

Spatial Ray Features for Real-Time Ego-Lane Extraction

Tobias Kühnl, Franz Kummert and Jannik Fritsch

Abstract—In order to support driver assistance systems in unconstrained environments, we propose to extend local appearance-based road classification with a spatial feature generation and classification. Therefore, a hierarchical approach consisting of multiple low level base classifiers, the novel spatial feature generation, as well as a final road terrain classification, is used. The system perceives a variety of local properties of the environment by means of base classifiers operating on patches extracted from monocular camera images, each represented in a metric confidence map. The core of the proposed approach is the computation of spatial ray features (SPRAY) from these confidence maps. With this, the road-terrain classifier can decide based on local visual properties and their spatial layout in the scene. In order to show the feasibility of the approach, the extraction and evaluation of the metric ego-lane driving corridor on an inner city stream is demonstrated. This is a challenging task because on a local appearance level, ego-lane is not distinguishable from other asphalt parts on the road. However, by incorporating the proposed SPRAY features the distinction is possible without requiring an explicit lane model. Due to the parallel structure of this bottom-up approach, the implemented system operates in real-time with approximately 25 Hz on a GPU.

I. INTRODUCTION

In order to decrease the number of traffic accidents accompanied by an increase of driving comfort for future cars, the topic of road terrain detection is of high interest for ADAS. Road terrain detection is beneficial for path planning and other kinds of object detection, because it creates knowledge about where the ego-vehicle and other traffic participants will probably move to and where other road users, e.g. cars and pedestrians, can potentially appear.

Due to lack of generality, commercial ADAS are often limited to specific scenarios. For example, Lane Keeping Assistant and Lane Departure Warning Systems are restricted to highway situations with certain conditions, e.g. a low curvature of the lane and good quality of the lane-markings. However, the robust recognition of the driving path on arbitrary roads will be needed for future ADAS operating in more complex traffic situations, especially in inner city and rural roads. If there are no explicit road boundaries (like curbstones / lane makers) detectable, e.g., because of parking cars on the side occluding them, current systems based on delimiter detection or explicit lane models (e.g. clothoids) are not working.

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 images showing the challenges in the urban dataset: Situations with bad lane-marking quality, curbstones as road delimiter, unmarked lanes, pot holes in the road and a strong curve.

Therefore, the proposed approach aims at improving appearance-based classification by incorporating the spatial layout of the scene. This allows handling a high variety of complex situations, which is demonstrated on an ego-lane detection scenario. For this task, classification decisions cannot be taken on a local appearance level without considering the spatial layout. In our dataset, geometrical constellations of the ego-lane and variations of asphalt color and texture on the road are manifold. This can be seen in Fig. 1, showing situations with bad lane-marking quality, curbstones delimiting the road, potholes (visually distinct from road), and unmarked road. The proposed system aims at detecting ego-lanes (see Fig. 2) in cases of both explicit (lane-markings or curbstones) and implicit (unmarked road) delimiters.

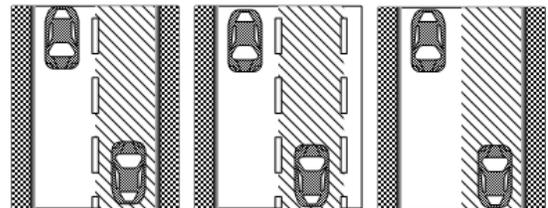


Fig. 2. Exemplary top view of road-scenes illustrating different scenarios for ego-lane detection that can be handled by the proposed system.

For that, the system represents visual properties of both, the road surface and delimiting elements in confidence maps based on analyzing local visual features. On such confidence maps, spatial ray features (SPRAY) that incorporate properties of the global environment are calculated. Only after this two-step extraction process, a decision for road terrain is taken, which implicitly reflects both, local visual properties and their spatial layout. Considering the spatial layout of properties helps for any classification task where there is a clear structural correspondence between properties at different spatial locations.

II. RELATED WORK

Vision-based road segmentation has been addressed in many papers in the last decade: One direction taken by a large number of researchers is the perception of road delimiting elements (e.g. curbstones, lane-markings) for the detection of actual driving space (see, e.g. [1], [2], [3]). Features for these models are extracted from longitudinal road structures like lane markings or road boundary obstacles (e.g. curbstones or barriers) by visual processing. This is mainly based on color and edge appearance (see e.g. [1]), 3D information from stereo processing (see e.g. [2]) or Structure From Motion (see e.g. [3]). From the extracted features, lane/road model parameters can be tracked using different road shape models (see e.g. [4]). However, especially for inner city the applicability of these approaches is limited because of violated model assumptions (intersections, parked cars occluding curbstones). In addition, using road delimiter information only is sometimes not sufficient, because lane-markings might be in a bad condition or not existing.

Alternatively, visual properties of the actual road surface, like, e.g., the mainly gray and untextured asphalt region, can also be beneficial in the detection process [5], [6], [7], [8]. In prior work, a learning approach [5] was presented that captures the typical visual properties of road-like area using monocular, patch-based classification. Pixel based classifications, using Conditional Random Fields (CRF), can be used to identify multiple scene elements in the field of view, including the road surface [7]. However, classifying the visual appearance on a local scale only, can lead to ambiguities. Therefore, Kang et al. [8] showed that incorporating a pixel's larger visual context by using multi-scale grid histograms increases the detection quality of all classes. This idea is similar to our approach, but our focus lies on capturing not spatial image statistics but rather geometrical constellations in a metric representation. This bottom-up spatial context allows detecting semantic categories like ego-lane, which is usually done with a model-based approach (e.g., [1]).

Complementary to such bottom-up approaches, top-down scene context (e.g., based on scene category, location of horizontal line and vanishing point) as prior to further enhance robustness of bottom-up classification decisions is proposed by Alvarez et al.[9].

III. SPRAY-BASED ROAD-TERRAIN DETECTION

The system (see Fig. 3) consists of three parts: base classification, SPRAY feature generation, and road terrain classification. The camera input is fed into each of the N base classifiers (M11) as can be seen in the right part of Fig. 3. Each base classifier (M11) obtains a confidence map for a specific visual property such as road, boundary and lane-marking appearance. On each confidence map, a spatial layout computation (M12-1) that captures spatial aspects of this confidence map's property is applied. All the individual features computed from the N different base classifiers are merged (M12-2) to obtain a spatial feature vector (I13) for each base point, before performing the road terrain classification (M13). The following subsections III A-C detail these

three system parts. Additional technical information on the system setup and its implementation are given in Sec. IV. Sec. V describes a metric driving corridor estimation which is used for evaluation.

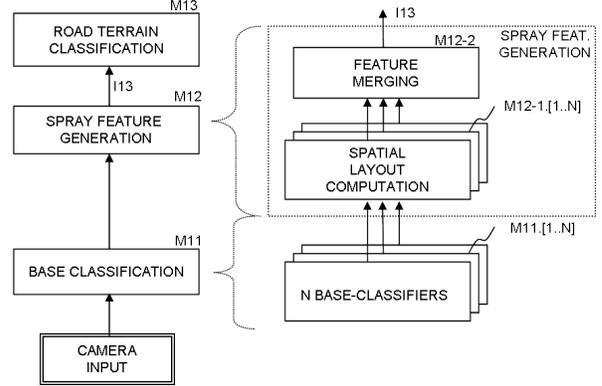


Fig. 3. System block diagram showing the main processing steps (left) and a more fine grained structure (right).

A. Base Classification

The block diagram in Fig. 4 shows the system setup with camera input and 3 base classifiers in module (M11).

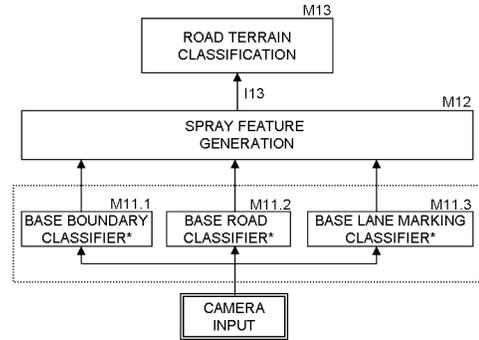


Fig. 4. System block diagram showing the setup for base classification in module M11. Note that all base classifiers (*) include preprocessing and inverse perspective mapping to provide metric confidence maps.

Each base classifier generates a map of confidence values, wherein each location corresponds to a certain location in the metric space which is internally obtained using inverse perspective mapping [10]. An entry of this confidence map contains confidence information about whether a corresponding cell in metric space has a certain property. The combination of all confidence maps builds a value-continuous spatial representation. We propose the use of 3 base classifiers (M11.1-3) which work on preprocessed camera images and result in three confidence maps in a metric representation. These are base boundary classifier (M11.1), base road classifier (M11.2) and base lane marking classifier (M11.3).

The base boundary classifier (M11.1) is specialized on detecting boundaries between the road-like area and adjacent regions like e.g. sidewalks, traffic islands, off-limits terrain, or non-road terrain based on the metric image representation. This base boundary classifier generates low confidences on road-like area and high confidences at locations that correspond to boundaries. The base road classifier (M11.2)

is specialized to generate high confidences on the road-like area and low confidences on non-road terrain. The two base classifiers for road and boundary are based on the work of [5]. This approach enables to learn the typical visual appearance of a given base class. Different training strategies are used to specialize each base classifier on its specific task.

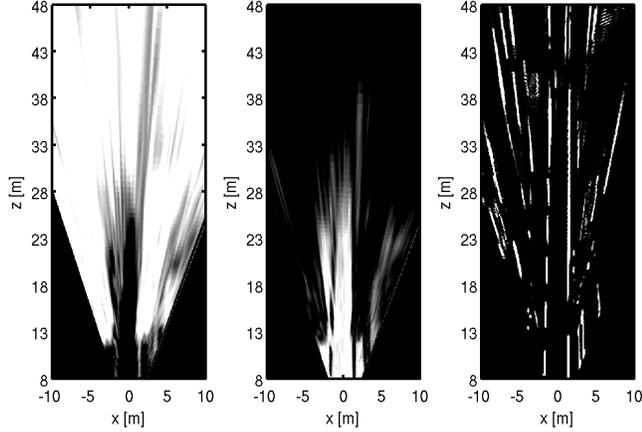


Fig. 5. Result of the base classification showing an illustration of the base boundary (left), base road (center) and base lane marking (right). The brightness denotes the confidence of the classification.

For the base lane marking classifier (M11.3), standard techniques (see, e.g., [11]), can be applied to generate a confidence map having a high confidence at locations corresponding to lane markings and low confidences on road terrain (e.g. road-like area, sidewalk, traffic island).

Fig. 5 shows exemplary results of the base classification in the metric space. These confidence maps are used as input for the SPRAY feature generation.

B. SPRAY Feature Generation

It has been shown that spatial features can be very beneficial for shape based classification [12]. Features extracted at different locations relative to a base point were also used for body part recognition [13]. In order to use these concepts for the task of road-terrain detection a ray-like feature approach inspired by [12] has been developed.

The SPRAY feature generation (M12) process is illustrated in Fig. 6. The left part of Fig. 6 shows the general processing steps for feature generation: Taking a confidence map from a base classifier as input (I12), for a defined number of base points (BP) in the metric representation, spatial feature vectors can be extracted.

The distribution of base points is in this example defined as a grid as shown in Fig. 7 (left). The spatial layout with respect to the confidence map is captured at each individual base point by radial vectors, which are called rays (see Fig. 6, right).

The example in Fig. 7 illustrates a confidence map of one base classifier in the metric space [10] which is used for explanation purpose. It represents a simplified confidence map (example for I12) for a two-lane road with lane markings in the center and curbstones on the left and right side. The simulated base classifier has the ability to

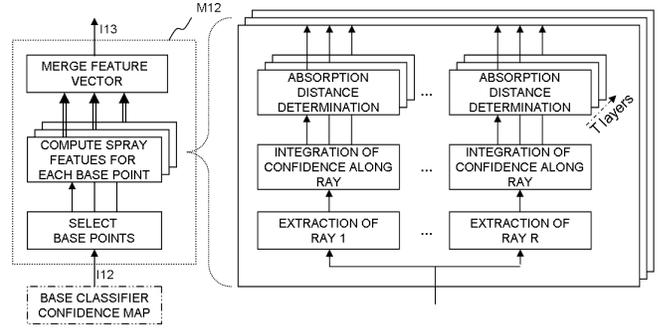


Fig. 6. System block diagram showing the general processing steps of spatial feature generation (left) and a fine-grained illustration of the spatial feature computation which is applied for each base point (right).

generate high confidences on curbstones and lane markings and low confidences on road terrain (dark color indicates high confidences) and is therefore comparable to the base boundary classifier.

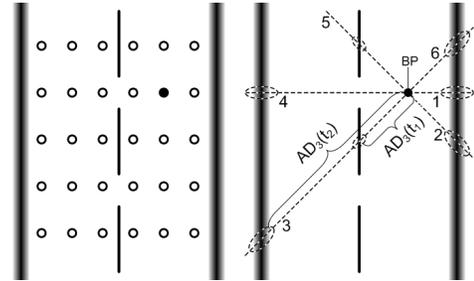


Fig. 7. Distribution of base points over the metric space (left) and the SPRAY feature generation procedure illustrated for one base point (right).

A ray-vector R_α includes all confidence values along a line, with a certain angular orientation α , starting from a specific base point and ending at the border of the metric representation. To convert this information into a defined number of feature values f , the integral of the confidence values $A_\alpha(\rho)$ along the ray R_α is computed (see Eq. 1). This integral can be interpreted as absorption of confidences along the ray.

$$A_\alpha(\rho) = \int_0^\rho R_\alpha(\gamma) d\gamma \quad (1)$$

$$AD_\alpha(t_i) = \underset{\rho}{\operatorname{argmin}} (\rho \mid A_\alpha(\rho) > t_i) \quad (2)$$

By defining a certain number T of absorption thresholds t_i , the absorption distances $AD_\alpha(t_i)$ (see Eq. 2), i.e. the locations where the integral value reaches a certain threshold t_i , are obtained as SPRAY features. The generation of the SPRAY features for a number of R rays with orientation α_r is performed for each base point. For one specific base point in a confidence map, R rays result in $R \cdot T$ absorption distances. The absorption distances serve as SPRAY features for the road terrain classification, because they describe the spatial layout of the environment, captured in the confidence maps of the properties, relative to the predefined base points. Notice that this has to be done independently on every confidence map. All SPRAY features are finally merged to a single feature vector.

In Fig. 7 (right part), an example of the spatial layout computation is shown for one base point. SPRAY features are extracted along six rays (1-6), numbered clockwise. For the third ray (3) additionally the absorption distances $AD_3(t_1)$ and $AD_3(t_2)$ for two thresholds are illustrated. The graph in Fig. 8 shows a principal sketch of the integral over the ray (3) given the confidence map mentioned above. The thresholds t_1 and t_2 lead to absorption distances that correspond to the distances from the base point to a lane marking $AD_3(t_1)$ and the left curbstone $AD_3(t_2)$. Obviously, the selection of 'good' thresholds is crucial for the method (cf. Sec. IV).

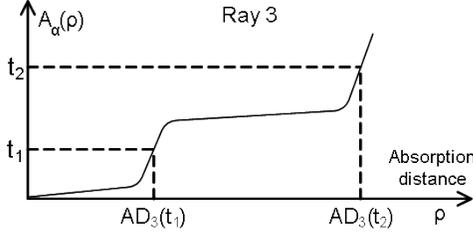


Fig. 8. Integral over the confidences (absorption) for a specific ray (here ray 3 in the scenario from Fig. 7). Two SPRAY features $AD_3(t_1)$ and $AD_3(t_2)$ are obtained, which reflect in this case the distance to the lane-marking and the left road border.

Combined SPRAY Features: In order to simplify the learning, a SPRAY feature f_α can be combined with its corresponding SPRAY feature pointing to the opposite side $f_{\alpha+180}$. This combined SPRAY feature $f_{cmb,\alpha}$, see Eq. 3, can be obtained for all ray orientations α that have a corresponding ray in the opposite side.

$$f_{cmb,\alpha}(t_i) = AD_\alpha(t_i) + AD_{\alpha+180}(t_i) \quad (3)$$

The presumption is that lanes (or even roads) have a fixed width which is captured by combined SPRAY features. For example, assuming ray number 5 and ray number 2 (example right part in Fig. 7) would have been combined: Then there will be a constant combined distance value for all base points BP that are located on the lane (assuming ray 5 hit the lane-marking). Even if this information could be captured in the learning process of the classifier by combining several features, the representation would be much more complex.

Ego SPRAY Features: To obtain an additional measure which indicates if a base point is located on the ego-lane, ego SPRAY features f_{ego} are used. The idea is to use the absorption value of the integral produced from a ray, send 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 value f_{ego} can be obtained with Eq. 1 after obtaining the angle α_{ego} (see Eq. 4) from BP to the ego-position.

$$\alpha_{ego} = \arctan\left(\frac{z_{BP} - z_{ego}}{x_{BP} - x_{ego}}\right) \quad (4)$$

In contrast to the standard SPRAY features the orientation α_{ego} is changing for different base points. This is beneficial for encoding ego-lane specific spatial properties.

C. Road Terrain Classification

For training a classifier based on the proposed SPRAY features a GentleBoost classifier [14] is used, because it has been shown that boosting is very successful in feature selection and classification [15]. The algorithm generates a sequentially weighted set of weak classifiers that build a strong classifier in combination. In every training iteration the method attempts to find an optimal classifier according to the current distribution of weights on the input signal.

Training of GentleBoost: We set up the weak classifiers with decision trees (6 tree split) and a maximum of 120 boosting iterations to get a classifier combining different SPRAY features. Ground truth is needed for training the road terrain classifier on a specific road terrain category. After training is finished, the classifier generates a confidence value for a given feature vector, indicating whether the corresponding base point is likely belonging to the trained category or not. For the proposed detection of ego-lane the classifier learns the distinction between ego-lane and non-ego-lane.

Processing: Once the road terrain classifier is trained, the system can process input images with the learned parameters. An exemplary classification result is given in the right part of Fig. 9. It is used as input for driving corridor extraction as described in Sec. V.

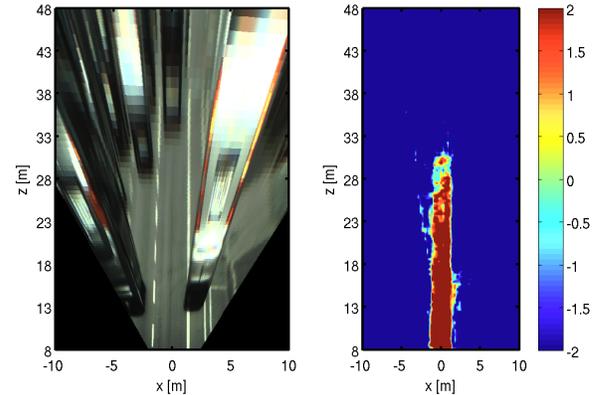


Fig. 9. Result of the road terrain classification showing the metric representation of the image (left) and the classification result for ego-lane as confidence map on the right. The ego-lane is unoccupied up to a distance of 33m. The corresponding base classification results can be found in Fig. 5.

IV. SYSTEM SETUP

For the proposed system RGB images with a resolution of 800×600 pixels are used. The metric representation is defined for a range of -10m to 10m in x direction (lateral) and 8m to 48m in z direction (see Fig. 9). With a resolution of 5cm , a representation with a block size of 800×400 is obtained. Because the GentleBoost road-terrain classifier has the ability to select the best out of a large variety of features it is proposed to use a feature setup, which is a trade-off between brute force (take all features one can get) and training duration. Here the SPRAY feature generation is set up to have 7 ray orientations $\phi = [-20, 0, 20, 160, 180, 200, 270]$ ($\phi = 0$ is rightwards, counted

clockwise positive). As mentioned above selecting good absorption thresholds is significant for the performance of the system. In the example of Fig. 8, thresholds are selected in an optimal way in order to encode the distance to relevant scene elements. This selection is of course not trivial in reality. However, by using a large number of thresholds one can use the GentleBoost to select the best. Based on experiments the absorption thresholds are set to $th = [3, 10, 30, 70, 120]$, although finer graduation and modification for every base-classifier could be more appropriate.

The whole system is implemented to run on a GPU using OpenCL. Due to the parallel system architecture a frame rate of $25 \frac{\text{frames}}{\text{s}}$ on a NVidia GTX 580 is obtained. Averaged timings of the main system parts measured in isolation can be found in Tab. I (note that several system parts run in parallel in the full system).

TABLE I
AVERAGE TIME OVER 200 ITERATIONS

	time [ms]
Base classification (bound./road/lane mark.)	11.2 / 8.8 / 9.9
SPRAY feature generation	13.4
Road terrain classification	2.8
Full system	38.4

V. DRIVING CORRIDOR ESTIMATION

In order to measure the quality of the ego-lane detection, a metric driving corridor is estimated. Therefore, multiple candidates of driving corridor elements with a certain width $w(z)$ are obtained. Starting at the approximate first visible ground element (8m from the rear axle of the ego-vehicle) the corridor is sampled at discrete distances with $\Delta_z = 0.5\text{m}$. In Fig. 5, it can be seen that the perception of the local visual properties by the base classifiers gets worse with increased distances, thus the estimation is applied up to a maximum distance of 28m (roughly 24m from the front bumper).

To represent the drivability of the driving corridor four categories are used: The *non-drivable corridor* has a width of $w(z) \leq 1\text{m}$, the *narrow corridor* width spans a range of $1\text{m} < w(z) \leq 2\text{m}$, the *drivable corridor* satisfies $2\text{m} < w(z) \leq 4\text{m}$ and for completeness the fourth category for $w(z) > 4\text{m}$ is named *oversized corridor*. For visualization purposes the metric corridor elements are mapped back into the perspective image as can be seen in Fig. 10.

VI. EVALUATION

The publicly available datasets of road scenes do not fit our requirements in terms of resolution or sequence length. Therefore, we evaluate the proposed system on an own but rather small dataset, consisting of images with manually annotated ground truth. The total stream length is approximately 4 minutes with 247 annotated frames (1 fps). We split the dataset into training and testing part by using N-fold cross validation with blocks of approximately 20 seconds (resulting in 12 blocks). Each of the 12 blocks is separately tested with a classification system trained on the remaining 11 blocks. The training blocks are split into those used for



Fig. 10. Visualization of driving corridor results. Examples I-VI corresponding to Fig. 1. Green indicating *drivable* and yellow *narrow corridor*. The visualization is limited to a distance of 28m.

training the base classifiers and those used for training the road terrain classifier (alternating blocks). This is important because first the base classifiers are trained on one half. Afterwards the road terrain classifier is trained on base classifier results generated on unseen images

The driving corridor visualization in Fig. 10 shows the results for challenging situations. There one can see that the major part of the ego-lane is captured by the system. Even in cases with bad-lane marking quality (ex. I and II), curbstones (ex. III), and missing lane-markings (ex. IV), the system works nicely. There are errors around road signs and close to lateral road markings (ex. VII and VIII). The system shows an error compensating behavior for appearance variances on the road surface caused by, e.g., shadows, asphalt joints, gully covers, and potholes (see ex. V).

For evaluating the driveability of the ego-lane, the metric driving corridor is used. Note that the results for every frame are obtained in isolation without any temporal information. From the four classes (cf. Sec. V), the corridor width is compared with ground-truth information (extracted from perspective polygonal-data). The histogram of all samples for each corridor class is illustrated in Fig. 11. One can see that the *narrow corridor* (orange) and the *oversized corridor* (blue) are very rare in the ground truth. However, these classes can be used to assess errors in the corridor estimation with a finer graduation.

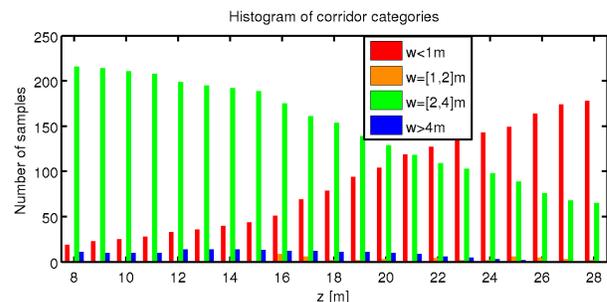


Fig. 11. Histogram of corridor categories over distance.

In Fig. 12, the results for the overall true positive rate TPR_{all} (purple) and false negative rate FNR as a function

of the metric distance to the ego-position are depicted. Here the false negative rate is split into its components corresponding to the four corridor categories, reflecting their missed detections. The feasibility of ego-lane detections with the presented approach can be seen by the TPR which varies from 90% to 80%.

False negative detections mainly occur in a construction site where the asphalt is milled off, and especially on road signs (e.g. arrows or letters, see ex. VII) on the ego-lane. Using a road sign detection method [16] in the base classification would allow to suppress these errors.

Additionally, the evaluation method is not adequate in the transition region between ego-lane and preceding cars because of ambiguities for the corridor category assignment.

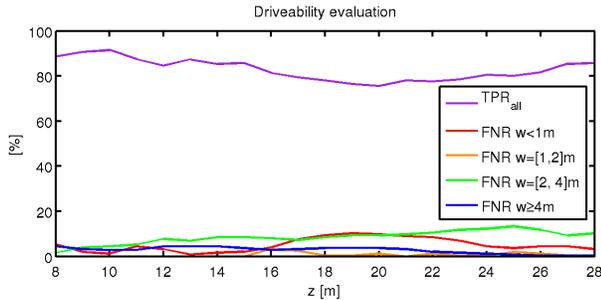


Fig. 12. Driveability evaluation of the driving corridor elements in distinct distances.

In Fig. 13, the quality (see Eq. 5, cf. [17]) of all classes (blue) and the drivable corridor class (green) as a function of the metric distance to the ego-position is shown.

$$Q = \frac{\sum_{i=1}^n TP}{\sum_{i=1}^n (TP + FN + FP)} \cdot 100\% \quad (5)$$

The quality measure is relevant because all errors, also including false positives FP , are considered. The decrease of quality for the *drivable corridor* (green) from 90% at 8m to 60% at 28m, can be explained by the low resolution in the image at far distances¹. This effect does not arise in the quality of all categories Q_{all} (purple) because in far distances mostly *non-drivable corridor* occurs. For *non-drivable corridor*, e.g., in case of a preceding vehicle occluding the corridor ahead of it, the detection of this category is based on the image part depicting the vehicle which has a higher resolution because it is closer to the camera.

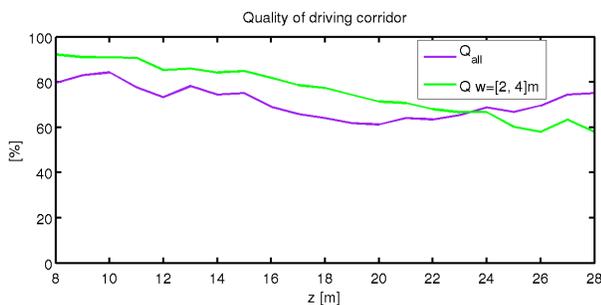


Fig. 13. Evaluation of the driving corridor detection quality in distinct metric distances for all (blue) and the drivable category (green).

¹This effect could be reduced using a multi-scale approach.

VII. CONCLUSION AND FUTURE WORKS

In this paper, SPRAY features that enhance local classification decisions are proposed. The approach enables to capture the geometric characteristics of man-made road. Embedded in the presented framework, learning the spatial layout of local visual properties for ego-lane detection is demonstrated. In a first evaluation, it was shown that this approach can handle various scenarios for ego-lane detection including parts without lane-markings and varying asphalt appearances.

In the future we want to improve the approach by finding optimal configurations for different scenarios (e.g., highway and inner-city). Furthermore, it is planned to train the system on a larger dataset.

REFERENCES

- [1] R. Danescu and S. Nedevschi, "New results in stereovision based lane tracking," in *Proc. IEEE Intelligent Vehicles Symp.*, 2011, pp. 230–235.
- [2] J. Siegemund, U. Franke, and W. Forstner, "A temporal filter approach for detection and reconstruction of curbs and road surfaces based on conditional random fields," in *Proc. IEEE Intelligent Vehicles Symp.*, 2011, pp. 637–642.
- [3] M. Darms, M. Komar, and S. Lueke, "Map based road boundary estimation," in *Proc. IEEE Intelligent Vehicles Symp.*, 2010, pp. 609–614.
- [4] M. Konrad, M. Szczot, and K. Dietmayer, "Road course estimation in occupancy grids," in *Proc. IEEE Intelligent Vehicles Symp.*, 2010, pp. 412–417.
- [5] T. Kuehn, F. Kummert, and J. Fritsch, "Monocular road segmentation using slow feature analysis," in *Proc. IEEE Intelligent Vehicles Symp.*, 2011, pp. 800–806.
- [6] T. Gumpp, D. Nienhuser, and J. M. Zollner, "Lane confidence fusion for visual occupancy estimation," in *Proc. IEEE Intelligent Vehicles Symp.*, 2011, pp. 1043–1048.
- [7] C. Wojek and B. Schiele, "A dynamic CRF model for joint labeling of object and scene classes," in *European Conference on Computer Vision*, vol. 5305, 2008, pp. 733–747.
- [8] Y. Kang, K. Yamaguchi, T. Naito, and Y. Ninomiya, "Multiband image segmentation and object recognition for understanding road scenes," *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 4, pp. 1423–1433, 2011.
- [9] J. M. Alvarez, T. Gevers, and A. M. Lopez, "3D scene priors for road detection," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2010, pp. 57–64.
- [10] 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.
- [11] T. Veit, J.-P. Tarel, P. Nicolle, and P. Charbonnier, "Evaluation of road marking feature extraction," in *Proc. 11th Int. IEEE Conf. Intelligent Transportation Systems ITSC 2008*, 2008, pp. 174–181.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] Y. Sha, X. Yu, and G. Zhang, "A feature selection algorithm based on boosting for road detection," in *Proc. Conf. Fuzzy Systems and Knowledge Discovery*, vol. 2, 2008, pp. 257–261.
- [16] T. Wu and A. Ranganathan, "A practical system for road marking detection and recognition," in *Proc. IEEE Intelligent Vehicles Symp.*, 2012, pp. 25–30.
- [17] T. Michalke, R. Kastner, M. Herbert, J. Fritsch, and C. Goerick, "Adaptive multi-cue fusion for robust detection of unmarked inner-city streets," in *Proc. IEEE Intelligent Vehicles Symp.*, 2009, pp. 1–8.