

Research Article

High-Level Interpretation of Urban Road Maps Fusing Deep Learning-Based Pixelwise Scene Segmentation and Digital Navigation Maps

Carlos Fernández,¹ Jesús Muñoz-Bulnes,² David Fernández-Llorca ,³ Ignacio Parra,³ Iván García-Daza,³ Rubén Izquierdo,³ and Miguel Á. Sotelo³

¹Karlsruhe Institute of Technology, 76131 Karlsruhe, Germany

²ALTRAN Innovation, Madrid, Spain

³Computer Engineering Department, University of Alcalá, Madrid, Spain

Correspondence should be addressed to David Fernández-Llorca; david.fernandezl@uah.es

Received 11 April 2018; Accepted 9 September 2018; Published 11 October 2018

Academic Editor: Jaeyoung Lee

Copyright © 2018 Carlos Fernández et al. This is an open access article distributed under the Creative Commons Attribution License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

This paper addresses the problem of high-level road modeling for urban environments. Current approaches are based on geometric models that fit well to the road shape for narrow roads. However, urban environments are more complex and those models are not suitable for inner city intersections or other urban situations. The approach presented in this paper generates a model based on the information provided by a digital navigation map and a vision-based sensing module. On the one hand, the digital map includes data about the road type (residential, highway, intersection, etc.), road shape, number of lanes, and other context information such as vegetation areas, parking slots, and railways. On the other hand, the sensing module provides a pixelwise segmentation of the road using a ResNet-101 CNN with random data augmentation, as well as other hand-crafted features such as curbs, road markings, and vegetation. The high-level interpretation module is designed to learn the best set of parameters of a function that maps all the available features to the actual parametric model of the urban road, using a weighted F-score as a cost function to be optimized. We show that the presented approach eases the maintenance of digital maps using crowd-sourcing, due to the small number of data to send, and adds important context information to traditional road detection systems.

1. Introduction

The high number of casualties on the road can be explained by many factors. As reported in the analytical report on road safety [1], more than 12.000 lives can be saved per year on European roads if everybody fasten their seat belt, respect speed limits, and do not drive under the influence of alcohol. Distraction is another factor since drivers need to keep their attention focused on surrounding traffic continuously, not just for their own safety but for the sake of their passengers and other road users too. In addition to the aforementioned driver factors, all road elements (vehicles, drivers, infrastructure, etc.) play an important role in the probability of crash or the final outcome. In order to reduce human errors and improve traffic efficiency and safety, Assistance

Intelligent Transportation Systems (AITS) [2], Advanced Driver Assistance System (ADAS), and Autonomous Driving (AD) were developed.

Self-driving cars require a precise and robust perception of the urban environment. It is a crucial point in the development of autonomous vehicles because the perception layer is the base for higher level systems, such as path planning and control algorithms. Most of the major car makers aim to produce fully autonomous vehicles by 2020. This has made the research on ADAS and AD a high priority research issue for both private and public agents.

Different levels of autonomy have been demonstrated in highways, urban scenarios, and cooperative environments [3–6]. However, in all cases a high definition 2D, or even 3D, map is required. Enhanced maps integrate precise

information of the environment such as road shape, lane markings, curbs, intersections, and buildings. The main drawbacks of this type of maps are their size ($\sim 2GB/km$), the complexity of measurements integration, and their maintenance.

On the other hand, the advent of deep learning techniques, namely, convolutional neural networks (CNNs), has involved a breakthrough in the field of artificial intelligence, with strong implications in a large variety of application domains. Thus, research on self-driving cars is experiencing a significant thrust due to the enhanced perception capabilities that the deployment of CNNs are making possible today. Powerful CNN architectures, such as AlexNet, VGG-16, or ResNet, are endowing self-driving cars with advanced capabilities to robustly and accurately interpret road scenes, even in complex urban scenarios with a great deal of clutter and complex road shapes. On top of that, the sensor costs of these road detection approaches are considerably low since they only involve the use of digital cameras. Thus, it eases the maintenance and it is more robust to environmental changes. Other approaches based on high cost LIDARs ($\sim 75K\$$) are not competitive for the car industry.

In this paper we present a high-level interpretation approach capable of estimating a parametric model of the road representing the real scenario appearing in front of the vehicle even in complex and cluttered environments. Based on our previous works [8, 9], a hybrid vision-map method is proposed. However, instead of using the best estimate of the road shape from an enhanced digital map as a feature of a hand-crafted road segmentation classifier [8], we propose a deep learning framework using ResNet-101 network with a fully convolutional architecture and multiple upscaling steps for image interpolation, to obtain an accurate estimation of the road, outperforming our previous results. We demonstrate that significant generalization gains in the learning process are attained by randomly generating augmented training data using several geometric transformations and pixelwise changes, such as affine and perspective transformations, image cropping, mirroring, blur, noise, distortions, and color changes. In addition, we show that the use of a 4-step upscaling strategy provides better learning results as compared to other similar techniques that perform data upscaling based on shallow layers with scarce representation of the scene data. This pixelwise segmentation process is combined with previously designed hand-crafted features extraction methods ([8, 10–12]) to provide a multiclass segmentation of the road and nearby regions. Finally, a high level interpretation module is trained to learn the best set of parameters of a mapping function that is capable of transforming the multiclass segmentation output to the actual parametric model of the urban road. The presented approach makes the maintenance of digital maps using crowd-sourcing easy, due to the small number of data to send, adding important context information to traditional road detection systems and increasing their robustness.

2. Related Work

2.1. Road Segmentation. Different vision-based methodologies to detect the road can be distinguished. Some of them

are based on road appearance learning, where the main hand-crafted features are texture and color information. The second approach focuses on road limits detection, assuming that the space between limits is the road surface. Finally, model-based approaches try to extract a compact high level representation of the road. In order to have a better overview of the different sensing technologies, road appearance and limits modeling techniques, geometrical models, and features integration we refer to [8].

It has been proved that CNNs can improve state-of-the-art results on image classification [13–16], and they have also been successfully applied to object detection [17–19] as well as monocular color image segmentation. There are several approaches to obtain a CNN-based pixelwise classification of an input image. First, the widely adopted fully convolutional networks (FCN) [7] adapt classifier networks, such as AlexNet [13] or VGG [14], to the segmentation task by replacing fully connected layers with convolutional ones and using a progressive interpolation approach. Other approaches follow this trend using other base networks, such as the ResNet [15]. In [20] dilated convolutions are introduced to reduce the downsampling performed by the convolutional layers, removing the need of the progressive interpolation stages. These kinds of dilated convolutions are further explored in [21] along with another upsampling method called *dense upsampling convolution*.

A more complex approach such as DeconvNet [22] learns a deep deconvolutional network on top of the convolutional one. SegNet [23] uses an encoder-decoder architecture. PSPNet [24] exploits pyramidal pooling to introduce global contextual priors in a dilated fully convolutional network. FRRN [25] is based on a novel architecture that keeps a stream with full-resolution features.

Other models do not directly produce a full image classification but work patchwise instead. In [26] the inputs to the classifier layers are enhanced with spatial information of the patch, and in [27] they use a deconvolution scheme. However, the current trend is to use fully convolutional, end-to-end models, from images to pixelwise classification masks [19].

In addition, there are specialized methods for road detection. One example is [28], where the goal is to optimize the models to speed-up inference and make them capable of being used in a real road detection scenario. In [29], MultiNet system is presented, which performs simultaneous street classification, vehicle detection, and road segmentation, all with the same CNN encoder and three different decoders.

In order to learn the huge number of parameters of CNNs architectures, thousands of images labeled with the categories to learn are required. There are different datasets available for autonomous vehicles. One of the most popular datasets is KITTI benchmark [30]. This dataset has labeled 289 images with road and nonroad labels. However, CityScapes dataset [31] is increasing the number of submissions due to the number of labeled categories (30) and the number of images (25K). The results of CNNs in the field of road detection have outperformed all the algorithms that use other approaches.




FIGURE 1: Map-based models: comparison between heavy 3D cloud-based models and light high-level interpretation maps.

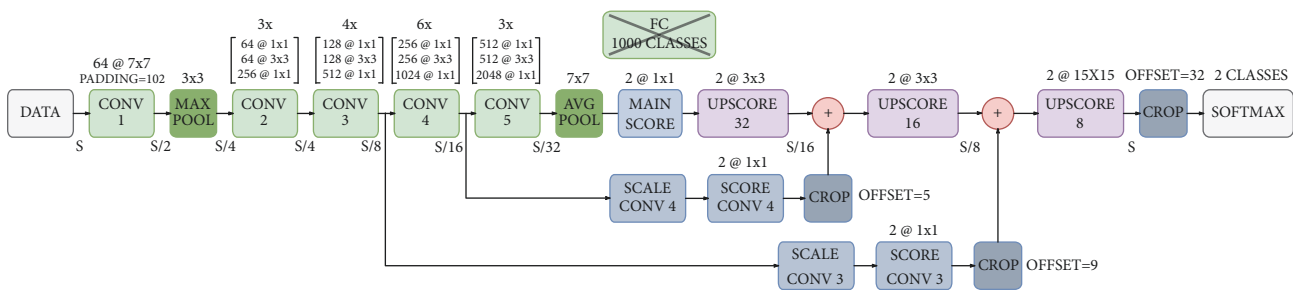


FIGURE 2: High-level schematic of the CNN-based road detector, implementing the FCN-8s architecture [7] on a ResNet-50.

2.2. Map Road Model Fitting. The simplest geometric models used for road boundaries are straight lines. Due to the pin hole camera model, straight parallel lines converge in a vanishing point. This principle is exploited in the state of the art to model the road using an edge descriptor extracted from the images [32]. More complex models are used to model curved roads, such as parabolic curves [33], clothoids [34], B-splines [35], or active contours (snakes) [36]. These parametric models improve the noisy bottom-up detections due to their constraints of width and curvature. However, urban environments are more difficult to model due to the presence of intersections and the variety of curvatures and width changes.

Nonparametric models are less common because they demand only that the line should be continuous. It provokes the model to be less robust than parametric models but more flexible to adapt to the irregular shapes present in urban environments [37] or rural paths [38].

Map-based models (see Figure 1(a)) are dominated by high definition maps. They have demonstrated to be a robust way to navigate [3–5] and they are usually built integrating several measurements of a multibeam LIDAR [39, 40] or multiple single beam LIDARs [41]. The drawback is their size ($\sim 2GB/km$), which is difficult to manage in a long trip or in a city. The ultimate goal is to drive everywhere with full functionality, but there are two main points of view on how to get to it. The first approach tries to drive in some places with full functionality using 3D detailed map ($\sim GB/km$) and low resolution sensing. On the other side, the goal is to drive everywhere with partial functionality using low resolution

maps and high resolution sensing. The problem of using a 3D highly detailed map is the scalability of the map and the updates. The problem of using low resolution images is that more intelligent and realtime perception algorithms are needed. The ultimate goal is to get cognitive perception as humans do.

3. CNN-Based Road Classifier

On the one hand, and given the generalized use of CNNs on the road detection problem, this paper proposes a road detector based on the ResNet network model [15] adapting its last stages to the fully convolutional architecture [7]. First, a ResNet-50 model already trained on the ImageNet dataset (1000 labels at image level) has been used. In contrast, our system is designed on the KITTI road detection dataset [30], which only defines 2 labels (road/nonroad) at pixel level. Accordingly, the original ResNet-50 architecture has to be transformed into a fully convolutional network to admit an input of an arbitrary size and to produce an output of the same size with pixelwise segmentation. In order to do that, the last inner-product fully connected layer (1000 outputs) is replaced with a new convolutional layer (two outputs) that will be learned from scratch. In addition, due to the fact that ResNet has an overall downsampling factor of 32, some upsampling stages are needed. As can be observed in Figure 2 the upsampling is performed in three interpolation stages:

- (1) UPSCORE 32: the main output scores are upsampled by a factor of 2. The output from the previous block

CONV 4 is added, since both scores have the same accumulated downsampling factor (16).

- (2) UPSAMPLE 16: the previous result is upsampled by a factor of 2. The output from the previous block CONV 3 (downsampling factor of 8) is then added.
- (3) UPSAMPLE 8: the result is upsampled by a factor of 8 to recover the scores in the original input size.

This process allows to recover pixelwise scores smoothly, with a high level of detail. The final output combines the coarser global features (MAIN SCORE) with some finer local features (SCORE CONV 4 and SCORE CONV 3).

Upsampling layers are initialized with bilinear interpolation kernels that do not need to be trained. The scores from the shallower layers are obtained with a two-output convolutional layer in the same manner as the main score. Also, a learnable scaling layer is placed before each one to help the network to adapt the different features to their addition. A large padding is added on the first stage (CONV 1) to compensate for the width and length reduction that pooling layers and convolutions combined with downsamplings can cause. Finally, some croppings have to be performed to align the score maps and match dimensions, with an offset which is calculated automatically during the architecture definition. The final output of our architecture consists of two channels that represent the probability maps for background and road respectively, obtained from the final SOFTMAX layer.

3.1. Data Augmentation. One of the main weaknesses of CNNs is their dependence on the training data. With data augmentation a better generalization can be achieved and different road conditions can be simulated, increasing the robustness of the network against illumination, color or texture changes, or variations in the orientation of the cameras. We adopt an online augmentation approach where modifications are performed at random each time. This way, the CNN never sees the same augmented image twice and this virtually infinite dataset does not require extra storage space on disk.

It can be distinguished between geometric transformations and pixel-value modifications. It is also possible to apply several augmentations on the same image or to apply different augmentations for each label (road or background) or to patches.

3.1.1. Geometric Transformations. These transformations are applied to both the image and the ground truth mask. Zero padding is added when needed to keep the original image size and the padding pixels are assigned to the “ignore label” of the classifier. The applied transformations include the following:

- (i) Random affine transformations: translations, rotations, scalings, and shearings are performed in order to change the positions of the points, while keeping lines parallel.
- (ii) Mirroring: apart from the affine transformations, horizontal flipping is performed independently to easily double the size of the training set.

- (iii) Cropping and scaling: the original image is cropped and scaled to the original size. Crops are defined by a random top left corner and also random size, within image limits.
- (iv) Distortion: random distortion parameters are applied to the image.
- (v) Perspective transformations: the original positions are selected empirically on road limits. The final positions are calculated adding Gaussian noise to the original ones with two restrictions. First, the shift of top points is the opposite of that of the bottom ones. Second, top points should not cross each other to prevent reflected images.

3.1.2. Pixel-Value Changes. These transformations are only applied to the image, since they produce changes only on pixel values.

- (i) Noise: random addition of Gaussian, speckle, salt & pepper noise, generation of an image with signal-dependent Poisson noise.
- (ii) Blur: the filters are applied independently to the image, creating a blur effect. The selected filters are: Gaussian, diagonal motion (left-to-right or right-to-left, at random), box, median and bilateral filtering.
- (iii) Color changes: three types of transformations are applied. The first one is casting, which consists in adding a random constant to each RGB channel [42]. The second one is an additive jitter, which is generated at random by means of exponentiation, multiplication and addition of random values. The last one is a PCA-based shift [13] which consists in adding to each pixel a linear combination of the found three principal components with magnitudes proportional to their corresponding eigenvalue times a random Gaussian variable (zero mean and standard deviation of 0.01).

3.2. Network Components and Training Variants. Regarding fully convolutional CNNs, there are several elements that can be optimized, such as the initialization of the score layers (with zeros, noise, etc.) and the initialization and training of the upsampling layers. We can also use more complex activation functions rather than the simple ReLU, such as parametric ReLUs (PReLU), which are recommended in combination with MSRA initialization [43]. Although it is not possible to change the original pretrained ResNet-50 structure, we can add PReLU to the new score layers. Training alternatives involve trying different learning rates and learning rate policies, such as decreasing the learning rate when the training stalls in previous trials or doing a *warmup* stage [15] at a reduced learning rate until error goes under 20%. Another alternative is to have a higher learning rate for score layers, which are learned from scratch and a lower rate for inherited layers. Moreover, in [7], several training schemes are defined: the standard accumulated learning (batch size of 20 and standard momentum of 0.9) or the *heavy learning* scheme, which uses a single image per gradient

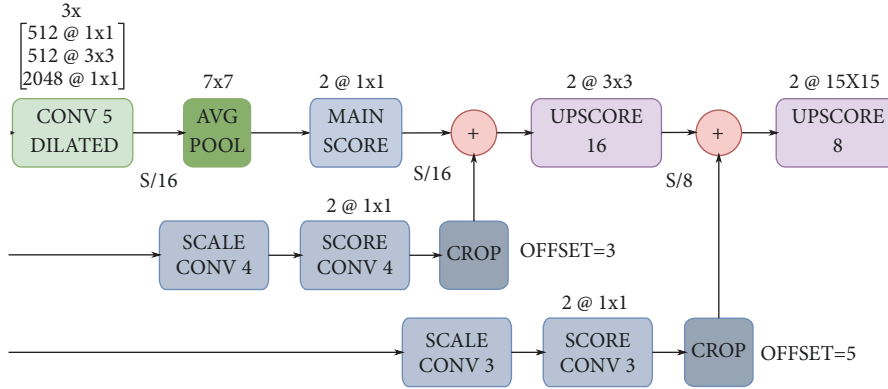


FIGURE 3: Detail of the scoring stage with dilated convolution in CONV 5. It removes the necessity of using the UPSCORE 32 layer.

actualization and a high momentum of 0.99, that simulates the gradient accumulation effect of the batch size. In [21] they use a variant of the accumulated learning (batch size of 12) with a polynomially decreasing learning rate which we try in the form $2.5 \cdot 10^{-4} \cdot (1 - \text{iter}/\text{max_iter})^{0.9}$.

3.3. Training in Bird's Eye View. The traditional training approach uses images in perspective view and obtains detections in this space. However, since KITTI benchmark evaluates its results with the F1-measure in Bird's Eye View (BEV) [44], we also train the network model directly in BEV. In this case, a less aggressive data augmentation strategy is used since geometric transformations in BEV creates important distortions.

3.4. Deeper Models. A ResNet-101 [15] model has been adapted using the same procedure applied to the ResNet-50, to test a deeper model in this problem. On the one hand, this model has an increased learning capacity. On the other hand, the risk of overfitting becomes more relevant.

3.5. Upsampling Variants. Apart from the schematics presented in Figure 2, the number of connections from shallower layers is modified. In order to obtain a more fine grained classification, both the full step-by-step upsampling (with additional connections from CONV 2 and CONV 1) and the four-step one (only additional connection from CONV 2) have been tested. Likewise, a two-step approach has also been evaluated to cover all possible cases, as well as the basic approach with no skip connections and just one large interpolation step.

There are other methods to increase the field of view of the deeper layers without downsampling the input features. The dilated convolution [20] and its improved version [21], which is claimed to avoid grid effects, are evaluated. This approach replaces the downsampling performed in one or more blocks with dilated convolutions in all of the subsequent layers. However, downsampling not only is necessary to enlarge the field of view but also plays an important role reducing the size of the input features to reduce the GPU memory consumption. If downsampling is completely removed, the

model would not fit in memory. For this reason, our tested method combines a dilated convolution in the deeper block of the ResNet-50, with two upsampling steps to achieve a tradeoff (see Figure 3).

4. High Level Interpretation Module

4.1. Digital Navigation Maps. There are two main types of maps. The first ones are navigation maps, which provide information about the steps to reach our destination. The second ones are high definition maps, which provide 3D information of the environment with centimeter precision. Most of the autonomous navigation vehicles are based on these types of maps [3]. In contrast to that, our approach is closer to the human way of driving. Human drivers do not need high definition maps. They drive using visual perception and local navigation methods. The only information they need are the indications and steps on the navigation map to reach the destination.

The digital navigation map used for our approach is Open Street Map. This collaborative map is created by a large community around the world and all the information stored in the map is editable and it is freely accessible. The map consists of a list of streets called ways. Every way is composed of a list of nodes with a location and its relations with the other nodes and ways. Thanks to the location and relation between the nodes, the shape of the current street and its surroundings can be estimated. In addition, digital maps include the number of lanes and road type.

4.2. Hand-Crafted Features. In order to make the paper self-contained, we provide a brief description of the main features used in the high-level interpretation module. See [8] for more details.

- (i) Vegetation areas: in order to segment the green areas of the image, multithresholding is applied over an area of the Hue channel of the HSV color space. A basic filtering is then applied on the resulted image.
- (ii) Obstacles: 3D points from the stereo cameras are processed to estimate normal (n_i) and curvature vectors (γ_i). Points having components $n_x \geq 0.5$,

TABLE 1: Model parameter range of values.

PARAM	START	COARSE		FINE	
		RANGE	STEP	RANGE	STEP
Lane Width	3.20 m	± 0.40 m	0.20 m	± 0.20 m	0.10 m
Lateral Displacement	0.0 m	± 8.00 m	0.50 m	± 0.50 m	0.10 m
Angular Offset	0.0°	$\pm 5.00^\circ$	0.25°	$\pm 0.20^\circ$	0.10°
Longitudinal Offset	0.0 m	± 8.00 m	0.50 m	± 0.50 m	0.10 m
Curvature Radius	7.00 m	± 3.00 m	1.0 m	± 1.0 m	0.50 m

$n_y \geq 0.5, \gamma_z \geq 0.5$ are considered as belonging to large obstacles. Every vector component is filtered by area and merged together to obtain the final result.

- (iii) Curbs: the curvature vectors variation of the 3D points is a good feature for the detection of curbs. However, depending on the curb height, curvature values differ in each scene. Five types of curvatures are segmented using five pairs of thresholds for different curb heights. The resulting clusters are filtered independently using morphological operators and contour analysis and finally merged.
- (iv) Road markings: a median filter is applied to the input image to remove white objects with horizontal size lower than the maximum horizontal size of markings. An adaptive window size of the median filter is needed due to perspective constraints. On the other hand, an adaptive thresholding with variable window size is applied to the input image to obtain white objects corresponding to markings. Both images are then subtracted to get the final result. This approach can be also carried out using the BEV image, with constant window sizes.

4.3. Road Map Modeling. The features extracted from the cameras (including the features obtained in [8]) are fused with digital navigation maps to obtain a high level interpretation of the urban scene. The map includes relevant information about the presence of railways, parking areas, buildings, or intersections, which are key points for a correct scene interpretation. Most of the elements in the map model are fixed (buildings, gardens, etc.). However, the road width should be adapted depending on the road type (residential, highway, intersection, etc.). Our proposed model has 6 degrees of freedom: number of lanes (l), lane width (w), lateral offset (y), longitudinal offset (x), angular offset (a), and curvature radius in intersections (r); see Figure 4.

It is assumed that the number of lanes is determined by the map. Nowadays, most of the maps indicate the number of lanes and the lane you should drive to reach your destination. The other parameters are evaluated in two steps, the first one for a coarse adjustment and the second one for a fine adaptation. Table 1 shows the range of every parameter, obtaining more than $\sim 800K$ and $\sim 15K$ combinations in the coarse and fine adjustment, respectively. The integration of the vehicle ego-motion with the previous models along the time creates a prior knowledge where the model should be. This prior knowledge removes the coarse adjustment and

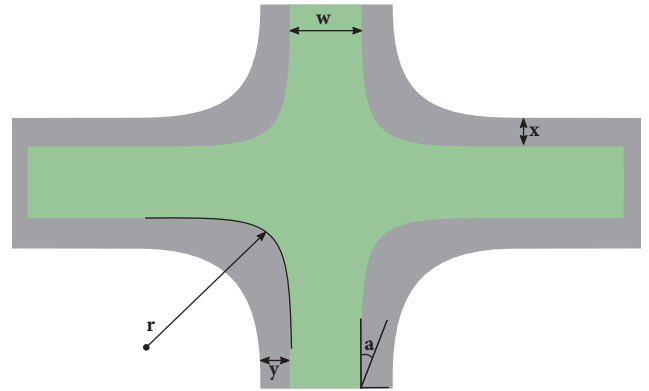


FIGURE 4: Parameters to adjust in the proposed model of the urban road. The real scene is painted in grey and the model to adjust is shown in green.

the fine adjustment could be extended. The selected option to reduce the number of combinations divides the process in three steps: the first step combines lane width and lateral displacement, the second one adjusts the angular offset, and finally the third step combines the longitudinal offset and the curvature radius. This process reduces the number of combinations from more than $\sim 815K$ to only 512.

The metric to evaluate the best adjustment is calculated using equations (1) and (2), where the precision and recall are computed by matching the model and the sensing of the environment. The matching is evaluated in 4 different groups (G). The first one compares vegetation areas (garden, grass, and forest) and obstacles (barriers, buildings and walls) [8]. The second one is the road provided by the CNN model. The third group is composed of curbs and road markings [8]. The last one only compares curbs [8] in order to reinforce the correct adjustment of the road boundaries. The weights of every group in the final score (S) are set after a training stage to optimize the correct adjustment. The combination of the parameters with the highest score (S) generates a map-based model which is the output of the algorithm. This map is then projected at the same space of the images obtained from the cameras.

$$F_{score} = \frac{2 \cdot precision \cdot recall}{precision + recall} \quad (1)$$

$$S = \sum_i w_i \cdot F_{score}(G_i) \quad (2)$$

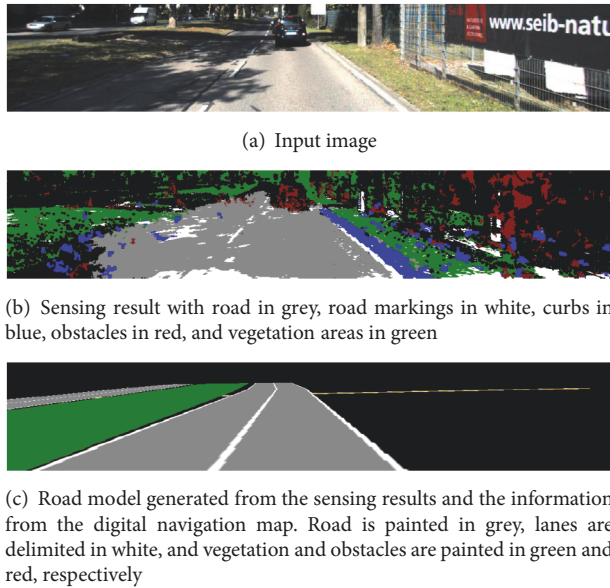


FIGURE 5: Results of the high-level interpretation module.

As an example, we depict in Figure 5 the results of the sensing stage. The map model is adapted to the current scenario taking into account every detected feature and their correspondence in the map model.

Note that the resulting model is based on static elements such as the road, vegetation areas, and buildings. This is a clear advantage in order to maintain an updated and enhanced version of the high-level structure of the digital navigation maps. When performing autonomous navigation, the drivable area is directly provided by the CNN-based road classifier which will not include dynamic objects such as pedestrians, bicycles, and cars. Previous hand-crafted [45–48] and more recent deep learning-based approaches [19] can be here adopted to detect dynamic obstacles.

5. Results

5.1. Experimental Setup. Our CNN-based road segmentation model was trained on the KITTI dataset, which is composed of 289 images manually labeled with two classes: road and nonroad. 50% of the images were used for training the net, and the remaining 50% are kept aside for validation.

More specifically, the ResNet-50 model previously trained on ImageNet was used for weight initialization, and then the full net is fine-tuned on the road detection task. This is performed with stochastic gradient descent at a constant learning rate of $5 \cdot 10^{-5}$ (except for the bilinear filters, which are fixed), weight decay of $5 \cdot 10^{-4}$, one image per iteration and high (0.99) momentum. This scheme is referred as *heavy learning* in [7]. The training is run for 24K iterations, with validation checkpoints every 4K iterations.

The Caffe framework [49] was used for the network prototype definition and the control of training and testing processes. In addition, a Python input data layer, adapted for KITTI images, is used for loading images and labels into

the net: each training image is randomly picked along with its corresponding label, and some minimal preprocessing must be done. The ImageNet per-channel pixel mean is subtracted, and the label images are converted into a $1 \times \text{height} \times \text{width}$ integer array of label indexes to be compatible with the loss function. As stated before, instead of passing the original image to the network, some data augmentation operations were applied to extend the training set, prevent overfitting, and make the net more robust to image changes. The data augmentation layer runs on CPU, and the rest of the processing can be done on GPU. It takes between 2 and 3 hours to complete the standard training on a single NVIDIA® Titan X GPU.

The parameters of the high-level interpretation module were estimated using the same images used to train the ResNet segmentation module from KITTI dataset. This module was implemented in C/C++. Training stage takes less than 1 minute and online estimation is performed in real time ($< 5\text{ms}$).

5.2. Road Segmentation Results. In this subsection we present the results obtained from the different variations on the network previously mentioned. As proposed in [44], quantitative results are calculated in terms of F1-measure, computed over the validation subset on Bird’s Eye View (converting from perspective view when the road detector is trained in this space). Namely, the MaxF is computed using the working point (confidence threshold) in the precision-recall curve that maximizes F1-measure.

5.2.1. Data Augmentation. We can see, in the training and validation losses curves (Figure 6), that the use of data augmentation prevents the network from overfitting, since the gap between training and validation losses disappears: training losses rise slightly whereas validation losses decrease. We have observed that geometric transformations introduce higher variability than pixel value changes, and using both kinds we obtain the best results. We transform the full image with a single random operation each time. This way, we get a high variability, as we note from the width of the losses curve in Figure 6(b), but with an acceptable level of noise.

With this method we can achieve an improvement of approximately 1% in F-measure when training in perspective space (from 94.59% to 95.76%), and 2% when training in Bird’s Eye View, which will be discussed later on.

Moreover, the trained model was tested on some sequences at the campus of the University of Alcalá, Madrid (Spain), to test the network in a different environment from that used in the training. Figure 7 demonstrates that data augmentation makes the model more robust against illumination, texture, perspective, and orientation changes.

5.2.2. Network Components and Training Variants. The upscore described in Figure 2 is composed of fixed bilinear kernels and score layers are initialized using the MSRA method because it is considered robust against symmetries in gradient propagation. Neither finetuning the bilinear filters (slower convergence and we get the same kernels in the end)

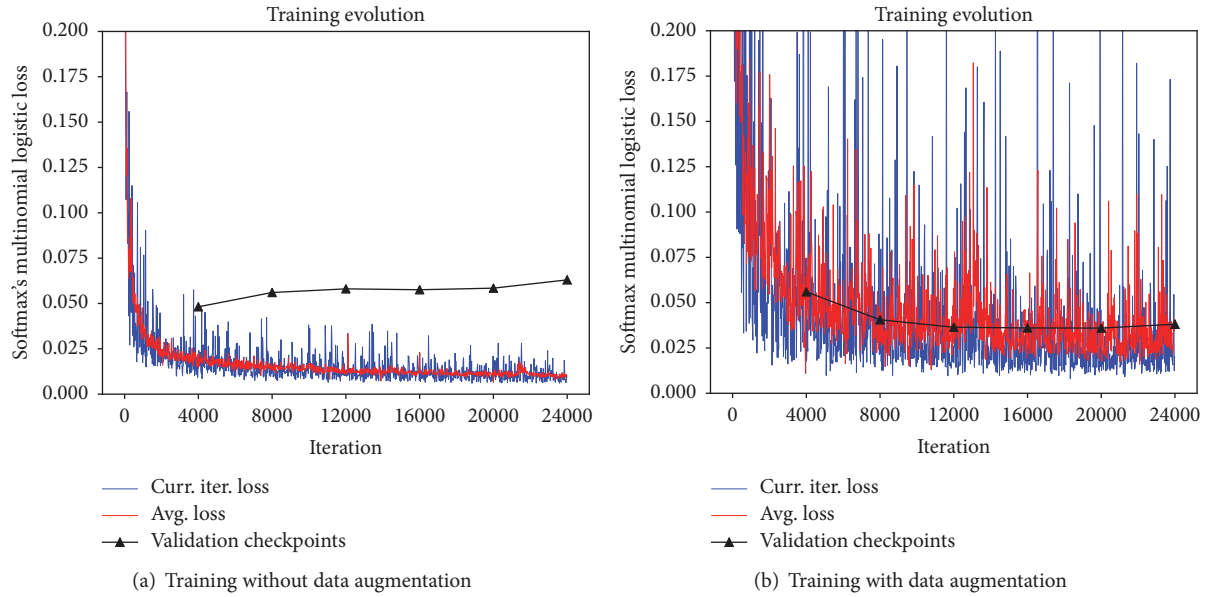


FIGURE 6: Comparison of the training and validation losses without data augmentation and with data augmentation.



FIGURE 7: The qualitative results of the proposed model are coloured as follows: TP in green, FP in blue, and FN in red. The scenarios with strong illumination changes and challenging road textures are better detected in the model trained with data augmentation.

nor learning them from scratch (smoother convergence and we get different interpolation kernels that use information from both classes' scores and scattered road limits) improves our previous performance. Regarding the learning rate, three different rates are compared ($1 \cdot 10^{-6}$, $5 \cdot 10^{-5}$, and $1 \cdot 10^{-4}$). The slower one ($1 \cdot 10^{-6}$) does not converge even with 40K iterations, the faster one ($1 \cdot 10^{-4}$) adds instability to the process, and the best results are obtained with the trade off between both approaches ($5 \cdot 10^{-5}$).

Since we are performing a fine-tuning with few iterations, changing the learning rate does not seem to have positive effects: decreasing policies can yield a monotonically decreasing validation losses curve, but the final losses and the F-measure are not better, and if the decrease is too abrupt, the training will become unstable, probably due to the high momentum; with the *warmup* scheme, there are no improvements either. It also appears to be better to let the whole net adapt to the new task of per-pixel road detection, instead of using a reduced ($0.1\times$) learning rate for inherited layers. Finally, accumulated learning policies lead to worse results and are also much slower (proportionally to the batch

size). The polynomially decreasing learning rate variant from [21] is better but still not superior to *heavy learning* in our trials.

5.2.3. Training in Bird's Eye View. In general, the model is able to learn better (less training losses) and also to generalize better (smaller gap with validation losses) during the training in perspective view because perspective images have more information about the scene, and more aggressive data augmentation recipes can be applied while maintaining the meaning of the image. Thus, without data augmentation, the model trained in BEV (94.08%) is worse than the one in perspective view (94.59%).

Data augmentation can significantly reduce the gap between training and validation losses and makes it worthwhile to train in BEV. Although the BEV approach with data augmentation is still worse at learning than the perspective one, the fact of learning in the same space as the evaluation obtains a better performance (96.06%). Analyzing in detail the performance, the model trained in BEV performs similarly at near and further pixels, whereas the perspective model

TABLE 2: Comparative results of deeper models performance.

Model	F-measure	Training Time	Iterations	Memory
ResNet-50	95.76	2h00	24K	7GB
ResNet-101	95.88	2h30	20K	10GB

TABLE 3: Quantitative results in F-measure on KITTI dataset for the main network variants evaluated.

Data aug.	BEV train	ResNet-101	4-step up.	F-measure
	✓			94.08%
				94.59%
✓				95.76%
✓			✓	95.80%
✓		✓		95.88%
✓	✓			96.06%
✓		✓	✓	96.09%
✓	✓	✓		96.13%
✓	✓	✓	✓	96.31%

has more problems with further pixels. Some problems of the BEV approach is that, in some cases, buildings at the end of the road or incoming tunnels can be confused with a continuation of the road.


5.2.4. Deeper Models. The training tests in perspective space over deeper models establish that the ResNet-101 achieves slightly better results over ResNet-50, which are obtained consistently with fewer iterations. As a drawback, the training takes slightly more time to complete than with ResNet-50 and more GPU memory is needed; see Table 2.

In the experiments evaluated in BEV, it is observed overfitting in the learning curve (training losses decrease while validation ones do not) because the deeper model has more learning capacity and needs a larger training set to generalize. Therefore, the training is stopped at 20K iterations to avoid the problem and the obtained F-measure is (96.13%). In conclusion, ResNet-101 offers a small but consistent improvement in detection performance, at the expense of needing more computing resources and time.

5.2.5. Upsampling Variants. Different upsampling variants are evaluated in a ResNet-50 trained in perspective view. As expected, the detections with the full step-by-step upsampling scheme have the highest resolution, but they are noisier and the F-measure is worse (95.49%), probably because the extracted features come from too shallow layers with little knowledge of the full scene. In the four-step case, the resolution is higher than in the original setup and the F-measure is slightly upraised (95.80%). The four-step approach has also been tried with a ResNet-50 trained in BEV and a ResNet-101 trained in both perspective and BEV spaces. Whereas the ResNet-50 gives similar results (95.97%), the ResNet-101 yields the best detections so far, with a F-measure of 96.09% and 96.31% in perspective and BEV spaces, respectively. The dilated convolution approaches yield also similar



(a) Single cyclist well segmented



(b) Group of cyclists considered a single obstacle

FIGURE 8: Results from the perspective-trained model in a scene with cyclists.

results. In particular, the method from [20] combined with two upsampling steps seems to be as good as the four-step approach in a ResNet-50 and less (20K) iterations, but it does not improve the results with the ResNet-101.

5.2.6. Global Road Segmentation Results. Table 3 summarizes the quantitative results in F-measure over our KITTI validation subset for the most interesting network variants. The baseline algorithm is the ResNet-50 model with the fully convolutional architecture and three-step upsampling shown in Figure 2.

The two best-performing methods, namely, the ResNet-101 with data augmentation and four-step interpolation, are trained with perspective and BEV images. Small obstacles such as pedestrians, cyclists, or cars are well differentiated from the road areas (Figure 8(a)), although two cyclists riding together are considered as a single obstacle (Figure 8(b)) since the space between them is not well segmented. This problem is also present when training in BEV and may be solved with higher resolution approaches.

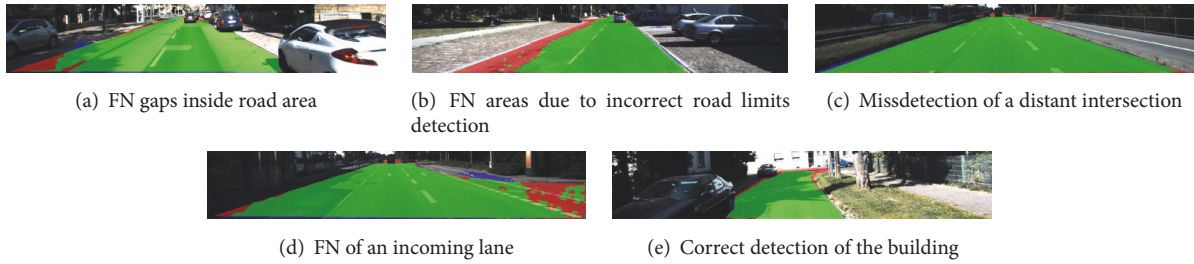


FIGURE 9: Results from the perspective-trained model.




FIGURE 10: Results from the BEV-trained model.

Both models sometimes leave FN gaps (Figures 9(a) and 10(a) on top-right corner), as well as FP patches outside road limits (Figure 9(d)) that could be filtered with some postprocessing methods. However, the model trained in BEV seems to be better delimiting road limits in the same image (Figure 10(b)) because in this representation they are straighter.

It can be seen that the BEV-trained model is better at detecting irregular road limits (Figures 10(c) and 10(d)) than the perspective-trained one (Figures 9(c) and 9(d)). However, the main problem of the BEV-approach is that, in some particular cases, the resulting image is so distorted and the net confuses buildings with the continuation of the road (Figure 10(e)). This would be very unlikely to happen if the image was analyzed in perspective space.

5.3. High-Level Interpretation Results. Considering the hand-crafted features of our scene interpretation module, we remark the following statements. The curb detection method was analyzed in detail in [10]. The algorithm was compared using different sources for the 3D cloud data (LIDAR and stereo), obtaining a lateral RMSE of 12cm in a range from 6m to 20m. The proposed algorithm can deal with curbs of different curvature and heights, from as low as 3cm, in

a range up to 20m whenever that curbs are connected in the curvature image. Finally, the use of fixed or empirical thresholds is avoided given that the proposed function is adapted automatically for different road scenes depending on the predominant curvature value. The boosting classifier was described and tested in [8, 11]. The weight of each feature in the final road classification reveals that 3D features (Y and Z coordinates) and its 2D representation (column and row) are the most discriminant features. However, some of the other 2D features are still important for the boosting classifier such as the grey value of (HSV) or the vegetation. Some of the road misclassifications are produced in sidewalks, where the only difference between the road and sidewalk is a small curb. Furthermore, in challenging urban scenarios the limit of a drivable area is difficult to distinguish from the nondrivable area, such as a cyclist lane. In some cases the limit is just a road marking or a variation in the pavement texture. It is remarkable that road markings have a low weight in the final response.

According to the results obtained in [8, 11], the weight of curbs and road markings in the road classification is very small because that type of features describes road limits instead of the road surface. The proposed algorithm follows a human reasoning. Therefore road markings and curbs should have a relevant role in the interpretation of the environment.

TABLE 4: Quantitative results on F-score in KITTI dataset.

Classifier	Image plane				Bird Eye View			
	UM	UMM	UU	All	UM	UMM	UU	All
Boosting	85.24	93.09	82.08	87.06	85.15	89.37	62.60	79.05
Map Model	89.35	88.56	90.20	89.37	85.17	84.76	86.66	85.53
CNN	96.64	95.36	94.89	95.50	95.09	94.94	93.39	94.59

TABLE 5: F-measure score for urban marked road, urban multiple marked lane, and urban unmarked road respectively of the KITTI test benchmark.

Type of road	UM	UMM	UU	All
F-measure	94.16%	95.45%	91.17%	93.98%

That is the reason to match the map model with 4 different groups of features: vegetation + obstacles, road, curbs + road markings, and curbs. After a weight training stage for each group, the qualitative results demonstrate the effectiveness of the proposed method to infer complex urban environments.

When evaluating our previous boosting-based classifier [8] with the new CNN-based approach, we can state that our CNN model with data augmentation clearly outperforms the classic hand-crafted features + classifier approach.

Table 4 shows the F-measure using the training set of the KITTI dataset (50%/50% training/validation) in both perspective view and BEV of urban marked road (UM), urban multiple marked lane road (UMM), and urban unmarked road (UU) and for all types of road (All). As can be observed, the CNN-based road segmentation provides an overall F-measure a 8.44% and a 15.54% better than the value provided by our previous approach in the perspective image and BEV, respectively. In addition, we have obtained the F-measure using the KITTI test dataset for the CNN-based approach (see Table 5). As can be observed, the performance decreased by 0.61% compared with the validation results, which clearly demonstrates that overfitting has been avoided.

We have also obtained the F-measure of the high-level interpretation module (see Table 4) for our validation dataset, yielding overall values of 89.37% and 85.53% in perspective view and BEV, respectively (note that we have not been able to obtain results from the KITTI test dataset since GPS positions are not available). These values are a 6.13% and 9.06% worse than the CNN classifier in the image plane and BEV, respectively. This is obviously an unfair comparison since the high-level road estimation module is not a pixelwise approach trained with the ground truth location of the road. It does not include dynamic obstacles, and, in some cases, it provides some regions of the road that are not even labeled in the ground truth (opposite lanes or not visible intersections). In general, pixelwise classifiers outperform model-based approaches because there is not any model that fits as close to the ground truth as the pixelwise classifiers. Figure 11 shows an example of a very complex intersection where the pixelwise classifier outperforms the map model




(a) Input image



(b) CNN classifier with data augmentation



(c) Map model



(d) Open street map diagram with aerial imagery layer

FIGURE 11: Example of a very complex intersection where the map-based model is not well fitted to the road limits. However, it provides richer high-level information to be used in autonomous navigation systems.




FIGURE 12: Example of a scene where the pixelwise classifier has a large number of false negatives and the use of a map-based model improves the result.

due to its inner architecture. The ways are usually centered with respect to the center of the real road, but this scene has ways that have different number of lanes before and after the intersection. In addition, the lanes for left turning are overlapped with each other, which is impossible to adjust to the map architecture. However, the map model includes relevant information to be supplied to any navigation module of an autonomous navigation system.

In some other scenarios, the map-based model improves pixelwise classifiers. Figure 12 shows an example of an urban scene where the CNN road classifier without data augmentation has many FN on the left and right boundaries of the road. The road model fills the missing pixels close to the road boundaries and it obtains a shape that fits better to the real scenario.

The qualitative results show the added value of the interpretative approach presented in this paper. Figure 13 shows an urban street with one lane for each direction separated by a fence. Furthermore, on the right side of each lane, there are slots for parking and buildings. Even with a coarse detection of the road (Figure 13(b)) or neither detect the road, the map model fits well to the scene and add high level information to the system. The use of different features for the map matching increases the robustness of the method because the absence of one of them does not make the system to fail. The scene represented in Figure 14 highlights the use of high level information extracted from the map. The use of map information reduces the possibilities to infer the cycle lane as a road lane and because of that the railways and the cycle lane are correctly labeled. Finally, Figure 15 shows an urban street where the map model fits well to the environment but the presence of parked vehicles creates a high number of FP. Dynamic obstacles, such as cyclist and vehicles, are effectively removed and not considered as drivable by the CNN-based road segmentation system. The high-interpretation module maintains the actual structure of the road besides the dynamic obstacles.

6. Conclusions and Future Work

A novel high-level interpretation approach that integrates pixelwise road segmentation, a set of hand-crafted features that describe road limits (vegetation, road markings and curbs), and enhanced data provided by digital navigation maps were presented. A ResNet-50 CNN model with a fully convolutional architecture and three interpolation steps, which are finetuned in perspective KITTI images, were used to obtain an accurate detection of the road, which represents the drivable area. Several variations were introduced to improve the training such as data augmentation, training in BEV space, training parameters tuning, using deeper models, and other upsampling architectures. Data augmentation offers a significant improvement within 2% in F-measure. The use of a ResNet-101 model with a four-step upsampling scheme, trained directly in BEV with data augmentation, improves the results up to a 96.31% in the validation subset. The proposed approach clearly outperforms our previous Boosting-based road detection approach [8].

The presented approach can be applied to update digital navigation maps using floating vehicles equipped with cameras. In parallel, an accurate estimation of the drivable area is supplied by our CNN-based road segmentation module.

Future works will be devoted to obtain smoother road detection results by adding a postprocessing layer into the system. More degrees of freedom can be considered to be capable of handling multilane roads with different lane widths. In addition, due to the new advances in semantic segmentation [19], the data provided by digital navigation maps and the output given by the high-level interpretation module will be enriched with new variables, including traffic lights, traffic signs, or even urban furniture (benches, bins, bus stops, etc.). Finally, different sensor ensembles will be tested to facilitate a quantitative evaluation of the high level interpretation module.

Data Availability

The data used to validate our approach corresponds to public dataset, in this case KITTI dataset.

Conflicts of Interest

The authors declare that they have no conflicts of interest.

Acknowledgments

This work was funded by Research Grants SEGVAUTO S2013/MIT-2713 (CAM), DPI2017-90035-R (Spanish Min. of Economy), BRAVE Project (H2020, Contract No. 723021), and UAH (CCGP2017-EXP/055). In addition, this project has received funding from the Electronic Component Systems for European Leadership Joint Undertaking under Grant Agreement no. 737469 (AutoDrive Project). This Joint Undertaking receives support from the European Unions Horizon 2020 Research and Innovation programme and Germany, Austria,




FIGURE 13: Results in a scene with parking slots on the right and vegetation on the left. Even with a coarse detection of the road (Figure 13(b)), the map model fits well to the scene and add high-level information to the system.




FIGURE 14: The use of a map model increases the robustness of a road classifier because the other boundary features compensate some road misclassifications.

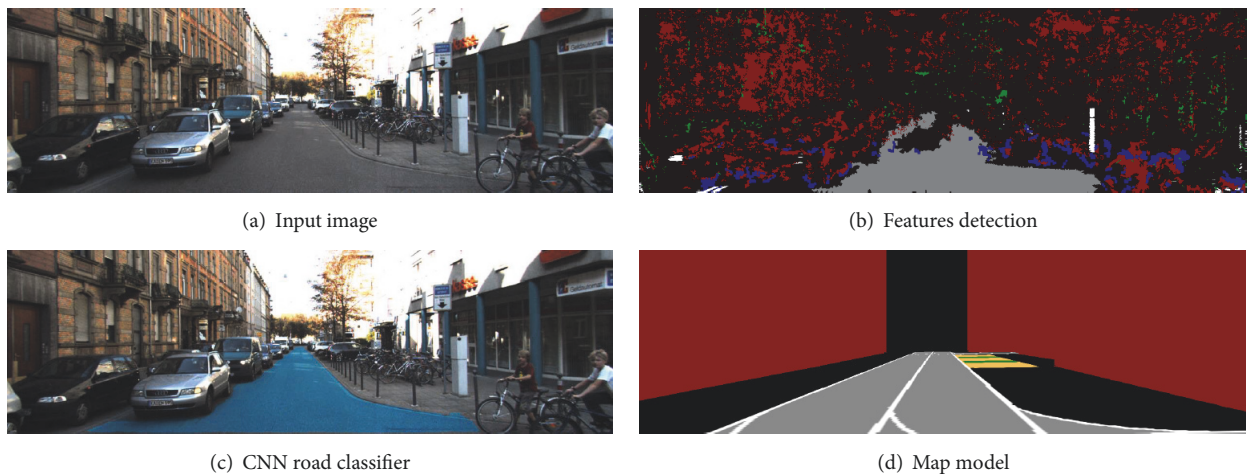


FIGURE 15: Dynamic obstacles (cyclist and vehicles) are effectively not considered as drivable by the CNN-based road segmentation system. The high-interpretation module maintains the actual structure of the road besides the dynamic obstacles.

Spain, Italy, Latvia, Belgium, Netherlands, Sweden, Finland, Lithuania, Czech Republic, Romania, and Norway.

References

- [1] Analytical report on road safety, “European Commission, Tech. Rep., 2010”.
- [2] D. Fernandez-Llorca, R. Q. Minguez, I. P. Alonso et al., “Assistive Intelligent Transportation Systems: The Need for User Localization and Anonymous Disability Identification,” *IEEE Intelligent Transportation Systems Magazine*, vol. 9, no. 2, pp. 25–40, 2017.
- [3] J. Ziegler, P. Bender, M. Schreiber et al., “Making bertha drive an autonomous journey on a historic route,” *IEEE Intelligent Transportation Systems Magazine*, vol. 6, no. 2, pp. 8–20, 2014.
- [4] P. Czerwionka, *A Three Dimensional Map Format for Autonomous Vehicles*, Freie Universität, Berlin, Germany, 2014.
- [5] E. Guizzo, “How google’s self-driving car works,” *IEEE Spectrum*, 2011, <http://spectrum.ieee.org/automaton/robotics/artificial-intelligence/how-google-self-driving-car-works>.
- [6] I. Parra Alonso, R. Izquierdo Gonzalo, J. Alonso, A. Garcia-Morcillo, D. Fernandez-Llorca, and M. A. Sotelo, “The experience of drivertive-driverless cooperative vehicle-team in the 2016 gcdc,” *IEEE Transactions on Intelligent Transportation Systems*, 13 pages, 2017.
- [7] E. Shelhamer, J. Long, and T. Darrell, “Fully Convolutional Networks for Semantic Segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 4, pp. 640–651, 2017.
- [8] C. Fernández, D. Fernández-Llorca, and M. A. Sotelo, “A Hybrid Vision-Map Method for Urban Road Detection,” *Journal of Advanced Transportation*, vol. 2017, Article ID 7090549, 21 pages, 2017.
- [9] J. Munoz-Bulnes, C. Fernandez, I. Parra, D. Fernandez-Llorca, and M. A. Sotelo, “Deep fully convolutional networks with random data augmentation for enhanced generalization in road detection,” in *Proceedings of the 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 366–371, Yokohama, October 2017.
- [10] C. Fernandez, D. F. Llorca, C. Stiller, and M. A. Sotelo, “Curvature-based curb detection method in urban environments using stereo and laser,” in *Proceedings of the IEEE Intelligent Vehicles Symposium, IV 2015*, pp. 579–584, July 2015.
- [11] C. Fernandez, R. Izquierdo, D. F. Llorca, and M. A. Sotelo, “A Comparative Analysis of Decision Trees Based Classifiers for Road Detection in Urban Environments,” in *Proceedings of the 18th IEEE International Conference on Intelligent Transportation Systems, ITSC 2015*, pp. 719–724, September 2015.
- [12] C. Fernandez, R. Izquierdo, D. F. Llorca, and M. A. Sotelo, “Road curb and lanes detection for autonomous driving on urban scenarios,” in *Proceedings of the 2014 IEEE 17th International Conference on Intelligent Transportation Systems (ITSC)*, pp. 1964–1969, Qingdao, China, October 2014.
- [13] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *Proceedings of the 26th Annual Conference on Neural Information Processing Systems (NIPS ’12)*, pp. 1097–1105, Lake Tahoe, Nev, USA, December 2012.
- [14] K. Simonyan and A. Zisserman, “Very Deep Convolutional Networks for Large-Scale Image Recognition,” *Computer Vision and Pattern Recognition*, 2014, <https://arxiv.org/abs/1409.1556>.
- [15] K. He, X. Zhang, S. Ren, and J. Sun, “Deep residual learning for image recognition,” in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016*, pp. 770–778, July 2016.
- [16] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. Alemi, “Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning,” *Computer Vision and Pattern Recognition*, 2016, <https://arxiv.org/abs/1602.07261>.
- [17] S. Ren, K. He, R. Girshick, and J. Sun, “Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 6, pp. 1137–1149, 2017.
- [18] J. Redmon and A. Farhadi, “YOLO9000: Better, Faster, Stronger,” in *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6517–6525, Honolulu, HI, July 2017.
- [19] K. He, G. Gkioxari, P. Dollár, and R. B. Girshick, “Mask R-CNN,” *CoRR*, vol. abs/1703, 2017, <http://arxiv.org/abs/1703.06870>.
- [20] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille, “DeepLab: Semantic Image Segmentation with Deep Convolutional Nets, Atrous Convolution, and Fully Connected CRFs,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 4, pp. 834–848, 2018.
- [21] P. Wang, P. Chen, Y. Yuan et al., “Understanding Convolution for Semantic Segmentation,” in *Proceedings of the 2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pp. 1451–1460, Lake Tahoe, NV, March 2018.
- [22] H. Noh, S. Hong, and B. Han, “Learning deconvolution network for semantic segmentation,” in *Proceedings of the 15th IEEE International Conference on Computer Vision, ICCV 2015*, pp. 1520–1528, Santiago, Chile, December 2015.
- [23] V. Badrinarayanan, A. Kendall, and R. Cipolla, “SegNet: A Deep Convolutional Encoder-Decoder Architecture for Image Segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 39, no. 12, pp. 2481–2495, 2017.
- [24] H. Zhao, J. Shi, X. Qi, X. Wang, and J. Jia, “Pyramid Scene Parsing Network,” in *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 6230–6239, Honolulu, HI, July 2017.
- [25] T. Pohlen, A. Hermans, M. Mathias, and B. Leibe, “Full-Resolution Residual Networks for Semantic Segmentation in Street Scenes,” in *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3309–3318, Honolulu, HI, July 2017.
- [26] C. Brust, S. Sickert, M. Simon, E. Rodner, and J. Denzler, “Convolutional Patch Networks with Spatial Prior for Road Detection and Urban Scene Understanding,” in *Proceedings of the International Conference on Computer Vision Theory and Applications*, pp. 510–517, Berlin, Germany, March 2015.
- [27] R. Mohan, “Deep deconvolutional networks for scene parsing,” *Machine Learning*, 2014, <https://arxiv.org/abs/1411.4101>.
- [28] G. L. Oliveira, W. Burgard, and T. Brox, “Efficient deep models for monocular road segmentation,” in *Proceedings of the 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4885–4891, Daejeon, South Korea, October 2016.
- [29] M. Teichmann, M. Weber, M. Zoellner, R. Cipolla, and R. Urtasun, “MultiNet: Real-time Joint Semantic Reasoning for Autonomous Driving,” *Computer Vision and Pattern Recognition*, 2016, <https://arxiv.org/abs/1612.07695>.
- [30] A. Geiger, P. Lenz, and R. Urtasun, “Are we ready for autonomous driving? the KITTI vision benchmark suite,” in

- Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR '12)*, pp. 3354–3361, Providence, RI, USA, June 2012.
- [31] M. Cordts, M. Omran, S. Ramos et al., “The Cityscapes dataset for semantic urban scene understanding,” in *Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016*, pp. 3213–3223, USA, July 2016.
- [32] H. Kong, J.-Y. Audibert, and J. Ponce, “General road detection from a single image,” *IEEE Transactions on Image Processing*, vol. 19, no. 8, pp. 2211–2220, 2010.
- [33] J. C. McCall and M. M. Trivedi, “Video-based lane estimation and tracking for driver assistance: survey, system, and evaluation,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 7, no. 1, pp. 20–37, 2006.
- [34] S. Nedeveschi, R. Schmidt, T. Graf et al., “3D lane detection system based on stereovision,” in *Proceedings of the Proceedings - 7th International IEEE Conference on Intelligent Transportation Systems, ITSC 2004*, pp. 161–166, October 2004.
- [35] A. Wedel, H. Badino, C. Rabe, H. Loose, U. Franke, and D. Cremers, “B-spline modeling of road surfaces with an application to free-space estimation,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 4, pp. 572–583, 2009.
- [36] L. Bentabet, S. Jodouin, D. Ziou, and J. Vaillancourt, “Road vectors update using SAR imagery: a snake-based method,” *IEEE Transactions on Geoscience and Remote Sensing*, vol. 41, no. 8, pp. 1785–1803, 2003.
- [37] J. Beck and C. Stiller, “Non-parametric lane estimation in urban environments,” in *Proceedings of the 25th IEEE Intelligent Vehicles Symposium, IV 2014*, pp. 43–48, June 2014.
- [38] A. V. Nefian and G. R. Bradski, “Detection of drivable corridors for off-road autonomous navigation,” in *Proceedings of the 2006 IEEE International Conference on Image Processing, ICIP 2006*, pp. 3025–3028, October 2006.
- [39] F. Moosmann and C. Stiller, “Velodyne SLAM,” in *Proceedings of the 2011 IEEE Intelligent Vehicles Symposium, IV'11*, pp. 393–398, June 2011.
- [40] R. Valencia, E. H. Teniente, E. Trulls, and J. Andrade-Cetto, “3D mapping for urban service robots,” in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '09)*, pp. 3076–3081, IEEE, St. Louis, Mo, USA, October 2009.
- [41] P. Pfaff, R. Triebel, C. Stachniss, P. Lamon, W. Burgard, and R. Siegwart, “Towards mapping of cities,” in *Proceedings of the 2007 IEEE International Conference on Robotics and Automation, ICRA'07*, pp. 4807–4813, April 2007.
- [42] J. Gao, T. Zhang, X. Yang, and C. Xu, “Deep relative tracking,” *IEEE Transactions on Image Processing*, vol. 26, no. 4, pp. 1845–1858, 2017.
- [43] K. He, X. Zhang, S. Ren, and J. Sun, “Delving deep into rectifiers: surpassing human-level performance on imagenet classification,” in *Proceedings of the 15th IEEE International Conference on Computer Vision (ICCV '15)*, pp. 1026–1034, IEEE, Santiago, Chile, December 2015.
- [44] J. Fritsch, T. Kuhn, and A. Geiger, “A new performance measure and evaluation benchmark for road detection algorithms,” in *Proceedings of the 16th International IEEE Conference on Intelligent Transportation Systems (ITSC '13)*, pp. 1693–1700, IEEE, The Hague, The Netherlands, October 2013.
- [45] M. Sotelo, J. Nuevo, L. Bergasa, M. Ocana, I. Parra, and D. Fernandez, “Road Vehicle Recognition in Monocular Images,” in *Proceedings of the IEEE International Symposium on Industrial Electronics*, pp. 1471–1476, Dubrovnik, Croatia, June 2005.
- [46] S. Álvarez, M. A. Sotelo, I. Parra, D. F. Llorca, and M. Gavilán, “Vehicle and pedestrian detection in esafety applications,” in *Proceedings of the World Congress on Engineering and Computer Science*, 2009.
- [47] D. F. Llorca, V. Milanés, I. Parra Alonso et al., “Autonomous pedestrian collision avoidance using a fuzzy steering controller,” *IEEE Transactions on Intelligent Transportation Systems*, vol. 12, no. 2, pp. 390–401, 2011.
- [48] V. Milanés, D. F. Llorca, J. Villagrà et al., “Intelligent automatic overtaking system using vision for vehicle detection,” *Expert Systems with Applications*, vol. 39, no. 3, pp. 3362–3373, 2012.
- [49] Y. Jia, E. Shelhamer, J. Donahue et al., “Caffe: convolutional architecture for fast feature embedding,” in *Proceedings of the ACM Conference on Multimedia (MM '14)*, pp. 675–678, ACM, Orlando, Fla, USA, November 2014.

