Honda Research Institute EU
Honda Research Institute Europe GmbH
https://www.honda-ri.de/

Visio-spatial Road Boundary Detection for Unmarked Urban and Rural Roads

Tobias Kühnl, Jannik Fritsch

2014

Preprint:

This is an accepted article published in Intelligent Vehicles Symposium (IV). The final authenticated version is available online at: https://doi.org/[DOI not available]

Visio-spatial Road Boundary Detection for Unmarked Urban and Rural Roads

Tobias Kühnl and Jannik Fritsch

***Abstract*— The robust detection of road boundaries is a prerequisite for Advanced Driver Assistant Systems (ADAS), such as Lane Departure Warning and Lane Keeping Assistant Systems. State-of-the-art ADAS rely on lane markings to draw inference about the extent of lanes or the road area. However, on many rural or urban roads markings are worn out or simply not existing. Therefore, this publication proposes a system for vision-based road boundary detection without requiring road markings at the outer side of the road. The fundamental approach is a boundary vicinity classification based on SPatial RAY (SPRAY) features which combines visual and spatial context information going beyond classical image patch analysis. More specifically, the inner road boundary vicinity (IBV), the outer road boundary vicinity (OBV), and the remaining part of the road area (RA) are detected. Because these classes occur in a defined sequence, i.e., a road boundary exhibits a transition from RA to IBV to OBV, this approach extracts the horizontal boundary transition pattern to make inference about possible locations of road boundaries and their direction (left or right road side). The implemented system was evaluated on unmarked urban and rural roads. The results show that the system effectively detects road boundaries such as curbstones and soft shoulders under challenging conditions.**

I. Introduction

Vision-based detection of driving space and road boundaries is an important source of information for ADAS, such as Lane Departure Warning Systems (LDWS). State-of-the-art ADAS rely on lane markings or elevated road delimiters (such as guardrails) to draw inference about the actual road course. These boundary types can be detected robustly because of their distinction in visual or physical properties compared to the driving space. However, not all roads are delimited by such scene elements and therefore it is important to detect other road boundaries such as curbstones or soft shoulders as well. Unfortunately, these are very challenging to detect because they exhibit a high variety in visual appearance properties and not necessarily a height difference which could be detected.

Research on visual discrimination of road and non-road area has been successfully applied to rural roads in order to detect soft shoulders. For discriminating road and non-road texture [1] and also color and edge features [2], [3] are extracted from locally restricted images regions, i.e., patches. Due to appearance changes of the road during runtime, some authors propose adapting the appearance models online

T. Kühnl is with the Research Institute for Cognition and Robotics, Bielefeld University, Bielefeld, Germany TKuehnl@cor-lab.uni-bielefeld.de

J. Fritsch and T. Kühnl are with the Honda Research Institute Europe GmbH, Offenbach am Main, Germany Jannik.Fritsch@honda-ri.de



Fig. 1. System scenarios: Urban and rural roads without outer markings.

[2], [4], [5]. Additionally, fusion of camera data with other sensors, such as radar is proposed to detect elevated obstacles such as guard rails [6]. Especially, in urban scenarios it appears that a high degree of the visual scene has a gray and untextured appearance similar to the road area (e.g., house walls, sidewalks, gray cars). Therefore, the discrimination of road and non-road based on local visual appearance information may suffer strongly from an ambiguity problem, i.e., the appearance of regions outside of the road area is visually very similar to those on the road area. This makes a classification based on local visual appearance very challenging. To this end, other approaches detect the road boundary based on physical elevation by using stereo vision [7], [8], [9] and even highly annotated digital map information [10]. However, many curbstones have an insufficient physical elevation to be robustly detected with depth-based approaches. In summary, existing vision-based approaches can not be fully applied for ADAS such as LDWS on urban roads yet.

In order to overcome the above mentioned drawbacks, we propose a new multi-class approach for detecting the inner and outer boundary vicinity of the road. The base system of this approach was discussed in previous work [11], [12], [13]. There it was shown that SPatial RAY (SPRAY) features capture visual and spatial characteristics of road environments by real-time image processing. Due to the incorporation of the spatial context, SPRAY features can for example distinguish visually similar but spatially different road regions such as different lanes [11]. The fundamental novelty of this paper is applying SPRAY for classifying image regions in the local vicinity of the road boundary whether they belong to road area (inner part) or non-road (outer part). More specifically, based on these unique visio-spatial features, a GentleBoost classifier can detect the inner road boundary vicinity (IBV), the outer road boundary vicinity (OBV), and remaining parts of road area (RA). Because these classes occur in a defined sequence, i.e., a road boundary

exhibits a transition from RA to IBV to OBV, an extraction procedure is used to identify left and right road boundary based on their horizontal boundary transition pattern.

This monocular pure vision-based approach shows high potential for detecting road boundaries on unmarked urban and rural roads. Qualitative and quantitative evaluations show that the system operates under challenging conditions and handles various types of road boundaries such as curbstones and soft shoulders (see Fig. 1).

The next Section gives an overview of the system architecture. The base system computing SPRAY features is explained in Sect. III. Subsequently, the boundary vicinity classification is detailed in Sect. IV. Then, Sect. V explains the extraction procedure for identifying left and right road boundary. In Sect. VI the application of partial linear regression on detected boundary positions is discussed. The evaluation results are presented in Sect. VII. The paper ends with a conclusion and future work in Sect. VIII.

II. System Overview

The overall vision system consists of four stages (see Fig. 2). Initally, the input image is processed to compute SPRAY features. Subsequently, processing is split into a left and right boundary pathway. In each pathway a boundary vicinity classification is carried out. Then, a procedure is applied to extract the metric course of the left or right boundary candidates in the direction of travel. Finally, for both boundaries a road boundary model is applied to obtain smoother results.

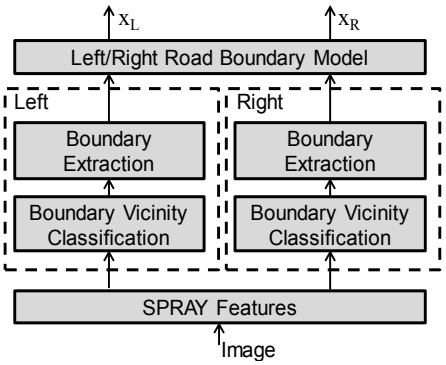

Fig. 2. Block diagram of the vision system showing the processing steps for road boundary detection. Extraction of boundary candidates is split into separate pathways for the left and right boundary.

III. SPRAY features

The basis for the presented approach is capturing the spatial layout by means of SPatial RAY (SPRAY) features [11]. In contrast to local feature gathering (e.g., inside a window with a fixed size), spatial features can combine feature information from multiple spatial locations and capture the extent of regions and shapes. The SPRAY feature generation process is split into two steps.

In a first step, three base classifiers are used to capture different local visual appearance properties from the image. In particular, these are road area appearance, road boundary appearance, and lane marking appearance. Each base classifier can be seen as an appearance model for the particular visual property (see [13]). Although, this paper focuses on detecting unmarked road boundaries, it is still useful to capture available lane markings, such as center lines. This is beneficial for finding road boundaries because the spatial relation of a particular detection position and surrounding lane markings is a useful feature. The output of each base classifier is a confidence map in the Birds Eye View (BEV, see [14]). This metric representation is beneficial for the subsequent processing because perspective changes of the road are compensated. For example, for straight roads the width of lanes or the road area stays approximately constant along the driving direction.

In the second step, SPRAY features are computed on each of the confidence maps. The SPRAY feature values encode the spatial layout of a particular base confidence map in the form of multiple distance values. These distance values are obtained by thresholding integral functions which are generated by integration from defined metric locations to the metric space boundary for multiple angular orientations. A SPRAY feature is defined as the distance resulting from a defined threshold applied on a particular integral function (each integral corresponds to one ray orientation). Further details on base classification and SPRAY features can be found in previous work [11], [13].

IV. Spatial classification of the boundary vicinity

Novelty of this paper is to train a classifier that discriminates regions of the inner and outer boundary vicinity of the road (see Fig. 3) and to suppress detections on the road area. Thus, we formulate a multi-class problem. If we assume road boundaries in the form of lines, the inner boundary vicinity (IBV) is defined as the part of the road area that contains all pixels for which the distance to the boundary line is lower than a given threshold. Accordingly, the outer boundary vicinity (OBV) corresponds to a region outside the road area. In [13] it was already shown that SPRAY features can be applied for detecting the road area. In order to suppress IBV and OBV detections on the remaining part of road area (RA) and therefore ease the extraction of road boundaries, RA is used as third class (The class RA comprises the road area excluding IBV). Note that don't care regions outside the road area (as depicted in Fig. 3) are not considered.

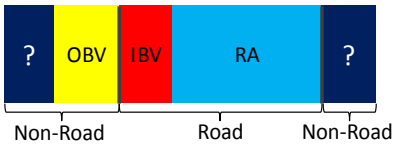

Fig. 3. Illustration of the boundary vicinity concept at the example of the left boundary. The boundary vicinity is separated into inner (IBV) and outer part (OBV). The road area (excluding IBV) is labeled RA. The region labeled with question marks are don't care regions.

The GentleBoost algorithm with decision trees as weak learners is used to automatically obtain class separation for this multi-class problem. The application of GentleBoost, which is originally intended for binary classification problems, to multi-class problems is discussed in [15].

For training the classifier, polyline annotations of the left and right road boundary define where training samples for IBV and OBV are extracted (for RA a polygon is used). Additionally, we make use of the mirror symmetry of SPRAY features for obtaining a more generic boundary vicinity classifier. This means that only a left road boundary vicinity classifier is trained based on SPRAY features for the left road boundary samples and mirrored SPRAY features (longitudinal axis of reflection) for the right road boundary samples. In the processing phase of the system, the trained left boundary vicinity classifier is used for both road boundaries. This implies that for right boundary vicinity classification the features have to be mirrored beforehand.

Fig. 4 depicts exemplary results for the left and right road boundary vicinity classification. Based on the confidences for each class, unconfident detections are suppressed or a decision on one of the classes is taken.

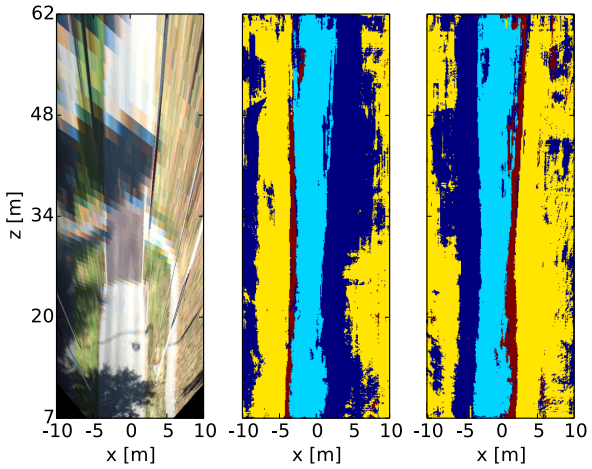

Fig. 4. Exemplary results of the spatial classification of the boundary vicinity on rural road (see Fig. 1 top). The 3 parts show a BEV of the scene (left), the classification result of the left road boundary (center), and the classification result of the right road boundary (right). The road area (light blue), inner boundary vicinity (red), outer boundary vicinity (yellow), and unconfident regions (dark blue) are depicted.

V. Boundary Extraction

The boundary extraction consists of two steps. Firstly, a filter is employed to emphasize left and right road boundary positions based on their horizontal boundary transition pattern. Secondly, an extraction procedure utilizing the Hough transform is applied to find road boundary candidates. These steps will be detailed in the following.

A. Horizontal Boundary Transition Pattern

Typically, each road boundary exhibits a horizontal boundary transition pattern (HBTP) from RA, followed by IBV, to OBV. Naturally, this transition is expected for a left road boundary in the following order OBV $\leftarrow$ IBV $\leftarrow$ RA ($\leftarrow$ denotes "left of"). Vice versa for a right road boundary we expect a HBTP in the form: RA $\leftarrow$ IBV $\leftarrow$ OBV.

The aim is to extract a score S for the HBTP at a particular location which helps us selecting good left or right road boundary positions. To this end, a filter and scoring procedure is proposed. For simplicity, this procedure is explained for left boundaries only. The score is computed by applying multiple filter operations. Each filter operation is designed to increase the score for a particular metric position if it conforms to a certain property of the HBTP and decrease otherwise.

Therefore, three binary matrices F_{RA}, F_{IBV}, and F_{OBV} are used, stating whether one of the three classes is detected (see regions in Fig. 4). Furthermore, three filter kernels f_{RA}, f_{IBV}, and f_{OBV} are defined (see Eq. 1-3). Each of the kernels f has size 1×14px and contains two different kernel values, each for one half of the kernel (left and right).

$$f_{\text{OBV}} = [1, ..., 1, -2, ..., -2] \quad (1)$$

$$f_{\text{IBV}} = [-1, ..., -1, 1, ..., 1] \quad (2)$$

$$f_{\text{RA}} = [-2, ..., -2, 1, ..., 1] \quad (3)$$

Finally, the score S (see Eq. 4) is the sum of the results from three subfilter operations with respect to F_{RA}, F_{IBV}, and F_{OBV}.

$$S = F_{\text{OBV}} * f_{\text{OBV}} + F_{\text{IBV}} * f_{\text{IBV}} + F_{\text{RA}} * f_{\text{RA}} \quad (4)$$

Exemplary results of this algorithm, i.e., the score maps for the left and right boundary are depicted in Fig. 5.

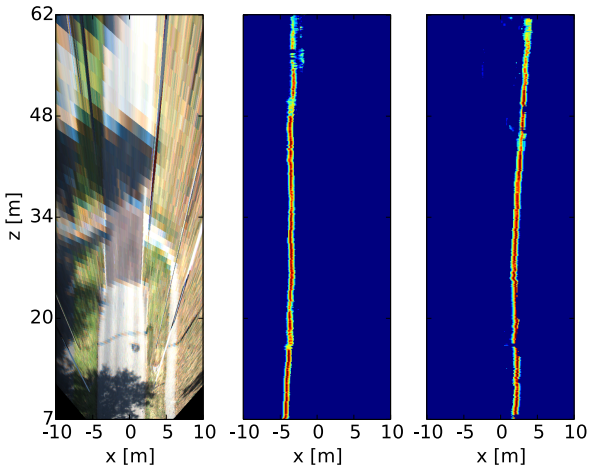

Fig. 5. HBTP filtering result, red values reflect high, blue values low score.

B. Extraction of road boundary candidates

In order to extract explicit road boundary candidates for each road side, an extraction procedure utilizing the Hough transform [16] is applied. Aim of using the Hough transform is emphasizing detections that form lines and to suppress solitary but confident false detections from the filtered results. The transform is applied locally at discrete distances from the ego-vehicle. For this purpose, the BEV is longitudinally sampled with $\Delta z = 1$m. For one metric distance, Fig. 6 shows the output matrix of the Hough algorithm in parameter space. Each element of the matrix corresponds to an integrated score of a line with orientation ϕ and lateral position x. Each integration is locally restricted, i.e., each line has a length l of 0.5m.

In order to extract definite lateral positions for a road boundary, a step-wise search algorithm is used. Starting from the bottom of the BEV, for each z_H an extraction of local

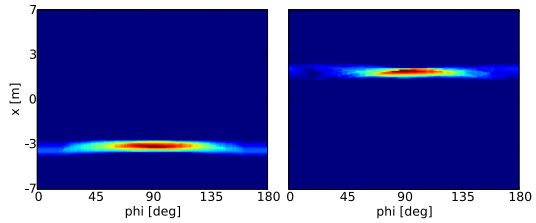

Fig. 6. Illustration of the Hough transform at $z_H = 20$m processing for the left and right road boundary. Red depicts high and blue low values.

maximas is applied to extract boundary candidates. In case of multiple local maximas, the lateral offset to the predecessor candidate is considered for the selection. Three exemplary scenes with detected boundary candidates for an unmarked road with soft shoulder, a partially marked road delimited by curbstones, and a curve, are depicted in Fig. 7.

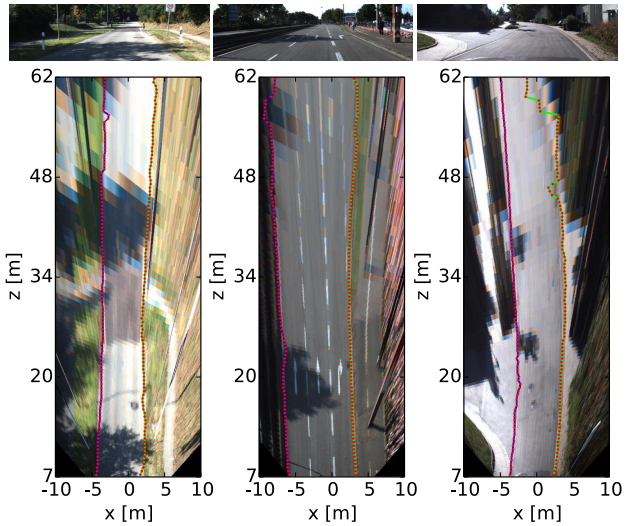

Fig. 7. Exemplary results of road boundary candidates (red dots).

VI. Filtering Lateral Boundary Positions using Partial Linear Regression

In order to smoothen results and filter outliers, partial linear regression is applied. This can be seen as a simple road boundary model. For this purpose, the metric space is separated into multiple ranges (see Fig. 8). The linear regression (least square approximation) is applied to all boundary candidates in a given range for each road side. Fig. 9 shows the results after applying the boundary model for the same scenes as in Fig. 7. In general, a smoother course of road boundaries can be observed.

VII. Evaluation

For assessing the capabilities of the proposed system, we conduct experiments and evaluations on urban and rural roads. Firstly, the approach is applied on data from the KITTI-ROAD vision benchmark containing urban roads [17]. Afterwards, the approach is tested with self recorded data containing rural roads. Furthermore, an application oriented experiment is conducted where the deviation of the estimated left and right road boundary position is evaluated. Finally, computational aspects are discussed.

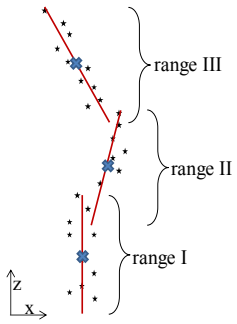

Fig. 8. Illustration of the Partial Linear Model for estimating the lateral position of road boundaries. The space is separated into multiple ranges. Black dots depict raw detections, red lines the fitted line approximation, blue crosses show the estimated road boundary position for each filter range.

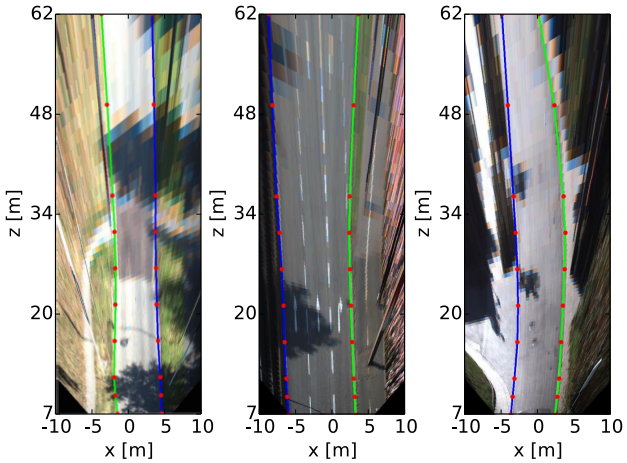

Fig. 9. Exemplary detection results after filtering. Red dots reflect the lateral position for different filter ranges.

A. Evaluation on urban roads (KITTI dataset)

For this experiment, images (1392 x 512 pixels) from the publicly available KITTI-ROAD dataset are used. For training the road boundary vicinity classifier, we annotated boundaries in 140 images from the different training subsets (UU,UM,UMM) depicting relevant boundaries. For testing we annotated 93 images from the UM test dataset. Note that the dataset consists of non-consecutive frames, i.e., a spatial and temporal offset of frames ensures a high decorrelation of the training and testing material.

Table I shows the confusion matrix of the boundary vicinity classification on the UM test dataset. As explained in Sect. IV, this classifier is used for left and right sides by making use of the mirror symmetry of the features. In total, the road boundary vicinity classifier obtains a recognition rate of 79.5%.

TABLE I
Confusion Matrix (KITTI Dataset).

	RA	OBV	IBV
RA	89.42%	3.30%	7.28%
OBV	7.14%	77.47%	15.39%
IBV	12.70%	13.10%	74.20%

Although, the recognition rates of the proposed classifier as shown in Table I imply an error prone process, due to the subsequent boundary extraction lateral boundary



Fig. 10. Final result for left (blue) and right (green) boundary estimation rendered as polyline into the perceptive image (KITTI raw dataset).

locations can be detected appropriately. Furthermore, the filtering reduces outliers and stabilizes noisy boundaries in far distances. This effect becomes visible if we compare the boundary candidates from Fig. 7 and the filtered results from Fig. 9. In addition, Fig. 10 shows more exemplary detections results after boundary extraction and application of the partial linear model. The illustrated results show the behavior of the system in challenging urban scenarios. Positive examples are depicted in rows A-C of Fig. 10. Different types of road boundary such as elevated and low curbstones, soft shoulders and transitions to other elevated obstacles, such as parked cars (see example B1) or hedgerows (see example B2) are depicted. Furthermore, challenging visual conditions in the dataset as, e.g., different types of asphalt, over exposure and shadows on the road are observable. In row D of Fig. 10 also erroneous detection examples are depicted.

A quantitative evaluation of the boundary quality on KITTI-ROAD reveals that this dataset focuses on urban conditions where boundaries are very diverse. Their shape often does not follow a smooth model but exhibits 'corners' (e.g., parked car occludes boundary), incoming roads interrupt the boundary, and opposing traffic occludes the boundary of the opposing lane. We provide the quantitative evaluation results (see Sect. VII-C) for completeness in Fig. 11 (left) but use a self-recorded dataset of rural roads as basis for carrying out a more appropriate quantitative evaluation.

B. Evaluation on rural roads

In the following evaluations, images recorded on rural roads with our test vehicle are used. The image data (1280x720 pixels) has a frame rate of 15 Hz. In total, the dataset contains 50 minutes of driving on rural roads. For evaluation frames with a temporal offset of at least 5 seconds are used.

The spatial classification of the boundary vicinity is evaluated on a sub-dataset containing 480 image frames using 3-fold cross validation. To assess the performance of the road boundary vicinity classifier, the recognition rate on the test folds is measured (see Table II). In contrast to the KITTI dataset, we obtain a higher recognition rate of 85.8% which is due to the larger amount of training frames.

TABLE II
CONFUSION MATRIX (RURAL ROAD DATASET).

	RA	OBV	IBV
RA	95.82%	0.52%	3.66%
OBV	3.29%	86.58%	10.12%
IBV	12.04%	13.36%	74.60%

C. Evaluation of lateral road boundary deviation

In this analysis we focus on unmarked road boundaries such as curbs and soft shoulders. For quantitatively evaluating the lateral deviation of the detection of such road boundaries in an application relevant manner, we selected a subset of the recorded rural road data (see Sect. VII-B) accordingly. In total, curbs and soft shoulders were annotated in 667 frames.

For evaluation, the estimated lateral left and right boundary positions can be compared with the ground truth polylines after mapping them into the metric BEV. In the conducted experiment no false negatives were observed. This means that in all cases where ground truth was annotated, the system delivered a valid lateral position estimation.

Fig. 11 (right) shows the metric mean absolute error for both road boundaries. For better understanding of the results a baseline was computed. For this purpose, we extract a static left and right boundary by averaging over the test ground truth. We see that the proposed algorithm clearly outperforms this baseline. The proposed system shows an overall mean absolute error of 17cm (varies from 12cm to 23cm) up to 40m distance. For comparison, lane marking detection combined with temporal integration was reported to have a mean absolute error of 10cm [18].

Fig. 12 (left) shows the distribution of variances of the proposed approach. In particular, the error over metric distance for the right road boundary with mean and standard deviation is shown. There, the magenta curve shows only those cases where the position was under estimated (within the road area). The blue curve depicts all over estimations (outside the road area). The green curve shows the combined error. We evaluate the baseline in the same way as our algorithm, results are depicted in the right part of Fig. 12.

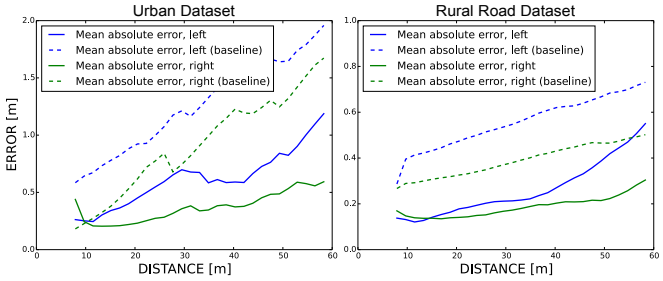

Fig. 11. Metric mean absolute error for the road boundaries (solid lines). For comparison baseline results are plotted (dashed lines).

We see lower standard deviations of the proposed method compared to the baseline. The overall standard deviation up to a distance of 40m is 30cm (varies from 26cm to 41cm). The above mentioned lane marking based approach [18] reports an overall standard deviation of 12cm.

While for lane markings the ground truth positions are clearly defined, ground truth of unmarked road boundaries is ambiguous, e.g., because of curb channels next to curbstones or grass reaching on the road. This results in much more ambiguities in the manual labeling. Therefore, higher variances of systems detecting unmarked road boundaries compared to lane marking based approaches can be expected.

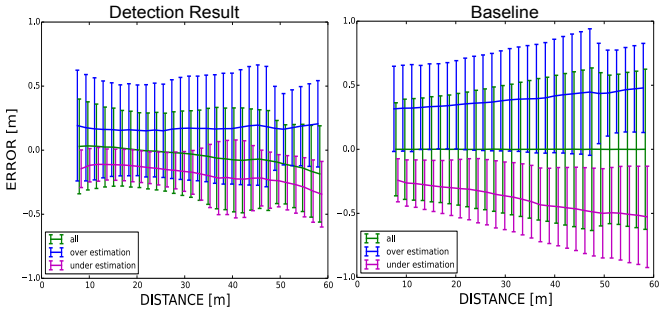

Fig. 12. Lateral position error of the proposed system (right road boundary) and the baseline measured over metric distance. The curves depict the mean error of over estimation, under estimation and all errors (vertical bars show standard deviations).

D. Computational aspects

Main factor affecting the processing time is the size of internal representations, i.e., the perspective image and the BEV. For the implemented prototype, the BEV is spanning a range of 7 to 62m long and -10 to 10m wide with a metric resolution of 5cm. The basis system, including all steps until the boundary vicinity classification, was already used in previous work [11], [12]. It is fully implemented in OpenCL, and runs with more than 12 Hz on a mobile GPU (Nvidia GTX680M). We can apply a subsampling and subsequent interpolation to the internal processing to reach 20 Hz. The modules on top of the basis system which are proposed in this paper are not yet optimized ($\sim$1Hz). Our objective is reaching real-time for the whole processing chain, i.e., at least 10 Hz on a mobile GPU by using OpenCL.

VIII. Conclusion and Future Work

This paper presented a method for road boundary detection without requiring roads to have lane markings at the outer sides. The presented multi-class approach is based on SPRAY features which represent the road environment not only by its local visual appearance but also the spatial layout. The monocular pure vision-based approach was tested on urban and unmarked rural roads and operates under challenging visual conditions and for any type of road delimiter. In the future, the real-time implementation and temporal integration of the results will be further investigated. The mean absolute error and standard deviations in the near range (up to 40m) obtained on single images show high potential for ADAS and autonomous driving.

References

[1] S. Strygulec, D. Mueller, M. Meuter, C. Nunn, S. Ghosh, and C. Woehler, "Road boundary detection and tracking using monochrome camera images," in *FUSION*, 2013, pp. 864–870.

[2] U. Franke, H. Loose, and C. Knoeppel, "Lane recognition on country roads," in *Proc. IEEE Intelligent Vehicles Symp.*, 2007, pp. 99–104.

[3] A. Seibert, M. Hahnel, A. Tewes, and R. Rojas, "Camera based detection and classification of soft shoulders, curbs and guardrails," in *Proc. IEEE Intelligent Vehicles Symp.*, 2013, pp. 853 – 858.

[4] T. Michalke, R. Kastner, J. Fritsch, and C. Goerick, "A generic temporal integration approach for enhancing feature-based road-detection systems," in *Proc. IEEE Intelligent Transportation Systems*, 2008, pp. 657–663.

[5] J. Álvarez, T. Gevers, Y. LeCun, and A. López, "Road scene segmentation from a single image," in *Computer Vision–ECCV 2012*. Springer, 2012, pp. 376–389.

[6] F. Janda, S. Pangerl, E. Lang, and E. Fuchs, "Road boundary detection for run-off road prevention based on the fusion of video and radar," in *Proc. IEEE Intelligent Vehicles Symp.*, 2013, pp. 1173 – 1178.

[7] D. Pfeiffer and U. Franke, "Efficient representation of traffic scenes by means of dynamic stixels," in *Proc. IEEE Intelligent Vehicles Symp.*, 2010, pp. 217–224.

[8] R. Danescu and S. Nedevschi, "New results in stereovision based lane tracking," in *Proc. IEEE Intelligent Vehicles Symp.*, 2011, pp. 230–235.

[9] C. Guo, S. Mita, and D. McAllester, "Robust road detection and tracking in challenging scenarios based on markov random fields with unsupervised learning," *IEEE Transactions on Intelligent Transportation Systems*, vol. 13, no. 3, pp. 1338–1354, 2012.

[10] M. Enzweiler, P. Greiner, C. Knoppel, and U. Franke, "Towards multi-cue urban curb recognition," in *Proc. IEEE Intelligent Vehicles Symp.*, 2013, pp. 902 – 907.

[11] J. Fritsch, T. Kuehnl, and F. Kummert, "Monocular road terrain detection by combining visual and spatial information," *IEEE Transactions on Intelligent Transportation Systems*, vol. PP, no. 99, pp. 1–11, 2014.

[12] T. Kuehnl, F. Kummert, and J. Fritsch, "Visual ego-vehicle lane assignment using spatial ray features," in *Proc. IEEE Intelligent Vehicles Symp.*, 2013, pp. 1101–1106.

[13] T. Kuehnl, "Road terrain detection for advanced driver assistance systems," Ph.D. dissertation, Bielefeld University, Bielefeld, 2013. [Online]. Available: http://pub.uni-bielefeld.de/publication/2633277

[14] H. A. Mallot, H. H. Bulthoff, J. Little, and S. Bohrer, "Inverse perspective mapping simplifies optical flow computation and obstacle detection," *Biological Cybernetics*, vol. 64, pp. 177–185, 1991.

[15] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: a statistical view of boosting," *The Annals of Statistics*, vol. 28, pp. 337–407, 2000.

[16] D. Ballard, "Generalizing the hough transform to detect arbitrary shapes," *Pattern recognition*, vol. 13, no. 2, pp. 111–122, 1981.

[17] J. Fritsch, T. Kuehnl, and A. Geiger, "A new performance measure and evaluation benchmark for road detection algorithms," in *Proc. IEEE Intelligent Transportation Systems*, 2013, pp. 1693–1700.

[18] J. C. McCall and M. M. Trivedi, "Performance evaluation of a vision based lane tracker designed for driver assistance systems," in *Proc. IEEE Intelligent Vehicles Symp.*, 2005, pp. 153–158.