

Vision-based path-planning in unstructured environments

Britta Hummel, Sören Kammel, Thao Dang, Christian Duchow and Christoph Stiller

Institut für Mess- und Regelungstechnik

Universität Karlsruhe (TH)

D-76131 Karlsruhe, Germany

Email: {name}@mrt.uka.de

Abstract—Autonomous driving in unstructured environments has attracted an unprecedented level of attention when the DARPA announced the Grand Challenge Competitions in 2004 and 2005. Autonomous driving involves (at least) three major subtasks: perception of the environment, path planning and subsequent vehicle control. Whereas the latter has proven a solved problem, the first two constituted, apart from hardware failures, the most prominent source of errors in both Grand Challenges.

This paper presents a system for real-time feature detection and subsequent path planning based on multiple stereoscopic and monoscopic vision cues. The algorithm is, in principle, suitable for arbitrary environments as the features are not tailored to a particular application.

A slightly modified version of the system described here has been successfully used in the Qualifications and the Final Race of the Grand Challenge 2005 within the *Desert Buceyes'* autonomous vehicle.

I. INTRODUCTION

Path planning is concerned with the problem of moving an entity from an initial configuration to a goal configuration. The resulting route may include intermediate tasks and assignments that must be completed before the entity reaches the goal configuration. External sensors provide input to the path planner, the most common of which include monoscopic and stereoscopic vision sensors and range finders, based on sonar, radar or laser light.

Path planning problems have been excessively studied within the robotics community, with applications ranging from robot manipulator navigation for medical or manufacturing applications to autonomous exploration in unknown environment, e.g. for military or planetary exploration purposes. [2] is an excellent survey on relevant work since the 1980's and proposes a classification scheme for path planning algorithms. [3] summarizes work on vision-based navigation and mapping.

The Grand Challenge turned out as an excellent testbed for comparison of different sensor data processing and path planning strategies (Fig. 1). It is an open contest for autonomous land vehicles that has been announced in 2004 and 2005 by the US Defense Advanced Research Projects Agency (DARPA). It requires autonomous robotic ground vehicles to successfully navigate a course of roughly 200 miles of off- and on-road terrain from Barstow, CA to Primm, NV. The course is roughly defined by GPS-waypoints that are published two hours before the race. Prior to the Grand Challenge race is a qualifying

event, testing the vehicles' abilities to autonomously navigate and avoid obstacles. In 2005, from 200 initial participants 40 were selected for the qualifying, 20 of whom were chosen for the final race.




Fig. 1. The Grand Challenge 2005 Qualifying event: The *Desert Buceyes* autonomous vehicle successfully navigating through a tunnel (left, onboard camera) and a narrow gate bounded by reflecting obstacles (right, spectator's camera).

This article introduces a fast path planning algorithm for unstructured environments based on multiple, complementary stereoscopic and monoscopic vision cues. Further incorporation of additional, arbitrary sensor information, e. g. lidar or radar, is straightforward.

The algorithm's path planning strategy is based on the following assumptions:

- A good path is maximally even.
- The transition between scene components involves a non-smooth change of color.
- The road exhibits a preferred texture orientation in driving direction due to antecedent vehicles passing (this can be observed even in the case of paved roads).

The algorithm is thus suitable to any environments where these assumptions at least roughly hold. Each assumption is tested for validity for each possible path by specially tailored feature extractors. A fast path planning strategy is introduced which is based upon a probabilistic grid map using an active testing model for efficient computation.

This document is structured as follows: First, we give a brief overview of the system architecture. Next, the used vision cues and the corresponding feature extraction schemes are examined. The next section focuses on the path planner, followed by real-time issues concerning the search strategy. We conclude by providing results from the system's intense

field test in the Grand Challenge Race and the Qualifying Event, thereby commenting on benefits and drawbacks.

II. SYSTEM ARCHITECTURE

Fig. 2 provides an overview of the system architecture. A monoscopic and a stereoscopic vision sensor, ultrasonic devices and four laser range finders serve as external sensors, while GPS plus Inertial Navigation System (INS) supply vehicle positioning data. The sensor data, including measurement uncertainties, and an internal database provide the input to an internal representation of the environment that is updated at 5Hz frequency. It serves as input to the path planning algorithm.

Whereas the global path planning is predominantly determined by map information and the GPS waypoints that are announced two hours prior to the race, the local, short-term path planner focuses on the information from positioning devices and external sensors.

The next section sketches the vision sensor's role in the generation of the environment representation and path planning.




Fig. 2. System architecture

III. IMAGE ANALYSIS

A robust path estimator is dependent on reliable, complementary environmental feature detectors. They serve to discriminate between (at least) two classes: *place to savely drive upon* and *obstacle of any kind*. A stereoscopic vision sensor, in principle, provides sufficient information for that task.

Three complementary, monoscopic and stereoscopic vision feature tests have been developed, exploiting information about depth, texture homogeneity and local texture orientation. They are considered suitable to a wide variety of environments. Each is presented in greater detail in the sequel.

A. Disparity feature

Stereo disparity is an appealing vision feature since it provides instantaneous depth measurements for most pixels of

the camera image and therefore enables identification of the road plane. The term disparity refers to the 1d displacement of corresponding pixels in a stereo image pair. It is inversely related to the distance of an observed object point and usually found via area based matching (see [6], [7] for reviews on stereo matching techniques in general and with an emphasis on automotive applications).

The Zero-Mean Sum of Squared Differences (ZSSD) is used for quantification of the dissimilarity between two image blocks:

$$\text{ZSSD}(R, L) = \sum_{\mathbf{x} \in B} [(R(\mathbf{x}) - \bar{R}) - (L(\mathbf{x}) - \bar{L})]^2, \quad (1)$$

where \bar{R} and \bar{L} denote the mean intensities of R , and L and B denotes the block pixels. It is closely related to the variance of the intensity difference between the two image regions: $\text{ZSSD}(R, L) \sim \text{Var}(R-L)$.

The ZSSD requires slightly more computational effort than e.g. the sum of absolute differences (SAD), but is invariant with respect to differing lighting conditions. Additionally, it induces a meaningful quality measure for disparity estimates:

$$\text{SNR}(R, L) = \frac{\text{Var}(R)}{\text{Var}(R-L)} = \frac{\sum_{\mathbf{x} \in B} (R(\mathbf{x}) - \bar{R})^2}{\text{ZSSD}(R, L)}. \quad (2)$$

This defines a signal-to-noise ratio, in terms of the ratio of the texture in the right block R with respect to the texture of the difference image $R-L$. Thus, Eq. (2) reflects the intuitive notion that corresponding image blocks with high texture may show larger dissimilarity than those with little texture.

For improved performance, and for generation of sparse disparity images, a matching with adaptive windows as introduced in [9] has been implemented. Figure 3 depicts disparity estimates for typical imagery. Only disparity measurements with a signal-to-noise ratio larger than 2.0 are displayed. The stereo matching procedure requires 80 ms for each run on a Intel Pentium M processor with 2 Ghz.

Following the work of [8], the ground plane can be detected from stereo disparities using the *v-disparity*. Each row in the *v-disparity* image is given by the histogram of the corresponding row in the disparity image. Each tilted plane in 3D space (with zero roll angle) becomes a straight line in the *v-disparity* image (see Fig. 4). In the current publication, the line parameters are estimated using the Radon transform, which is implemented efficiently using the central slice theorem (cf. [10]), and readily provides a quality measure, namely the number of pixels in the image that are consistent with the detected ground plane. The advantages of the *v-disparity* method are its simplicity and its robustness, whereas a drawback is the difficult determination of the ground plane's roll angle.

After estimation of the ground plane, the height above/below ground (*parallax*) $d_{\Delta}(\mathbf{x})$ of an arbitrary image point $\mathbf{x} \in \mathcal{R}^2$ is given by the difference between measured disparity and the disparity of the ground plane at that position. The parallax serves as input to the feature test.

Figure 4 shows the estimated parallax for the stereo images from Figure 3 as a bird's eye view.




Fig. 3. (a) and (b) Raw stereo image pair, (c) disparity image, (d) right image with partial overlay of the ground plane, estimated using the v-disparity. The matching block size was set to 9×9 pixels. Only measurements with a signal-to-noise ratio larger than 2.0 are displayed.




Fig. 4. (a) The height above/below ground in brightness-encoded birds' eye view. Light coloring denotes positive, dark denotes negative heights, respectively. Unobservable parts, either due to occlusion by obstacles or due to a limited field of view, are colored in black. (b) V-disparity image. The white dashed line is the estimated ground plane line. For illustration purpose, the horizon is visualized in the original image.

B. Texture feature

Depth information cannot be computed for regions without dominant gradients in both directions due to the aperture problem. In addition, computed depth values become less reliable at greater distances. However, a human can discriminate quite reliably between roads and obstacles even for homogeneously textured or far-distance regions because he extrapolates in various ways from reliably classified points. A completely unsupervised segmentation, i. e. not relying on any known classification results, would not be feasible for the scenes considered here, because of the abundance of possible road and obstacle textures.

The texture feature described below has been designed to adopt this extrapolation strategy with respect to color information, starting at reliably classified *seed points* which are provided by a simple disparity based segmentation into road plane and obstacle. In the sequel this problem is formulated

as a graph theoretic approach and solved by computing some shortest paths within the graph.

The path cost $c(\mathbf{x}, \mathbf{y})$ between two arbitrary points \mathbf{x} and \mathbf{y} is defined as the minimum sum of the pixel transition costs from start to goal over all possible paths

$$c(\mathbf{x}, \mathbf{y}) = \min_{(\mathbf{x}_1, \dots, \mathbf{x}_n), \mathbf{x}_1 = \mathbf{x}, \mathbf{x}_n = \mathbf{y}} \sum_{i=1}^{n-1} \text{cost}(\mathbf{x}_i, \mathbf{x}_{i+1}), \quad (3)$$

where \mathbf{x}_i and \mathbf{x}_{i+1} are 4-neighboring pixels. The transition cost between neighboring pixels is defined as the squared, normalized difference of the color median in a rather large environment:

$$\text{cost}(\mathbf{x}, \mathbf{y}) = \left(\frac{\| \text{med}(g(\mathbf{x})) - \text{med}(g(\mathbf{y})) \|}{\| g_{\min} - g_{\max} \|} \right)^2. \quad (4)$$

The path cost is thus a measure for the smoothness (w.r.t. color) of the best path between two image regions. Note that, in this formulation, neither short nor smooth paths are generally preferred. While this measure works more robustly in the RGB color space it has successfully been used here with grayvalue cameras.

By interpreting image pixels and their 4-neighbors as the vertices and edges of an undirected graph, the shortest path problem can be solved using either Kruskal's or Dijkstra's single source shortest path algorithms (see [1]).

The input to the feature test is the difference between the minimum path cost to any obstacle seed point $\mathbf{s}_{0,i}$ and the minimum path cost to any road seed point $\mathbf{s}_{1,i}$:

$$t(\mathbf{x}) = \min_i (c(\mathbf{x}, \mathbf{s}_{0,i})) - \min_i (c(\mathbf{x}, \mathbf{s}_{1,i})). \quad (5)$$

This quantifies degree of certainty that the pixel in question belongs to an obstacle ($t(\mathbf{x}) = 1$) or the road ($t(\mathbf{x}) = -1$). A low value for $|t(\mathbf{x})|$, on the contrary, indicates an unreliable estimate.

While more seed points increase classification robustness, the algorithm can principally cope with an arbitrary number of given seed points.

The classifier is illustrated in Fig. 5.

Previous downsampling of the rectified input image eases computational load without loss of classification robustness, as the feature has been designed to only consider the rough color impression of an image patch. The thereby suppressed high frequency information is exploited in the orientation feature described in the next section.

C. Orientation feature

The feature of local orientation is based on the assumption that the path in question features structures that are oriented in the direction of the path. These structures could be tire marks within the path or the edge of the path, for example. Subsequently, a probabilistic model for the local orientation at a given image point is formulated. The probability of the point to belong to the path is then computed using Bayes' formula. This forms the basis for computing the probability for the patch in question to belong to the path. First, the




Fig. 5. Texture-only based classification. Results for an image without apparent change in texture between obstacle (the upper left part of the field) and driveable area (the road and the lower left part of the field, where the grass is cut). Displays the original image after 5×5 median filtering and (a) parts classified as obstacles and (b) parts classified as flat area, based on the given road and driveable area seeds. Correctly, the classifier does *not* take notice of the texture transition between road and field on the left, because of the driveable area seeds present in both parts.



Fig. 6. Brightness-encoded output of the local orientation feature, tested for an orientation of $\mu = 0$, i. e. straight ahead. High grayvalues correspond to small orientation differences. Note how the tire marks in the sand cause, in average, significantly larger outputs than the other parts of the image.

local orientation is computed from the grey value gradient covariance matrix (cf. [4]):

$$M(\mathbf{x}) = \begin{bmatrix} \left(\frac{\partial g(\mathbf{x})}{\partial x_1}\right)^2 & \left(\frac{\partial g(\mathbf{x})}{\partial x_1}\right)\left(\frac{\partial g(\mathbf{x})}{\partial x_2}\right) \\ \left(\frac{\partial g(\mathbf{x})}{\partial x_1}\right)\left(\frac{\partial g(\mathbf{x})}{\partial x_2}\right) & \left(\frac{\partial g(\mathbf{x})}{\partial x_2}\right)^2 \end{bmatrix} \quad (6)$$

with $g(\mathbf{x}) = g(x_1, x_2)$ being the grey value. We are interested in image points with one dominant orientation, therefore, we consider those points \mathbf{x}_i in the image with a high ratio $\frac{\lambda_1}{\lambda_2}$ of the two eigenvalues $\lambda_1 > \lambda_2$ of $M(\mathbf{x})$. Then, the orientation $\alpha(\mathbf{x})$ is computed as

$$\alpha(\mathbf{x}) = \frac{1}{2} \arctan \left(\frac{2 \left(\frac{\partial g(\mathbf{x})}{\partial x_1}\right) \left(\frac{\partial g(\mathbf{x})}{\partial x_2}\right)}{\left(\frac{\partial g(\mathbf{x})}{\partial x_1}\right)^2 - \left(\frac{\partial g(\mathbf{x})}{\partial x_2}\right)^2} \right) + \frac{\pi}{2} \quad (7)$$

Fig. 6 shows the results of the orientation-only classifier output with the direction straight upwards in the image being tested for.

IV. PATH PLANNING

We represent a path by a sequence $\mathbf{p} = (\mathbf{p}_1, \dots, \mathbf{p}_N)$ of N patches \mathbf{p}_n . Each patch is described with a small set of parameters. In our experiments we have used rectangular patches of length l and width b . As depicted in Figure 7, each patch may then be described by its center vector $\mathbf{p}_n \in \mathcal{R}^2$ where a seed patch $\mathbf{p}_0 = (l, 0)^T$ has been predefined to start the path directly in front of our vehicle.




Fig. 7. Definitions of coordinate system and path

We comprise all available features into a vector \mathbf{y} . Considering \mathbf{p} and \mathbf{y} as realizations of random vectors \mathbf{P} and \mathbf{Y} , respectively, one may formulate the *a posteriori* distribution for the path given the features with the Bayesian paradigm as

$$p(\mathbf{P} = \mathbf{p} | \mathbf{Y} = \mathbf{y}) = \frac{p(\mathbf{Y} = \mathbf{y} | \mathbf{P} = \mathbf{p}) \cdot p(\mathbf{P} = \mathbf{p})}{\sum_{\mathbf{p}} p(\mathbf{Y} = \mathbf{y} | \mathbf{P} = \mathbf{p}) \cdot p(\mathbf{P} = \mathbf{p})}. \quad (8)$$

The conditional entropy $H(\mathbf{P} | \mathbf{Y} = \mathbf{y})$ quantifies our remaining uncertainty about the path. We continue to reduce this uncertainty by successively gathering features from the images until the distribution (8) concentrates at one or a few paths that are overwhelmingly likely. This procedure belongs to the class of *active testing* (cf. [5]). The two distributions on the right side of (8) are formulated in the following subsections.

A. Path Model

The *prior* $p(\mathbf{P} = \mathbf{p})$ expresses our expectations on typical paths. We have selected a first-order Markov chain that favors a smooth sequence of patches.

$$\begin{aligned} p(\mathbf{P} = \mathbf{p}) &= \prod_{n=1}^N p(\mathbf{P}_n = \mathbf{p}_n | \mathbf{P}_{n-1} = \mathbf{p}_{n-1}) \\ &= k e^{-\lambda \sum_{n=1}^N \mathbf{p}_n^T \cdot \mathbf{p}_{n-1}}, \end{aligned} \quad (9)$$

where the partition function k is a normalization factor. It is worth noting that maximum probability is associated to a straight path. Any difference in orientation between subsequent patches is punished with weight λ . Extensions that favour constant curvature or constant curvature rate can easily be formulated with higher-order Markov chains. However, it is emphasized that this model only weakly constrains path geometry. In contrast parametric models that impose hard constraints any complex path may be estimated if sufficiently supported by the image data.

B. Observation Model

The likelihood $p(\mathbf{Y} = \mathbf{y} | \mathbf{P} = \mathbf{p})$ in (8) constitutes the relationship between the image information \mathbf{y} to a given path \mathbf{p} . We gather our feature vector as a sequence of active tests $\mathbf{y} = (y_1, \dots, y_K)$. In this sequence, each test y_k is influenced by only a single patch \mathbf{p}_i . Let $p_1(y_k)$ and $p_0(y_k)$ denote the distribution of the test on and beyond the path, respectively. Assuming conditional independency of the test sequence we can write

$$\begin{aligned} p(\mathbf{Y} = \mathbf{y} | \mathbf{P} = \mathbf{p}) &= \prod_{k=1}^K p(Y_k = y_k | \mathbf{P} = \mathbf{p}) \\ &= \prod_{\substack{k=1 \\ \mathbf{p} \notin C_k}}^K p_0(y_k) \cdot \prod_{\substack{k=1 \\ \mathbf{p} \in C_k}}^K p_1(y_k), \end{aligned} \quad (10)$$

where C_k denotes the set of all paths that include the patch that is investigated by the k^{th} test.

The next step consists in defining the distributions p_1 and p_0 for the described image features, reflecting our expectations concerning the feature output for a given on- or off-road patch.

1) *Disparity feature test*: The distribution of the disparity parallax d_Δ is modelled as white gaussian:

$$p_{i,D}(y_k) = p_{i,D}(d_\Delta) = \prod_{\mathbf{x} \in \mathcal{G}_k} \frac{1}{\sqrt{2\pi\sigma_i^2}} e^{-\frac{d_\Delta(\mathbf{x})^2}{2\sigma_i^2}}, \quad i \in \{0, 1\}. \quad (11)$$

\mathcal{G}_k denotes the set of points belonging to the patch that is tested for in the k^{th} test. The on-road variance is significantly smaller than the off-road variance, i. e. $\sigma_1 \ll \sigma_0$. Hence, the disparity test favours planar patches lying in the road plane. Each parallax estimate exceeding a quality threshold, defined in terms of Equation 2, constitutes one feature test.

2) *Texture homogeneity feature test*: The texture feature $t(\mathbf{i})$ quantifies the belief that a certain image pixel belongs to an obstacle ($t = 1$) or the road ($t = -1$). A low value for $|t_{\mathbf{i}}|$ indicates an unreliable estimate. The distributions are therefore modelled as:

$$\begin{aligned} p_{1,T}(y_k) &= p_{1,T}(t) \\ &= \prod_{\mathbf{x} \in \mathcal{G}_k} \frac{2}{\sigma\sqrt{2\pi}} \cdot e^{-\frac{(1-t(\mathbf{x}))^2}{2\cdot\sigma^2}} \cdot \sigma(1 - t(\mathbf{i})) \quad (12) \\ p_{0,T}(y_k) &= p_{0,T}(t) \\ &= \prod_{\mathbf{x} \in \mathcal{G}_k} \frac{2}{\sigma\sqrt{2\pi}} \cdot e^{-\frac{(-1-t(\mathbf{x}))^2}{2\cdot\sigma^2}} \cdot \sigma(t(\mathbf{i}) + 1) \quad (13) \end{aligned}$$

Hence, this feature test favors patches with a smooth path (w.r.t. color) to a known road part compared to the smoothness of a path to a known obstacle.

3) *Texture orientation feature test*: The distribution for the local orientation $\alpha \in [0, \pi]$ is modeled as:

$$\begin{aligned} p_{1,O}(y_k) &= p_{1,O}(\alpha) \\ &= \prod_{\mathbf{x} \in \mathcal{G}_k} \frac{C_1}{\sigma\sqrt{2\pi}} \cdot e^{-\frac{(\mu - \alpha(\mathbf{x}))^2}{2\cdot\sigma^2}} + \frac{1 - C_1}{\pi} \quad (14) \end{aligned}$$

$$\begin{aligned} p_{0,O}(y_k) &= p_{0,O}(\alpha) \\ &= \prod_{\mathbf{x} \in \mathcal{G}_k} \frac{1}{\pi}. \quad (15) \end{aligned}$$

Eq. 14 is a mixture of a gaussian, with a mean angle μ equalling the direction of the tested patch, and a uniform distribution, with $0 < C_1 < 1$ being the mixing parameter (however, the algorithm is rather robust towards the choice of C). The gaussian assigns a high probability to regions that are oriented in patch direction. The uniform distribution increases robustness towards gross outliers.

Eq. 15 reflects the assumption that no preferred orientations are assumed off-path.

C. Search tree

Insertion of (9) and (10) into (8) yields the *a posteriori* distribution of the path for the given features. Since its complexity increases exponentially with path length N an exhaustive search is prohibitive even for moderate N . We have implemented a search tree that guarantees to find the path of length N with maximum *a posteriori* probability. The length N increases over time and we stop our search when either a sufficient length is reached or a preset time interval dedicated to path computation has expired.

V. EXPERIMENTAL RESULTS

A. Hardware setup

A custom stereo vision sensor was build for the special requirements in the Mojave desert. We used Flea cameras from PointGrey Research which offer a VGA resolution with a good SNR and an automatic synchronization between the cameras without disabling the AGC feature. They were mounted with a base width of 37cm. The cameras were mounted inside of two protection cases, additionally featuring a large sun shield that prevented from direct sunlight on the CCD sensors in most cases. Attached wipers provided the possibility of regularly cleaning the cases' windows.

Future developments include the use of a (probably adaptive) optical filter, for further reduction of the possibility of sensor blooming under the extreme light conditions in the desert. Additionally CMOS chips, offering a logarithmic sensitivity to incoming light, could be used, if attaching custom external trigger hardware.

B. Results

Fig. 8 displays the computed path for a stereo image from a preliminary test sequence taken on a German dirt road. It clearly illustrates the reduction of measurement uncertainty due to inclusion of complementary feature tests, compared to single feature detectors. Fig. 9 provides examples on path computation for other terrains.

A slightly modified version of the algorithm was used within the DARPA Grand Challenge Qualifying and the Final Race. Additional information was provided by the external sensors described in section II and by a positioning device. The *Desert Buckyeyes* successfully avoided all obstacles in the Qualifying and autonomously drove a distance of about 50km in the final race in completely unknown terrain.




Fig. 8. Calculated path. Obstacles are the sign to the right and the unreaped field in the upper-left. (a) Path, based on disparity information only. A 30° orientation difference between camera and vehicle was assumed in the camera calibration unit, resulting in the corresponding orientation of the starting patch. In the overexposed spot in the middle of the road disparity could not be computed. Therefore, and as a smooth path is generally favoured, the path direction remains constant until approaching the unreaped field. (b) Path, after inclusion of the texture and orientation information. The local orientation of the tire marks lead to a path correction towards the road, despite its higher smoothness cost.




Fig. 9. Calculated path for varying terrains. In all observed cases, obstacles have been clearly avoided.

VI. CONCLUSION AND OUTLOOK

We have introduced a fast and robust vision-based path-planning algorithm that has been successfully used within the DARPA Grand Challenge 2005 within the *Desert Buckyeyes* vehicle, cooperatively developed by the Ohio State University and the University of Karlsruhe.

The system features a robust real-time feature detection that has been tested in a wide variety of environments including desert, paved roads, unpaved tracks and meadows. The underlying holistic probabilistic framework allows a reliable obstacle detection and subsequent path planning.

The active testing scheme allows for a straightforward integration of scene domain (disparity) and image domain (texture) features. This prevents from those errors that are caused by imperfect coordinate transformations and projections.

Whereas the disparity feature clearly outperforms the 2D feature tests in standard situations, the latter prove important in regions where disparity information cannot be reliably

computed, i. e. in homogeneously textured and in far distance regions.

The system can easily be extended to a stand-alone global path planning solution by additional inclusion of positioning information. The position of the next GPS waypoint can constitute an additional feature test, favoring a vehicle orientation towards the mission goal.

Current work on the system includes temporal tracking of the detected path, which further reduces the impact of noise within the feature outputs.

One drawback is, that the current implementation is strongly dependent on the quality of the ground plane estimation. Furthermore, vision sensing in general faces problems with reflecting obstacles. To cope with these, additional inclusion of further complementary feature tests, based for example on radar or lidar measurements, are recommended.

VII. ACKNOWLEDGEMENT

We would like to thank the *Ohio State University* and all the members from the *Desert Buckyeyes* team for the inspiring partnership and the fruitful discussions.

REFERENCES

- [1] R. Sedgewick, *Algorithms in C++*, Addison-Wesley, Boston, MA, 1992.
- [2] Y. Hwang and N. Ahuja, *Gross Motion Planning – A Survey*, ACM Computing Surveys 24(3), September 1992, 221-291.
- [3] G. DeSouza and A. Kak, *Vision for Mobile Robot Navigation: A Survey*, IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 24(2) February 2002, 237-267.
- [4] Jinyou Zhang; Nagel, H.-H., *Texture-based segmentation of road images*, Intelligent Vehicles '94 Symposium, Proceedings of the; 24-26 Oct. 1994, 260-265.
- [5] D. Geman and B. Jedynak, *An active testing model for tracking roads in satellite images*, IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), 18(1) January 1996, 1-14.
- [6] D. Scharstein, and R. Szeliski, *A Taxonomy and Evaluation of Dense Two-Frame Stereo Correspondence Algorithms*, International Journal of Computer Vision, Vol. 47, pp. 7-42, 2002.
- [7] H. Sunyoto, W. van der Mark and D. M. Gavrilu. A Comparative Study of Fast Dense Stereo Vision Algorithms, *IEEE Intelligent Vehicles Symposium*, Parma, Italy, pp. 319-324, 2004.
- [8] R. Labayrade, D. Aubert, and J. Tarel, Real Time Obstacle Detection in Stereovision on Non Flat Road Geometry Through "v-Disparity" Representation, *IEEE Intelligent Vehicle Symposium*, pp. 646-651, 2002.
- [9] H. Hirschmüller, P.R. Innocent, and J.M. Garibaldi, Real-Time Correlation-Based Stereo Vision with Reduced Border Errors, *IJCV*, Vol. 47(1/2/3), pp. 229-246, 2002.
- [10] J. Beyerer and F. Puente Leon, Die Radontransformation in der Digitalen Bildverarbeitung, *Automatisierungstechnik*, vol. 50(10), pp. 472-480, 2002.