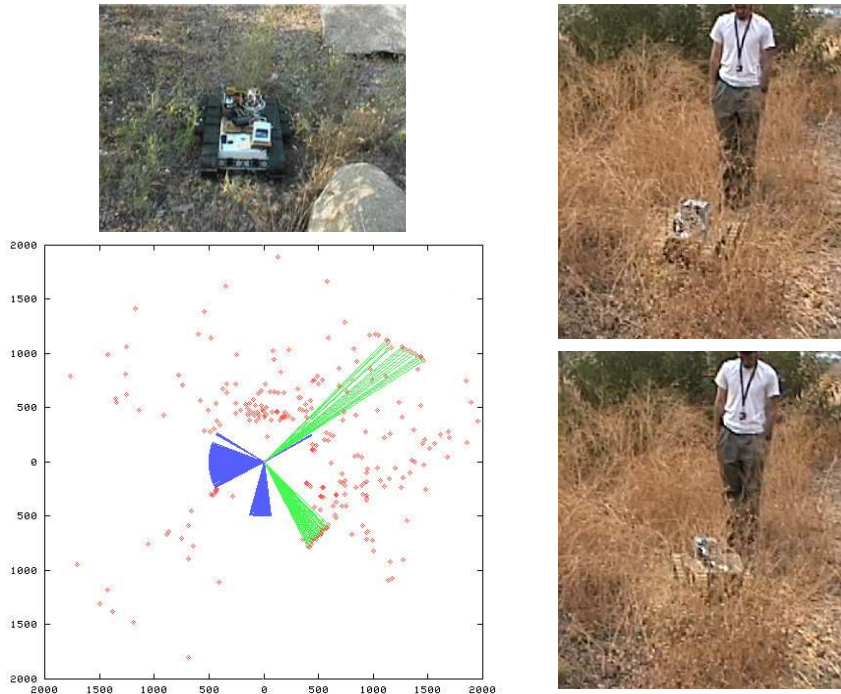


Fig. 1. Upper left: urban robot carrying a 360-degree, single-axis scanning ladar. Lower left: range data from one scan of the ladar for scene in upper left. Red dots are range measurements, blue cones mark structure on the robot that obstructs the beam, and green lines show detected obstacles (ie. the two rocks). Right column: two frames from a movie of obstacle avoidance in dense, tall grass, in this case using a 180-degree SICK ladar.

2 Ladar

Figure 1 illustrates the situation we address and shows results we have demonstrated for ladar-based obstacle avoidance in dense vegetation. We have developed several algorithms for ladar-based obstacle detection in vegetation. The most recent [15] essentially uses a window of three pixels in single-axis scan ladar data to estimate the second derivative of the range data. Where this is high, the central pixel is probably a blade of grass; where it is low, the central pixel may be an obstacle. Further stages of the algorithm attempt to suppress false alarms and to find the full extent of obstacles by reasoning about nearby pixels that are labelled as obstacles. As shown in Figure 1, this algorithm is quite effective at recognizing grass and can detect obstacles a short distance into grass. The question now is: how far into grass can obstacles be detected with ladar?

We address this question here with a simple statistical model that provides sufficient insight for our current engineering purposes. We model the foliage as plant stems with a constant diameter d that occur according to a Poisson

Table 1. Plant density and stem diameter data for established crops, drawn from the crop insurance industry.

Crop	Density (m^{-2})	Diameter (mm)	λd (m^{-1})
Wheat	250	4	1
Sweet corn	5	20	0.1
Kentucky bluegrass	2100	3	6.3
Timothy grass	1050	2	2.1
Orchard grass	840	2	1.68
Alfalfa	430	3	1.29

distribution with mean λ [8]. We approximate the ladar as having an infinitesimal, non-divergent beam. We also ignore uncertainty in the range measurement for now. Under these conditions, the probability density function (pdf) of the measured range will have an exponential distribution with mean λd , that is:

$$f(r) = \lambda d e^{-\lambda d r} .$$

A similar model was used in [16] for two-axis ladar data in a forest. Although real ladar beams are slightly divergent, we found in [8] that the exponential distribution was close to the true distribution for the divergent case and that it modelled experimental data quite well.

If an obstacle is at distance r_0 in the field of grass modelled above, then the probability that a ladar beam will reach the obstacle is

$$g(r_0) = \Pr(r \geq r_0) = \int_{r_0}^{\infty} f(r) dr = e^{-\lambda d r_0} .$$

To estimate realistic values of the parameters λ and d , we obtained published data from the crop insurance industry for a variety of crops [17]; Table 1 shows a selection of this data. Note that the density refers to the number of plants, which would not account for multiple leaves on each plant. For comparison, in [8] we obtained a value of $\lambda d = 3.7/m$ in fitting an exponential distribution to scans taken with the robot in Figure 1 of a fairly dense field of tall grass. Estimating the grass blade diameter at 3 mm, this gives $\lambda = 1233 m^2$, which falls comfortably in the range of agricultural grass data in Table 1. Figure 2 plots $g(r)$ for a few of the plants in the table.

One could go further to develop models of detection and false alarm probabilities as a function of range for a given obstacle detection algorithm. For our current needs, we take a simpler approach with the following heuristic argument. Robust line-fitting algorithms can have a break-down point of up to 50% outliers (eg. grass in front of an obstacle); if such were the basis of an obstacle detection algorithm, then $g(r) = 0.5$ would give a plausible estimate of the maximum range for reliable detection. For the examples in Figure 2, this gives

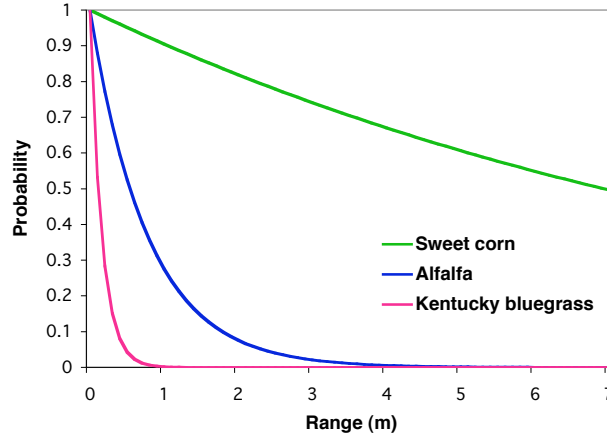


Fig. 2. Probability $g(r)$ that a range measurement will reach an obstacle as a function of range, r , for selected crop data from Table 1

$r = 0.11, 0.54,$ and 6.9 m for Kentucky bluegrass, alfalfa, and sweet corn, respectively. This gives a useful guide to the “penetration depth” for lidar-based obstacle detection in plants. For many plants, the answer will be well under one meter, which is not bad for robots in the man-portable size class but is quite limited for large vehicles.


The above analysis assumed an infinitesimal, non-diverging beam, whereas real beams have a finite initial diameter and diverge with increasing range; for example, a typical beam spread is around 10 milliradians (mr). This means that the beam cross-section is already at least 1 cm in diameter at 1 m from the robot. Given the thinness of much vegetation, this means that mixed pixels will be ubiquitous in off-road navigation; that is, pixels where the beam footprint spans a range discontinuity. With AMCW lidar, it is well known that mixed pixels can produce range measurements literally anywhere within the dynamic range of the sensor [18]. Anecdotally, pulsed time-of-flight lidars like the SICK produce mixed pixels with range measurements that are between the foreground and background surfaces; however, we have not seen that documented in print.

To begin to explore how mixed pixels might affect interpretation of vegetation, we did some experiments to characterize mixed pixel behavior of the SICK lidar. Figure 3 shows the experimental set-up. Figure 4 shows scans taken with the backstop at two different distances from the rods. Mixed pixels were clearly present when the distance between the foreground and background was 60 cm, but were not present when the distance was 200 cm. Trials were conducted with foreground-background distances varying in 20 cm steps between these extremes; mixed pixels did not occur for separations beyond 160 cm. We interpret this as reflecting the ability of the pulse detection circuitry to discriminate the returns from the foreground and background, which depends on the separation and on the laser pulse duration. When the object separation is large enough, the two reflections are sufficiently separated in time for the circuitry to detect the first

reflection; otherwise, the two reflections overlap and the measured range is intermediate between foreground and background. Similar behavior occurs with multi-return, airborne terrain mapping ladars, which cannot distinguish consecutive returns within a threshold separation distance (eg. also about 160 cm).

In ladar scans of vegetation with sensors mounted on a ground vehicle, the foreground-background distance within mixed pixels will often be less than the minimum discriminable return separation, particularly at close ranges. We suspect that this will tend to make vegetation appear flatter than it is in reality – which may be convenient. Future experiments to confirm this would be valuable.

Fig. 4. Set-up for experiments to characterize mixed pixel behavior of the SICK ladar. The ladar is at a variable distance in front of a row of vertical rods that mimic stems, which in turn are at a variable distance in front of a flat vertical backstop. The red wedge is a graphic overlay illustrating the plane of the scan.



3 Radar

Radar-based object detection in foliage has been studied for remote sensing of crops and forests [12] and for detecting military targets in trees [11]. We conducted experiments on radar-based obstacle detection through foliage using a radar that was available to us through the DARPA Tactical Mobile Robotics program. In this section, we very briefly summarize issues pertinent to radar remote sensing through foliage, then briefly describe the sensor we used, the data processing we applied, and the results we achieved.

Two key issues are penetration depth of the radar energy through the foliage and angular (or spatial) resolution of the radar system. Penetration depth is a function of the dielectric constant of the material being penetrated; for natural materials, including vegetation, this depends primarily on the moisture content of the material. Formally, penetration depth is defined as the depth at which the signal is attenuated to $1/e$ (37%) of its original value; typically, this happens within a few wavelengths. Wavelengths used in remote sensing of vegetation vary from about 70 cm (~ 0.5 GHz) to 1 cm (~ 30 GHz). The radar conveniently available to us was a micro-impulse radar (MIR) from Lawrence Livermore National Laboratory (Figure 5a), which was developed with buried mine detection as a primary application. This unit has a center frequency of 2.2 GHz (wavelength of 13.6 cm), frequency content from about 1-4 GHz, a pulse width of 0.5 ns, average power of 0.1 mW, and a beam width of 90x90 degrees with an

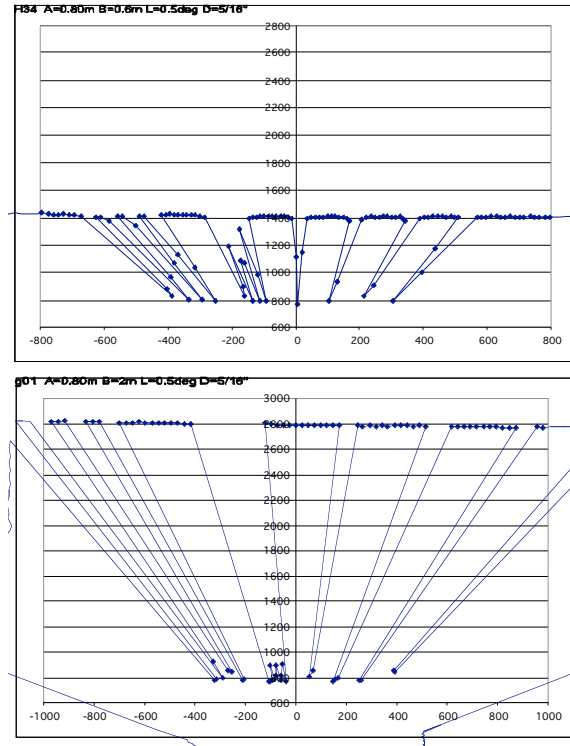


Fig. 4. Results from mixed pixel characterization experiments from the set-up shown in Figure 4. Both graphs show plots of a single scan of range data, with blue dots representing range pixels and line segments drawn between them to clarify scan order. In both graphs, the upper row of dots are pixels on the backstop and the bottom-most dots are pixels on the rods. Intermediate dots in the top graph are mixed pixels. In both graphs, the ladar was 80 cm from the rods. Top: 60 cm between rods and backstop; bottom: 200 cm between rods and backstop

approximate antenna size of 11x7 cm. These parameters made it reasonably well suited to the initial feasibility study we desired to conduct, with the exception of the very wide beam width.

Beam width is an inverse function of antenna width; hence, to get a narrow beam requires either a wide real aperture antenna or synthetic aperture algorithms applied to multiple measurements from a moving sensor. Since we desired at least a crude image of material behind foliage, rather than just one range measurement, we also needed either a scanning radar or multiple transceivers. The technique we explored was to take multiple scans by translating one radar perpendicular to the beam axis, then to estimate a 2-D map of scattering surfaces using a radar signal processing algorithm based on diffraction tomography using content from multiple frequencies [14]. The concept was that the single, translating radar would be replaced by multiple transceivers in a deployed system on a robotic vehicle.

Figure 5b-5c shows a large tree trunk obscured by foliage that was used for testing. In this data set, we acquired 45 scans, with the sensor moved 4 cm per scan. Figure 6a shows the amplitude data for all 45 scans concatenated together (the X axis is the sensor translation axis). Deep blue represents low amplitude; bright red represents high amplitude. The reconstruction algorithm [14] transforms this data to the frequency domain, then applies a line-to-line backward propagation algorithm to estimate the spatial frequency map of scatterers one depth (Z value) at time, and finally transforms back to the spatial domain to obtain the intensity map in Figure 6b. The tree trunk was very clearly detected, behind over 2.5 m of branches and leaves.




Fig. 5. (a) Micro-impulse radar, shown with a quarter for scale. (b) Side view of a tree that was imaged with the radar, showing a large trunk over 2.5 m behind thick foliage. (c) Front view of the tree; the red line shows the path along which the sensor was translated to acquire scans.

5 Summary and Conclusions

We demonstrated obstacle detection and avoidance behavior in dense, tall grass with a SICK ladar mounted on a portable robot. To gain insight into how far into grass and other vegetation obstacles can be detected, we presented a simple probabilistic model, based on assuming a Poisson distribution for the vegetation, parameterized in terms of the stem diameter and frequency. We instantiated these parameters for various plant types with data drawn from the crop insurance industry. This led to predictions that the reliable obstacle detection distance could range from on the order of 10 cm for very dense grass to about 7 m for a field of sweet corn. We observed that mixed pixels will be ubiquitous in ladar scans of vegetation and speculated that this may cause the vegetation to appear flatter than it really is, particularly close to the robot; this could be helpful to driving control algorithms, because it would smooth out the apparent bumps in the terrain. This prediction needs to be tested experimentally.

We then presented results for obstacle detection through foliage with a 2.2 GHz micro-impulse radar. The radar was translated perpendicular to the beam axis to acquire 45 scans, 4 cm apart, simulating a dense linear array of transceivers. Applying a multi-frequency diffraction tomography algorithm to this data produced a high resolution, 2-D map of scattering surfaces, which revealed a large tree trunk behind 2.5 m of thick foliage. We have not had the opportunity yet to

build an array with multiple transceivers, but these results suggest that this has great potential for allowing robots to detect obstacles a few meters through vegetation that would be too dense for lidar.




Fig. 6. (a) Concatenated amplitude data from 45 scans, spaced 4 cm apart, of the tree in Figure 5b-5c. (b) 2-D density map obtained by applying a multi-frequency diffraction tomography algorithm [14] to the data in (a). The bright red spot is the tree trunk

Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the DARPA Tactical Mobile Robotics, Mobile Autonomous Robot Software, and Perceptor programs through agreements with the National Aeronautics and Space Administration. Reference herein to any specific commercial product, process, or service by trade name, 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. Matthies, L., Bellutta, P., McHenry, M. (2003) Detecting Water Hazards for Autonomous Off-Road Navigation. Proc. SPIE Symp. on Unmanned Ground Vehicle Technology V, Orlando, Florida
2. Sarwal, A., Rajagopalan, V., Simon, D., Rosenblum, M., Nett, J. (2003) Terrain Classification. Proc. Collaborative Technology Alliances Conference on Robotics. U.S. Army Research Laboratory, Adelphi, Maryland
3. Matthies, L., Rankin, A. (2003) Negative Obstacle Detection by Thermal Signature. Proc. IEEE Int. Conf. on Robotics and Automation. Taipei, Taiwan
4. Olin, K.E., Tseng, D.Y. (1991) Autonomous Cross-Country Navigation. IEEE Expert, **6(4)**
5. Elachi, C. (1987) Introduction to the Physics and Techniques of Remote Sensing. John Wiley and Sons, New York
6. Matthies, L., Kelly, A., Litwin, T., Tharp, G. (1996) Obstacle Detection for Unmanned Ground Vehicles: a Progress Report. In: Giralt, G., Hirzinger, G. (eds.): Robotics Research: Proceedings of the 7th International Symposium. Springer-Verlag
7. Jensen, J. R. (2000) Remote Sensing of the Environment, Prentice-Hall
8. Macedo, J., Manduchi, R., Matthies, L. (2000) Ladar-based Discrimination of Grass from Obstacles for Autonomous Navigation. Proc. 7th Int'l. Symp. of Experimental Robotics, Honolulu, Hawaii.
9. Lacaze, A., Murphy, K., DelGiorno, M. (2002) Autonomous Mobility for the Demo III Experimental Unmanned Vehicles. Annual Symp. of the Association for Unmanned Vehicle Systems International (AUUVSI), Orlando, Florida
10. Hebert, M., Vandapel, N., Keller, S., Donamukkala, R.R. (2002) Evaluation and Comparison of Terrain Classification Techniques from LADAR Data for Autonomous Navigation. Army Science Conference, Orlando, Florida
11. Xu, X., Narayanan, R.M. (2001) FOPEN SAR Imaging using UWB Step-frequency and Random Noise Waveforms. IEEE Transactions on Aerospace and Electronic Systems, **37(4)**
12. Henderson, F.M., Lewis, A.J. (eds.) (1998) Manual of Remote Sensing, Vol. 2: Principles and Applications of Imaging Radar. John Wiley and Sons, New York
13. Talukder, A., Manduchi, R., Castano, R., Owens, K., Matthies, L., Castano, A., Hogg, R. (2002) Autonomous Terrain Characterization and Modelling for Dynamic Control of Unmanned Vehicles. Proc. IEEE/RSJ Conf. on Intelligent Robots and Systems, Lausanne, Switzerland
14. Mast, J.E., Johansson, E.M. (1994) Three-dimensional Ground Penetrating Radar Imaging Using Multi-frequency Diffraction Tomography. Proc. SPIE Vol. 2275: Symp. on Advanced Microwave and Millimeter Wave Detectors
15. Castano, A., Matthies, L. (2003) Foliage Discrimination Using a Rotating Ladar. IEEE Int'l Conf. on Robotics and Automation, Taipei, Taiwan
16. Huang, J., Lee, A.B., Mumford, D. (2000) Statistics of Range Images. IEEE Conference on Computer Vision and Pattern Recognition, Hilton Head, South Carolina
17. //www.agric.gov.ab.ca/agdex/100/100_22-1.html
18. Hebert, M., Krotkov, E. (1992) 3-D Measurements from Imaging Laser Radars: How Good are They? Int'l J. of Image and Vision Computing, **10(30)**