Terrain Perception for DEMO III

P. Bellutta, R. Manduchi, L Matthies, K. Owens, A. Rankin

Jet Propulsion Laboratory
California Institute of Technology
Pasadena, CA 91109
Tel. (818)354-1897 – Fax (818)393-4085

[manduchi,lhm,bellutta,kowens,arankin]@telerobotics.jpl.nasa.gov

Abstract

The Demo III program has as its primary focus the development of autonomous mobility for a small rugged cross country vehicle. Enabling vision based terrain perception technology for classification of scene geometry and material is currently under development at JPL. In this paper we report recent progress on both stereo-based obstacle detection and terrain cover color-based classification. Our experiments show that the integration of geometric description and terrain cover characterization may be the key to enabling successful autonomous navigation in cross-country vegetated terrain.

Keywords: Autonomous navigation, obstacle detection, stereo vision, color classification.

1 Introduction

The ability to navigate autonomously in vegetated, off-road terrain may be the most critical technology needed for Unmanned Ground Vehicles (UGV) today. The DEMO III program [7] has as primary focus the development of autonomous mobility technology for a small ground vehicle over cross-country, rugged terrain. The Experimental Unmanned Vehicle (XUV) must be able to drive autonomously at speeds of up to 40 mph on roads, 20 mph off road by day, and 10 mph off road by night or in foul weather conditions. The Autonomous Mobility (AM) sensor suite of the DEMO III XUV includes a LADAR, a radar, and color and infrared (FLIR) cameras (visible on the front of the vehicle in Figure 1). The Jet Propulsion Laboratory (JPL) is developing technology for the vision-based *terrain perception* part of the DEMO III system, which is a key component to enabling autonomous mobility in cross-country vegetated terrain. The JPL terrain perception system interfaces with the world model system developed by NIST, which integrates data from the other sensors (radar and laser rangefinder).

One of the challenges posed by DEMO III is the integration of geometric terrain description with terrain cover characterization. Geometric information is used to recover the 3-D scene structure. It allows the detection of positive obstacles (that is, non-traversable areas that

should be avoided), negative obstacles (such as ditches or holes), and terrain surface characteristics to help determine the most effective yet safe velocity for the traversal. While in typical urban environments the geometry description alone is sufficient to characterize the traversability of a path, terrain cover characteristics should also be taken into account for off-road navigation. For example, assume that the range sensors detect an obstacle 40 cm high in front of the vehicle. This obstacle should be steered around if it is a rock, a log, or some other impenetrable material. However, if this obstacle is just a small bush or a tuft of grass, the vehicle may safely run over it, with obvious advantages in terms of driving efficiency. An extreme case is given by navigation in a field of tall grass. The elevation map will represent the scene as a basically horizontal surface above the ground level (where the vehicle lies); that is, as a big obstacle in front of the vehicle. Indeed, planning a path through tall grass is possible only if the navigation system realizes that the *load-bearing surface* lies below the top grass level. It is apparent that only by integrating the geometry description with terrain cover characterization will a robot be able to navigate in such critical conditions.

Figure 1: The XUV DEMO III.

In this paper we report recent progress on the development and implementation of algorithms for terrain perception at JPL. Only passive sensors (color and FLIR cameras) are considered here. In the area of geometric terrain analysis, we demonstrate the performance of simple real-time algorithms for positive and negative obstacle detection both by day and night. These techniques work directly on the range image, computed by real-time stereo triangulation. Due to the extreme terrain conditions considered in the DEMO III program, it is not advisable to analyze stereo images with respect to a horizontal *horop-*

ter, i.e., to a plane fitted to the ground surface, as often proposed for navigation on paved roads. Indeed, fitting an horizontal horopter on a very rugged or vegetated terrain surface with large variations of the surface height is a questionable operation. Furthermore, tall obstacles (such as a hanging branch) will produce large disparity values in stereo pairs warped with respect to a ground plane horopter, and may not be detected if a reduced size search window is used.

Terrain cover characterization is performed on color images using pattern recognition techniques. The proposed color-based classification algorithm is very fast and robust to changing illumination conditions. The classification is completely registered to the range data; subsequent reasoning may determine the actual traversability of obstacles detected by range sensing, toward the final goal of estimating the actual load bearing surface in front of the vehicle.

Figure 2. *The JPL HMMWV.*

Our algorithms are tested on the JPL HMMWV (High Mobility Multi-Wheeled Vehicle) shown in Figure 2. This testbed vehicle is provided with color camera pairs (Hitachi HV-C20 and Sony DXC 9000), a Pulnix TM9701 monochrome camera pair installed on a pan-tilt head, and an Amber Radiance 3-5 micron FLIR camera pair. The computing system includes a number of Motorola CPU cards (MVME172, 2400, 2700) running VxWorks, which are used for iris control, image acquisition, stereo vision, obstacle detection, velocity control, gaze control and terrain cover classification. The vehicle is equipped with a NovAtel RT20 differential global positioning system (DGPS) receiver, which yields a 20 cm horizontal circular error probable (CEP) accurate positioning solution at 10 Hz. It also contains a Honeywell Modular Azimuth Positioning System (MAPS), an inertial navigation system (INS) that produces position, orientation, and velocity data at 25 Hz. The INS and DGPS solutions are integrated with an external Kalman filter. The vehicle has robotically actuated throttle, brake, steering and gear selection.

This paper is organized as follows. In Section 2, we describe the application of our real-time stereo system to geometric terrain representation and obstacle detection. In Section 3 we introduce our color-based terrain cover classification system.

2 Geometric Terrain Representation

Resolution requirements for obstacle detection in DEMO III have been derived in [1] and [5], based on the procedure described in [5]. The nominal obstacle sizes for the DEMO III XUV are 35 cm for "positive" obstacles and 70 cm for "negative" obstacles. The study in [5] has shown the need for a stereo system mounted on a pan/tilt device, which allows coverage of a wide field of regard (around 75 degrees) while maintaining a narrow instantaneous field of view (IFOV) of about 1 mrad/pixel. For what concerns FLIR cameras to be used for night vision, [6] derives the requirements in terms of signal-to-noise ratio and exposure time (to avoid motion-induced blurring). It is maintained in [6] that in order to fulfil the DEMO III requirements, a cooled FLIR stereo pair should be used. The current XUV vehicle has a pair of 3-5 micron Merlin cooled FLIR cameras, produced by Indigo Systems.

The JPL stereo system is able to produce disparity maps at a rate of 6 Hz on a single PowerPC 750 microprocessor with 7x7 pixel correlation window and search range of 40 pixels on images at resolution level 1. (Resolution level 0 corresponds to the full 640x480 pixel image, level 1 corresponds to half size, and so on). We also implemented a "dual resolution" scheme where disparities are computed on a subwindow of attention at resolution level 0. Currently, we are converting the system to run on a dual PowerPC G4 processor, which includes 128-bit vector instructions, with an expected performance increment of 4 times.

2.1 Obstacle detection

Once a range image is produced, a set of range-based obstacle detectors are executed on each image column, looking for gaps and discontinuities in the range data that indicate impassable regions. Impassable regions are classified as either negative or positive obstacles. Negative obstacles include ditches, holes, sudden drop-offs, and steep down grades (see examples in Figure 5-7). Positive obstacles are any upward protrusion out of the ground plane steep enough to be impassable or cause a tip-over hazard. Our system has separate detectors for negative and positive obstacles.

Negative obstacles. The height profile of one column of elevation data from Figure 5 is shown in Figure 3. A characteristic gap in the range data, followed by a vertical edge is usually present where a ditch or hole exists. The length of the vertical edge at the far side of a ditch or hole is a function of the obstacle's range from the sensor (see
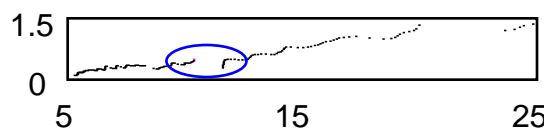

Figure 3: Elevation profile along one column of the image of Figure 5, showing a 60 cm wide negative obstacle (encircled). All units are in me-

Figure 4). If a gap threshold and a range-dependent vertical edge threshold are exceeded, that image pixel is labelled a negative obstacle pixel.
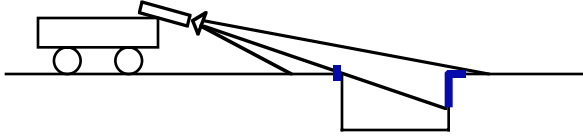

Figure 4: Negative obstacle detection. Highlighted are the visible front edge and back wall of the negative obstacle.
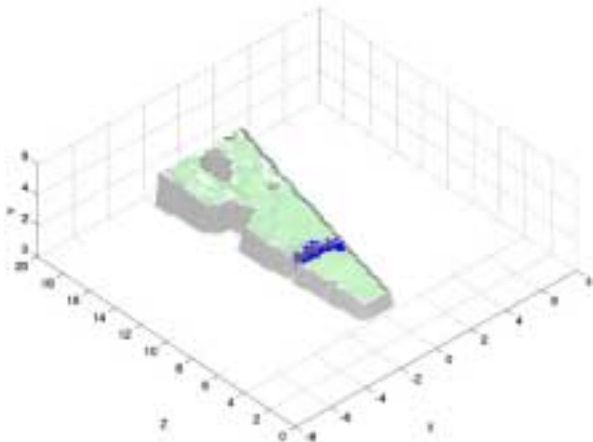

Figure 5. Top: left CCD image of a stereo pair acquired at the Perryman facility, Aberdeen Proving Ground. A 60 cm wide ditch is visible at a distance of 10 meters (showing as a horizontal feature in the image). Bottom: elevation map computed via stereo triangulation. Points corresponding to actual range measurements are drawn in pale green; surface patches drawn in shaded white have been interpolated from such measurements. All units are in meters; the vehicle is located at position (0,0,0)). The edges of the ditch (detected as a negative obstacle by the algorithm of Section 2.1) are marked in blue; the occluded surface behind the front edge of the negative obstacle is marked in purple.

Examples of negative obstacle detection with CCD and FLIR cameras are shown in Figure 5 to 7. Note that only the front edge and part of the back edge of a negative obstacle are visible, and they have been marked in blue in the elevation maps. The area marked purple behind the front edge of a negative obstacle corresponds to occluded points for which range measurements are not available.
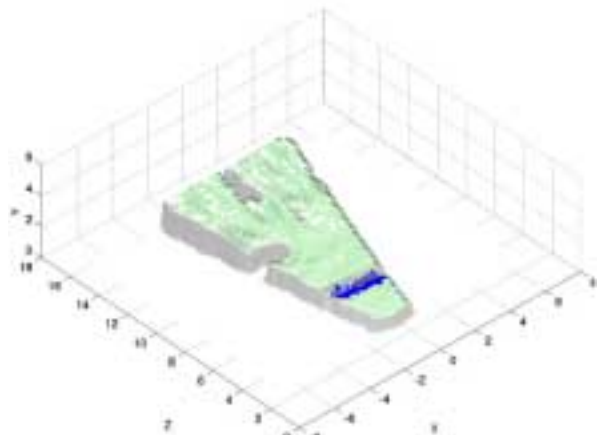

Figure 6. See caption of Figure 5. In this case, a 60 cm wide ditch was at a distance of 6 meters in front of the vehicle, and the image is from a pair of FLIR cameras.

Positive obstacles. Positive obstacles are detected by checking for upward slanted edges in the range data. If the height of an upwards-slanted edge exceeds a height threshold and the slope of the edge exceeds a slope threshold, the image pixel is labelled as a positive obstacle. Once all columns are processed for negative and positive obstacles, separate blob filters are passed over the obstacle image to remove isolated detection that is likely to be false alarms. An obstacle list is then passed to a map management module that populates an obstacle map. Examples of positive obstacle detection are shown in Figure 8 to 10 for both CCD and FLIR images.

2.2 Sensor pointing

A pointing subsystem is used to control the pan/tilt head holding the narrow FOV stereo cameras. The user can specify a goal position up to 50 meters away. An obstacle-avoiding planner searches the obstacle map for the shortest distance two-dimensional path constructed primarily of clothoid segments. This path is projected onto the range image to determine the elevation of each way-

point. A gaze control algorithm filters and smooths the resultant 3D path and points the pan/tilt cameras at the path the vehicle will be commanded to travel.

A position on the path is chosen by applying a look-ahead metric that is dependent on the vehicle's current speed. The look-ahead distance is larger than the predicted distance required to bring the vehicle to a stop on cross-country terrain. The goal of the current gaze control algorithm is to keep the look-ahead position centered in the vision sensor's field of view. The gaze control algorithm also sends a set of points to a low-level image processing module that bound the path near the stopping distance. These points are used to determine where to locate a high-resolution window of attention within which negative obstacles will be searched for.

Figure 7. Top: left image of a stereo pair acquired in the area surrounding JPL. The road terminates abruptly on a crag over a field of tall grass. Bottom: elevation map computed via stereo triangulation (see caption of Figure 1 for the description of the color code). All units are in meters; the vehicle is located at position (0,0,0)). The steep drop-off has been detected in spite of the difficult light conditions.

3 Terrain Cover Classification

While the elevation and obstacle maps give a complete geometric representation of the scene, they tell us only one part of the story. Once an obstacle has been detected on the grounds of geometric analysis, the next step is to

characterize it in terms of its traversability properties. Our approach to this problem is to classify each point of the scene as a "material class" in a predetermined family, and to infer its traversability characteristics from the combination of this information and range data. This information fusion step may take the form of a rule-based system. For example, one simple set of rules could be the following: "if the detected obstacle is up to 25 cm tall, drive through it with unchanged velocity; if it is between 25 and 35 cm and it is classified as a rock or other impenetrable material, slow down to a safer velocity; if it is more than 35 cm tall, reduce velocity and drive through it only if it is classified as penetrable vegetation".

Figure 8. Top: left image of a FLIR cameras stereo pair acquired at the Aberdeen Proving Ground. The white areas are rocks lying on the ground. Bottom: elevation map computed via stereo triangulation. Points corresponding to actual range measurements are drawn in pale green; surface patches drawn in shaded white have been interpolated from such measurements. Positive obstacles, detected by the algorithm of Section 2.1, are marked in red. All units are in meters; the vehicle is located at position (0,0,0)).

The choice of the terrain cover class taxonomy depends mainly on two factors: which classes are useful for autonomous navigation, and which classes can actually be detected with the available sensor information. The right balance must be struck between the desire for high descriptiveness (which would lead to a dense taxonomy)

and the need for robust classification (which favours fewer classes).

In our implementation we perform terrain classification based on color. Other visual features (such as shape and texture) are the object of current research [5]. A basic set of classes, used for the experiments in this paper, are: green vegetation, dry vegetation, soil/rock, and outliers (i.e., anything that cannot be safely classified into any of the previous classes). The classification algorithm is based on Bayesian assignment. The class likelihoods are represented using a mixture-of-Gaussian model, and the parameters of the models are estimated from training data using the Expectation Maximization algorithm. A support region for the model likelihood is estimated, and all color vectors that are outside such support are classified as outliers. An advantage of the Bayesian approach is that the classification is expressed in terms of posterior class probabilities, which facilitates fusion with data from other sensors.

A well-known problem of color classification is that the measured color has a spectrum which depends both on the spectrum of the illuminant and on the reflectivity characteristics of the viewed surface [2]. Ideally, only the last term is present; in practice, we have to deal with the changing illumination conditions. A typical preprocessing step is lightness normalization. While this procedure is known to work well for indoor imagery, for outdoor images we have not found experimental evidence supporting its utility. One possible explanation is that a surface in the shadow is actually illuminated by a different spectrum than a surface in direct sunlight. In other words, not only the lightness is different for the two surfaces, but also the spectral composition of the reflected light. In fact, we have verified that most of the chromatic variations of light reflected from any of the classes of the system is well represented by the mixture-of-Gaussian class-likelihood profile. However, in order to deal with large chromatic variation (corresponding to different times of the day and of the year and to changing weather), it is advisable to use an external reference onto which to calibrate the color classifier.

The second row of Figure 9 and 10 shows the results of color classification for those pixels which correspond to points less than 50 meters from the camera. Note in passing that we didn't provide a specific "sky" class; therefore, all pixels corresponding to the sky are (correctly) classified as outliers. Also, note that the metal pipes visible in Figure 10 have been (incorrectly) classified as soil/rock. This is because in the set used to train the classifier there were a number of images with rocks having a very "metallic" color. Since metal objects are typically non-traversable, there was no point in refining the material class taxonomy to account for such a specimen.

For each stereo pair, positive obstacles have also been detected, and the corresponding class color code overlayed in the last row of the figure. In the case of Figure 10 (right), the obstacles have been classified as rock/soil, and therefore are not traversable. In the case of Figure 9 (left), obstacles on the ground have been classi-fied as dry vegetation, and in fact they are just small (traversable) bushes

Acknowledgements

References

[1] Demo III Experimental Umanned Vehicle (XUV) Program: Autonomous Mobility Requirements Analysis. Technical Report, Robotics Systems Technologies, 1998.

[2] G.H. Healey, S.A. Shafer. KL.B. Wolff, (eds.), *Physics-based Vision: Principles and Practice*, Jones and Bartlett Publishers, London, England, 1992.

[3] R. Manduchi. Bayesian Fusion of Color and Texture Segmentations. *Proceedings of the IEEE International Conference on Computer Vision*, Kerkyra, September 1999, 956-962.

[4] L. Matthies, P. Grandjean. Stochastic Performance Modeling and Evaluation of Obstacle Detectability with Imaging Range Sensor*. IEEE Transactions on Robotics and Automation, Special Issue on Perception-based Real World Navigation*, 10(6), December 1994.

[5] L. Matthies, T. Litwin, K. Owens, A. Rankin, K. Murphy, D. Coombs, J. Gilsinn, T. Hong, S. Legowik, M. Nashman, B. Yoshimi. Performance Evaluation of UGV Obstacle Detection with CCD/FLIR Stereo Vision and LADAR. *Proceedings of the IEEE Workshop on Perception for Mobile Agents*, Santa Barbara, CA, June 1998.

[6] K. Owens, L. Matthies. Passive Night Vision Sensor Comparison For Unmanned Ground Vehicle Stereo Vision Navigation. *Proceedings of the IEEE Workshop on Computer Vision Beyond the Visible*. Fort Collins, CO, June 1999.

[7] C.M. Shoemaker, J.A. Bornstein. The Demo III UGV Program: A Testbed for Autonomous Navigation Research", *Proceedings of the IEEE International Symposium on Intelligent Control*, Gaithersburg, MD, September 1998.
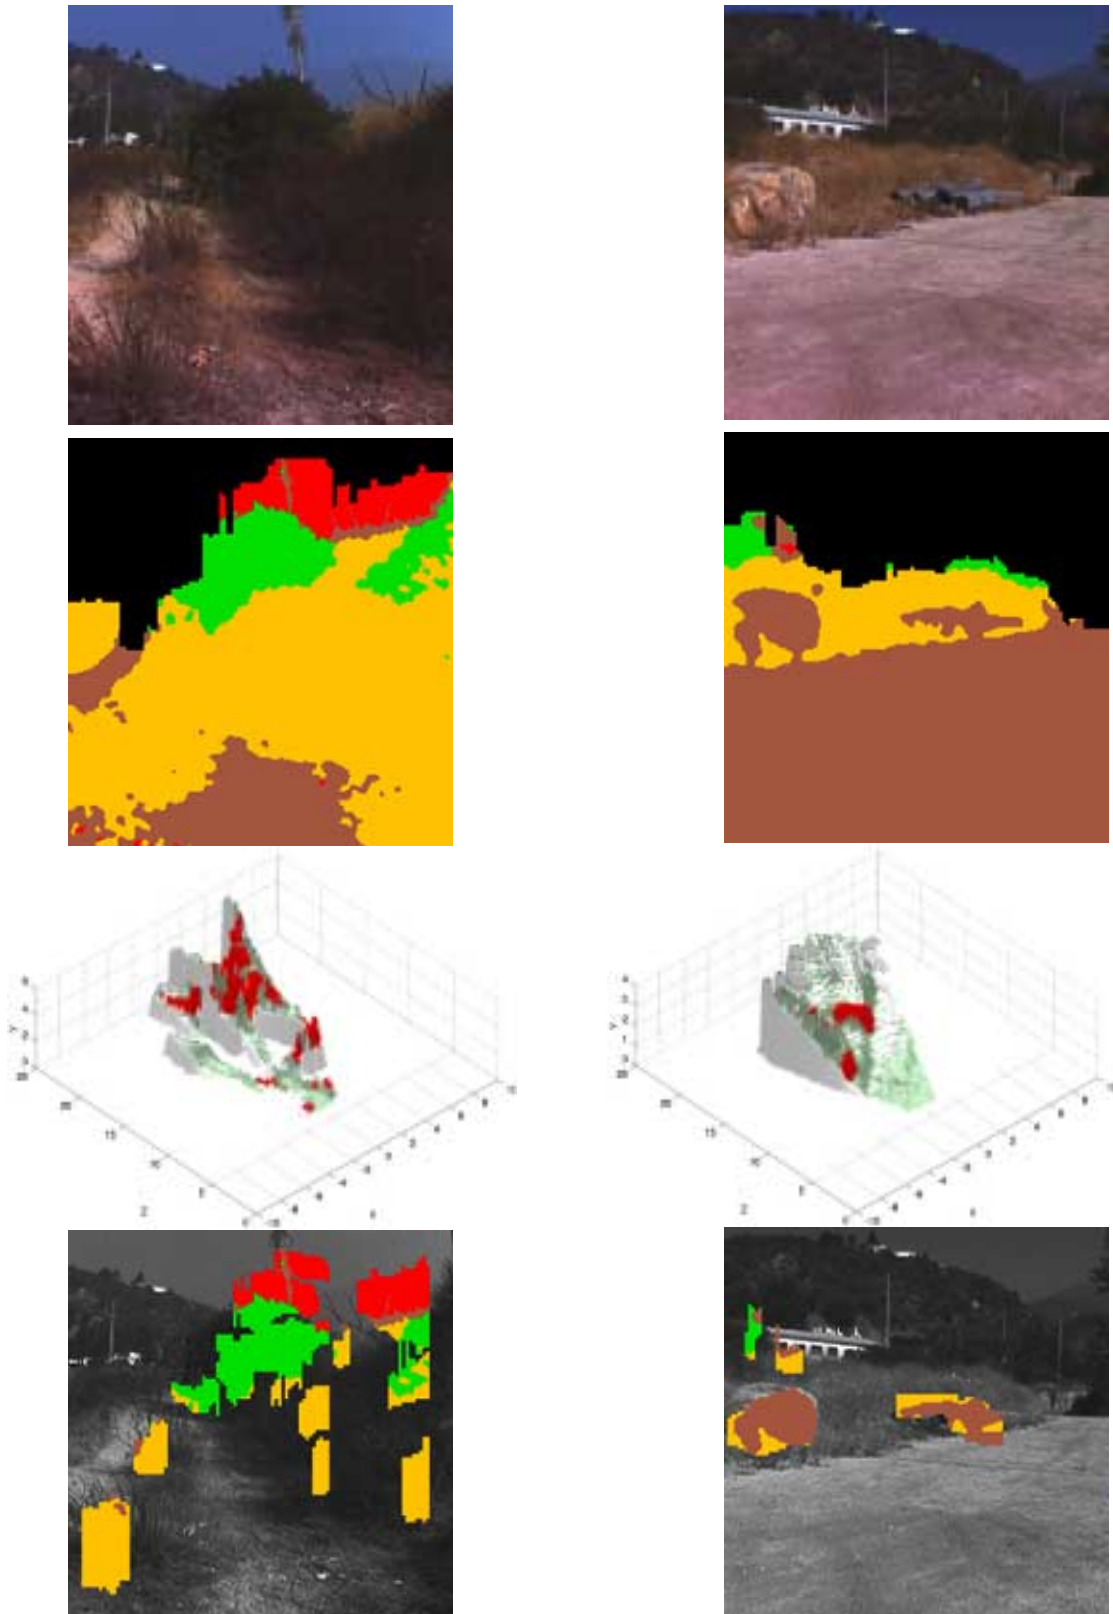
Figure 9 (left) and 10 (right)*: Top row: left image from a stereo pair acquired at JPL. Second row: terrain cover classification, color-coded as follows: brown: soil/rock; yellow: dry vegetation; green: green vegetation; red: outlier. Third row: elevation map computed via stereo matching (see caption of Figure 8 for the description of the color code). All units are in meters; the vehicle is located at position (0,0,0)). Fourth row: obstacle map overlay on the image plane, where the obstacles are color-coded according to the corresponding terrain cover class.*