

A Sensor for Urban Driving Assistance Systems Based on Dense Stereovision

Sergiu Nedevschi, Radu Danescu, Tiberiu Marita, Florin Oniga, Ciprian Pocol, Stefan Sobol, Corneliu Tomiuc, Cristian Vancea

Marc Michael Meinecke, Thorsten Graf, Thanh Binh To, Marian Andrzej Obojski

Technical University of Cluj-Napoca, ROMANIA
e-mail: Sergiu.Nedevschi@cs.utcluj.ro

Volkswagen A.G. GERMANY
email: marc-michael.meinecke@volkswagen.de

Abstract—The urban driving environment is a complex and demanding one, requiring increasingly complex sensors for the driving assistance systems. These sensors must be able to analyze the complex scene and extract all the relevant information, while keeping the response time as low as possible. The sensor presented in this paper answers to the requirements of the urban scenario through a multitude of detection modules, built on top of a hybrid (hardware plus software) dense stereo reconstruction engine. The sensor is able to detect and track clothoid and non-clothoid lanes, cars, pedestrians (classified as such), and drivable areas in the absence of lane markings. The hybrid stereovision engine and the proposed detection algorithms allow accurate sensing of the demanding urban scenario at a high frame rate.

I. INTRODUCTION

Urban Driving Assistance Systems have little in common with the systems targeted for highways. Not only is the problem itself more complex, with difficult and atypical road geometries, crowded traffic, the presence of pedestrians and other type of traffic participants which are not normally found on the highway, but the expectations from such a system are considerably higher. The detection errors are not tolerated, the measurement errors need to be considerably smaller, and the field of view must be considerably increased, to account for the lateral objects. The thematic network ADASE (Advanced Driving Assistance Systems in Europe, www.adase2.org) helps define the driving assistance requirements by harmonizing and communicating active safety functions, identifying technological needs and focusing on essentials and preparing architectures, roadmaps and standards.

A realistic analysis of the requirements and of the possibilities of the urban traffic environment leads to the establishment of several goals for urban traffic assistance, to be implemented in the near future:

- Follow-to-Stop
- Stop to fixed/ non-moving Obstacles/ Vehicles
- Go Inhibit if Objects are in front of the own Vehicle
- Lateral Support in narrow road conditions
- Set of Max-Velocity depending on lanes width

Manuscript received 15.01.2007. The work was funded by Volkswagen A.G. in the framework of the DESBOR (Dense Stereo-Based Object Recognition) project.

- Go (automatically/ driver initiated)

In order to achieve these urban ACC applications, a vision sensor must provide the following functions:

- Lane Detection / Lane Parameters Estimation
- Navigable channel detection and channel parameters estimation in crowded environments
- Vehicle Detection and tracking
- Detection of fixed (non-moving) Obstacles
- Classification of Pedestrians

The sensorial systems for driving assistance (highway and urban) are today the focus of large, joint research projects, which combine active and passive sensors, GPS navigation, and telematics. Projects such as INVENT (www.invent-online.de), PREVENT (www.prevent-ip.org), CARSENSE (www.carsense.org) bring together car manufacturers and research partners for the common goal of solving the driving assistance problem. Sensing in urban environments is also a long-time effort of the DaimlerChrysler research department [1].

Many new research papers deal with problems that are present in the urban driving environments. The researchers at Toyota [2] present a stereovision-based system combined with a near infrared projector for road and obstacle detection in any environment. The unstructured scenario driving is approached in [3], [4], and a combined, structured plus non-structured system is presented in [5]. A dense stereo system for obstacle detection for go inhibit, with possible applications in urban environments is presented in [6]. New methods for lane detection, suitable for urban environments, are presented in [7] and [8].

The research team of the Technical University of Cluj Napoca, in cooperation with Volkswagen AG, has already implemented a stereovision-based sensor for the highway environment [9]. This sensor was able to detect the road geometry and the obstacle position, size and speed, from a pair of synchronized grayscale images, using edge-based software stereo reconstruction with general geometry.

The urban scenario required important changes in the detection algorithms, which in turn required more stereo information. Thus, the edge-based stereo engine was discarded, and replaced with a dense stereo system. A software dense stereo system being time consuming, a hybrid solution was chosen: software rectification and down sampling, followed by hardware correspondence search. The

time gained by the hardware part compensated the increase in complexity of the new algorithms.

The dense stereo information is vital for the new obstacle reconstruction module, which extracts oriented objects even in serious clutter, and also allows better shape segmentation for recognition of pedestrians. Dense stereo information allows us to compute and track an unstructured elevation map, which provides drivable areas in the case when no lane markings or any other road delimiting features are present or visible.

Lane detection requires edges, but the vertical profile is better computed from dense stereo information. The edge based lane detection algorithms are completely changed, adapted to the limited and variable viewing distance of the urban environment. A freeform lane detection module was added, in order to solve the problem of the non-standard geometry roads.

The dense stereovision based sensor presented in this paper provides complex and accurate functionality on a conventional PC architecture, covering many of the problems presented by the urban traffic environment, and promising to be a valuable addition to a driving assistance system.

II. STEREOVISION SYSTEM ARCHITECTURE

A. Image Acquisition

The hardware acquisition system (fig. 1) includes two JAI CV-M4+CL cameras with 2/3" (1380x1030) CCD sensors and 6.5 mm fixed focal length lenses, allowing a horizontal field of view (HFOV) of 72 [deg]. The cameras are mounted on a rigid rig with a baseline of 320 [mm] (fig. 2). The images are acquired at full resolution with a microEnable3-XXL acquisition board with a maximum frame rate of 24 fps.

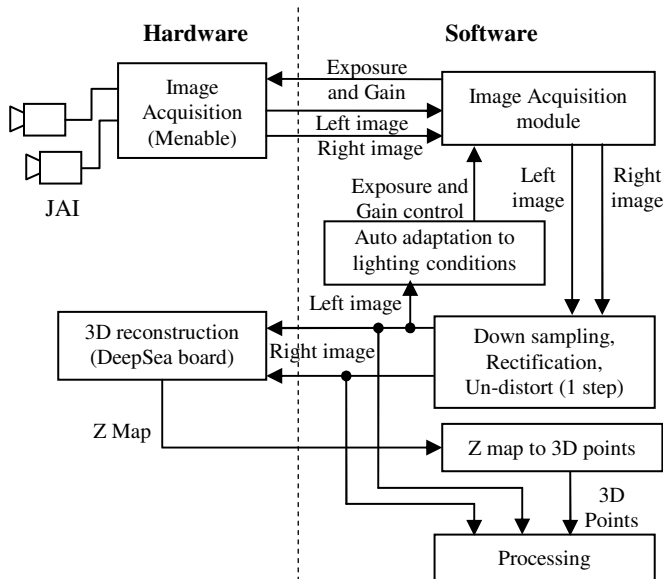


Fig. 1. The stereovision system architecture.

The camera parameters are calibrated using a dedicated method optimized for high accuracy stereovision [10] using the full resolution images.

The images are further enhanced by lens distortion correction and rectified in order to fulfill the dense stereo reconstruction requirements (canonical images). A down-sampling step is used to adapt the image size to the DeepSea board parameters (53 pixels width) and to minimize the noise introduced by the digital rectification and image correction. The whole process is reduced to an image warping approach performed in a single step (fig. 1) using reverse mapping and bilinear interpolation [11]. An optimized implementation using MMX instructions and lookup tables was used in order to minimize the processing time.

B. 3D Reconstruction

The 3D reconstruction of the scene is performed using the DeepSea hardware board provided by TYZX [12]. The input of the board consists in two rectified images and the output can be either a disparity or a Z map (expressed in the left camera coordinate system). Our system uses 3D points set for scene representation; therefore the preferred output is the Z map. Using the Z coordinate value, the X and Y coordinate can be computed and then transformed into the car coordinate system using the extrinsic camera parameters [10].

C. Detection Range

With the current system setup a detection range optimally suited for the urban environments is obtained (fig. 2):

- minimum distance: 0.5 m in front of the ego car (approximately 2.5 m in front of the cameras) – the near range distance is limited by the baseline, focal length and maximum disparity allowed by the DeepSea board.
- delimiters of the current lane (considered approximately of 3,5 m wide) are visible at 1.0 m;
- reliable detection range: 0.5 ... 35 m, with a maximum detection range (up to which 3D points can be reconstructed) of 50 m;




Fig. 2. Detection range of the current stereo system setup.

III. LANE DETECTION

The urban environment brings new requirements for the lane detection system. The scenes are more complex, the presence of nearby vehicles and other obstacles shorten the visibility range, and the geometry of the lanes does not always comply with the clothoid model. Therefore, a serious method overhaul was needed.

The first step is to detect the vertical profile and to separate the road obstacle points. The pitch angle and the vertical curvature are detected using histograms, in a manner similar to the Hough transform [13].

The highway lane detection approach required little information about lane markings, because it could rely greatly on the lane model. For the urban environment, however, we require a fast and robust lane marking extraction algorithm.

The lane marking extraction method relies on the well-known dark-light-dark transition detection [14]. We have to search for pairs of gradients of opposing sign and equal magnitude. We have improved the method by using a variable filter for computing the horizontal gradient. The size of the filter is the size of a standard width lane marking projected in the image space, and varies because of the perspective effect. Applying the variable width filter we preserve the level of detail in the distance while filtering the noise in the near areas.

The gradient maxima and minima are paired and the DLD pairs are extracted as lane markings. The complete technique is described in [15].




Fig. 3. Lane marking detection – top left, original image; top right, results of the adaptive gradient filter; bottom, lane marking results

Although the clothoid model is not always accurate for the urban scenario, it has several benefits, such as good results when the lane is delimited by simple edges (unmarked roads). Due to the short visibility range, we have decided to avoid matching the whole clothoid model on the image data, but to match pairs of line segments instead, in two zones: near and far.

First, we make an attempt for the near zone (2m to 5 m). Hough transform is performed on the image edges

corresponding to the road points, and line segments are extracted. Lane markings will have a higher weight in the Hough bins, and therefore they will have a higher priority. We divide then the line segments in two sets – left and right. The segments on the left are paired with the segments on the right, and the best pair is selected as our lane measurement. The linear measurement will update the clothoidal model parameters using the Extended Kalman Filter. If the linear fit for the near zone is successful, the same is done for the far zone, and the model parameters are updated again.



Fig. 4. Updating the clothoidal lane profile from linear segments using the Extended Kalman Filter

Sometimes the clothoid lane model is not suited for the road we wish to observe, and the detection will be incorrect. For these cases, a freeform lane detection system has been implemented. Because we don't have strong models to guide our search, we have to discard the non-marking delimiters, and work with lane markings only. The markings are projected onto a top view image, and then distance transform is performed, to facilitate the fitting of a lane border. The left and the right lane borders are represented as Catmull-Rom splines with four control points. The lateral coordinates of the four control points are the model parameters, and they are found using a simulated annealing search strategy in the model space.




Fig. 5. Top view of the lane markings and the distance transform image used for freeform lane matching




Fig. 6. Freeform lane detection succeeds in situations where the clothoid model fails.




Fig. 7. Freeform detection versus model-based detection, bird-eye view comparison

IV. DRIVABLE AND NON-DRIVABLE AREAS DETECTION

There are some urban scenarios where the 3D lane cannot be detected, especially when not enough lane delimiters exist (ex. road crossing). An alternative method must be used to detect elevated areas (obstacles), regions where the ego vehicle cannot be driven. Complementary, the obstacle-free road areas can be considered as drivable.

The dense stereo engine usually reconstructs most of the road surface points even if lane markings are not present. Thus, the surface of the road can be computed by fitting a geometric model to the 3D data. The fitting must be performed in a least-square fashion (LSQ), or, more robustly, using a statistical approach (ex. RANSAC). The model used for the road is a planar one, allowing for non-zero pitch and roll angles of the ego car. The algorithm can be extended to fit complex surfaces, such as quadratic or cubic.

The 3D data available is a set of 3D points (80,000 to 30,000). Fitting the road surface to this set, in real-time, is not possible because it has a high computational complexity. A (bird-eye rectangular, 3x35 meters) region of interest of the 3D space can be represented similar to a digital elevation map. An image of elevations is formed, with each pixel (cell) having the intensity proportional to the 3D height (fig. 8.b). If a cell has more than one 3D point, then the greatest height is used. Morphological dilation is used to fill voids and compensate for the perspective effect (the 3D space gets sparser with the depth) (fig. 8.b).

Since the LSQ fitting is sensitive to noise in the data set, a pre-selection of candidate road points is required. Most of the 3D points in front of the ego car are road points if no nearby large obstacle is occluding the road. A histogram of intensities from the lower part of the elevation image is computed (fig. 9) and the dominant intensity is selected. Candidate road points are those having the intensity close to

the dominant one (in the range depicted by arrows in figure 9.b).




Fig. 8. The urban scenario (a) and the elevation image (b. initial, c. with morpho-dilation). Darker means more elevated.




Fig. 9. The road level is selected from the patch in a, using the histogram in b. Candidate points are shown in c. The final result is presented in d (white means road and gray means elevated).

The road model is fitted to the selected road points in a LSQ fashion by minimizing the error along the vertical direction (vertical axis in the 3D space). A better approach can be used, the RANSAC algorithm, if the image quality is poor and a lot of noise exists in the 3D data.

For the final classification into drivable/non-drivable areas, the depth uncertainty model from [16] was extended to a height uncertainty model (1). The expected uncertainty Y_{err} is a function of the height Y and the depth Z of the 3D point, height of the camera H_{cam} , and the estimated depth uncertainty Z_{err} . Z_{err} is also a function of the system parameters and the expected disparity uncertainty. The disparity uncertainty was chosen experimentally as 1.5 pixels, although a more complex model for estimating the correlation's accuracy can be developed.

$$Y_{err} = \left| \frac{(Y - H_{cam}) * Z_{err}}{Z} \right| \quad (1)$$

Each point is labeled as road if it is closer to the road surface than its estimated height uncertainty. The result for the scenario in figure 8.a is shown in figure 9.d. Another result is presented in figure 10, re-projected as a grid onto the left image.




Fig. 10. The result for two scenes re-projected back as a grid onto the left image: pedestrians, sidewalks and vehicles (dark gray) are separated from the road surface (white).

V. OBSTACLE DETECTION

The road/obstacle separation is done using the lane's vertical profile. The 3D points situated in a convenient height range above the road (for example 0 .. 1.5 meters) are used by the obstacle detection algorithm (fig. 11.b).

Ideally, an obstacle detection algorithm should meet two opposite requirements:

- do not split a real obstacle into smaller parts;
- do not merge more distinct obstacles into a single one.

Previous experience has shown that these requirements are hard to be fulfilled mainly due to the reconstruction errors of the 3D points.

The current approach prefers to merge first more distinct obstacles into a single one, named *occupied area*, and then, to fragment it upon its components.

It is supposed that the obstacles do not overlap each other; no obstacle is hanged up in the air above other one. In other words, in a top view (fig. 11.c) the obstacles are disjoint.

Consequently, in what follows, the Y coordinate (the elevation) of the 3D points will be ignored and all the processing is done using only the X and Z coordinates (top view).

To compensate the problem of the points' density variation with the depth (less points for higher distances), a method used to divide the top view Cartesian space into smaller tiles with constant density is proposed (fig. 11.c). The horizontal field of view of the camera is divided into slices with constant aperture in order to compensate the sparseness on the X-axis. The depth range is also divided into intervals having bigger lengths as the distance grows in order to compensate the sparseness on the Z-axis.

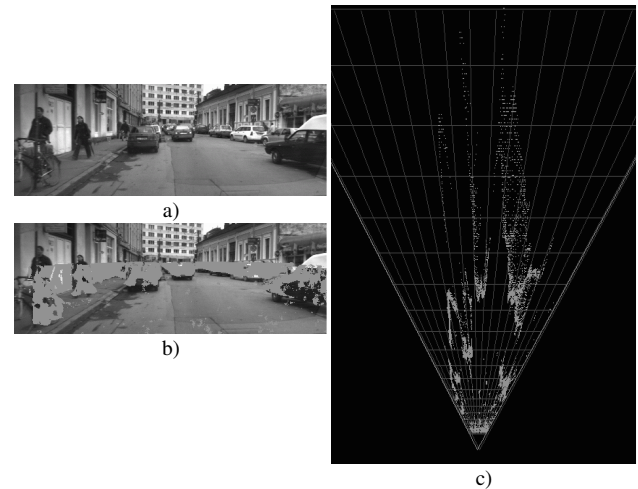


Fig. 11. Tiles division. a) Gray scale image, b) 3D points – perspective view, c) 3D points – top view; the tiles are here considerably larger for visibility purpose. Wrong reconstructed points can be observed in random places. Reconstruction error is visible as well (spread points for further obstacles).

A specially *compressed space* is created (fig. 3.a). The cells of the compressed space correspond to the trapezoidal tiles of the Cartesian space from figure 11.b. The compressed space is, in fact, a bi-dimensional histogram, each cell counting the number of 3D points found in the corresponding trapezoidal tile (fig. 3.a). The cells without any point represent free space. The cells with few points are also considered free (most probably those points were wrong reconstructed). The other cells, with enough points are signaling the existence of obstacles.




Fig. 3. The compressed space (for scene from figure 11): a) bi-dimensional histogram counting 3D points, b) labeled adjacent cells.

On these cells of the compressed space (having many points), a labeling algorithm is applied: it groups adjacent

high density cells (fig. 3.b). Each group of cells, obtained by the labeling algorithm, represents an *occupied area*. The small groups are filtered out.

Better obstacle detection requires the fragmentation of occupied areas into primitive obstacles. A *primitive obstacle* is defined as an obstacle without concavities. An obstacle with concavities must be fragmented into primitive obstacles.

The steps of the fragmentation process are exemplified in figure 13 for an inward corner of a building.

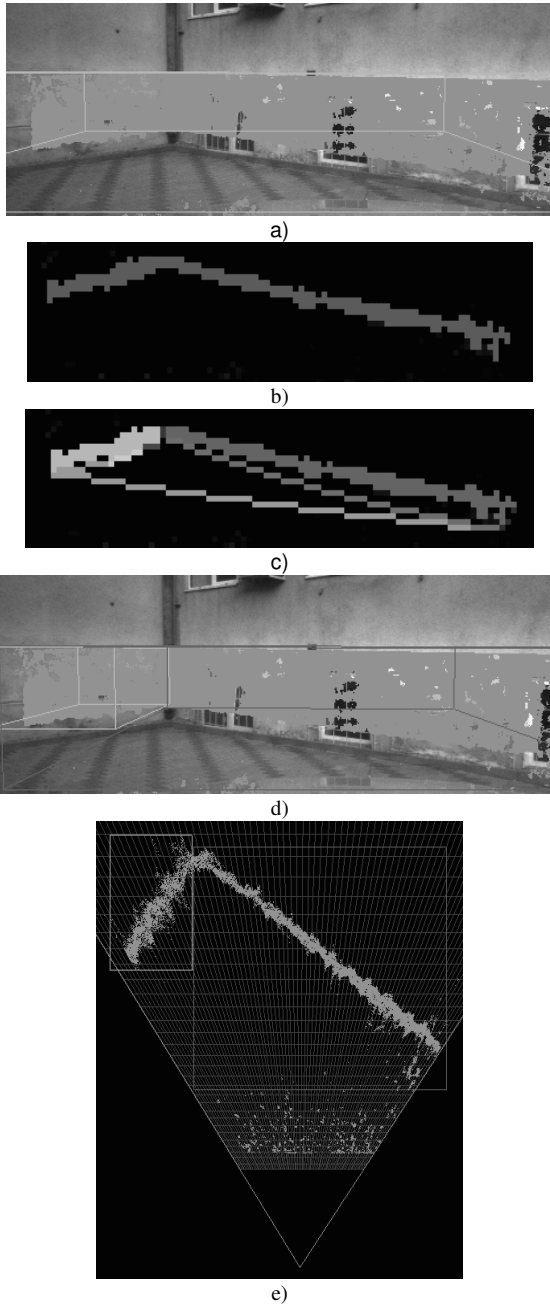


Fig. 13. Fragmentation of occupied areas into primitive obstacles: a) an occupied area, b) the labeling in the compressed space, c) sides of the envelope and the two primitive obstacles of the occupied area – compressed space, d) the two primitive obstacles – perspective view, e) the two primitive obstacles – top view

The idea is to determine the envelope of the cells of an occupied area. Then for each side of the envelope, the concavity between the side and the occupied cells is determined. If it is big enough, there are in fact two primitive obstacles (or occupied areas) and the deepest point of the concavity gives the column where the division will be done. The two sub-parts are subject to be divided again and again as long as concavities are found.

In figure 13.c the bottom side of the envelope for the cells in figure 13.b delimits a big concavity. For each new sub-part, the envelope of the cells has been computed again (and painted as well), but without revealing big concavities for new division steps.

By reconsidering the coordinates (including Y) of the 3D points that have filled the cells of an obstacle, the limits of the circumscribing box are determined. Boxes are shown in figure 13.d (perspective view) and figure 13.e (top view).

As observable in figure 13.e, the bounding boxes, which are parallel with the coordinate system's axes, still encompass free space. The real obstacle is oblique oriented in the coordinates system. The "silhouette" of the 3D points of this obstacle reveals a straight boundary. The envelope of the points can model this silhouette, and then, an analysis of its visible sides can determine the orientation. The visible sides are those in the front of the obstacle, toward the camera. If the analysis cannot determine a preponderant orientation of these sides, the box remains parallel with the axes of the coordinates system (un-oriented box).

In figure 14.a such an envelope is shown. The algorithm searches chains of consecutive sides having components with similar slopes. If the length of the longest chain is at least 70% from the length of the visible envelope, the obstacle orientation is the weighted average of the slopes of the chain's sides. The weights used are the lengths of the sides.

Along the found orientation, a rectangle is fitted on the visible sides. It gives the base of the oriented box of the obstacle. The result is shown in figure 14.b.




Fig. 14. Obstacle orientation: a) The computed envelope of the 3D points of one obstacle is shown with light gray. b). The oriented obstacles in perspective view




Fig. 15. Obstacle detection results with oriented 3D boxes in the perspective view of the traffic scenario from figure 11.a.

As future work the tracking of obstacle fragmentation decisions and obstacle orientation is needed. This would stabilize the detection over frames. Criteria based on optical flow/motion field would also help to merge and/or divide obstacles. Big obstacles that have a convex top view shape (envelope) but do not look like a rectangle could be divided also in sub-obstacles or could be described by a free shape polygonal model.

VI. OBJECT CLASSIFICATION

The object classification is performed using a new approach to pattern matching exploiting both 2D image information and 3D dense stereo information. Because the 3D information accuracy does not allow the direct classification of the 3D shape, a combined 3D-2D method is used. The 3D object information consists of 3D position, size and aspect ratio, obtained from dense stereo data through the obstacle detection's grouping process. The 3D data is used for model selection and scale estimation. Based on the 3D orientations of the object, the appropriate 2D views of the model are generated from a database of 3D models. The corresponding 2D image window is selected based also on the 3D object information. The scaled models are matched against the features found in the selected window using an elastic high speed matching based on chamfer distances.

An important advantage in the use of dense 3D information is the possibility to make a better separation between the object and its background. This separation is done based on the difference in depth between the object of interest and the background behind it. A clear separation of the features belonging to the object allows the use of pattern matching classification techniques with a higher success rate since no additional noise derived from the object's background will be considered in the feature set for the classification process.




Fig 16. Object segmentation

To be able to handle a wide variety of objects the model database can contain a mixture of 2D and 3D models. For simple objects with a rigid structure, the model database can

contain 3D models. The orientation of the object relative to the camera is used in the classification process to transform the 3D models in 2D models using a projection according to the specific viewing angle. This allows the classification of objects that have a general orientation related to the observer. For objects that have a more complex shape or do not have a rigid structure, 2D models representing different points of view and different states are stored directly in the database. The orientation of the object can still be used to select the correct 2D model for these complex objects. The direction of movement associated with the 3D object is used to determine if the object is seen from the front or from the rear, or for pedestrians to determine which way they are facing.

The model selection is done based on the 3D size and aspect ratio of the object (car models, truck models, pedestrian models, etc).

The 3D sensorial information provides the exact positions and dimensions of the objects in the scene. By projection, the size and position of the objects in the 2D image can be evaluated and the model scaling factor can be deduced.




Fig. 17. 2D model determination: a) Top view of 3D object with orientation b) Object features c) 3D model with point of view d) Model projection

The use of dense stereo information in the classification process provides significant improvements both in speed and success rate for the classification process. The 3D information allows the direct determination of the parameters for the model (type, scale factor, approximate position) which leads to a reduction in computational time. The selection of relevant features for the object (background separation) improves the success rate of the classification system.

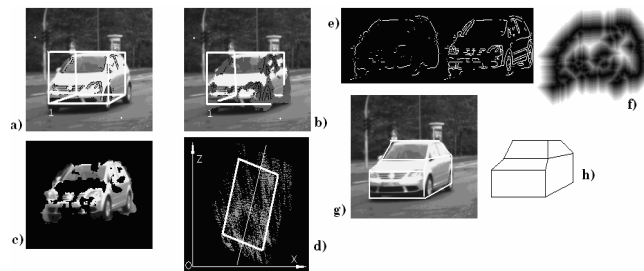


Fig 18. Results: a) Object selection (ROI selection) b) Used 3D information c) Image feature selection d) Object orientation e) Edge selection f) distance transform g) Result of pattern matching h) Model used

The same approach has been successfully used in the case of pedestrian classification. The only modification in the process is the use of a hierarchical model database to reduce the number of pedestrian models used in the pattern matching stage.



Fig.19 Pedestrian classification.

VII. EXPERIMENTAL RESULTS

The dense stereovision based sensorial system was subjected to extensive test situations, targeted for each algorithm. The tests included recorded urban images from Wolfsburg (Germany) and Cluj Napoca (Romania), in light, medium and heavy traffic, in different times of the day and different weather conditions. The system was also tested online (on board of the vehicle, while driving) in the same situations. The lane detection, object detection and pedestrian detection algorithms performed correctly under a vast majority of the test situations, and the failed test results provided valuable information for future algorithm improvement. The detection range was limited by the stereo reconstruction capability of 6.5 mm focal length cameras to 35 m, a tradeoff for an increase in the field of view. The time performance of the whole processing cycle on Pentium Core 2 Duo E6600 architecture (2.4 GHz) is 20 frames per second.

VIII. CONCLUSIONS

We have proposed and implemented a sensor system dedicated to urban driving assistance applications, built on the infrastructure of a dense stereovision engine. The dense stereo information provides the foundation for a new object reconstruction algorithm, dedicated to identify objects in the city clutter, and also allows the identification of drivable/non drivable areas in the case when no lane delimiting information is visible. Lane detection takes advantage of a better vertical profile provided by dense stereo, but its model matching algorithms were also changed to cope with the urban situations. Freeform lane border estimation was introduced to handle the non-clothoidal road geometries.

The results give reasons to believe that dense stereovision is a valuable tool for the driving assistance sensorial systems, and draw the directions of the future work: the lane detection system will be enhanced by adding a curb extraction algorithm, and by combining the freeform and model based results using a fusion module; the object detection will benefit from orientation-based tracking; the drivable/non drivable area results will be fused with the obstacle detection and with the lane detection results.

REFERENCES

- [1] U. Franke, D. M. Gavrila, S. Götzig, F. Lindner, F. Paetzold and C. Wöhler, "Autonomous Driving Approaches Downtown", *IEEE Intelligent Systems*, vol.13, nr.6, 1998, pp. 40-48.
- [2] Yasuhiro Takagi, Setsuo Tokoro, Masayuki Usami, Tatsuya Shiraishi, "The Development of Stereo Image Obstacle Recognition System", *World Automotive Congress 2006 (FISITA 2006)*, Yokohama, Japan.
- [3] S. Thrun, M. Montemerlo, A. Aron, "Probabilistic Terrain Analysis for High-Speed Desert Driving", *Robotics Science and Systems Online Proceedings*, www.roboticsproceedings.org.
- [4] B. Hummel, S. Kammel, T. Dang, C. Duchow, C. Stiller, Vision-based Path Planning in Unstructured Environments, in *Proceedings of IEEE Intelligent Vehicles Symposium 2006*, Tokyo, Japan, June 13-15, 2006, pp. 176-181.
- [5] S. Kolski, D. Ferguson, M. Bellino, R. Siegwart, Autonomous Driving in Structured and Unstructured Environments, *IEEE Intelligent Vehicles Symposium 2006*, Tokyo, Japan, June 13-15, 2006, pp. 558-563.
- [6] M. Bertozzi, A. Broggi, P. Medici, P.P. Porta, A. Sjogren, "Stereo Vision-Based Start-Inhibit for Heavy Goods Vehicles", *IEEE Intelligent Vehicles Symposium 2006*, Tokyo, Japan, June 13-15, 2006, pp. 350-355.
- [7] C. Duchow, "A Novel, Signal Model Based Approach to Lane Detection for Use in Intersection Assistance", *IEEE Intelligent Transportation Systems Conference (ITSC 2006)*, Toronto, Canada, September, 2006, pp. 1162-1167.
- [8] Labayrade, Raphaël, Aubert, Didier, Douret, Jérôme, "A Multi-Model Lane Detector That Handles Road Singularities", *IEEE Intelligent Transportation Systems Conference (ITSC 2006)*, Toronto, Canada, September, 2006, pp. 1143-1148.
- [9] S. Nedevschi, R. Danescu, T. Marita, F. Oniga, C. Pocol, S. Sobol, T. Graf, R. Schmidt, "Driving Environment Perception Using Stereovision", *Proceedings of IEEE Intelligent Vehicles Symposium, (IV2005)*, June 2005, Las Vegas, USA, pp.331-336.
- [10] T. Marita, F. Oniga, S. Nedevschi, T. Graf, R. Schmidt, Camera Calibration Method for Far Range Stereovision Sensors Used in Vehicles, *Proceedings of IEEE Intelligent Vehicles Symposium, (IV2006)*, June 13-15, 2006, Tokyo, Japan, pp. 356-363.
- [11] C. Vancea, S. Nedevschi. Analysis of different image rectification approaches for binocular stereovision systems. *IEEE 2nd International Conference on Intelligent Computer Communication and Processing (ICCP 2006)*, vol. 1, 1-2 September 2006, Cluj-Napoca, Romania, pp. 135-142.
- [12] J.I. Woodfill, G. Gorden, R. Buck, Tyx DeepSea High Speed Stereo Vision System, *IEEE Conference on Computer Vision and Pattern Recognition*, Washington, D.C., 2004, pp. 41-45
- [13] S. Nedevschi, 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)*, 2004, Washington, USA, pp. 161-166.
- [14] J. Goldbeck, B. Huertgen, "Lane Detection and Tracking by Video Sensors", in *Proc. IEEE International Conference on Intelligent Transportation Systems*, October 5-8, 1999, Tokyo Japan, pp. 74-79
- [15] R. Danescu, S. Nedevschi, "Robust Real-Time Lane Delimiting Features Extraction", in *proc. of International Conference on Intelligent Computer Communication and Processing, (ICCP) 2006*, Cluj Napoca, Romania, pp. 77-82.
- [16] S. Nedevschi, F. Oniga, R. Danescu, T. Graf, R. Schmidt, Increased Accuracy Stereo Approach for 3D Lane Detection, *Proceedings of IEEE Intelligent Vehicles Symposium (IV2006)*, June 13-15, 2006, Tokyo, Japan, pp. 42-49.