

Road Surface and Obstacle Detection Based on Elevation Maps from Dense Stereo

Florin Oniga, Sergiu Nedevschi, Marc Michael Meinecke, and Thanh Binh To

Abstract— A new approach for the detection of the road surface and obstacles is presented. The 3D data from dense stereo is transformed into a rectangular elevation map. A quadratic road surface model is first fitted, by a RANSAC approach, to the region in front of the ego vehicle. This primary solution is then refined by a region growing-like process, driven by the 3D resolution and uncertainty model of the stereo sensor. An optimal global solution for the road surface is obtained. The road surface is used for a rough discrimination between road and above-road points. Above-road points are grouped based on vicinity and false areas are rejected. Each above-road area is classified into obstacles (cars, pedestrians etc.) or traffic isles (road-parallel patches) by using criteria related to the density of the 3D points. The proposed real-time algorithm was evaluated in an urban scenario and can be used in complex applications, from ego-pose estimation to path planning.

I. INTRODUCTION

PROCESSING 3D data from stereo (dense or sparse) is a challenging task. A robust approach can prove of great value for a variety of applications in urban driving assistance.

There are two main algorithm classes, depending on the space where processing is performed: disparity space-based and 3D space-based.

Disparity space-based algorithms are more popular because they work directly on the result of stereo reconstruction: the disparity map. The “v-disparity” [1] approach is well known and used to detect the road surface in a variety of applications [2]. Unfortunately, it is not a natural way to represent 3D data and has some drawbacks: it assumes that the road should occupy most of the image, and it is sensitive to roll angle changes. If the roll angle of the ego car changes from the initial calibration, then the road profile becomes blurry and harder to detect on the “v-disparity” image.

3D space-based algorithms are mainly used for ego-pose estimation [3], [4], but also for lane and obstacle detection [5], [6].

In [3] the ego-pose is estimated with respect to the road

plane. The road plane is fitted (not in real-time) by a RANSAC-approach to the whole set of dense 3D points (after filtering non-road points). A constant band around the road is used to select inliers and outliers, even though this is against the fact that the 3D uncertainty from stereo increases with the depth. The assumption that most of the 3D points are road points is made again (if the ego car is close to a sidewalk with more 3D points than the road, it is likely to fail).

A planar road surface is estimated from tracking in [4]. The method provides robust numerical results, but fails if occlusions (obstacles) are in front of the ego car. Lack of high-gradient road features also leads to failure (lane markings, borders, etc.).

Obstacles are detected as clusters of image edge points reconstructed in the 3D space ([5]). Road features as lane-markings can also be detected as obstacles. Later in [6] a 3D lane model is proposed and used for obstacle/road points’ separation. Again the method requires high-gradient road features (edges) to be present and uses a constant band to select road inliers and outliers.

The algorithm that will be presented in this paper manages to solve most of the drawbacks presented above.

It takes as input dense 3D reconstructed points, overcoming the lack of road edge features (high-gradient). To achieve real-time processing, the 3D set of points is transformed into a digital elevation map. The road is modeled as a quadratic surface to allow vertical curvatures, often present in urban scenarios. The road surface is fitted in a way which does not require road points to occupy most of the image. The 3D uncertainty increasing with the depth is taken into account.

First we will present the basic math needed to model and fit the quadratic road surface. The depth resolution and uncertainty are also modeled from the stereo system’s geometry. After that the algorithm will be presented and evaluated.

II. THE QUADRATIC ROAD SURFACE MODEL

The planar road model was first tested. This model proved to be less robust for obstacle/road separation, especially for the detection of road delimiters such as curbs (ex. border of a sidewalk). This happens because the road surface usually presents small longitudinal and lateral curvatures, and the planar assumption cannot cope with these curvatures.

We used a different model that allows quadratic variations

Manuscript received April 2, 2007.

F. Oniga and S. Nedevschi are with the Computer Science Department, Technical University of Cluj-Napoca, Romania, email: florin.oniga, sergiu.nedevschi@cs.utcluj.ro (corresponding author’s phone: +40264401219; e-mail: sergiu.nedevschi@cs.utcluj.ro).

M. M. Meinecke and T. Binh To are with the Research Department of Volkswagen AG, Wolfsburg, Germany, e-mail: marc-michael.meinecke, thanh-binh.to@volkswagen.de.

of the height (the Y coordinate) with the horizontal displacement (the lateral X coordinate) and the depth (the longitudinal Z coordinate). Figure 1 emphasizes the differences between the planar geometry and the quadratic one.

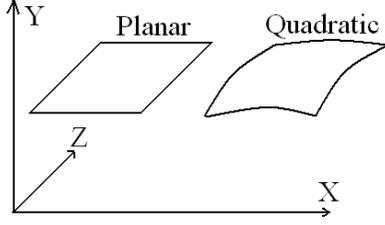


Fig. 1. Planar vs quadratic geometry for the road surface.

Equation (1) shows the algebraic form of the road model, by defining the height value Y with respect to the depth Z and the horizontal displacement X.

$$Y = -a \cdot X - a' \cdot X^2 - b \cdot Z - b' \cdot Z^2 - c. \quad (1)$$

Fitting the quadratic surface to a set of 3D points involves minimizing an error function. The error function S represents the sum of squared errors along the height:

$$S = \sum_{i=1}^n (Y_i - \bar{Y}_i)^2, \quad (2)$$

Where Y_i is the height of the 3D point i and \bar{Y}_i is the height of the surface at coordinates (X_i, Z_i) . Minimizing only along the Y-axis is enough because, even for curved roads, the normal of the surface is close to the Y-axis (only few degrees difference). The computational complexity is highly reduced by avoiding minimization against the normal of the surface.

By replacing (1) into (2) the following equation is obtained, where the unknowns are a , a' , b , b' , and c :

$$S = \sum_{i=1}^n (Y_i + a \cdot X_i + a' \cdot X_i^2 + b \cdot Z_i + b' \cdot Z_i^2 + c)^2. \quad (3)$$

The minimum of this function is 0 for perfect fitting. For S to have a minimum value, its partial derivatives with respect to the unknowns must be 0. The following system of equations must be solved:

$$\left\{ \frac{\partial S}{\partial a} = 0, \frac{\partial S}{\partial a'} = 0, \frac{\partial S}{\partial b} = 0, \frac{\partial S}{\partial b'} = 0, \frac{\partial S}{\partial c} = 0. \right. \quad (4)$$

After writing explicitly each equation, the system (4) becomes (matrix form):

$$\begin{bmatrix} S_{X^2} & S_{X^3} & S_{XZ} & S_{XZ^2} & S_X \\ S_{X^3} & S_{X^4} & S_{X^2Z} & S_{X^2Z^2} & S_{X^2} \\ S_{XZ} & S_{X^2Z} & S_{Z^2} & S_{Z^3} & S_Z \\ S_{XZ^2} & S_{X^2Z^2} & S_{Z^3} & S_{Z^4} & S_{Z^2} \\ S_X & S_{X^2} & S_Z & S_{Z^2} & n \end{bmatrix} \begin{bmatrix} a \\ a' \\ b \\ b' \\ c \end{bmatrix} = \begin{bmatrix} -S_{XY} \\ -S_{X^2Y} \\ -S_{ZY} \\ -S_{Z^2Y} \\ -S_Y \end{bmatrix}, \quad (5)$$

Where n is the number of points, and generically

$$S_\alpha = \sum_{i=1}^n \alpha_i \quad (\text{for example } S_{XZ} = \sum_{i=1}^n X_i \cdot Z_i). \quad (6)$$

This system has 5 linear equation and 5 unknowns, therefore solving it is a trivial algebra problem.

III. MODELING THE DEPTH RESOLUTION AND UNCERTAINTY

In order to keep road points connected in the Euclidian 3D space, the depth (or range) resolution must be modeled. Connectivity is required by the step of road surface growing.

The 3D surface fitting will involve selection of inliers and outliers of the road surface (road/non-road points). Because the 3D uncertainty increases with the depth, using a constant band is quite inefficient when working with a large range of depths (false inliers close and outliers far).

A. The depth resolution from stereo

The set of 3D points from dense stereo will be transformed into a rectangular elevation map, each cell containing a height value. Due to the perspective projection, 3D points reconstructed by stereovision are sparser with the depth. This effect is more visible for the road surface, because it is quasi-parallel to the optical axis of the camera.

A 3D space compressed with the depth was used in [5] to achieve connectivity for obstacle points in order to detect obstacles. We will use the opposite approach, by applying a morphological dilation of heights on the elevation map. The amplitude of dilation will be computed for each cell from the stereo geometry.

The effect of decreasing depth resolution for road points is shown in figure 2. The distance between image-adjacent 3D road points increases with the depth.

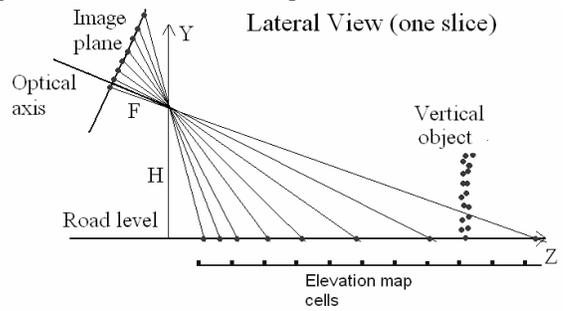


Fig. 2. A lateral view of the depth resolution problem (F - focal length, H - height of the camera relative to the road level). Image-adjacent road points are not always adjacent in the elevation map. Points from vertical objects have higher densities in cells.

Considering the canonical geometry of the stereo system and a planar road, the Z -value of a road point placed at the vertical coordinate Y_{im} (in pixels) in the image can be computed using (7). H is the height of the camera in the world reference frame, F the focal length of the stereo system, and α is the pitch angle of the camera.

$$Z(Y_{im}) = H \cdot \frac{F \cdot \tan\left(\frac{\pi}{2} - \alpha\right) + Y_{im}}{Y_{im} \cdot \tan\left(\frac{\pi}{2} - \alpha\right) - F}. \quad (7)$$

Furthermore the depth resolution ΔZ can be computed at any depth Z by (8).

$$\Delta Z = \left| Z(Y_{im}) - Z(Y_{im} + 1) \right|. \quad (8)$$

Later in this paper it will be explained how to use this basic model for computing the depth resolution even for non-planar road surfaces.

B. The depth uncertainty from stereo

The 3D (localization) uncertainty is caused by a low accuracy computation of the disparity value and is mainly visible in the depth value. The height is also influenced and modeling the uncertainty will help for a robust detection of the road.

In [7] a simple model for the uncertainty of the depth was proposed for a canonical stereo system. The depth uncertainty Z_{err} for a point at depth Z , was modeled as a function (9) of the system's parameters (baseline B and focal F known from calibration) and of the disparity uncertainty D_{err} .

$$Z_{err} = \left| \frac{Z^2 \cdot D_{err}}{B \cdot F - Z \cdot D_{err}} \right|. \quad (9)$$

By using the standard canonical stereo equations, we extended this model to compute the uncertainty Y_{err} of a point with height Y and depth Z (10). H is the height of the camera in the world reference frame.

$$Y_{err} = \left| \frac{(Y - H) \cdot Z_{err}}{Z} \right|. \quad (10)$$

Figure 3 shows an example: the region for inliers computed for a simple planar road surface assuming a disparity uncertainty D_{err} of 1 pixel. This value is not exaggerated because the road surface usually presents poor texture (except for markings), resulting in a larger disparity uncertainty compared to good-textured objects.

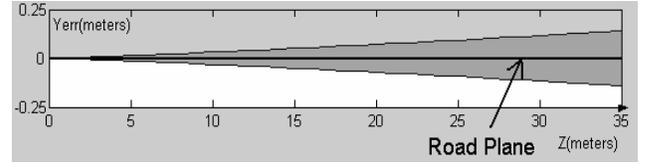


Fig. 3. Lateral view of the inliers (darker gray) region around a planar road. The region is computed using the proposed model for the height uncertainty.

IV. ALGORITHM DESCRIPTION

The road and obstacles detection algorithm presented in this paper takes as input dense 3D reconstructed points. The output is multiple: the quadratic surface of the road in parametric form, 3D points classified as road/traffic isles/obstacles, and cell clusters representing individual traffic isles or obstacles. Traffic isles (elevated surfaces parallel to the road, such as sidewalks) are classified distinctly because they should have a different significance for higher-level applications, compared to relevant obstacles (cars, pedestrians, buildings, etc.).

To achieve real-time processing, the 3D set of points is transformed into a digital elevation map. All processing steps will be performed on the elevation map. The road surface is fitted using a RANSAC approach to a small patch in front of the ego vehicle. This primary solution is then refined through a region growing-like process to insure the optimality of the global solution.

A. Building the elevation map

The elevation map representation is usually used for relief representation. An intensity map (image) is superimposed, from the top, to the 3D environment. Each pixel has a gray value proportional to the height of the underlying 3D location.

Computing the elevation map (Fig. 4) from the set of dense 3D points is straightforward. A 3D space of interest (40m x 12m from bird-eye view) in front of the car is considered. The longitudinal Z and lateral X coordinates of each 3D point are scaled into the (image) coordinates space of the elevation map. The elevation map image has the same aspect ratio as the 3D space of interest. A cell has a size of 7.5 cm x 7.5 cm in the XZ world plane.

Each cell (32 bits) of the elevation map will contain the highest height of the 3D points contained in the cell, with a resolution of 4 mm. We centered the world zero level (road level at system calibration) at 128 in the elevation map for a better visualization. For all the images in this paper only heights around the road (1 m band) are displayed correctly due to a limited number of gray values (256).

3D points higher than 2 meters from the ground (zero level from calibration) will not be stored because they are out of interest. Empty cells are flagged and not used further.

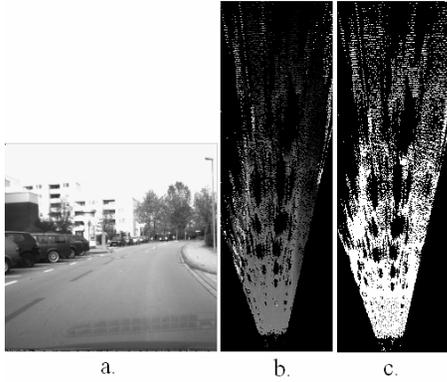


Fig. 4. The elevation map (b) for the scene in a. In c valid cells (non-empty, with 3D data) are highlighted.

As seen in fig. 4, the elevation map presents poor connectivity between road points at far depths. The amount of dilation needed for each cell is computed by applying (8) in the cells space. To compensate for possible vertical road curvatures the amount of dilation obtained will be increased by 50%. This is necessary to avoid under-connectivity. Figure 5.a shows the dilated elevation map.

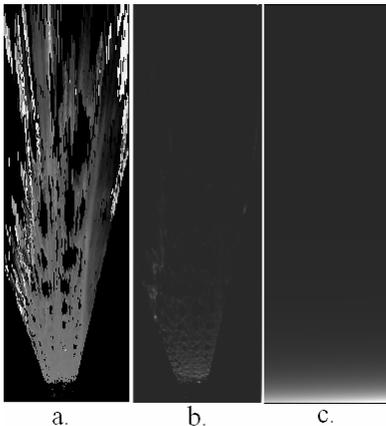


Fig. 5. The dilated elevation map in a. The averaged measured density of 3D points per cell is shown in b. The road expected density of 3D points per cell in c.

Another two features, related to the density of 3D points, are computed for each cell: *expected road density* and average *measured density*. These two features can be used for a rough discrimination between road and non-road features.

By using the depth resolution model presented above, the expected density of road points in each cell can be computed (Fig. 5.c). This density has a large range of values from about 200 3D points/cell near the ego-vehicle, down to 0.1 points/cell at 35 meters depth.

When the elevation map is formed, a counter is stored for each cell. This counter shows how many 3D points are contained in each cell. For road points at far depths this counter is not equivalent with the local measured density. This happens because of the decrease in depth resolution. The local measured density can be estimated by averaging the map of counters with an adaptive mask, with a size equal

to the amount of dilation used for connectivity. The average measured density is obtained (Fig. 5.b).

As a conclusion, instead of a set of 3D points without explicit connectivity, we obtained a discrete 3D space as an elevation map. Full connectivity of road points is available. For the sake of simplicity, in the next subsections (B to F) we will use the term *point* (described 3-dimensionally as the map coordinates and height) for a cell of the elevation map. All formulas presented in section II and III can be applied in the elevation map space (it is still a 3D space, but with discrete locations).

B. Selection of points for initial surface fitting

The RANSAC approach [8] is a robust method for fitting a model to a data set containing noise. Instead of fitting the model to the whole set in a least square fashion (due to noise the solution will not be optimal), the RANSAC approach chooses a number of samples (subsets of the data set). For each sample the model is fitted and a score is computed. The sample with the highest score is selected.

A rectangular patch from the elevation map is selected in front of the ego car. Valid points from this patch will represent the data set. Not all the points should be selected because there are two situations when even RANSAC will fail (Fig. 6): when traffic isles or obstacles are dominant in the selected patch.

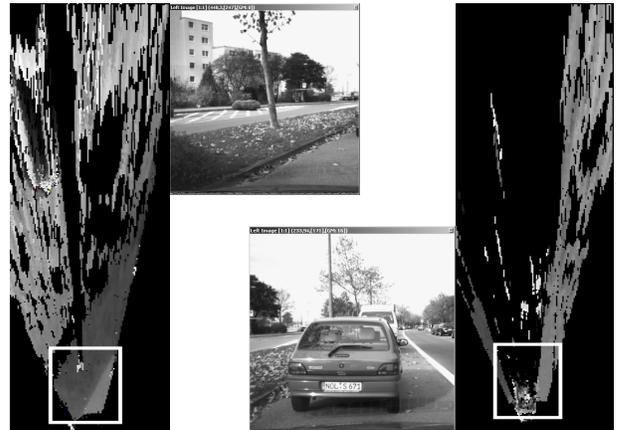


Fig. 6. Two situations where RANSAC can fail. The patch selected for initial fitting is shown as a white rectangle.

Two constraints are used to filter (result shown in Fig. 7.d) the data set before applying RANSAC:

- Curbs are detected in the rectangular patch. Only points placed on the same side as the ego car, relative to the detected curbs, are considered.
- Each considered point must have a computed density no more than 150% of the estimated road density for that point (most of object points should be discarded).

To detected curbs in the interest patch, edge points are detected on the elevation image, and the Hough transform [9] is used to select relevant lines. Five lines are selected, having the highest Hough score. Each line is analyzed by

counting how many of its points have a height variation between 5 cm and 35 cm (normal range for curbs). Lines with a score higher than 40% of the total number of line points are considered valid. Two of the valid lines, with the highest scores, are selected as curbs. This method for detecting nearby curbs has a good detection rate of about 1% false positives (evaluated on 120 urban scenes with curbs).

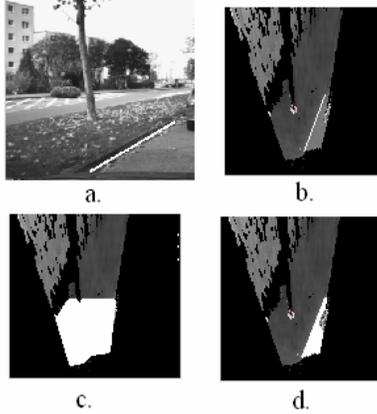


Fig. 7. One curb is detected on the elevation map in b and projected onto the left image in a. From the valid points in c only highly probable road points are selected in d for RANSAC.

The RANSAC method is applied to the filtered set of points. We used a number of 200 samples of 5 points each. The quadratic surface was computed for each sample. The score for each sample was considered as the number of inliers from the whole set. The uncertainty model from section III.B was used to classify a point as part of the road surface. The sample with the highest number of inliers is selected as the primary road surface. If the total surface of inliers is less than (equivalent in 3D) 1 m^2 , then the detection of the road surface is aborted and only the density-based method proposed in section IV.E is used for obstacle detection.

C. Uncertainty model-driven surface growing

The primary road surface is detected optimally for the points in front of the vehicle. Optimality is not granted for the whole scene because vertical curvatures of the road surface can be computed correctly only using large road patches.

The primary solution can be refined through a region growing process (Fig. 8) where the initial region is the set of inliers from the initial rectangular patch. A new point can be added to the current region if it fulfills the following conditions:

- To be adjacent to inliers from the current region,
- To verify the current road surface equation according to the uncertainty model described by equations (8) and (9) with $D_{err}=1$ pixel.

The surface is re-computed, in a least square-fashion, each time the region has expanded its border with 1-2 pixels (about 100 new points). This insures that the surface is

refined gradually. On average the surface is recomputed about 150-200 times. This can be very time consuming since it involves computing the sums defined by (6). A real-time implementation is possible by using the partial sums between two consecutive re-computations (only the amounts for the new points are added).

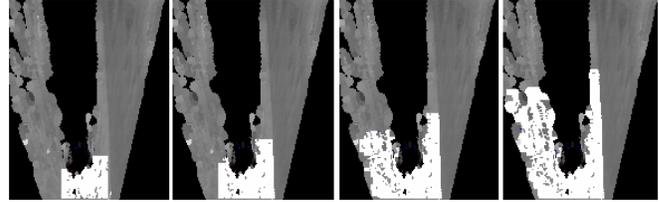


Fig. 8. Road inliers detected in the initial patch (as white). Intermediate regions, growing from left to right.

A global optimal solution for the road surface is obtained after this step.

D. Rough classification of points

Elevation map points are roughly classified (Fig. 9.c), with respect to the road surface, considering the following rules:

- If the height of the point P relative to the road is below the estimated height uncertainty (according to the defined model, with $D_{err}=1.5$ pixels) the point is considered **Road**,
- Otherwise, let us consider $Quotient = ExpectedDensity(P) / MeasuredDensity(P)$:
 - If $Quotient$ is higher than 1 and the height of P from the road is below 45 cm, the point is considered **Traffic Isle**,
 - Otherwise ($Quotient$ less than 1):
 - P is considered **Obstacle** if its height from the road is higher than $Quotient * 60$ cm (in other words the lower its density is, the higher it should be from the road in order to be object point),
 - Otherwise P is considered **Traffic Isle**.

The rough classification provides good positives but also some false positives and negatives (fig. 9.c).

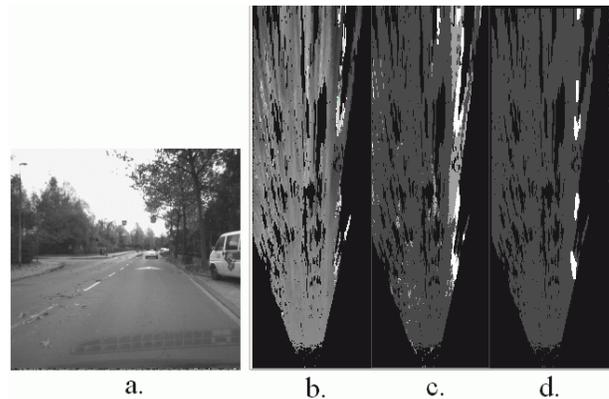


Fig. 9. The elevation map (b) of a scene (a) is roughly classified (c) (dark gray - road, gray - traffic isle, white - obstacles). The result of the density-based classification is shown in d (dark gray - road, white - obstacles).

E. Density-based obstacle classification

When the elevation map is built, two features are also computed for each point (cell): the expected road density (the theoretical density of 3D points in the cell) and the measured density (what actually exists).

The expected road density is computed for a planar road surface but it can be easily used for vertically curved roads (curvatures are small). This is done by simply using a larger value for the expected density (50% more than computed).

Obstacles points will have much larger densities than the road for the same depth. We use a standard approach called double thresholding, which is frequently used in edge detection [10].

The following steps are performed to detect obstacle points on the elevation map (result shown in Fig. 9.d):

- Points are flagged as **Obstacle** if the measured density of the point is higher by TH (=6) times than the estimated road density,
- Other points are flagged (recursively) as **Obstacle** if they are adjacent to an **Obstacle** point and the measured density of the point is higher by TL (=3) times than the estimated road density.

The threshold values TH and TL were chosen experimentally, considering the following constraints:

- The rate of false positives is below 1%,
- Small obstacles such as near road poles (7-8 cm thin and 20 cm tall) should also be detected.

This detector has a very low rate of false positives, but some objects (about 10%) are only partially detected due to poor dense reconstruction. The method is used as a stand-alone detector (if the road surface cannot be computed) or it is robustly combined with the rough classification described previously.

F. Obstacle filtering and classification

This step performs fusion between the result of rough classification (Fig. 9.c) and the density-based obstacle classification (Fig. 9.d).

The fusion and error filtering are performed using the following rules:

- Small **Traffic Isle** areas from the rough classification are discarded (less than 0.5 m^2),
- **Obstacles** areas from the rough classification are rejected if they do not overlap obstacle points from the more robust density-based classification,
- Our experiments shown that for depths higher than 25 meters, **Traffic Isle** detection is less reliable and provides lots of false positives. For depths higher than 25 meters only the result of the density-based classification is used.

False elevations (dark gray spikes in Fig. 10.a) due to wrong dense 3D data are correctly classified as drivable (road).

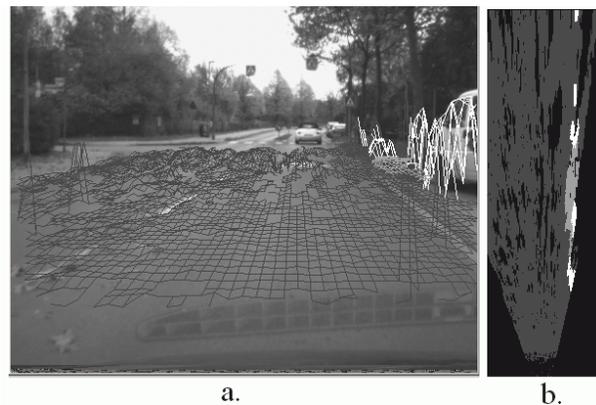


Fig. 10. The final result after error filtering and classification is shown on the left image (a) for the classified elevation map (b). Cells from the elevation map are the vertices of the image projected grid (dark gray - road, gray - traffic isle, white - obstacles).

V. RESULTS

The algorithm was implemented in C++. The dense 3D information was generated using a calibrated stereo rig with grayscale cameras and a commercial dense stereo board [11].

Due to the use of software-specific C optimizations and the elevation map representation, a processing time of 20 ms was achieved for the algorithm itself (on Pentium 4 at 2.6 Ghz). Overall, with the image acquisition and the dense hardware reconstruction, a sustained processing frame-rate of 23 fps is obtained. By using a more recent processor (Pentium Dual Core) a processing frame-rate of 35 fps can be achieved (for our system the cameras have a limited acquisition frame-rate of 25 fps).

Regarding the robustness of the algorithm we performed only a simple evaluation which will be extended in the future:

- A number of 40 stereo images of different (random) scenes were selected (out of several hours of stored stereo images, recorded while driving the ego-car).
- Results were analyzed for each frame in terms of: missed obstacles, partially detected obstacles, missed traffic isles, false traffic isles.

The numbers obtained proved the robustness of the algorithm:

- Missed obstacles: 3 out of 153 - very small objects,
- Partial obstacles: 8 out of 153 - due to the lack of dense 3D data,
- Missed traffic isles: 2 out of 28 - the height was too small (4-6 cm),
- False traffic isles: 1 - on one of the scenes (road junction) the quadratic road model was not good because the road surface was made of 2 main quadratic surfaces with different parameters.



Fig. 11. Results for various scenes. The algorithm performs well even for scenes with sparse 3D road reconstruction and noisy 3d data (the grid is projected only where dense 3D data is provided by the dense stereo engine). The result for the first scene is displayed using a virtual camera. Grid colors represent dark gray - road, gray - traffic isle, and white - obstacles.

VI. CONCLUSIONS

A new road and obstacle detection algorithm was presented. It transforms the 3D dense data from stereovision into an elevation map. The depth uncertainty and resolution are modeled and used to fit the road quadratic model to the elevation map. An RANSAC-approach, combined with region growing, is used for the detection of the optimal road surface. Obstacles and traffic isles are detected by using the road surface and the density of points.

The algorithm proved to be robust. Most of the erroneous detections are caused by bad 3D data from dense stereo. The current version copes well with most (but not all) types of errors in the 3D data set.

The proposed algorithm works in real-time and provides robust results. It can be used in a variety of applications, from ego-pose estimation to complex 3D path planning.

Future development and evaluation are required:

- Further and more complex evaluation of the proposed algorithm,
- Using tracking can greatly improved the robustness of the method,
- A more complex road model will be proposed and tested (cubic or 4th degree, or a cubic-spline surface - might be robust even for off-road driving),
- Using MMX/SSE processor features can reduce the processing time to half,
- Enhancing the algorithm to cope better with various types of noise in the 3D data set.

REFERENCES

- [1] 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.
- [2] A. Broggi, C. Caraffi, P. Paolo Porta, and P. Zani, "The Single Frame Stereo Vision System for Reliable Obstacle Detection used during the 2005 DARPA Grand Challenge on TerraMax", in *IEEE Intelligent Transportation Systems Conference*, Toronto, Canada, September 17-20, 2006, pp. 745-752.
- [3] A. Sappa, D. Gerónimo, F. Dornaika, and A. López, "Real Time Vehicle Pose Using On-Board Stereo Vision System", *Int. Conf. on Image Analysis and Recognition*, LNCS Vol. 4142, Springer Verlag, Póvoa de Varzim, Portugal, September 18-20, 2006, pp. 205-216.
- [4] M. Cech, W. Niem, S. Abraham, and C. Stiller, "Dynamic ego-pose estimation for driver assistance in urban environments". In *IEEE Intelligent Vehicles Symposium*, Parma, Italy, 2004, pp. 43-48.
- [5] S. Nedeveschi, R. Danescu, D. Frentiu, T. Marita, F. Oniga, C. Pocol, R. Schmidt, T. Graf, "High Accuracy Stereo Vision System for Far Distance Obstacle Detection", *IEEE Intelligent Vehicles Symposium (IV2004)*, Parma, Italy, pp. 292-297, 2004.
- [6] S. Nedeveschi, R. Schmidt, T. Graf, R. Danescu, D. Frentiu, T. Marita, F. Oniga, C. Pocol, "3D Lane Detection System Based on Stereovision", *IEEE Intelligent Transportation Systems Conference (ITSC)*, Washington, USA, pp.161-166, 2004.
- [7] S. Nedeveschi, F. Oniga, R. Danescu, T. Graf, R. Schmidt, "Increased Accuracy Stereo Approach for 3D Lane Detection", *IEEE Intelligent Vehicles Symposium*, (IV2006), June 13-15, 2006, Tokyo, Japan, pp. 42-49.
- [8] M. Fischler, R. Bolles, "Random sample consensus: A paradigm for model fitting with applications to image analysis and automated cartography", *Graphics and Image Processing* 24(6) (1981), pp. 381-395.
- [9] D. H. Ballard, "Generalizing the Hough Transform to Detect Arbitrary Shapes", *Pattern Recognition*, 13(2):111-122, 1981.
- [10] J. Canny, "A computational approach to edge detection", *IEEE Trans. Pattern Analysis and Machine Intelligence*, pp. 679-698, Nov. 1983.
- [11] J. Woodfill, G. Gordon, D. Jurasek, T. Brown, and R. Buck, "The Tyx DeepSea G2 Vision System, A Taskable, Embedded Stereo Camera", *IEEE Computer Society Workshop on Embedded Computer Vision, Conference on Computer Vision and Pattern Recognition*, (New York, NY), June 2006.