

Vegetation Detection for Driving in Complex Environments

David M. Bradley, Ranjith Unnikrishnan, and James Bagnell

Abstract—A key challenge for autonomous navigation in cluttered outdoor environments is the reliable discrimination between obstacles that must be avoided at all costs, and lesser obstacles which the robot can drive over if necessary. Chlorophyll-rich vegetation in particular is often not an obstacle to a capable off-road vehicle, and it has long been recognized in the satellite imaging community that a simple comparison of the red and near-infrared (NIR) reflectance of a material provides a reliable technique for measuring chlorophyll content in natural scenes. This paper evaluates the effectiveness of using this chlorophyll-detection technique to improve autonomous navigation in natural, off-road environments. We demonstrate through extensive experiments that this feature has properties complementary to the color and shape descriptors traditionally used for point cloud analysis, and show significant improvement in classification performance for tasks relevant to outdoor navigation. Results are shown from field testing onboard a robot operating in off-road terrain.

I. INTRODUCTION

Current autonomous navigation techniques work well for environments such as hallways and on roads, where obstacles are static and usually rigid. In these cases, size and shape are sufficient for determining which obstacles can be driven over and which need to be avoided. In off-road driving, however, the assumption that every obstacle is rigid and would be lethal to the robot quickly presents problems. In situations such as a field of tall grass, there may be dense geometric obstacles on all sides of the robot. In order to plan safe, efficient paths the robot must be able to reliably discriminate between vegetation that it can drive through if necessary, and rigid obstacles such as tree trunks and rocks that can cause damage (Figure 1). For safe high-speed operation performing this discrimination at range becomes increasingly important.

Methods have been developed to detect vegetation from 3-D point clouds [1], [2], but there is still significant room for improvement, particularly at longer ranges where the limited viewpoint of onboard sensors, reflection of the laser pulses away from the scanner, laser beam divergence, and partial occlusion by other objects make it difficult to obtain point clouds of sufficient quality and density.

Fortunately there are well-established techniques for measuring chlorophyll content using a multi-spectral camera [3], [4], [5], [6], [7], [8], [9], [10] that have been developed for satellite-based remote sensing. A simple pixel-by-pixel comparison between red and Near-InfraRed (NIR) reflectance,

This work was sponsored by DARPA under contracts "UPI" and "Perception for Off-Road Mobility (PerceptOR)" (contract number MDA972-01-9-0016) and by the Army Research Office through the NDSEG fellowship program. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing official policies or endorsements, expressed or implied, of the U.S. Government.

The authors are with the Carnegie Mellon University Robotics Institute, 5000 Forbes Avenue, Pittsburgh, PA 15213, USA dbradley@cs.cmu.edu



Fig. 1. Autonomous mobile robot used for this evaluation. The addition of NIR data helps in discrimination between the rock on the left and the bush on the right.

normally referred to as a vegetation index or a band-ratio, provides a powerful and robust way to detect vegetation. Further, most CCDs have significant spectral response all the way out to wavelengths of around 1000nm, meaning that a standard monochrome CCD can be converted into a NIR camera simply by covering it with a longpass filter. Although the viewpoint of a satellite is drastically different from that of a mobile robot, we show that the technique is still effective despite additional complications such as views of the sky, or shadowed areas that are lit by both light reflected off of other surfaces and light reflected from the sky.

In the mobile robotics community, surprisingly little attention has been given to the use of multi-spectral information for ground-based terrain classification for navigation. An early attempt, [11], used data from a filter wheel camera to label pixels as chlorophyll-rich vegetation or soil based on thresholding the ratio of the NIR and red values. Later work in [12] mentioned the usefulness of NIR in detecting photosynthetic vegetation, but described the use of a Gaussian mixture model-based classifier with only RGB features. Aside from the overall speed of the classifier, no quantitative performance analysis was given, and the role played by the features in the larger system was largely presented as anecdotal. This work is an extension of [13], which provides an extensive qualitative evaluation of several vegetation indices across different environments and geographic locations, but did not incorporate a lidar or stereo system and so could not evaluate the utility of combining vegetation indices with 3-D data.

To bridge this gap in understanding, we choose a broad suite of classification tasks relevant to off-road navigation using combinations of popularly used shape, density and

color features and show the effect of augmenting these features with NIR and NDVI information. In particular for each task we analyze the reliability of different feature sets with respect to sensing range. We demonstrate through a set of experiments that this feature has properties complementary to the color and shape descriptors traditionally used for point cloud analysis, and show significant improvement in classification performance for tasks relevant to outdoor navigation.

The paper is organized as follows. Section II gives an overview of existing approaches to vegetation detection using spectral and geometric information, and its application to navigation. Section III describes the overall system implementation on our robot platform including the sensing modalities available and the design choices made for internal representation and path planning. Section IV describes the datasets used and presents results on voxel classification experiments. Section V describes how improved voxel classification translates into system level performance gains. Finally we conclude in Section VI and discuss future work.

II. APPROACHES TO VEGETATION DETECTION

This paper combines two complementary approaches to vegetation detection, vegetation indices (also referred to as band-ratio techniques) that have long been used in the remote sensing community, and more recent 3-D structure based techniques from the mobile robotics community.

A. Vegetation Indices

The spectral properties of chlorophyll-rich vegetation are primarily determined by the absorption spectra of water and chlorophyll, and the refraction of light at cell walls [14]. The water present in cells absorbs light with wavelengths longer than 1400 nm. Chlorophyll strongly absorbs visible light, especially red and blue wavelengths [3]. The remaining light is efficiently scattered by the critical internal reflection caused by the change in refractive index from water to air at the cell wall. As a result, those wavelengths between 700 nm and 1400 nm that escape both water and chlorophyll are strongly reflected in all directions.

The sharp difference between the reflectance of vegetation at 645nm (red) and at 780nm (NIR) has long been exploited in the field of satellite remote sensing. Kauth and Thomas [7] noticed that plotting NIR reflectance against Red reflectance for satellite images produced a scatter diagram with a line of points formed by pixels containing bare soil, and a cluster of points from pixels completely covered with vegetation. Points with a mixture of vegetation and soil appear between the soil line and the vegetation point. Figure 2 shows this scatter plot created from one of our images. Because our camera also includes a view of the sky, our scatter plot contains a blue sky region (bottom image, marked in blue) as well as the soil region (middle image, marked in red) and the vegetation region (top image, marked in green). Clouds blend into the soil line, but are still very distinct from vegetation. Pixels containing vegetation and blue sky are remarkably well separated from everything else in a natural scene.

One of the most popular ways to use the information contained in the red and NIR bands for remote sensing applications is to compute a quantity known as the Normalized Difference Vegetation Index (NDVI) which varies from -1 (blue sky) to 1 (chlorophyll-rich vegetation) [8].

$$\text{NDVI} = \frac{\rho_{\text{NIR}} - \rho_{\text{RED}}}{\rho_{\text{NIR}} + \rho_{\text{RED}}} \quad (1)$$

An informative interpretation of the NDVI is given in [15], where it is shown to be a measurement of the slope from the origin to the location of the pixel in a 45 degree rotation of the red-NIR space. Several attempts have been made in the remote sensing literature to correct deficiencies in this index [10], [5], [16], [15], particularly in shadows and underexposed areas. Since the NDVI measures the slope from the origin, sensor noise and errors in the radiometric calibration of the red and NIR sensors have a much greater effect in underexposed areas. Shadows pose a more challenging problem, since the reflected light that illuminates shadowed regions can have a spectral distribution that is significantly different from that of sunlight, usually shifted towards blue wavelengths because of atmospheric scattering.

The typical use of NDVI in remote sensing is to measure the Leaf Area Index (LAI), the percentage of the ground surface that is covered by vegetation. However, when the goal is linear classification into vegetation and non-vegetation categories, it is useful to provide the raw NIR and red values as well as the NDVI to the classifier. This is because a linear decision boundary based solely on the NDVI corresponds to a line intersecting the origin of the red-nir space, whereas a linear classifier operating on the raw pixel values and a constant bias feature can produce a decision boundary with an arbitrary intercept. In this paper the NDVI is discussed to tie this vegetation detection technique back to its origins in the remote sensing community, but the actual classifiers evaluated in section IV use the raw pixel values as well as the NDVI, since for a linear classifier they may be more informative than using many of the variants of the NDVI.

1) *Validation of Approach:* The USGS digital spectral library [17], provides a useful tool for selecting an appropriate NIR filter, and verifying the effectiveness of using NIR in conjunction with red across a variety of different types of vegetation and non-vegetation. This library contains the spectral signatures of over 800 different materials. Many of these spectra are from rare minerals that are not relevant to our application. The relevant subset of the library used for testing contained 105 common types of vegetation and 169 common soil mixtures, artificial materials, and coatings. The optimal linear separator on this dataset classifies 258 of 274 materials correctly. A yellow aspen leaf and dry grass are misclassified as non-vegetation, and various man-made materials and natural mixtures involving hematite are misclassified as vegetation (figure 3).

B. 3-D Point Distribution

A complementary approach to vegetation detection, presented in [1], uses the spatial distribution of the local lidar point cloud to classify the region into surfaces, linear structures, and a class referred to as scatter, which includes tree

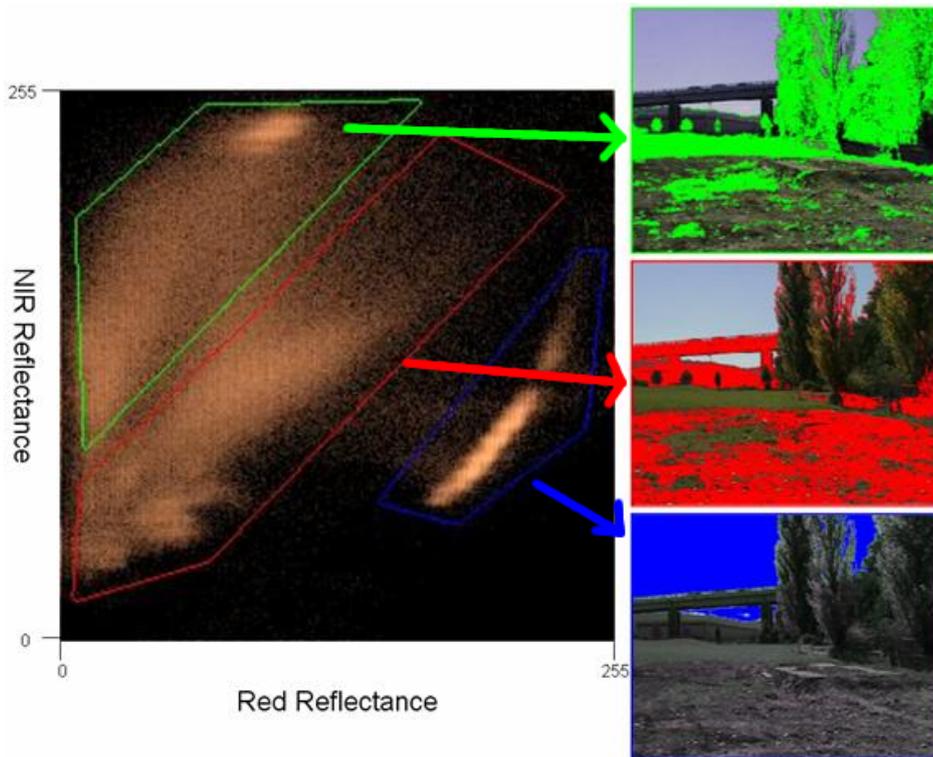


Fig. 2. Scatter plot of NIR reflectance vs. red reflectance for all pixels in a typical image. Different regions in the scatterplot correspond to different types of materials in the image. Pixels in the green region correspond to vegetation (top image), pixels in the red region are mainly soil and man-made structures (middle image), and pixels in the blue region correspond to sky (bottom image).



Fig. 4. Small images show the RGB and color appearance of two scenes. Large images are the RGB image with areas of high NDVI highlighted in bright green. While generally reliable in natural scenes, NDVI can fail on synthetic materials such as the paint on the truck in the right image. In this case 3-D methods would easily classify the side of the truck as a rigid obstacle instead of vegetation. Note that NDVI correctly classifies the brown dead grass in the left image as vegetation.

canopy and porous vegetation. This method first computes the eigenvalues of the covariance matrix of the local point cloud (defined as all points within a certain distance of the point of interest), and then classifies the point based on the relative magnitudes of those eigenvalues. Linear structures have one dominant eigenvalue. Surfaces have two large eigenvalues, and an area is declared to be scatter when the third largest eigenvalue is a significant fraction of the largest eigenvalue. In addition, the estimated surface normal of the local area, another useful feature of the local point cloud, is recovered by this computation as it is simply the eigenvector corresponding to the smallest eigenvalue of the covariance matrix.

This 3-D method performs particularly well on certain man-made structures where the NDVI approach is known to fail. For instance, certain types of vehicle paint give off a vegetation-like NDVI signature (figure 4), but the flat sides of vehicles are easily detected as surface 3-D structures (figure 5). However, the 3-D method does require a relatively dense, high-quality point cloud, which limits its application to areas closer to the robot.

III. SYSTEM OVERVIEW

The vegetation index technique for detecting chlorophyll-rich vegetation can (and has) been implemented in many different ways, but since this is an experimental paper which

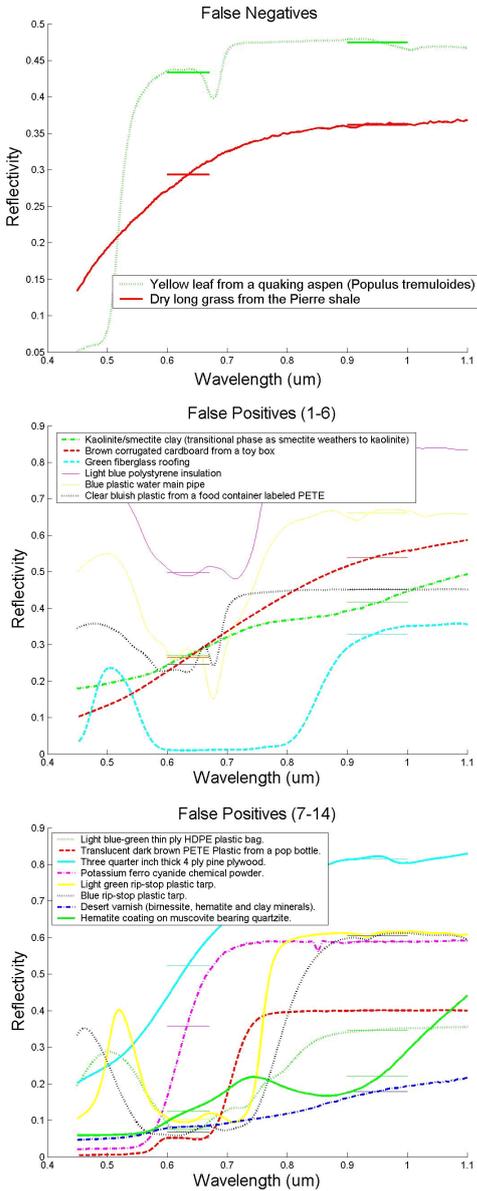


Fig. 3. Material spectra misclassified by an optimal linear discriminator using the red and NIR bands (600-670nm and 900-1000nm). Horizontal line segments indicate the simulated camera responses used for classification. 258 of 274 common materials are classified correctly. The failures are mainly dead vegetation where the chlorophyll has had time to degrade, man-made materials such as plastic, and natural mixtures involving hematite.

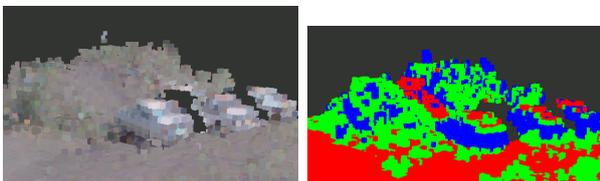


Fig. 5. Left: average RGB values of each voxel in a scene containing cars. Right: after 3-D classification the voxels containing the flat sides of the cars are classified as obstacles (blue). Voxels containing curved car surfaces have more of a vegetation-like signature (green), and voxels corresponding to ground are marked in red.

addresses system-level effects we include a brief description of the specific implementation to aid in analyzing the results.

The robotic system used for this evaluation extends the approach outlined in [18]. A set of laser scanners and cameras provide the input to the perception system, which consists of 3-D points that have been projected into camera images, and tagged with local properties of the image such as color and the NDVI value (Figure 6). The local perception system then discretizes the space surrounding the robot into a 3-D grid of voxels, and accumulates summary statistics of the tagged laser points over each voxel. The summary statistics include the averages of the point tags, eigenvalues of the local point cloud scatter matrix, the surface normal (3rd eigenvector of the scatter matrix), and the probability and strength of laser pulse reflection from the voxel. The perception system is then responsible for condensing this 3-D grid of voxels into a 2-D grid of cost values so that the planner (a variant of A*) can then plan a minimum cost path. Various interpretations have been proposed for the meaning of the cost values, such as mobility risk, but due to the tightly coupled nature of mobile robot systems they have no fundamental interpretation apart from the paths that they cause the planner to produce through the environment.

Costs are produced from the voxels by applying a set of learned classifiers in conjunction with several hand-tuned rules for combining the classification results in each vertical column. Linear maximum entropy (multi-class logistic regression) classifiers [19], are used in order to meet the strict real-time requirements imposed by continuous motion at several meters per second. These classifiers find a conditional distribution $P(c|d)$ for the class $c(d)$ of training example $d \in D$, that has maximum entropy (i.e. makes as few assumptions as possible) subject to the constraint that the expected value of each feature $f_i(d, c)$ of each example matches its average value over the training set (2).

$$\frac{1}{|D|} \sum_{d \in D} f_i(d, c(d)) = \sum_d P(d) \sum_c P(c|d) f_i(d, c) \quad (2)$$

In practice the empirical distribution of d in the training set is used to approximate the true distribution $P(d)$, and a weak Gaussian prior is added to control overfitting. The resulting maximum entropy conditional distribution takes the form of a linear classifier over the features (3), that is normalized to form a probability distribution (4).

$$P(c|d) = \frac{1}{Z(d)} e^{\sum_i \lambda_i f_i(d, c)} \quad (3)$$

$$Z(d) = \sum_c e^{\sum_i \lambda_i f_i(d, c)} \quad (4)$$

The classifiers are trained from labeled data sets that are gathered by either driving the robot over areas of different terrain type (and labelling everything that passes under the robot), or in the case of obstacles, by hand-labelling their locations in a set of point clouds using a simple paint-like interface. Currently there are two ways that classifiers are used in the system. The classification of a voxel as obstacle, vegetation, or ground influences how the data it contains is

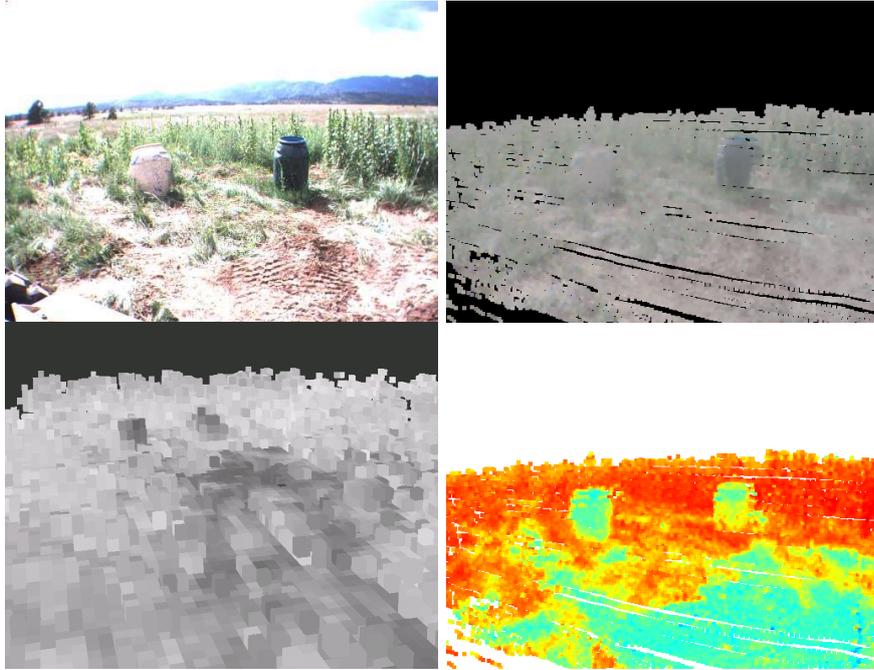


Fig. 6. Top left: A typical camera image. Inputs to the perception system consist of 3-D points that have been projected into the camera images and tagged with local image properties such as color (top right) and the NDVI value (bottom right). The points are then accumulated in a discretized 3-D grid of voxels, with the average NDVI of each voxel shown in the bottom left image.

used in estimating the ground plane. Also, the rigid/non-rigid obstacle classifications of the voxels in each column are used to compute an overall rigid/non-rigid flag for the column, which allows for the assignment of different costs to rigid and non-rigid obstacles in the 2-D costmap.

Because the beamwidth of the ladar we use is approximately 10x the angular resolution of the cameras, ladar points tagged with the center point of the ladar beam may be incorrectly tagged with pixels from the sky. Because of this effect voxels straddling vegetation/sky image boundaries can have average NDVI values that look like non-vegetation (blue sky is particularly bad because of its exceptionally low NDVI value). To reduce this problem, the exposure on the cameras is controlled to correctly expose only those portions of the image where ladar returns were received, which generally means that sky pixels are marked as overexposed and their NDVI tags do not contribute to the voxel's classification.

IV. VOXEL CLASSIFICATION RESULTS

A. Data Set

The system was trained and tested using data from two significantly different physical environments. Training examples of rocks, bushes, grass, trees, and man-made structures such as cars, telephone poles, and barrels were gathered from a site in Western Pennsylvania. More examples of rocks and bushes were collected several weeks later in natural terrain in the foothills of the Colorado Rockies. Voxels were labeled by either hand-labelling the point cloud, or in some cases labeling everything the vehicle drives over. In both cases the features (density, surface normal, scatter matrix eigenvalue, RGB, NIR, and NDVI) of each labeled voxel were recorded every time a new laser point in that voxel was observed.

B. Classification Results

We start our evaluation with quantitative results on the effects of using NDVI for several important discriminative tasks. We first compare the performance of the feature sets for similar (but physically separate) environments and lighting conditions. Figure 7 compares classifiers trained for the binary task of discriminating between non-rigid voxels (grass, bushes, etc...) and rigid voxels (Tree trunks, rocks, cars, telephone poles, bare ground, etc...). As discussed in Section III, this task is crucial to the way the robot generates a costmap for path planning. Both the RGB and the NDVI features can accomplish this task successfully, showing roughly a 20% improvement in classification accuracy over ladar features alone. The ladar features become crucial, however, in the more complicated three-way classification task used by the robot to estimate the true ground plane of the scene. This task is similar to the previous task, with the exception that the rigid voxel class is divided into an obstacle class and a road class (horizontally oriented bare ground surfaces). In Figure 8 we see that combining the camera features with the ladar features boosts the total classification accuracy by approximately 10%. This performance boost is almost entirely from improvements in the ability to discriminate between the obstacle and road classes, as shown in Tables II & III. Again in this task we see that the camera features are very helpful in discriminating between non-rigid voxels and the other two classes (Table I).

Finally we investigate the generalization ability of the different feature sets across different geographic environments on the task of discriminating between rocks and bushes. For this test the training set is from Pennsylvania, and the

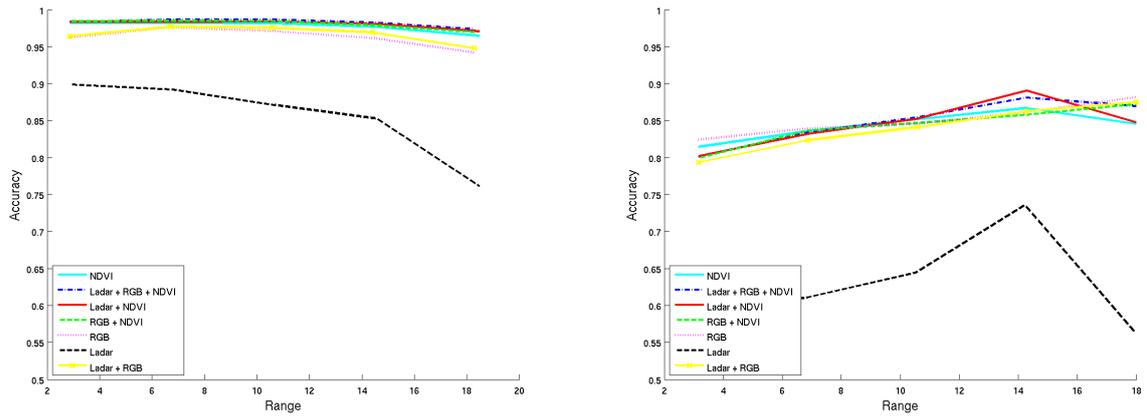


Fig. 7. Voxel classification accuracy vs range for telling rigid obstacles (rocks, tree trunks, telephone poles, cars, etc...) apart from non-rigid obstacles (grass and bushes). Left: training set performance. Right: Test set performance. Inclusion of RGB or NDVI features provides a significant performance boost over ladar features alone.

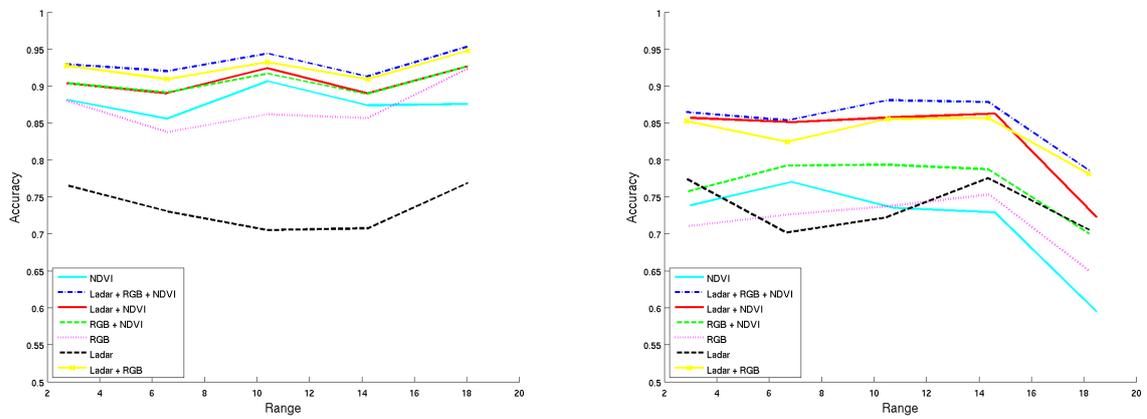


Fig. 8. Voxel classification accuracy vs range for discriminating between rigid obstacles, non-rigid vegetation, and ground. Left: training set performance. Right: Test set performance. This task, which is used in ground plane estimation, benefits significantly from combining ladar and RGB or NDVI features.

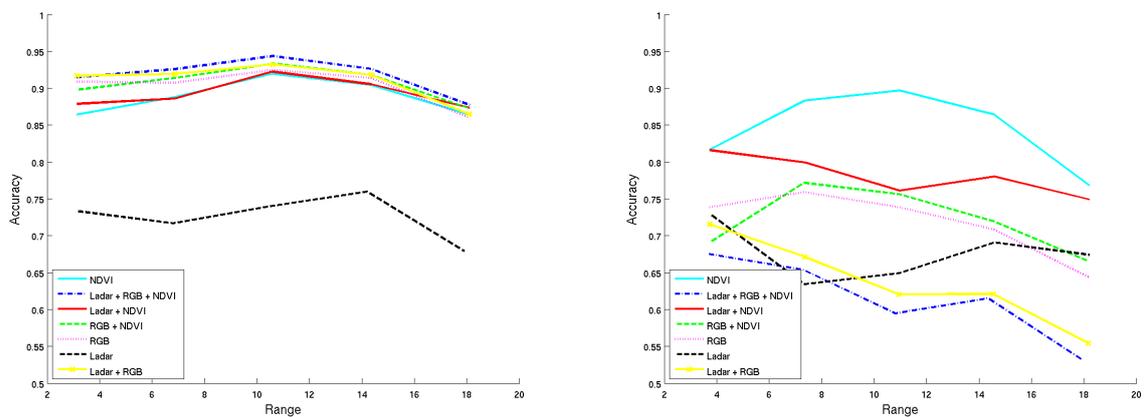


Fig. 9. Voxel classification accuracy vs range for discriminating between rocks and bushes. Left: training set performance. Right: Test set performance. The training set for this task was collected in Pennsylvania, and the test set was collected in Colorado. The NDVI features show superior generalization to the novel environment.

Predicted Class	True Class		
	obstacle	vegetation	ground
obstacle	88.7	11.4	3.1
vegetation	9.5	64.0	27.9
ground	1.8	24.6	68.9

TABLE I

CONFUSION MATRIX FOR A CLASSIFIER TRAINED ONLY WITH LADAR FEATURES

Predicted Class	True Class		
	obstacle	vegetation	ground
obstacle	88.1	0.2	4.7
vegetation	6.4	95.1	21.2
ground	5.5	4.7	74.1

TABLE II

CONFUSION MATRIX FOR A CLASSIFIER TRAINED WITH ALL FEATURES

test set is from Colorado. As shown in figure 9, the NDVI features are hardly affected by the change in environment, and their performance degrades only slightly. The RGB features, on the other hand, allow the classifiers to overfit to the specific lighting and flora of the training set, leading to vastly degraded performance in the novel environment.

Fortunately, it is not necessary to have many NIR sensors to benefit from the generalization ability of the NDVI features. As long as a representative sample of RGB-NIR tuples are available for the current environment and lighting condition, a RGB-based classifier can be trained to predict the NDVI values. This representative sample might come from a single NIR camera whose field of view overlaps with that of an existing RGB camera on the robot. Figure 10 shows the results of predicting the NDVI value at each pixel in an image using simple features computed from the local RGB values consisting of the RGB, HSV, and Lab values of the pixel when blurred with a gaussian at four scales (sigma of 1, 2, 4, and 8 pixels). The 100-node regression tree used for the prediction was trained on a random sample from a 4-color (RGB and NIR) image with a field of view covering a separate portion of the same scene. Table IV gives the relative error rates of predicting NDVI over a sequence of 1646 4-color images from the same environment. Numbers given have been normalized by the error produced by simply predicting the mean NDVI value of all the pixels in the image sequence. Classifiers were re-trained every 25 images using an image that was 20 seconds old (which corresponds to approximately 30m of travel in this sequence, enough to ensure a significant viewpoint difference). The regression tree produced less than 8% of the error obtained by predicting

Predicted Class	True Class		
	obstacle	vegetation	ground
obstacle	78.7	0.2	21.1
vegetation	6.7	95.1	21.4
ground	14.6	4.7	57.5

TABLE III

CONFUSION MATRIX FOR A CLASSIFIER TRAINED WITH ONLY CAMERA FEATURES



Fig. 10. NDVI can be predicted from RGB for a given environment and illumination condition, meaning the benefits of NDVI can be captured even if the NIR sensors on the robot have a much more limited field of view than the RGB sensors. The 4-color image (RGB & NIR, only RGB shown) at top left is used to train a classifier that predicts the NDVI image (bottom right) for the RGB image at top right. The ground truth NDVI image is shown at bottom left. Note how the vegetation behind the cars is detected against the dark green portion of the building.

100-node regression tree	linear least squares
0.078	0.117

TABLE IV

MEAN NDVI PREDICTION ERROR OVER A SEQUENCE OF 1646 4-COLOR IMAGES CAPTURED ONBOARD A MOBILE ROBOT WHILE DRIVING IN A GRASSY URBAN AREA.

the mean NDVI value.

V. SYSTEM-LEVEL PERFORMANCE GAINS FROM IMPROVED VEGETATION DETECTION

The increase in voxel classification accuracy from adding RGB and NDVI features improves the robot's overall ability to autonomously avoid rigid obstacles without being overly afraid of bushes or tall grasses. Figure 6 shows the robot's view of two barrels that it is avoiding during a 2 m/s autonomous run. Figure 11 shows the classification maps the robot generates from this position with and without the NDVI features. The tall grass behind the barrels creates many false obstacles for the lidar-only perception system. With NDVI information the barrels stand out clearly as the two large blue blobs, and there is only one small false-positive obstacle. To show generalization across environments, this run took place in Colorado using classifiers that were trained exclusively on data collected in Pennsylvania.

VI. CONCLUSIONS AND FUTURE WORK

The near-infrared and red color bands provide a robust method for discriminating between rigid obstacles and non-rigid vegetation that thus far has been underutilized on mobile robotics systems. By showing how it contributed to performance improvements in a field tested robotic system, we hope to encourage more widespread use of this technique

More work is needed in the area of color constancy, or compensating for the effect of differences in the illuminants

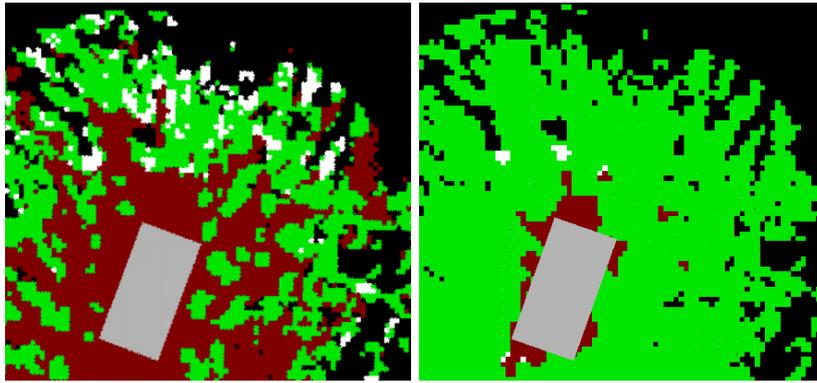


Fig. 11. Left: classification map using lidar features only. Right: adding NDVI information allows the system to be more selective, and display the barrels from Figure 6 clearly in the classification map without also picking up non-rigid vegetation. The barrels are the two large white blobs. White indicates obstacle classification (high cost), green indicates vegetation classification (low cost), and red indicates road classification (very low cost). The position of the robot is represented as a gray rectangle.

encountered. The general problem of color constancy is underdetermined, however robotic systems that use ladars in conjunction with NIR and color cameras have a crucial advantage in that they often have access to information such as laser remission, surface normal, and surface shape that can be used to help deduce the approximate spectral distribution of the illumination.

VII. ACKNOWLEDGMENTS

The Authors would like to thank the Defense Advanced Research Projects Agency and the Army Research Office for their support of this research through the UPI program and the NDSEG graduate fellowship program. The authors are especially grateful to the UPI team for making this research possible, and in particular to Carl Wellington and Cris Dima for all of their hard work and advice. The authors would also like to thank the reviewers for their time and helpful comments.

REFERENCES

- [1] N. Vandapel and M. Hebert, "Natural terrain classification using 3-d lidar data," in *IEEE Intl. Conf. on Robotics and Automation (ICRA)*, 2004.
- [2] J. Macedo, R. Manduchi, and L. Matthies, "Ladar-based discrimination of grass from obstacles for autonomous navigation," in *Proc. of the International Symposium on Experimental Robotics*, Honolulu, HA, December 2000.
- [3] R. N. Clark, G. A. Swayze, K. E. Livo, R. F. Kokaly, S. J. Sutley, J. B. Dalton, R. R. McDougal, and C. Gent, "Imaging spectroscopy: Earth and planetary remote sensing with the usgs tetracorder and expert systems," *J. Geophys. Res.*, December 2003.
- [4] R. E. Crippen, "Calculating the vegetation index faster," *Remote Sensing of Environment*, vol. 34, pp. 71–73, 1990.
- [5] A. R. Huete, "A soil adjusted vegetation index (savi)," *Remote Sensing of Environment*, vol. 25, pp. 295–309, 1988.
- [6] C. F. Jordan, "Derivation of leaf area index from quality measurements of light on the forest floor," *Ecology*, vol. 50, pp. 663–666, 1969.
- [7] R. J. Kauth and G. S. Thomas, "The tasseled cap - a graphic description of the spectral-temporal development of agricultural crops as seen by landsat," in *LARS: Proceedings of the Symposium on Machine Processing of Remotely Sensed Data*, West Lafayette, IN, 1976, pp. 4B–14–4B–51.
- [8] F. J. Kreigler, W. A. Malila, R. Nalepka, and W. Richardson, "Preprocessing transformations and their effects on multispectral recognition," in *Proc. of the Sixth International Symposium on Remote Sensing of Environment*, Ann Arbor, MI, 1969, pp. 97–131.
- [9] R. Willstatter and A. Stoll, *Untersuchungen über Chlorophyll*. Berlin: Springer, 1913.
- [10] A. J. Richardson and C. L. Wiegand, "Distinguishing vegetation from soil background information," *Photogrammetric Engineering and Remote Sensing*, vol. 43, pp. 1541–1552, 1977.
- [11] L. Matthies, A. Kelly, T. Litwin, and G. Tharp, "Obstacle detection for unmanned ground vehicles: A progress report," in *Proc. of the International Symposium on Robotics Research (ISRR)*, 1996, pp. 475–486. [Online]. Available: citeseer.ist.psu.edu/matthies95obstacle.html
- [12] R. Manduchi, A. Castano, A. Talukder, and L. Matthies, "Obstacle detection and terrain classification for autonomous off-road navigation," *Autonomous Robots*, vol. 18, pp. 81–102, 2005. [Online]. Available: citeseer.ist.psu.edu/manduchi04obstacle.html
- [13] D. Bradley, S. Thayer, A. T. Stentz, and P. Rander, "Vegetation detection for mobile robot navigation," Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-04-12, February 2004.
- [14] C. A. Shull, "A spectrophotometric study of the reflection of light from leaf surfaces," *Bot. Gazette*, vol. 87, no. 5, pp. 583–607, 1929.
- [15] C. Ünsalan and K. L. Boyer, "Linearized vegetation indices based on a formal statistical framework," *IEEE Trans. on Geoscience and Remote Sensing*, vol. 42, no. 7, pp. 1575–1585, 2004.
- [16] J. Qi, A. Chehbouni, A. R. Huete, Y. H. Kerr, and S. Sorooshian, "A modified soil adjusted vegetation index: Msavi," *Remote Sensing of Environment*, vol. 48, pp. 119–126, 1994.
- [17] R. N. Clark, G. A. Swayze, R. Wise, K. E. Livo, T. M. Hoefen, R. F. Kokaly, and S. J. Sutley, "USGS digital spectral library splib05a, USGS open file report 03-395," 2003.
- [18] A. Kelly, A. Stentz, O. Amidi, M. Bode, D. Bradley, A. Diaz-Calderon, M. Happold, H. Herman, R. Mandelbaum, T. Pilarski, P. Rander, S. Thayer, N. Vallidis, and R. Warner, "Toward reliable off road autonomous vehicles operating in challenging environments," *International Journal of Robotics Research*, vol. 25, no. 5-6, pp. 449–483, 2006.
- [19] K. Nigam, J. Lafferty, and A. McCallum, "Using maximum entropy for text classification," in *IJCAI-99 Workshop on Machine Learning for Information Filtering*, 1999, pp. 61–67. [Online]. Available: citeseer.ist.psu.edu/nigam99using.html