

The Single Frame Stereo Vision System for Reliable Obstacle Detection used during the 2005 DARPA Grand Challenge on TerraMax™

Alberto Broggi, Claudio Caraffi, Pier Paolo Porta, Paolo Zani
VisLab-Dipartimento di Ingegneria dell'Informazione,
Università di Parma, Parma I-43100, Italy
{broggi, caraffi, portap, zani}@ce.unipr.it

Abstract—Autonomous driving in off-road environments requires an exceptionally capable sensor system, especially given that the unstructured environment does not provide many of the cues available in on-road environments. This paper presents a variable-width-baseline (up to 1.5 meters) single-frame stereo vision system for obstacle detection that can meet the needs of autonomous navigation in extreme environments. Efforts to maximize computational speed—both in the attention given to accurate and stable calibration and the exploitation of the processors MMX and SSE instruction sets—allow a guaranteed 15 fps rate. Along with the assured speed, the system proves very robust against false positives. The system has been field tested on the TerraMax™ vehicle, one of only five vehicles to complete the 2005 DARPA Grand Challenge course and the only one to do so using a vision system for obstacle detection.



Fig. 1. The TerraMax™ vehicle

I. INTRODUCTION

In recent years the DARPA Grand Challenges (<http://www.grandchallenge.org/>) have given a good deal of attention to autonomous vehicle operation in off-road environments. One key component of any autonomous system is its perception system: the vehicle must first sense its surrounding environment before it can determine how to safely drive within that environment.

Performance in the inaugural 2004 Challenge gave a clear indication of how “Grand” a Challenge DARPA had proposed. No vehicle could travel even 10 miles of a 142-mile course. Autonomous operation is quite a challenge, but the unstructured off-road environment confounds basic environmental perception. Color has no clearly defined meanings or interpretations. Slopes frequently change. Roads are not flat. No lane markings exist, in fact, boundaries can be defined by severe negative obstacles. All of this gets coupled with high vehicle speeds (up to 45 mph): the Challenge is not simply to complete the course but to do it faster than anyone else.

This paper presents the stereo vision obstacle detection system that the TerraMax™ autonomous vehicle (Fig. 1), used in the 2005 DARPA Challenge. Oshkosh Truck Corporation, Rockwell Collins, and University of Parma joined forces to form Team TerraMax™ [22]. Parma developed the vision system to meet the following requirements:

- supply valuable information to the path planner;
- reliably detect obstacles and estimate their positions;
- avoid false detections.

This work was supported by Oshkosh Truck Corporation

TerraMax™ used laserscanners as well as the vision system, however the limited testing time did not allow development of algorithms to fuse vision and laserscanner data at a low level. Instead, each sensing system perceived obstacles and built its own real world coordinate map to send to the path planner system (like in [8]). GPS, inertial sensors, and map databases rounded out the tools TerraMax™ used to sense and understand its environment.

The paper is organized as follows. In section II we present the hardware architecture implemented on TerraMax™ vision system. In section III we introduce the principles that motivated our choices during the algorithm design. Section IV presents the ground information extraction capabilities of the V-disparity image, while section V describes the core of our obstacle detection algorithm. In section VI we show some results of our algorithm. Conclusions are drawn in section VII.

II. SYSTEM SETUP

Human beings rely only partially on stereo vision while driving; using motion analysis and accumulated experience, we can easily drive with a closed eye. This capability appears to be portable on a computer platform in urban environments, where it is easier to find known image patterns to be tracked (as in [17]), but the same task seems more difficult in less structured environments.

In some cases it is then appropriate to give some “advantage” to the computer platform and exploit its computational abilities. For example, we can give it better eyes, e.g. using

infrared cameras when the conditions suggest it. Or, as done in this case, we can use a wide baseline¹ on the stereoscopic system.

Without any a priori knowledge (apart from the calibration measurements) wide baselines allow accurate distance estimation even for very distant obstacles. However, at close range wide baselines give very different viewing angles and make stereo matching difficult. TerraMaxTM used a three-camera variable-baseline system to allow for good precision and efficient computation at a wide range of viewing distances. The cameras are installed on a rigid bar over the vehicle hood (see Figure 2).



Fig. 2. TerraMaxTM cameras setup

Selecting two cameras at a time, the system can switch through the different baselines (0.5, 1 and 1.5 m). During the DARPA Grand Challenge TerraMaxTM selected the viewing baseline based on vehicle speed. Higher speeds required greater sensing distances and thus wider baselines. The graph of Figure 3 shows the growth of the distance estimation error with a growing distance for the different baselines. No sub-pixel disparity analysis is considered.

The 3 cameras feature a Bayer pattern output and mount low distortion 6mm optics. In the DARPA Grand Challenge environments, not much information is carried by color so this feature is not taken into consideration during stereo matching: the images are grayscaled after acquisition. Synchronization is supplied directly by the cameras through the firewire bus.

¹The baseline is the distance between the two cameras that form the stereoscopic system




Fig. 3. Estimation of the mean error of the distance measure depending on the distance and on the chosen baseline (BL), plotted considering cameras parameters and the image resolution.

During calibration, every effort is made to obtain a standard form, i.e. a perfect alignment and co-planarity of the image planes. The goal is to reduce the number of software transformation needed to rectify the images², avoiding problem complexity growth. In their work, Nedevschi and al. [2] comprehensively explain the importance of this step. Once the best achievable alignment is reached, the remaining calibration error (e.g. misalignment) can be measured. Subsection IV-B will show how V-disparity image analysis may help in this task. Small calibration errors can be compensated by (small) image shifting, rotation, and magnifying.

III. ALGORITHM OVERVIEW

The hardware setup of our system requires the camera pairs to have a horizontal baseline in order to perceive vertical obstacles. As [2] states, in these conditions it is reasonable to perform localization of homologous points on vertical edges obtained by means of a Sobel filter. Furthermore, a Sobel filtered image does not suffer from stereo matching problems caused by brightness differences between cameras. Unfortunately image areas with constant appearance (low contrast or lack of texture) do not have many (vertical) edges and thus do not provide sufficient information for stereo reconstruction via a simple DSI (Disparity Space Image) computation. In such cases, image segmentation can still provide information for stereo reconstruction, but it does so at a high computational cost (e.g. in [20] it takes between 1 and 7 seconds per image). This problem is simply one specific instance of the general phenomena noted in [8]: it is difficult to detect obstacles of any kind with a single approach. The solution to the general problem is to use more than one sensor system. In this specific example, the laserscanners mounted on TerraMaxTM can easily detect wide textureless obstacles, so the vision system design need not consider obstacles without edges. Exploiting the complementary strengths of vision and laserscanners can provide a complete real-time sensing solution.

Full 3D stereo reconstruction is currently receiving a great deal of attention by the scientific community; nevertheless, computational time aspects are not always addressed (for example, the algorithm described in [19], that reaches very good performances in term of 3D reconstruction, takes 40 seconds to process a single pair of images). Obviously, this clashes with our real time requirements.

As a real-time alternative, this paper proposes a two-step approach to full 3D stereo reconstruction. The first step is a global one, using V-disparity image [10], [11] analysis to obtain basic information about the position and shape of the ground (taking slope changes into consideration). This information helps to provide some structure—which off-road environments would otherwise lack. The second step detects obstacles with a DSI algorithm that uses information from the first step to both reduce the region of interest prior to calculation and to filter the results after calculation.

²In this context, "to rectify images" means geometrically transform them as they were taken by cameras in standard form

In off-road environments vehicle oscillations caused by terrain bumps cannot be neglected; to overcome this problem, image stabilization is necessary. In [1] we introduced a system that relies on V-disparity image properties to stabilize images. In such environments other approaches are also available [18], but we did not take into consideration time-correlation-based stabilization: currently we are aiming to extract as much information as possible from a single stereo image pair.

Using the information about the ground, the second step of our algorithm computes a DSI and searches for obstacles. The obstacles, mapped in real world coordinates, are then grouped following a neighbor rule. In their work [3], [4], Dang and others make a further step dividing adjacent obstacles by means of optical flow. DARPA Grand Challenge context is demanding, but can be considered static: that is the reason why we did not focus on obstacles motion. Obstacle tracking is delegated to an external system that relies on inertial data. Obstacle classification is left for further development as well.

The similarity measurements are made via a SAD (Sum of Absolute Differences) computation³. This allowed code optimizations to exploit processors MMX and SSE instruction set [21], as seen also in [2], [7], [12]. The right image is chosen as reference, and homologous windows are searched in the left image. In this way, disparities grow as the search window shifts from left to right on the left image.

IV. V-DISPARITY IMAGE ANALYSIS

A. Computation

Figure 4 shows an example of V-disparity image. V-disparity images are 3D graphical representations of the similarity measures between left and right image rows (v coordinate) depending on the disparities (d coordinate) used to compare them. Brightness, that is used as third dimension, is directly proportional to the similarity measured. In the example of Fig. 4, the slanted segment in the V-disparity image is caused by the linearly changing maximum similarity disparity among ground components. This segment, called the “ground correlation line” in this study, contains the information about the position and the shape of the ground in the images.




Fig. 4. A V-disparity image example on synthetic images. The frame of reference is plotted. In this case the similarity measures are computed on edges.

³For a complete discussion about matching methods and similarity measurements see [14], [15].




Fig. 5. Ternarized images and their V-disparity image. In the former, red is negative Sobel phase and green positive Sobel phase. In the latter, the red line lays on the zero disparity value.

The similarity measure of image rows can be computed in many ways. In their works, Labayrade [10], [11] and others compute a fast DSI; then they obtain the similarity collecting, for each image row, the number of matches for each disparity. In [1] we introduced a slightly different approach to V-disparity image computation: we consider comparison windows that correspond to the image lines. That is, we measure similarity among windows 1 pixel high and of the same width of the image. Making this choice we considered that ground has different disparities at different heights in the image, so it is preferable to maintain windows as short as possible.

In this step we aim to extract ground information from images. In off-road environments, a good method to do this is to apply a vertical Sobel mask to the image, and to map the obtained values in a ternary domain (-1,0,+1) obtaining the so-called ternarized image. This approach is thresholdless and takes advantage of a property of the ground that is usually verified: it takes up the largest part of the image. Fig. 5 shows an example of ternarized images and of the derived V-disparity image.

Code optimization allows to compute the whole V-disparity image in less than 6 ms on a Pentium IV @ 2.8GHz processor system.

Considering a flat ground, the behavior of the ground correlation line in case of vehicle pitch variations is to oscillate parallel to itself. Knowing cameras static calibration the algorithm can compute a set of candidate ground correlation lines by varying the pitch angle. By collecting the V-disparity image values along these lines and choosing the best one, it is possible to estimate the pitch at time of acquisition (see Fig. 6).

The information obtained at this step is used in the subsequent obstacle detection to locate a region of interest in the image where to compute the DSI. It can also be forwarded to other sensors (like laserscanner [23] and vision path detector [24]) in order to improve their environment

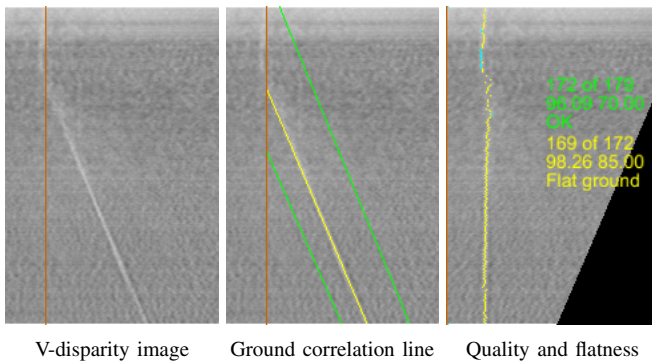


Fig. 6. V-disparity image analysis. In the center image, the ground correlation line is plotted in yellow. Green lines bound the search space. In this case, the ground correlation line is closer to the left bound, so that the horizon height in the image (which corresponds to the point where red and yellow line intersect) is below its position when the vehicle is static. The algorithm can then state that the vehicle is pitching up. In the right image, set A points (see section IV-B) are plotted in yellow and set B points in cyan. Green numbers describe V-disparity image quality analysis. 179 is the number of maxima below the horizon, that are then taken into consideration during the subsequent analysis. 172 is the sum of set A and set B cardinalities. 96.09 is the percentage ratio between 172 and 179 and 70.00 is the threshold applied to distinguish between a good quality image and a bad one. Yellow numbers describe flatness analysis. Again, 172 is the sum of set A and set B cardinalities, while 169 is set A cardinality. 98.26 is the percentage ratio between 169 and 172 and 85.00 is the threshold applied to decide about ground flatness. In this case quality is high and road is flat.

knowledge. The horizon estimation can also be useful to select an appropriate image area to be used to tune the cameras automatic gain control.

B. Quality and road flatness measurement

A verification of the pitch estimation result would be appreciable because the V-disparity image quality may be insufficient to reliably detect the ground correlation line. Possible causes of low quality V-disparity images are underexposed and overexposed images, cameras misalignments, textureless terrains, and strong roll angles. Furthermore, in unstructured environments the flat road assumption is too strong; anyway V-disparity image contains information about slope changes, since slope changes bend the ground correlation line.

In the following step the positions of the V-disparity image maxima in each row are considered. A maximum close to the ground correlation line increases the probability of facing a flat ground. Maxima are classified in three categories. Set A is defined as the set of maxima falling within a limited range centered onto the ground correlation line. All the other maxima not belonging to set A, are either rejected or included in set B. In particular, a maxima is rejected when it is isolated (i.e. the number of maxima that can be found in a given small neighborhood is lower than a threshold). To ease the computation of set A and set B, we warp [16] the V-disparity image using the ground correlation line disparities, so that the points of set A are placed in a vertical strip.

The quality of the V-disparity image is estimated in inverse proportion to the number of isolated best values. The ground flatness is estimated as $\frac{|set A|}{|set A|+|set B|}$, where $|set A|$ is the

cardinality of set A. Figure 6 shows an example of flat ground, while Figure 7 shows a non flat ground.

A complete road profile reconstruction through estimation of the shape of the ground correlation line is currently under development. Nonetheless, information about the number of maxima for each disparity allows to implement a slope filter used in the obstacle detection step, to get rid of false positives caused by increasing slopes. In structured environments, [2], [13] rely on slightly changing slopes and on a small number of wrong matches in the DSI image to perform this task. In particular, the triangulation of the wrong matches must not lead to 3D points whose coordinates lie underground. Both conditions (smooth slopes and small number of wrong matches) can fail in off-road environments. A reliable road profile reconstruction through V-disparity image analysis will have to use time correlation in order to overcome the low quantity of information present in unstructured environments.

Ternarized images are usually a good method to extract ground information. Nonetheless, in some cases the terrain texture is not enough appreciable (e.g. smooth asphalt roads) and even the use of ternarized images cannot provide sufficient information to the V-disparity image. Anyway, when the quality of the V-disparity image is low, it could be recomputed using information other than ternarized images (e.g. road markings extraction in the case of asphalt roads).

V-disparity image quality analysis allows to estimate also the precision of calibration: in particular regarding the cameras alignment. Fig. 8 shows a sequence of 8 V-disparity images simulating a growing cameras misalignment, obtained shifting downward the left image of Figure 5 of 1 pixel steps.

During the processing of an image sequence V-disparity image quality may change; if this drift remains for a given amount of time, then it can be inferred this is due to a cameras misalignment, caused by a physical movement of the cameras. Although not yet implemented, a software re-alignment (through images shift) can be applied to every incoming image for small angle corrections.

Strong roll angles cause the ground correlation line to appear blurred. Luckily, the TerraMaxTM vehicle characteristics limited lateral oscillation, so that we did not register strong roll angles variations. We are currently considering to try quality analysis on different V-disparity images computed on images rotated at different degrees, to discover if the original images were acquired with a non null roll angle. Other approaches [13] are to be considered.

C. Limits

A few cases are not addressed yet. In general the approach fails when the terrain does not feature a prevailing slope, and when the assumption that ground should occupy the largest image part is not met (i.e. in front of wide obstacles). A possible solution to this problem should be to compute the V-disparity image only considering image regions classified as road by a former algorithm.




Fig. 7. V-disparity image analysis in a non flat ground scene. Set A points are plotted in yellow and set B points in cyan.

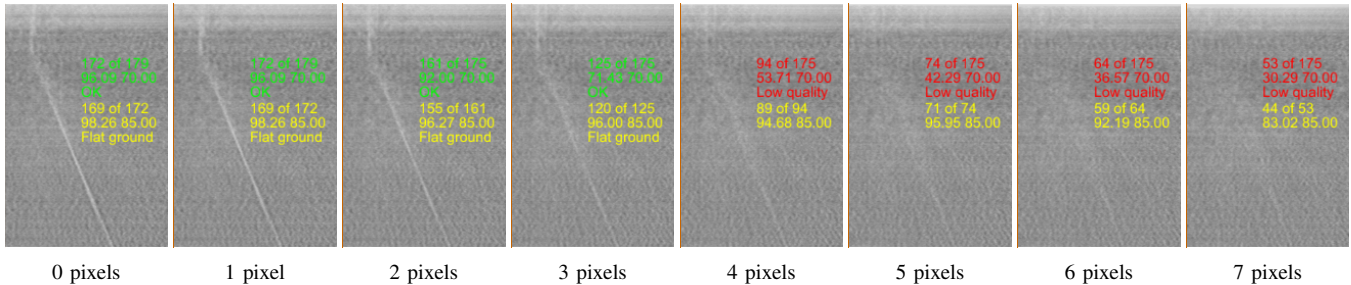


Fig. 8. V-disparity image quality analysis for different grades of artificial misalignment of images of Fig. 5. The V-disparity images computed in correspondence to shifts up to 3 pixels present a sufficient quality to allow a reliable detection of a flat ground.

V. OBSTACLE DETECTION

A. DSI computation

Once the global analysis of the images has finished, the algorithm proceeds computing a region of interest using pitch estimation. That is, in order to reduce the computation cost, real world regions too close to the vehicle for a successful stereo matching and image areas well above the estimated horizon are not processed for obstacle detection.

A DSI is then computed, using a $4 \text{ rows} \times 3 \text{ columns}$ confrontation window. The chosen window is narrow in order to allow the detection of thin obstacles and short to ease matching of ground features (as stated in subsection IV-A). The DSI is not computed in textureless regions.

The disparity search range is centered considering the expected ground disparity depending on the v coordinate in the image. Thanks to the code optimization, and the resulting reduction of the computational time, we have been able to increase our DSI resolution from 106×80 to 320×240 pixels. The DSI is computed within 15 ms. In the mean-time 3D world coordinates are computed via stereo triangulation, so that the DSI encodes also a range information. The wide stereo baselines provide a sufficient accuracy in depth estimation, thus not requiring sub-pixel disparity measure.

B. Obstacles search in the DSI

Many methods have been studied to extract obstacles from a DSI/Range Image. In road environments, [2], [13] exploit their knowledge about ground surface in order to delete ground matches in the range image: the remaining points are grouped into obstacles. Anyway, as discussed earlier, the

definition of a “ground plane” is not straightforward in off-road environments.

An exhaustive method, suitable for off-road environments, has been introduced in [5] and applied also in [6], [7]. This approach evaluates connectivity of points in the range image, builds candidate obstacles, and estimates if they have a sufficiently high angle with the ground plane. A tree search allows then to connect close candidate obstacles and obtain 3D points clouds, whose densities have to be measured in order to classify them as wrong matches or as obstacles. Unfortunately, applications of this method have not passed 1.5 Hz throughput yet. Furthermore, this method seems to be inadequate for thin obstacles detection, one of the main targets of our study.

[3] uses a standard flood filling algorithm to find connected regions (clusters) of similar disparity in the DSI, and evaluates clusters size to decide if they represent obstacles.

In the intent of keeping computational load under control, in our approach we developed a series of fast filters to be applied to the DSI, in order to detect disparity concentrations that are eligible to be obstacles. A full example of the filter series is shown in Fig. 9.

Farther obstacles have smaller sizes in the image. To locate obstacles, [2] maps the Range Image data in a bird’s eye view map and looks for points concentration. To address the fact that 3D information is more and more sparse as distance grows, it compresses the space in direct proportion to the distance. As we look for obstacles directly in the images, we keep the threshold on obstacles size (measured in pixels) in inverse proportion to the distance.

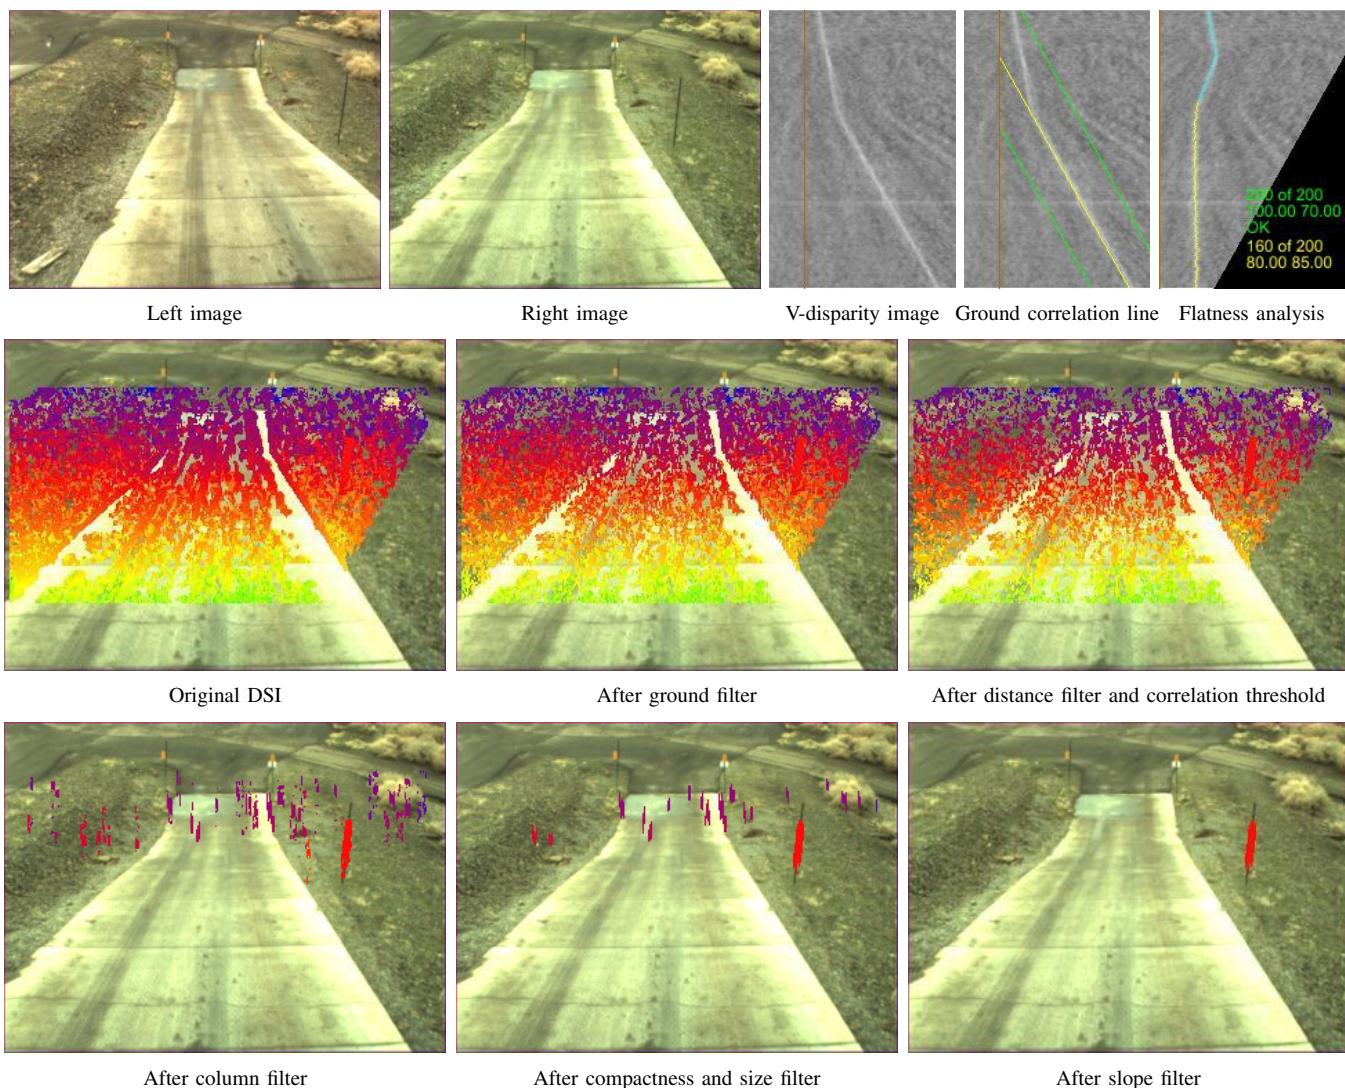


Fig. 9. Example of frame processing in a situation in which the vehicle was running downhill, facing a flat road after a steep descent. After the V-disparity analysis, obstacles are found via a series of filters applied to the DSI. The slope filter helps to remove the false positives that are mainly present in the region featuring a slope change.

C. Real world coordinates mapping

At present, the obstacles found in the image are not framed with bounding boxes because the algorithm does not need any further step of obstacle classification. The obstacles are mapped in real world coordinates, and set of obstacles close to each other (closer than the vehicle width) are grouped into a single obstacle using a convex hull algorithm. Figure 10 shows the bird's eye view map of obstacles obtained from the example of Fig. 9. The whole filtering and real world coordinates mapping are accomplished within 5 ms. Other output examples in different scenarios are shown in Fig. 11.

VI. RESULTS AND PERFORMANCE

The algorithm thresholds were tuned in order to avoid false detections. Preparatory testing in the desert near Barstow, California, showed the vision obstacle detection system to be very effectively complemented by the other sensor systems. Vision could detect obstacles which the others sensors missed



Fig. 10. Obstacles bird's eye view map for example of Fig. 9. The field of view is shown in green.

—namely tall thin ones— and ran with a guaranteed 15 Hz throughput. The obstacle detection system can perform the entire computation within 30 ms on a Pentium IV @ 2.8GHz processor system, allowing other needed image processing routines to run on the same machine while still meeting the

15 Hz throughput.

During NQE (National Qualification Event prior to DARPA Grand Challenge 2005), the system detected all the bounding cones with only one false detection in four runs. On October, 8th and 9th 2005 the TerraMaxTM vehicle was one of only five vehicles to complete the entire DARPA Grand Challenge course and was the only finisher exploiting vision for obstacle detection.

Data from the high precision laserscanner helped tune the distance estimates from the vision system to meet all requirements. The pitch detection system, provided by the V-disparity image quality analysis, can detect its own failure in all ground correlation line detection errors.




Fig. 11. Obstacle detection for different scenarios. The color of the obstacles varies with distance. In the bird's eye view map the field of view varies because of different baselines and camera setup. In the bottom image, a far gate in the left side is detected (in purple).

VII. CONCLUSIONS

We have presented a fast and reliable algorithm for off-road obstacle detection. A certain number of possible cases (e.g. textureless obstacles) are dealt with other sensors. The system is robust to false detection, even in complex environments, with slopes, shadows and different quality of textures.

An accurate calibration is of paramount importance. The paper also presented a method to detect misalignment through V-disparity analysis. Our system may serve as a general feature for higher-level object detection and classification.

VIII. ACKNOWLEDGEMENT

Authors would like to thank the members of Team TerraMaxTM for the great time spent together developing, testing, and finally seeing our vehicle rolling across the finish line. A special thank to Chris Hubert for his valuable suggestions for improving this paper.

REFERENCES

- [1] A. Broggi, C. Caraffi, R. I. Fedriga, and P. Grisleri, "Obstacle Detection with Stereo Vision for off-road Vehicle Navigation," in *Intl. IEEE Wks. on Machine Vision for Intelligent Vehicles*, San Diego, USA, June 2005.
- [2] S. Nedeveschi, R. Danescu, D. Frentiu, T. Marita, F. Oniga, C. Pocol, T. Graf, and R. Schmidt, "High accuracy stereovision approach for obstacle detection on non-planar roads," in *IEEE Intelligent Engineering Systems (INES)*, Cluj Napoca, Romania, 2004.
- [3] T. Dang and C. Hoffmann, "Fast object hypotheses generation using 3D position and 3D motion," in *Intl. IEEE Wks. on Machine Vision for Intelligent Vehicles*, San Diego, USA, June 2005.
- [4] T. Dang, C. Hoffmann, and C. Stiller, "Fusing optical flow and stereo disparity for object tracking," in *Intelligent Transportation Systems Conference*, Singapore, Sept. 2002, pp. 112–117.
- [5] A. Talukder, R. Manduchi, L. Matthies, and A. Rankin, "Fast and reliable obstacle detection and segmentation for cross country navigation," in *IEEE Intelligent Vehicles Symposium*, Versailles, France, 2002.
- [6] R. Manduchi, A. Castano, A. Talukder, and L. Matthies, "Obstacle detection and terrain classification for autonomous off-road navigation," *Autonomous Robot*, 18, pp. 81–102, 2003.
- [7] J. van den Hauvel, J. Kleijweg, W. van der Mark, M. Lievers, and L. Kester, "Obstacle detection for people movers using vision and radar," in *10 th World Conference on Intelligent Transport Systems and Services*, Madrid, Spain, 2003.
- [8] R. Arturo, A. Huertas, and L. Matthies, "Evaluation of stereo vision obstacle detection algorithms for off-road autonomous navigation," in *32nd AUVSI Symposium on Unmanned Systems*, June 2005.
- [9] P. Bellutta, R. Manduchi, L. Matthies, K. Owens, and A. Rankin, "Terrain perception for DEMO III," in *Intelligent Vehicles Conference*, 2000.
- [10] R. Labayrade, D. Aubert, and J.-P. Tarel, "Real time obstacle detection on non flat road geometry through V-disparity representation," in *IEEE Intelligent Vehicles Symposium, Versailles*, June 2002, pp. 646–651.
- [11] R. Labayrade and D. Aubert, "A single framework for vehicle roll, pitch, yaw estimation and obstacles detection by stereovision," in *Intelligent Vehicles Symposium Proceedings, Columbus*, June 2003.
- [12] Z. Hu, F. Lamosa, and K. Uchimura, "A complete U-V-disparity study for stereovision based 3D driving environment analysis," in *3-D Digital Imaging and Modeling (3DIM'05)*, 2005.
- [13] V. Lemonde and M. Devy, "Obstacle detection with stereovision," in *Mechatronics & Robotics (MECHROB'04)*, vol. 3, Aachen, Germany, Sept. 2004, pp. 919–924.
- [14] K. Konolige, "Small vision systems: Hardware and implementation," in *Eighth Intl. Symposium on Robotics Research, (Hayama, Japan)*, Oct. 1997, pp. 111–116.

- [15] D. Scharstein, R. Szeliski, and R. Zabih, "A taxonomy and evaluation of dense two-frame stereo correspondence algorithms," in *IEEE Workshop on Stereo and Multi-Baseline Vision, Kauai, HI*, Dec. 2001.
- [16] T. A. Williamson, "A high-performance stereo vision system for obstacle detection," Ph.D. dissertation, Carnegie Mellon University, Sept. 1998.
- [17] T. Kato, Y. Ninomiya, and I. Masaki, "An obstacle detection method by fusion of radar and motion stereo," in *IEEE Trans. Intell. Transport. Syst.*, vol. 3, Sept. 2002, pp. 182–188.
- [18] D. Nister, O. Naroditsky, and J. Bergen, "Visual odometry," in *Computer Vision and Pattern Recognition (CVPR) 2004*, vol. 1, June 2004, pp. 652–659.
- [19] J. Sun, Y. Li, S.-B. Kang, and H.-Y. Shum., "Symmetric stereo matching for occlusion handling," in *CVPR*, vol. 2, San Diego, USA, June 2005, pp. 399–406.
- [20] F. Estrada and A. Jepson, "Quantitative evaluation of a novel image segmentation algorithm," in *CVPR*, vol. 2, San Diego, USA, June 2005, pp. 1132 – 1139.
- [21] *IA32-Intel Architecture Optimization Reference Manual*, Intel, June 2005, order Number 248966-012 <http://developer.intel.com>.
- [22] Oshkosh Trucks Corporation, Rockwell Collins, and University of Parma, "Team TerraMax 2005," available <http://www.terramax.com>.
- [23] A. Broggi, S. Cattani, P. P. Porta, and P. Zani, "A Laserscanner–Vision Fusion System Implemented on the TerraMax Autonomous Vehicle," in *IEEE Intl. Conf. on Intelligent Robots and Systems (IROS06)*, Beijing, China, Oct. 2006, in press.
- [24] A. Broggi, C. Caraffi, S. Cattani, and R. I. Fedriga, "A Decision Network Based Frame-work for Visual Off-Road Path Detection Problem," in *IEEE Intelligent Transportation System Conference (ITSC06)*, Toronto, Canada, Sept. 2006.