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

Image-based Lane Level Positioning using Spatial Ray Features

Tobias Kühnl, Franz Kummert, Jannik Fritsch

2013

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]

Visual Ego-Vehicle Lane Assignment using Spatial Ray Features

Tobias Kühnl, Franz Kummert and Jannik Fritsch

Abstract— **Assigning the ego-vehicle to a lane is not only beneficial for navigation but will be an essential element in future Advanced Driver Assistance Systems. This paper describes an approach for ego-lane index estimation using only a monocular camera and no additional sensing equipment like, e.g., the typically employed GPS and Inertial Measurement Unit. Key aspect of the approach are SPatial RAY (SPRAY) features which represent the spatial layout of the road in the visual scene. The proposed method perceives a variety of local visual properties of the scene by means of base classifiers operating on patches extracted from camera images. The spatial arrangement of these local visual properties are captured using SPRAY features. With a boosting classifier trained on these features the ego-lane index is obtained. The system is evaluated on low traffic density and complementary to an object-based approach suitable for heavy traffic. In the conducted experiments, the proposed approach reaches recognition rates of 93% to 97% on individual highway images without applying any kind of temporal filtering.**

I. Introduction

In order to increase traffic safety and driving comfort, research on Advanced Driver Assistance Systems (ADAS) is shifting from processing of single entities to more complete understanding of the environment. One highly relevant aspect is the information about how many lanes there are and on which the ego-vehicle is currently driving.

Current vision-based ADAS perform lane departure warning and lane keeping assistance by only using information about the ego-lane. However, information on the relative location of the ego-lane on a multi-lane road (the ego-lane index) is not available from these systems. Identifying the ego-lane index is not straightforward. Without information from a high-precision GPS signal, as used, e.g., in the DARPA Challenges, many approaches rely on combining image processing with detailed digital maps [1], [2].

The proposed system has the target to assign the ego-vehicle to a lane without having any prior knowledge, e.g., on the total number of lanes. Therefore, it can be seen as a complementary contribution to the related work. Instead of estimating the total number of lanes, we separately identify left and right ego-lane index. A combined system output is illustrated in Fig. 1, the boxes in the top right corner show the assigned ego-lane indexes and implicitly the total number of lanes (black box denotes ego-lane). Here the left ego-lane index is 2, meaning that there are two more lanes to the left. Because of the emergency lane the right ego-lane index is 1.

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

F. Kummert is with the Faculty of Technology, Bielefeld University, Bielefeld, Germany Franz@techfak.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. Demo image showing a typical scene on a German highway. Here the left ego-lane index is 2, and the right ego-lane index is 1.

An easily realizable benefit of the proposed method is to provide the ego-vehicle's lane position to a navigation system which allows for more adequate routing instructions. More importantly, for future ADAS operating on highways, information about other lanes provides context to predict behavior of other drivers. For example, knowing the ego-lane index and encountering an entrance lane allows to predict cut-in maneuvers by entering road users [3].

The proposed system captures the spatial layout of visual categories (road, boundary, lane-marking) using SPatial RAY (SPRAY) features. Due to the holistic nature of SPRAY features we increase the likeliness of measuring the full extent of the road even in case of, e.g., overtaking vehicles. The spatial analysis has been used earlier for road terrain classification [4]. The novelty of this contribution is to apply this algorithmic concept for the classification of lane indexes.

The next Section will outline related work in more detail. Subsequently, Section III gives an overview of the system architecture. The base classification is covered in Section IV. SPRAY features are explained in Section V. The application of SPRAY features for training a ego-lane index classifier is detailed in Section VI. The evaluation results are presented in Section VII. The paper ends with a conclusion and future work in Section VIII.

II. Related Work

In previous work a number of sensing techniques have been employed to achieve ego-vehicle positioning. A classical approach that was also extensively used in the DARPA challenges is the localization on a highly detailed digital map using costly high-precision GPS receivers [5]. There is also promising research on affordable and accurate GPS techniques [6]. However, for safety reasons additional environment sensing, e.g., by camera, will be needed for future ADAS. As alternative sensor technologies, inter-vehicle communication has been used for lane positioning [7] as well as a high-precision inertial measurement unit (IMU) in combination with a terrain map [8]. Furthermore, radio-frequency-based landmarks were applied to urban intersections [9].

Avoiding the need for extra/costly sensor hardware, Mattern et al. [1] apply an extended digital map that contains information about road borders. These landmarks are used in a generative approach where a hypothetical sensor image is generated from the landmarks and subsequently matched to the actual camera data to obtain precise positioning.

Popescu et al. [2] also use GPS and an extended digital map containing road infrastructure information (e.g., number and width of lanes, lane delimiter types). A Bayesian network is used for positioning on the lane level where the network parameters are adapted to the current map information. A visual detection of the delimiter types of the ego-lane serves as input data together with information about the dynamic environment (position and movement of other vehicles).

Using a navigation map and a fixed road width, Konrad et al. [10] convert this map data into virtual road border information. Road borders are extracted in the form of lane markings, color differences, or obstacles and matched to the virtual road borders to perform the positioning. The approach relies on successful border detection and seems to be limited to two-lane roads unless an extended digital map is used.

III. System Overview

The overall system depicted in Fig. 2 consists of three stages: Firstly, base classifiers capturing local visual properties. Secondly, SPRAY feature generation capturing spatial properties. And lastly, a lane index classification.

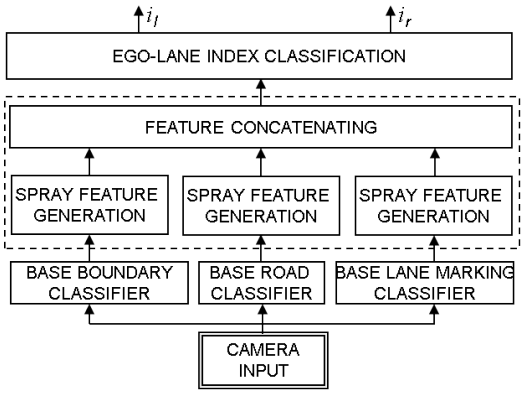

Fig. 2. Blockdiagram showing the main three stages of the proposed system: Base classification, SPRAY feature generation and ego-lane index classification. Output is the left and right ego-lane index (i_l, i_r).

The RGB images from the camera are fed into the base classifiers for analyzing specific local visual properties. We propose using 3 base classifiers: base boundary classifier, base road classifier, and base lane marking classifier. These classifiers can be seen as appearance models which generate three confidence maps in a metric representation (Section IV). For each confidence map, SPRAY features capture spatial aspects of a base property relative to a manually selected spatial position which we call base point (Section V). All the individual SPRAY features computed from the different base classifiers are concatenated to obtain a single feature vector. These features have been previously used for road-terrain classification [4]. The main contribution of this paper is the application of SPRAY features to pure visual ego-lane index classification (Section VI).

IV. Base Classifiers

The task of each base classifier is to analyze the input image with respect to a specific appearance property and to generate a metric map of confidence values. In order to be more robust to lighting changes in the visual scene, global image normalization is applied as a preprocessing step in order to transform the RGB-image to have zero mean and unit variance. Each base classifier produces output in form of the so-called birds-eye-view (BEV), a metric representation of the image which can be obtained using Inverse Perspective Mapping (IPM) [11].

The base boundary classifier is specialized on detecting boundaries between road area and adjacent regions like guardrails or non-road terrain, generating high confidences at locations that correspond to boundaries. The base road classifier is specialized to generate high confidences on road-like area (e.g., asphalt) and low confidences on non-road terrain. The base lane marking classifier aims to detect lane markings as these are visually very specific, forming a special group of road delimiters that can best be captured by a separate classifier. In order to extract lane markings a simple dark-light-dark transition detection is applied to the luminance channel of the BEV. Then, based on the dark-light luminance margin, a confidence is computed (see Fig. 3).

For base boundary and base road classifier a classical patch-based procedure with the following four steps is used: Patch extraction, feature extraction, classification, and mapping into the BEV. After a patch ($21\times21px$) is extracted from the input image and features are calculated on this patch, a classifier assigns a confidence to this patch. Note that the systems for base boundary and base road are structurally identical. However, different internal parameters result from the dissimilar datasets used for training the classifiers. A base classifier outputs a perspective confidence map, which is obtained by interpolating patch confidences, indicating for every pixel position how well it matches the respective appearance model. After IPM, an entry of the metric confidence map contains confidence information about whether a corresponding cell in metric space has a certain appearance property or not.

In the following, three different features types that are used in the base boundary and base road classification are introduced.

A. Feature Computation

Slow Feature Analysis (SFA) is a learning technique which enables to find useful and invariant representations by using unsupervised learning [12]. During training, the algorithm performs an optimization in order to obtain a static transformation from a highly varying multidimensional temporal input signal. Result of the optimization is a slowly-varying output signal (for details see [13]). For the given task, 20 slow features are a good trade-off between performance and computational effort.

For each patch we obtain 6 color features by calculating mean and variance for each channel of the RGB image. In addition, each patch is divided into four rectangles and the mean values for all color channels of these subregions are calculated. The difference of left to right and top to bottom in mean and variance for every channel results in 12 additional color features.

In order to obtain texture features we apply the 2D Walsh Hadamard transform on greyscale images [14]. In total 64 texture features are obtained here.

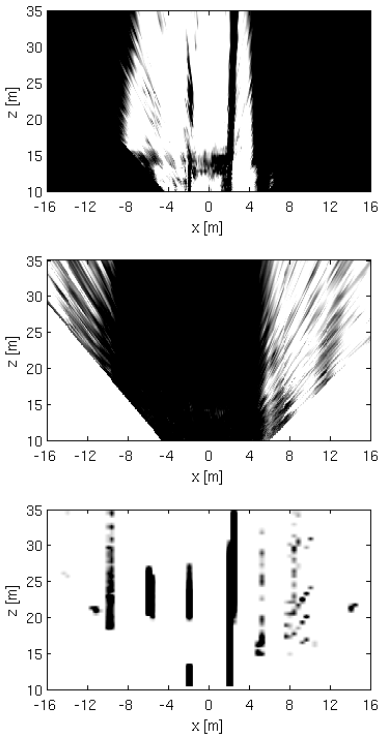

Fig. 3. Output of the base classification: Confidence maps for base boundary (top), base road (center), and base lane marking (bottom) in the BEV. Dark points denote high confidence of the classification. Note, for road and boundary only the positive values of the confidence maps are shown. The input image was shown in Fig. 1.

B. Base Classification

We use the GentleBoost classification method [15] for patch classification. For each image patch the feature extraction generates a 102-dimensional vector $x_{\text{Class},1}$ consisting of slow, color, and texture features. The method has to be trained with data samples and corresponding target classes. The algorithm generates a sequentially weighted set of weak learners that build a strong classifier in combination. In every iteration of the procedure, according to the current distribution of weights on the input samples, the method attempts to find another weak learner in order to optimize the overall classification.

After training, in the processing phase, patches are extracted all over the image using a regular grid with a 4 pixel offset. Similar to training, we compute $x_{\text{Class},1}$ for every patch and obtain patch-based confidences $y_{\text{Class},1}$ with the trained classifier. Based on the patch center position c_i, we map the confidences $y_{\text{Class},1}$ into a metric confidence map $conf_1^{\text{BEV}}$ by applying IPM and 2D bi-linear interpolation. These confidence maps contain negative and positive values indicating low and high probability for the respective base class. For further processing it is necessary to split the confidence maps into their positive and negative parts. Exemplary results of the base classification in the metric space are depicted in Fig. 3. Note, that only the positive parts of the confidence maps are shown.

V. SPRAY Features

This section describes our approach for capturing the spatial layout of local visual properties in order to classify the ego-lane index. In [16] it has been shown that spatial features can be very effective for classifying biological cell shapes. Features extracted at different locations relative to a base point have also been used for body part recognition [17]. In order to apply these concepts to the given task a ray-like feature approach inspired by [16] has been developed.

The SPRAY feature generation process is depicted in Fig. 4. A base classifier provides a metric confidence representation $conf_1^{\text{BEV}}(x, z)$ on which the spatial feature vectors are extracted. The example, shown in Fig. 5, illustrates a confidence map of an idealized base classifier in the metric space for a three-lane road with lane markings. The simulated high confidences at lane markings or other road delimiters and low confidences on road terrain are comparable to the base boundary or the base lane marking classifier (see Fig. 3).

The spatial layout with respect to a confidence map is captured at a base point BP by radial vectors, which are called rays. A ray vector R_α includes all confidence values along a line, with a certain angular orientation α, starting from a base point and ending at the border of the metric representation. The example in Fig. 5 shows six rays enumerated clockwise.

To capture spatial information in a certain direction, the integral of the confidence values $A_\alpha(\rho)$ along the ray vector R_α is computed (1). This integral can be interpreted as

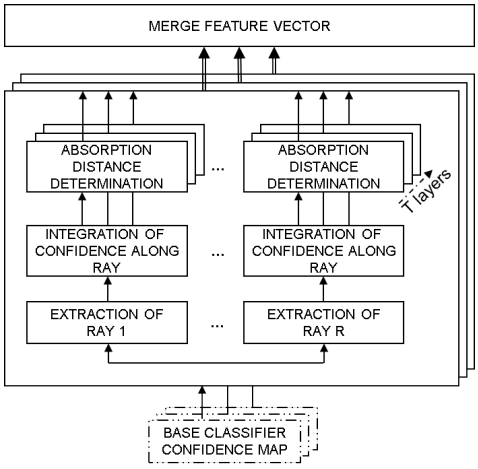

Fig. 4. System block diagram showing the processing steps for spatial feature generation which is applied for each base confidence map.

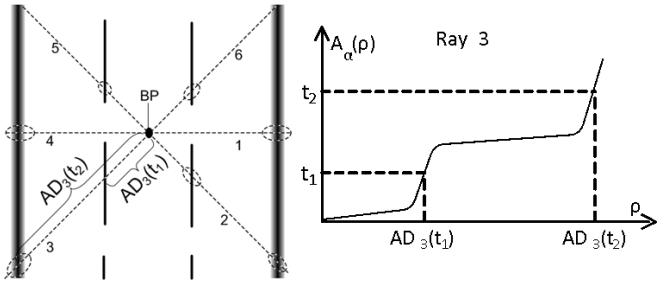

Fig. 5. Illustration of SPRAY feature generation procedure (left) and the integral over the confidences (absorption) for the third ray (right). Two SPRAY features $AD_3(t_1)$ and $AD_3(t_2)$ are obtained, which reflect in this case the distance to a lane-marking and the left road border. For the confidence map (left) high confidences are depicted with dark colors.

absorption of confidences along the ray.

$$A_\alpha(\rho) = \int_0^\rho R_\alpha(\gamma)\,\mathrm{d}\gamma \tag{1}$$

$$f_{\text{SPRAY},\alpha}(t_i) = AD_\alpha(t_i) = \underset{\rho}{\operatorname{argmin}}\,(\rho \,|\, A_\alpha(\rho) > t_i) \tag{2}$$

By defining a certain number T of absorption thresholds t_i, the absorption distances $AD_\alpha(t_i)$, i.e., the locations where the integral value reaches a certain threshold t_i, are obtained as SPRAY features f_{SPRAY} (2). The generation of SPRAY features for a number of R rays with orientation α_r is performed for a specific base point.

For the third ray in Fig. 5 the absorption distances $AD_3(t_1)$ and $AD_3(t_2)$ for two thresholds are illustrated. Fig. 5 (right) shows a principal sketch of the resulting integral for the third ray. The thresholds t_1 and t_2 lead to absorption distances that correspond to the distances from the base point to a dashed lane marking $AD_3(t_1)$ and the leftmost delimiter $AD_3(t_2)$.

For one specific base point in a confidence map, R rays result in $N_{AD} = R{\cdot}T$ absorption distances. These absorption distances serve as SPRAY features for ego-lane index classification f_{SPRAY}, because they describe the spatial layout of local visual properties relative to the predefined base point. Note that this feature generation is done independently for every confidence map $conf_1^{\text{BEV}}(x, z)$. As mentioned in Section IV, in total 5 confidence maps are given (2 $\times$ road, 2 $\times$ boundary, 1 $\times$ lane marking). All SPRAY features are finally concatenated to a single feature vector.

In addition to the SPRAY features outlined above, we use the absorption value of the integral produced by a ray, sent from a base point (x_{BP}, z_{BP}) to the ego-position (x_{ego}, z_{ego}) in the metric representation, as a feature. The feature serves as a good indication for noise level and whether the base point is located on the ego-lane.

VI. Lane Level Classification

The problem of finding the ego-lane index and the number of lanes can be reformulated into an ego-lane index classification to the left and respectively to the right. The ego-lane index $i_{l,r} = [0..n]$ is given by a discrete number, where 0 indicates the ego-lane is the left- or rightmost lane. Consequently, 1 indicates the second most lane and so on.

For the conducted experiments a maximum ego-lane index of 3 (both sides) was considered. In order to capture the extent of a four-lane highway even if the ego-vehicle is on the left- or right-most lane our system uses a metric representation for computing SPRAY features with a lateral range from $-16m$ to $16m$. In longitudinal direction a range from $5m$ to $30m$ (measured from the front bumper of the car) was analyzed. With a discretization of $0.1m$ per pixel we obtain a BEV-resolution of $250 \times 320px$.

For base classification, images from an in-car-mounted camera with a resolution of $271 \times 1280px$ are used. We use pre-learned parameters, obtained from an inner-city benchmark dataset with 750 frames. This benchmark comprises multiple weather conditions (overcast, sunny and rainy weather) and is therefore suitable for general practice.

The base point for SPRAY feature generation was set to a distance of $11m$ centered in front of the vehicle. For driving on highway or motorways it is assumed that this metric location is always unoccupied.

Obviously, for ego-lane index classification laterally oriented SPRAY features of the road scene are more relevant than those measured along the driving direction. Thus, SPRAY features are configured with 6 rays, 3 to the right $[-30°, 0°, 30°]$ and 3 to the left $[150°, 180°, 210°]$. The $\pm 30°$ rays are useful in order to have a higher likelihood to measure the full lateral extent of the road even if cars are driving on neighboring lanes. Furthermore, five thresholds are used $[3, 10, 30, 70, 120]$ which were empirically chosen. Consequently the SPRAY features encode the scene with 30 distance values per confidence map. We obtain a feature vector length 150 plus 5 ego-rays from the 5 confidence maps (2$\times$ base boundary, 2$\times$ base road, and 1$\times$ base lane marking, cf. Section IV).

A GentleBoost classifier is used to automatically select good features and obtain class separation. The basic problem is a multiclass classification. Note that GentleBoost is originally designed for binary classification problems. The strategy for applying it to the multi-class problem to obtain a confidence for each class is described in [18]. Finally a decision on the ego-lane index can be taken based on the class with the maximum confidence. In Fig. 6 two examples for the ego-lane index classification are depicted. The upper image shows the correct estimation on a five-lane highway. Note, that the system is not able to distinguish an emergency lane from an usual traffic lane. The lower image shows a challenging situation because of hardly visible lane markings (between leftmost and second leftmost lane) and cast shadow caused by the fence. The proposed system can handle these situations, while a pure lane-marking based approach would probably fail.

VII. Evaluation

Two highway dataset recorded with 20 fps are used to evaluate the performance of lane index classification. The datasets were recorded at daytime with mostly sunny weather conditions. Especially cast shadows (see lower image in Fig. 6) can cause noisy behavior of the base classifiers.



Fig. 6. Exemplary classification results where estimated ego-lane indexes are illustrated with white boxes. The black box denotes the ego-lane.

Furthermore, various asphalt types require to have very general base classifiers.

The first dataset "Day 1", capturing approximately 30 km of driving is evaluated using four-fold cross validation. In the first experiment, for every fold $\frac{3}{4}$ of the dataset was used for training ($\frac{1}{4}$ for testing) the GentleBoost classifier (decision trees with 5 splits as weak learners and 100 iterations). In order to show the generalization capabilities of the approach in a second experiment, dataset "Day 2" (11 km) is used for testing. Consequently, no additional training is applied. The learned classifiers from "Day 1" are simply reused. Note that "Day 2" contains more parts with higher traffic density than "Day 1" but both exhibit mostly low traffic density.

For both datasets the left and right ego-lane index was manually annotated for every frame. In situations of unclear class membership, e.g., lane changes or other unclear situations, the frame was not used for evaluation and training. The variety of different situations in the dataset go from ramps ($i_{l,r} = 0$) up to four-lane highways ($i_{l,r} = [0..3]$). Note that traffic, entrance, and emergency lanes are handled equally.

We separately evaluate left and right ego-lane index classification. In all experiments SPRAY features were only computed for a single base point as we did not notice a significant improvement when using multiple base points.

A. Left and Right Ego-Lane Index Classification (Day 1)

The evaluation for ego-lane index classification on "Day 1" was carried out on 19884 frames for left and 18522 frames for right using 4-fold cross validation. The distribution of the ego-lane index classes is given in Table I.

TABLE I

CLASS DISTRIBUTION OF EGO-LANE INDEXES

	Day 1		Day 2	
class	left	right	left	right
0	9803	7358	2290	3661
1	6922	7776	3111	3813
2	2756	3287	920	400
3	403	101	34	0
$\sum$	19884	18522	6355	7874

The confusion matrices for the two classifiers are depicted in Table II. The proposed method obtains similar results for both sides with overall recognition rates of $R_{l,1} = 97.6\%$ (left) and $R_{r,1} = 97.6\%$ (right).

TABLE II

CONFUSION MATRIXES FOR EGO-LANE INDEX CLASSIFICATION

A. 4-fold cross validation on "Day 1"

Left				
	0	1	2	3
0	9639 (98.3 %)	137 (1.4 %)	21 (0.2 %)	6 (0.1 %)
1	58 (0.8 %)	6786 (98.0 %)	75 (1.1 %)	3 (0.0 %)
2	44 (1.6 %)	78 (2.8 %)	2631 (95.5 %)	3 (0.1 %)
3	0 (0.0 %)	19 (4.7 %)	25 (6.2 %)	359 (89.1 %)
Right				
	0	1	2	3
0	7302 (99.2 %)	34 (0.5 %)	22 (0.3 %)	0 (0.0 %)
1	123 (1.6 %)	7567 (97.3 %)	86 (1.1 %)	0 (0.0 %)
2	107 (3.3 %)	69 (2.1 %)	3111 (94.6 %)	0 (0.0 %)
3	0 (0.0 %)	0 (0.0 %)	0 (0.0 %)	101 (100 %)

B. Generalization experiments on "Day 2"

Left				
	0	1	2	3
0	2272 (99.2 %)	17 (0.7 %)	1 (0.0 %)	0 (0.0 %)
1	101 (3.2 %)	3008 (96.7 %)	2 (0.1 %)	0 (0.0 %)
2	12 (1.3 %)	147 (16.0 %)	760 (82.6 %)	1 (0.1 %)
3	0 (0.0 %)	0 (0.0 %)	34 (100 %)	0 (0.0 %)
Right				
	0	1	2	
0	3591 (98.1 %)	29 (0.8 %)	41 (1.1 %)	
1	215 (5.6 %)	3500 (91.8 %)	98 (2.6 %)	
2	21 (5.2 %)	84 (21.0 %)	295 (73.8 %)	

B. Generalization Experiments (Day 2)

In order to analyze the generalization capabilities of the classification system a second dataset, recorded on a different day with a different route, was used. Instead of using one classifier from a specific fold from the previous step, a majority decision of all four classifiers is computed. The ego-lane index result is obtained by summing up the per-class confidence, and select the class with the maximum argument. As shown in the confusion matrices in Table II a high generalization performance for the first two ego-lane indexes (class 0 and class 1) is obtained. The overall recognition rates on the "Day 2" dataset are $R_{l,2} = 95.0\%$ to the left and $R_{r,2} = 93.8\%$ to the right. One cause for the worse performance for ego-lane index class 2 and 3 is the higher amount of traffic on "Day 2". Additionally we see in Table I that the ego-lane index class 2 and 3 are underrepresented, so that the classifiers do not generalize that well for these classes. We would need to use a more balanced dataset in terms of class occurrence for better training.

C. Discussion

Eighteen exemplary classification results are depicted in Fig. 7 with I-XVI showing correct classification results. Even though, very high recognition rates are obtained, there are several reasons for false detections: Although, the SPRAY-feature ray casting strategy increases the likeliness of capturing the full extent of the road, in some cases long vehicles occluding a significant part of the road lead to false detections. This can be seen in example XVIII of Fig. 7 where the ego-lane index is underestimated. This

could be compensated by incorporating vehicle detections, e.g., using radar or camera, into the system as proposed by [2]. A situation where the ego-lane index is overestimated is shown in Fig. 7 XVII, this effect is caused by the road-like appearance of the barrier on the right. This problem might be solvable with an approach capturing the height over ground in the scene (e.g., by stereo).

We tried applying a median for the last 5 detections as a simple temporal integration but this did not significantly improve the results. This is because our errors occur in larger temporal blocks, e.g., during an overtaking maneuver of a truck. A noisy behavior of the lane index classifier is not observed. Therefore, a more sophisticated temporal integration is needed to obtain a performance increase.

The whole system[1] runs with a frame rate of 16 fps on a GPU (NVidia GTX 580) using OpenCL. The base classification turns out to be the system part with the highest averaged iteration time of 53.8 ms (in 500 iterations). The SPRAY computation, with 3.6 ms, and ego-lane index classification, with 1.2 ms, are compareply fast. In total, the whole system takes 59.9 ms for one iteration.

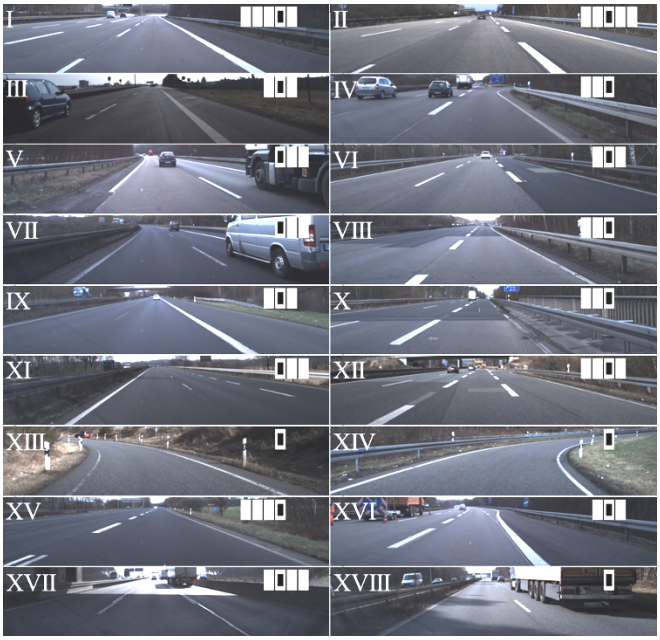

Fig. 7. Exemplaric classification results in multiple scenes.

VIII. Conclusion and Future Works

In this contribution we have demonstrated that SPRAY features provide an effective representation of spatial environments like highway roads. Using a purely vision-based approach operating on individual images, the resulting ego-lane index detection provides recognition rates of 93%-97%. Towards the application in real ADAS, a combination of the presented approach with other information sources will be beneficial. Most importantly, having vehicle detections from stereo, radar, or object recognition will allow to reduce the miss-detections occuring during take-over maneuvers and to ensure correct operation also in very dense highway traffic. In addition, incorporating temporal filtering and lane change notifications from lane tracking, turning lights, a.s.o., will ensure robustness during continous driving. Although there is, therefore, still some work to integrate this approach into an ADAS like, e.g., a prediction of cut-in maneuvers [3], the simplicity of the feature calculation and lane classification are promising aspects of our approach.

References

[1] N. Mattern, R. Schubert, and G. Wanielik, "High-accurate vehicle localization using digital maps and coherency images," in *Proc. IEEE Intelligent Vehicles Symp.*, 2010, pp. 462–469.

[2] V. Popescu, R. Danescu, and S. Nedevschi, "On-road position estimation by probabilistic integration of visual cues," in *Proc. IEEE Intelligent Vehicles Symp.*, 2012, pp. 583–589.

[3] S. Bonnin, F. Kummert, and J. Schmudderich, "A generic concept of a system for predicting driving behaviors," in *Proc. IEEE Intelligent Transportation Systems*, 2012, pp. 1803–1808.

[4] T. Kuehnl, F. Kummert, and J. Fritsch, "Spatial ray features for real-time ego-lane extraction," in *Proc. IEEE Intelligent Transportation Systems*, 2012, pp. 288–293.

[5] M. Buehler, K. Iagnemma, and S. Singh, Eds., *The DARPA Urban Challenge: Autonomous Vehicles in City Traffic, George Air Force Base, Victorville, California, USA*, ser. Springer Tracts in Advanced Robotics, vol. 56. Springer, 2009.

[6] V. Knoop, P. Buist, C. Tiberius, and B. van Arem, "Automated lane identification using precise point positioning: An affordable and accurate gps technique," in *Proc. IEEE Intelligent Transportation Systems*, 2012, pp. 939–944.

[7] T.-S. Dao, K. Y. K. Leung, C. M. Clark, and J. P. Huissoon, "Markov-based lane positioning using intervehicle communication," *IEEE Transactions on Intelligent Transportation Systems*, vol. 8, no. 4, pp. 641–650, 2007.

[8] A. J. Dean and S. N. Brennan, "Terrain-based road vehicle localization on multi-lane highways," in *Proc. American Control Conf.*, 2009, pp. 707–712.

[9] H. Kloeden, D. Schwarz, E. Biebl, and R. Rasshofer, "Vehicle localization using cooperative RF-based landmarks," in *Proc. IEEE Intelligent Vehicles Symp.*, 2011, pp. 387–392.

[10] M. Konrad, D. Nuss, and K. Dietmayer, "Localization in digital maps for road course estimation using grid maps," in *Proc. IEEE Intelligent Vehicles Symp.*, 2012, pp. 87–92.

[11] 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.

[12] L. Wiskott and T. Sejnowski, "Slow feature analysis: Unsupervised learning of invariances," *Neural Computation*, vol. 14, no. 4, pp. 715–770, 2002.

[13] T. Kuehnl, F. Kummert, and J. Fritsch, "Monocular road segmentation using slow feature analysis," in *Proc. IEEE Intelligent Vehicles Symp.*, 2011, pp. 800–806.

[14] Y. Alon, A. Ferencz, and A. Shashua, "Off-road path following using region classification and geometric projection constraints," in *Conf. on Computer Vision and Pattern Recognition*, 2006, pp. 689–696.

[15] Y. Freund and R. Schapire, "A decision-theoretic generalization of on-line learning and an application to boosting," *Journal of Computer and System Sciences*, vol. 55, pp. 119–139, 1997.

[16] K. Smith, A. Carleton, and V. Lepetit, "Fast ray features for learning irregular shapes," in *Proc. IEEE Int Computer Vision Conf*, 2009, pp. 397–404.

[17] J. Shotton, A. Fitzgibbon, M. Cook, T. Sharp, M. Finocchio, R. Moore, A. Kipman, and A. Blake, "Real-time human pose recognition in parts from single depth images," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2011, pp. 1297–1304.

[18] 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.

[1] Experiment setup from "Day 1" with one lane index classifier