



Bergische Universität Wuppertal

Fakultät für Mathematik und Naturwissenschaften

Institute of Mathematical Modelling, Analysis and Computational  
Mathematics (IMACM)

Preprint BUW-IMACM 20/43

K. Kowol, M. Rottmann, S. Bracke, H. Gottschalk

**YOdar: Uncertainty-based Sensor Fusion for  
Vehicle Detection with Camera and Radar  
Sensors**

October 12, 2020

<http://www.imacm.uni-wuppertal.de>

# YOdar: Uncertainty-based Sensor Fusion for Vehicle Detection with Camera and Radar Sensors

Kamil Kowol,<sup>1</sup> Matthias Rottmann,<sup>1</sup> Stefan Bracke,<sup>2</sup> Hanno Gottschalk<sup>1</sup>

<sup>1</sup>University of Wuppertal, School of Mathematics and Natural Sciences / IZMD / IMACM,

<sup>2</sup>University of Wuppertal, Chair of Reliability Engineering and Risk Analytics / IZMD  
{kowol, rottmann, bracke, hgottsch}@uni-wuppertal.de.

## Abstract

*In this work, we present an uncertainty-based method for sensor fusion with camera and radar data. The outputs of two neural networks, one processing camera and the other one radar data, are combined in an uncertainty aware manner. To this end, we gather the outputs and corresponding meta information for both networks. For each predicted object, the gathered information is post-processed by a gradient boosting method to produce a joint prediction of both networks. In our experiments we combine the YOLOv3 object detection network with a customized 1D radar segmentation network and evaluate our method on the nuScenes dataset. In particular we focus on night scenes, where the capability of object detection networks based on camera data is potentially handicapped. Our experiments show, that this approach of uncertainty aware fusion, which is also of very modular nature, significantly gains performance compared to single sensor baselines and is in range of specifically tailored deep learning based fusion approaches.*

## 1 Introduction

One of the biggest challenges for computer vision systems in automated driving is to recognize the cars environment appropriately in any given situation. The data of different sensors has to be interpreted correctly while operating with limited amount of computing resources. Besides dense traffic, different weather conditions like sun, heavy rain, fog and snow are challenging for state-of-the-art computer vision systems. Previous studies show that the use of more than one sensor, the so called sensor fusion, leads to an improvement in object detection accuracy, e.g. when combining camera and lidar sensors [1–5]. Up to now, datasets containing real street scenes using radar data and another sensor are rather the exception, although synthetic data with different sensors, such as camera and radar sensors, can be generated using simulators like CARLA [6] or LSGVL [7]. With the publication of the nuScenes dataset [8], the scientific community obtained access to

real street scenes recorded with different sensors including radar. In total there are 5 radar sensors distributed around the car that generate the data. Ever since, the number of published papers dealing with sensor fusion combining radar with other sensors with the help of the nuScenes dataset has increased, see e.g. [9–11]. We now briefly review these approaches.

The authors of [9] propose a convolutional neural network (CNN) for object detection named RVNet which is equipped with two input branches and two output branches. One input branch processes image data, the other one radar data. Similarly to YOLOv3 [12], the network utilizes two output branches to provide bounding box predictions, i.e., one branch is supposed to detect smaller obstacles, the other one larger obstacles. The authors conclude that radar features are useful for detecting on-road obstacles in a binary classification framework. On the other hand the features extracted from radar data seem not to be useful in a multiclass classification framework due to the sparsity of the data.

Another deep-learning-based radar and camera sensor fusion for object detection is the CRF-Net (CameraRadarFusionNet) [10], which automatically learns at which level the fusion of both sensor data is most beneficial for object detection. The CRF-Net uses a so-called BlackIn training strategy and combines a RetinaNet (VGG backbone), a custom-designed radar network and a Feature Pyramid Network (FPN) for classification and regression problems. The main branch is composed of five VGG-Blocks, every block receives pre-processed radar and image data for further processing which is forwarded to the FPN-Blocks. The network is tested on the nuScenes dataset and a self-build one. The authors provide evidence that the BackIn training strategy leverages the detection score of a state-of-the-art object detection network.

Furthermore, a fusion approach for lidar and radar is introduced in [11]. This approach is designed for multi-class object detection of pedestrian, cyclist, car and noise (empty region of interest) classes. To this end, lidar and radar data are first processed individually. The lidar branch detects objects and tracks them over

time. On the other hand, the radar branch provides the object classification, where three independent fast Fourier transforms (FFTs) are applied on the range-Doppler-angle spectrum. After time synchronisation the two branches are merged, resulting in regions of interest. These regions are fed to the CNN, based on the VGGNet architecture, which computes the classes probabilities. Since this approach works well for vehicle and noise classification but has problems with pedestrians and cyclist classes the network was improved by applying a tracking filter on top of the classifier. They used a Bayes filter which improved the classification performance for the two challenging classes.

In summary, the works presented [9,10] aim at simultaneously fusing and interpreting image and radar data within a CNN. In [11], lidar and radar data are first fused and afterwards a CNN processes the fused input. While these approaches are beneficial with respect to maximizing performance, they require additional fallback solutions in case that a sensor drops out. Also in contrast to other sensor fusion solutions, our approach preserves the option to use both networks redundantly. In this way, indications of only one of the networks can be used for scenario constructions that are alternative to the main scenario provided by the fusion approach.

It also seems inevitable that sensor fusion approaches require additional uncertainty measures to verify the quality of the developed methods and networks. A tool for semantic segmentation called MetaSeg that estimates prediction quality on segment level was introduced in [13] and extended in [14–16]. It learns to predict whether predicted components intersects with the ground truth or not, which can be viewed as meta classifying between two classes ( $IoU = 0$  and  $IoU > 0$ ). To this end, metrics are derived from the CNN’s output and pass them on to another meta-classifier. This work of false positive detection was extended in [17] where the number of overlooked objects was reduced by only paying with a few additional false positives. The overproduction of false positives is suppressed by MetaSeg. Following these approaches for uncertainty quantification, we use metrics from the output of two CNNs. We pass them through a gradient boosting classifier, which reduces the number of false positive predictions. In addition, by reducing the score threshold for object detection, we are able to improve over the performance of the respective single sensor networks.

In our tests, we utilize a YOLOv3 [12] as a state-of-the-art object detection network to process the camera data and complement this with a custom-designed CNN that performs a 1D binary segmentation which is supposed to detect obstacles. Further downstream of our computer vision pipeline we introduce a very general uncertainty-based fusion algorithm. Based on the predictions of both CNNs and their uncertainties as well as other geometrical meta-information, the fusion algorithm learns to provide a prediction by means of a structured dataset. In our experiments we demonstrate

for the case of street scenes recorded at night, that this approach significantly leverages the object detection accuracy. Furthermore, both networks only show moderate correlation which further supports our safety argument.

**Outline** The remainder of this work is organized as follows. In section 2 we briefly describe the characteristics of different sensors used for perception in automated driving. In section 3 we introduce our 1D-segmentation network for the detection of vehicles by radar data. We describe the preprocessing, the network architecture and the loss function we used for our method. This is followed by section 4, where the sensor fusion approach is presented and in section 5 our choice of metrics are introduced. In section 6 we discuss the numerical results. Finally, we present the conclusion and outlook in section 7.

Specifications		Camera	Radar	Lidar
Distance	Range	++	+++	+++
	Resolution	++	+++	++
Angle	Range	+++	++	+++
	Resolution	+++	+	++
Classification	Velocity Resolution	+	+++	++
	Object Categorization	+++	+	++
Environment	Night Time	+	+++	+++
	Rainy/Cloudy Weather	+	+++	++

+ = Good, ++ = Better, +++ = Best

Table 1: Overview of the advantages and disadvantages of the most common sensors for autonomous driving [18].

## 2 Sensor Characteristics

Today’s vehicles are equipped with sensors for recording driving dynamics, which register movements of the vehicle in three axes, as well as sensors for detecting the environment. The latter try to map the vehicle environment as accurately as possible to promote automated driving. This section briefly describes the characteristics of the three most used sensors for the perception of the environment for automated driving, i.e., camera, radar and lidar, including their advantages and disadvantages.

Camera sensors take two-dimensional images of light by electrical means. They are accurate in measuring edges, contour, texture and coloring. Furthermore they are easily integrated into the design of a modern vehicle. However, 3D localization from images is challenging and weather-related visual impairment can lead to higher uncertainties in object detections [19].

Radar sensors use radio waves to determine the range, angle and relative velocity of objects. Long-range radar sensors have a high range capability up to 200 – 250m [9, 20, 21] and are cheaper than Lidar sensors [22]. Compared to cameras, radar sensors are

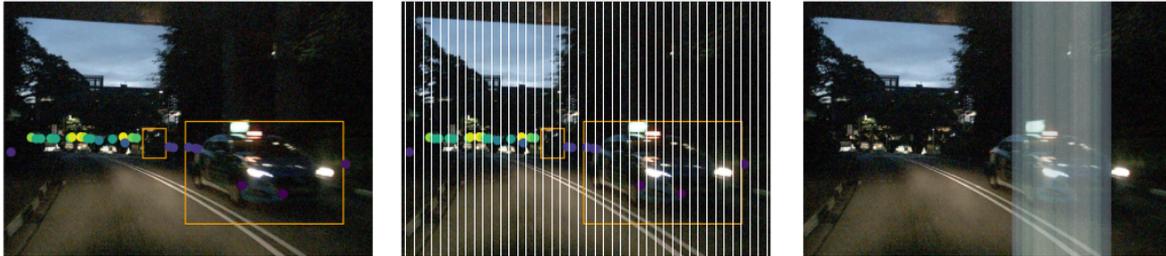


Figure 1: Preprocessing and prediction of the radar network. On the left we see the radar points projected to the front view camera image. The center image is divided into a certain number of slices  $N_s$  from which we generate the input matrix for the training of the CNN. The right hand image corresponds to the output of the radar network.

less affected by environmental conditions and pollution [21–23]. On the other hand, radar data is sparse and does not delineate the shape of the obstacles [9,23].

Lidar sensors use a light beam, emitted from a laser, to determine the distances and shapes of objects. Lidar sensors are highly accurate in 3D localization and surface measurements as well as a long-range view up to 300m [24]. They are expensive to buy and bad weather conditions like rain, fog or dust reduce the performance [21,22].

Each sensor has its advantages and disadvantages (summarized in table 1), so that a sensor fusion with at least two sensors makes sense in order to provide a better safety standard.

### 3 Object Detection via Radar

In this section we introduce the 1D segmentation network that we equip for detecting vehicles. First we explain the pre-processing method, then we describe the network architecture, the loss function and the network output.

#### 3.1 Preprocessing

In a global 3D coordinate system, the radar data is situated in a 2D horizontal plane. Hence, before training a neural network, we pre-process the radar data for two reasons. First of all, in order to simplify the fusion after processing each sensor with a neural network separately, we project the given radar data into the same 2D perspective as given by the front view camera. Secondly, one can observe that after this projection, the remaining section of the radar sensor modality is close to 1D. Figure 1 depicts radar points projected into the front view camera image. Darker colors indicate closer objects and brighter colors indicate more distant objects. Consequently, we build and train a neural network to perform a 1D segmentation.

To be more specific, we pre-process the ground truth for training the radar network as illustrated in figure 2. That is, we divide the given front view image into a

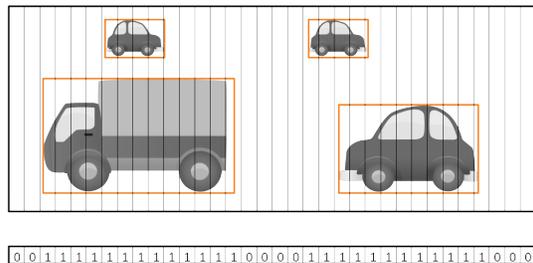


Figure 2: Ground truth vector for radar data. Every slice that overlaps with a bounding box in the front view camera image obtains the value 1, otherwise 0.

chosen number  $N_s$  of slices and generate an occupancy array of length  $N_s$ . The  $i$ th entry of this array is equal to 1 if there is a ground truth object intersecting with the  $i$ th slice of the image and 0 else. This ground truth construction defines the desired prediction for the radar network.

The radar data is pre-processed similarly to the ground truth as we aim at providing the neural network with an input tensor of size  $N_s \times N_t \times N_f$  where  $N_t$  denotes the number of considered time steps and  $N_f$  denotes the number of features. By assigning radar points to image slices we can drop the  $x$ -coordinate (which is implied by the array index up to an quantization error). For each slice  $i = 1, \dots, N_s$  which contains at least one radar point we store the following  $N_f$  features in the input matrix:  $y$ -coordinates (indicating the distance of the reflection point), the height coordinate with respect to the front view image that the radar point obtains by projection into the image plane as well as the relative lateral and the longitudinal velocity.

#### 3.2 Network Architecture

The network architecture for the radar network is based on a FCN-8-network [25] and is depicted in figure 3. The hidden layers consist of three convolution

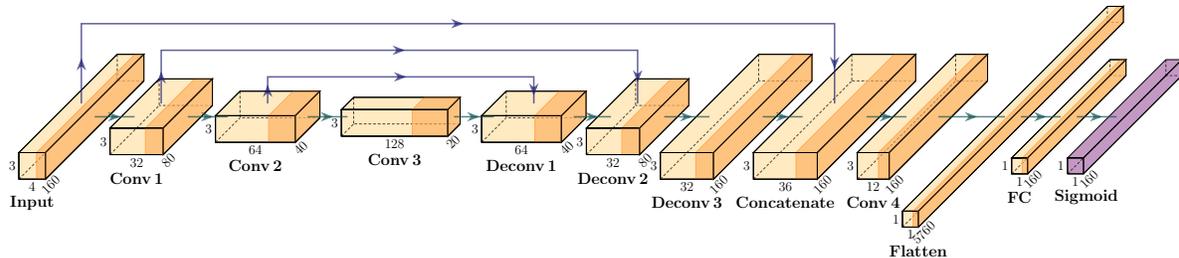


Figure 3: CNN architecture of our custom FCN-8 inspired radar network.

blocks, three deconvolution blocks followed by a concatenate layer, a fourth convolution block, a flatten layer and one fully-connected block. Finally, the network contains a sigmoid layer from which we get in  $[0, 1]$ . Each convolution and deconvolution block includes a (de-)convolution layer, a batch normalization and a leaky ReLU as activation function, respectively. The convolution blocks capture context information while losing spatial information whereas the deconvolution blocks restore these spatial information. Using bypasses, context information can be linked with spatial information. Furthermore, each fully connected block consists of a dense layer followed by a leaky ReLU activation function (except for the final layer where we use a sigmoid activation function).

### 3.3 Loss function

Let  $D = \{(r^{(i)}, t^{(i)}) : i = 1, \dots, n\}$  denote a dataset of tuples containing radar data  $r^{(i)} \in \mathbb{R}^{N_s \times N_t \times N_f}$  and ground truth  $t^{(i)} \in \{0, 1\}^{N_s}$ . The radar network  $g$  provides an array of estimated probabilities indicating whether a given slice  $s$  is occupied or not. We denote  $y^{(i)} = g(r^{(i)})$ . For training the neural network we use the binary cross-entropy for each array entry  $s = 1, \dots, N_s$ , i.e., for a single data sample  $(r^{(i)}, t^{(i)})$  we have

$$\ell(t_s^{(i)}, y_s^{(i)}) = -\alpha t_s^{(i)} \log(y_s^{(i)}) - (1 - t_s^{(i)}) \log(1 - y_s^{(i)}) \quad (1)$$

where  $\alpha$  is a tunable parameter. We introduced this parameter in order to account for the imbalance of zeros and ones in the ground truth. When training with stochastic batch gradient descent, the loss function is summed over all  $s = 1, \dots, N_s$  and then the mean is computed over all indices  $i$  in the batch.

### 3.4 Output

The predictions of the radar network result in a vector consisting of values in  $[0, 1]$  that estimate the probability of occupancy. Neighboring slices whose predicted probabilities are above a certain threshold  $T_g$  are recognized as one coherent object, also called slice bundle. Figure 4 shows for example an image with four slice bundles. The more a slice bundle fills in a bounding box, the higher the  $1D$ -IoU gets.



Figure 4: Image of a radar detection example with four predicted slice bundles

## 4 Object Detection via Sensor Fusion

After describing the object detection method using radar sensor data in the previous section, this section deals with the image detection method and the fusion of both methods. Various object detection networks have been developed in recent years, whereby the YOLOv3 network has become a very good choice when fast and accurate real-time detection is desired [12, 26]. It has been observed that YOLOv3 works very well under good weather and visibility conditions but has problems with object recognition in bad visibility like hazy weather [27, 28] or darkness, like Xiao et al. have investigated for object detection with RFB-Net [29]. In our experiments, we focus on object detection by camera and radar at night. To this end, we first use each method separately in order to connect both outputs with gradient boosting [30], see figure 5. Similarly to [13] we derive metrics from each CNN output and pass them through a gradient boosting classifier to increase the number of detected vehicles. The data and metrics used for the classifier are explained in section 5. The YOLOv3 threshold  $T_f$  for vehicle detection is set to a low value such that we get a higher number of bounding box predictions. From the many predicted bounding boxes, gradient boosting select those boxes that are likely to contain an object according to the output of both networks. On the one hand, the radar sensor should detect vehicles not recognized by

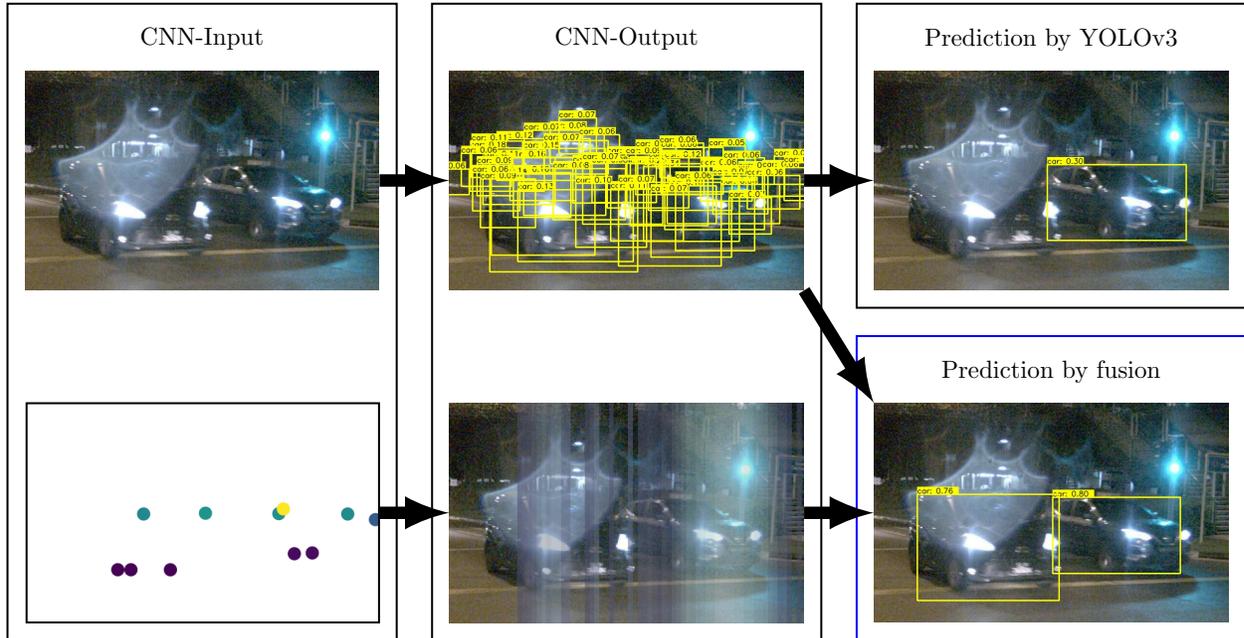


Figure 5: Illustration of our YODar method. In the top branch, the YOLOv3 produces many candidate bounding boxes, but after score thresholding and non maximum suppression, only one of both cars is detected. By lowering the object detection threshold and fusing the obtained boxes with the radar prediction, both cars are detected by YODar as shown in the bottom right panel.

the YOLOv3 network. On the other hand, gradient boosting should support the decision making process by additional information in case the YOLOv3 network is uncertain.

## 5 Fusion Metrics and Methods

The fusion method that we introduce in this section is of generic nature. Therefore, we denote by  $f$  the an arbitrary camera network and by  $g$  a 1D segmentation radar network. Given an input scene  $(x, r)$ , we obtain two network outputs, one for the image input  $x$ , one for the radar input  $r$ . Each prediction obtained by  $f$  consists of a set  $B = \{b_1, \dots, b_k\}$  containing  $k$  boxes where  $k$  depends on  $x$ . Each  $b_i$  is a tuple that contains an objectness score value  $z_i$ , a probability  $f(v|x, b_i)$  that the box  $b_i$  contains a vehicle  $v$ , a center point with its x-coordinate  $c_i^x \in \mathbb{R}$  and y-coordinate  $c_i^y \in \mathbb{R}$  as well as width  $w_i \in \mathbb{R}$  and height  $h_i \in \mathbb{R}$  of the box, i.e.,

$$b_i = (z_i, f(v|x, b_i), c_i^x, c_i^y, w_i, h_i). \quad (2)$$

For the radar network  $g$  we obtain a 1D output of probabilities,  $g_s(v|r)$  for each of the slices  $s = 1, \dots, n$ , that this slice  $s$  belongs to a vehicle  $v$ , recall figure 2. As depicted in figure 1 we identify slices  $s$  and bounding boxes  $b_i$ . In order to aggregate slices  $s$  from the radar network over bounding boxes  $b_i$  obtained by the camera network, let  $S_i$  denote the set of all slices  $s$  that

intersect with the box  $b_i$ . We denote by

$$\mu_i = \frac{1}{|S_i|} \sum_{s \in S_i} g_s(v|r) \quad (3)$$

the average probability of observing a vehicle in the box  $b_i$  according to the radar network’s probabilities. The standard deviation corresponding to equation (3) is termed  $\sigma_i$ . As a set of metrics, by which we compute a fused prediction, we consider

$$M_i(x, r) = (b_i, A_i, \mu_i, \sigma_i) \quad (4)$$

where  $A_i = w_i \cdot h_i$  denotes the size of the box  $b_i$ . In summary, we use these nine metrics  $M_i(x)$  for all scenes  $x$  and boxes  $b_i$  that are visible with respect to the front view camera. To perform the fusion of the camera based network prediction and the radar based network prediction we proceed in two steps. First we compute the ground truth which states for each box predicted by the camera network whether it is a true positive (TP) or a false positive (FP). Afterwards we train a model to discriminate by means of  $M_i$  whether  $b_i$  is a TP or an FP.

More precisely, for the sake of computing ground truth, we define TP and FP in the given context as follows: For a predicted box  $b_i$  and a ground truth box  $a$ , which has the biggest intersection  $|a \cap b_i|$  of all ground truth boxes of the same class, the intersection over union is defined as follows:

$$IoU(b_i) = \max_a \frac{|a \cap b_i|}{|a \cup b_i|}. \quad (5)$$

Oftentimes we omit the argument  $b_i$  if it is clear from the context. Given a chosen threshold  $T \in [0, 1)$  we define that  $b_i$  is a TP if  $IoU(b_i) > T$  and an FP if  $IoU(b_i) \leq T$ . For the sake of completeness, we define that a false negative is a ground truth box  $a$  that fulfills  $\max_{b_i} \frac{|a \cap b_i|}{|a \cup b_i|} \leq T$ . Note that in the latter definition, the ground truth element is fixed while the left hand side of this expression is maximized over all predicted boxes.

After computing the ground truths, i.e., whether a box  $b_i$  predicted by the camera network yields a TP or an FP, we train a model. The gathered metrics  $M_i$  yield a structured dataset where the columns are given by the different metrics and the rows are given by all predicted boxes  $b_i$  for each input scene  $x$ . By means of this dataset and the corresponding TP/FP annotation, we train a classifier to predict whether a box predicted by the camera network is a TP or an FP.

The  $IoU$  can be calculated in different dimensions  $D = 1, 2$ . In this work the  $1D-IoU$  is used for the radar network. As soon as a low threshold  $T_g$  has been reached, a predicted object is considered as TP, otherwise as FP. The  $2D-IoU$  is used for the YOLOv3 network, analogously we speak of TP and FP according to a threshold  $T_f$ . An illustration is given in figure 6.

The mean average precision ( $mAP$ ) is a popular metric used to measure the performance of models. The  $mAP$  is calculated by taking the average precision (area under precision as a function of recall, a.k.a. precision recall curve) over one class.

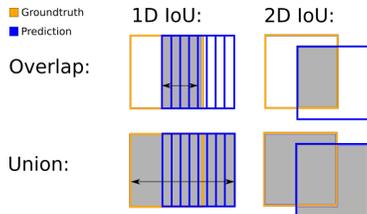


Figure 6:  $IoU$  calculation for 1- and 2- $D$  bounding boxes

## 6 Numerical Experiments

As explained in detail in the previous section, we use a custom FCN-8-like network for processing the radar data from the nuScenes dataset [8]. It contains urban driving situations in Boston and Singapore. The dataset has a high variability of scenes, i.e., different locations, weather conditions, daytime, recorded with left- or right-hand traffic. In total, the dataset contains 1,000 scenes, each of 20 seconds duration and each frame is fully annotated with 3D bounding boxes. The vehicle used, a Renault Zoe, was set up with 6 cameras at 12Hz capture frequency, 5 long-range radar sensors (FMCW) with 13Hz capture frequency, 1 spinning lidar with 20Hz capture frequency, 1 global positioning system module (GPS) and 1 inertial measurement unit

(IMU). Each scene is divided into several time frames for which each sensor provides a suitable signal.

splitting	number of images/frames		night images/frames [%]	
	YOLOv3	Radar	YOLOv3	Radar
train	10,000	29,853	7.08	9.04
val	3,289	3,289	8.57	8.57
test	1,006	1,006	100	100

Table 2: Data used for training, validation and testing.

For our experiments, the YOLOv3 network was pre-trained with day images from the COCO dataset [31] and afterwards with 249 randomly selected scenes from the nuScenes dataset containing 10,000 images with different weather conditions and times of day. Furthermore we have trained a CNN with radar data from the nuScenes dataset to detect vehicles. 743 of the scenes (29,853 frames) were used as training data, 82 scenes (3,289 frames) as validation data and 25 scenes (1,006 frames) as test data, see table 2. For the test set, we consider exclusively all frames recorded at night that are not part of the training data in order to create a perception-wise challenging test situation. For the training and validation sets we used a natural split of day and night scenes pre-defined by the frequencies in the nuScenes dataset. Due to the resolution of the radar data, we focus on the category vehicle in our evaluation. This includes the semantic categories car, bus, truck, bicycle, motorcycle and construction vehicle.

### 6.1 Training

As input, the radar network obtains a tensor with the dimensions  $160 \times 3 \times 4$  that contains for each of the 160 considered slices the current frame (i.e., the current time step) and two previous frames. Each frame contains four features, i.e., x-, and y-coordinates, lateral and longitudinal velocity. From the radar data we removed all ground truth bounding boxes that do not contain any radar points with valid velocity vectors.

The radar network is implemented Keras [32] with TensorFlow [33] backend. Training on one NVIDIA Quattro GPU P6000 takes 229 seconds training time. The network structure is shown in figure 3 and the training parameters are shown in table 3. We have trained the neural network three times with three different learning rates i.e., the first 20 epochs with a learning rate of  $10^{-3}$ , 10 epochs with  $10^{-4}$  and 10 epochs with  $10^{-5}$ . The networks output vector has the same dimension ( $160 \times 1$ ) like the ground truth vector, where each entry contains a probability value, whether there is a vehicle in the respective area or not. If the probability value of a single slice is equal or higher than the threshold  $T_g = 0.5$ , then the network predicts a vehicle. The higher the probability value, the brighter the slice is displayed in figure 1.

The YOLOv3 network was trained with 10,000 images consisting of different scenes. Therefore we converted the 3D bounding boxes in the nuScenes dataset



Figure 9 shows the relative amount of objects recognized by the YOLOv3 network. It shows that mainly objects closer to the ego car and straight ahead are recognized, while objects farther away or located on the very left or very right end of the image remain often unrecognized. Figure 10 states in absolute numbers how many additional objects are recognized in each particular area when using YOdor instead of YOLO. The increase is clear and also mostly in the relevant areas close to the ego car and straight ahead. This is in line with the idea of focusing with the radar on objects in motion (by considering objects that carry velocities). These results show that an uncertainty-based fusion approach like YOdor is indeed able to increase the performance significantly. This finding is also confirmed by the  $mAP$  and accuracy values stated in table 5. While YOLO3v achieves 31.36%  $mAP$ , Yodor achieves 39.40% which is also close to state of the art deep learning based fusion results for the nuScenes dataset with the natural split of day and night scenes as reported in [10].

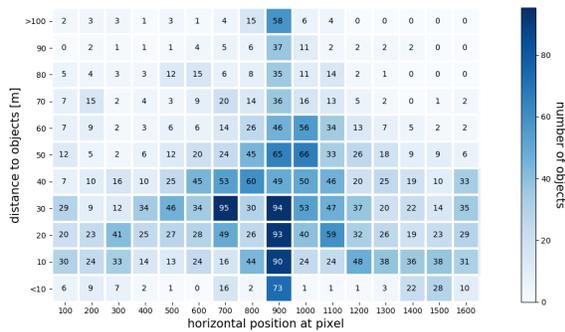


Figure 8: Ground truth heat map displaying the spatial distribution of the test data, broken down by distance (vertical axis) and the pixel intervals of width 100 of the front view input image with a total width of 1600 pixels (horizontal axis). More precisely, each intervals denoted by  $i$  on the horizontal axis represents the pixels (row, column) with  $column \in [i - 99, i]$ .

	Radar	YOLOv3	YOdor
Accuracy [%]	14.41	33.90	43.08
mAP [%]	7.93	31.36	39.40

Table 5: Accuracies and  $mAP$  scores of all three networks.

## 7 Conclusion and Outlook

In this paper we have introduced the method YOdor which detects vehicles with camera and radar sensors. With this uncertainty-based sensor fusion approach the camera and radar data are first processed individually. Each branch detects objects, the camera

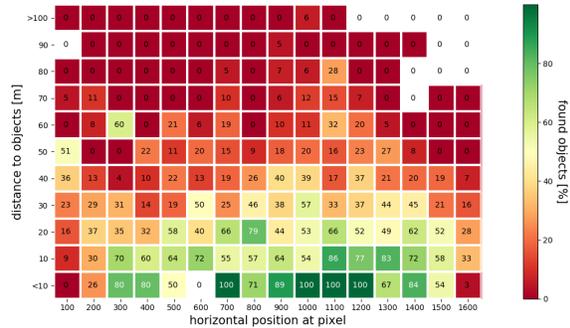


Figure 9: Relative amount of objects recognized by the YOLOv3 network evaluated on the test data. The underlying geometry is the same as in figure 8.

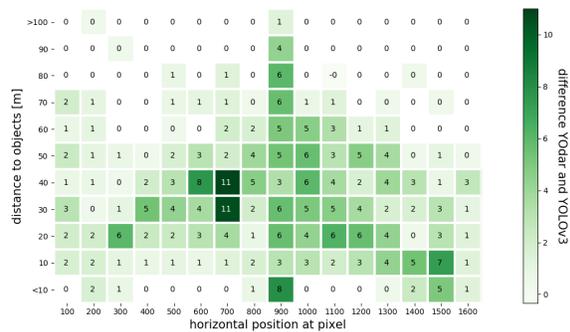


Figure 10: Number of objects recognized by YOdor minus the number of objects recognized YOLOv3. The underlying geometry is the same as in figure 8.

branch uses a YOLOv3 network trained with day and night scenes and the radar branch uses a custom-based radar network. The outputs of every branch are aggregated and then passed through a post processing classifier that again learns the same vehicle detection task. Compared to the YOLOv3 network, the YOdor fusion method detects at night a significant additional amount of vehicles in total. While YOLO3v achieves 31.36%  $mAP$ , YOdor achieves 39.40%  $mAP$  which is also close to state of the art deep learning based fusion results for the nuScenes dataset with the natural split of day and night scenes. In future work, additional sensors will be added for training and evaluation, since this approach uses only the front camera and the front radar sensor. Furthermore, we plan to optimize the radar network such that more objects can be detected by YOdor. Moreover, we plan to extend this approach also the detection of pedestrians as more dense radar data becomes available.

## Acknowledgements

K.K. acknowledges financial support through the research consortium bergisch.smart.mobility funded by



the ministry for economy, innovation, digitalization and energy (MWIDE) of the state North Rhine Westphalia under the grant-no. DMR-1-2.

## References

- [1] Varuna De Silva, Jamie Roche, and Ahmet Konoz. Robust Fusion of LiDAR and Wide-Angle Camera Data for Autonomous Mobile Robots. oct 2017.
- [2] Tai-En Wu, Chia-Chi Tsai, and Jiun-In Guo. LiDAR/camera sensor fusion technology for pedestrian detection. In *2017 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 1675–1678, 2017.
- [3] Ziyi Liu, Siyu Yu, Xiao Wang, and Nanning Zheng. Detecting Drivable Area for Self-driving Cars: An Unsupervised Approach. may 2017.
- [4] Mikkel Kragh and James Underwood. Multi-Modal Obstacle Detection in Unstructured Environments with Conditional Random Fields. jun 2017.
- [5] Shuo Gu, Yigong Zhang, Jinhui Tang, Jian Yang, and Hui Kong. Road Detection through CRF based LiDAR-Camera Fusion. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 3832–3838, 2019.
- [6] Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. Carla: An open urban driving simulator. *arXiv preprint arXiv:1711.03938*, 2017.
- [7] Guodong Rong, Byung Hyun Shin, Hadi Tabatabaee, Qiang Lu, Steve Lemke, Mārtiņš Možeiko, Eric Boise, Geehoon Uhm, Mark Gerow, Shalin Mehta, et al. Lgsvl simulator: A high fidelity simulator for autonomous driving. *arXiv preprint arXiv:2005.03778*, 2020.
- [8] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for autonomous driving. mar 2019.
- [9] Vijay John and Seiichi Mita. RVNet: Deep Sensor Fusion of Monocular Camera and Radar for Image-Based Obstacle Detection in Challenging Environments. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 11854 LNCS, pages 351–364, 2019.
- [10] Felix Nobis, Maximilian Geisslinger, Markus Weber, Johannes Betz, and Markus Lienkamp. A Deep Learning-based Radar and Camera Sensor Fusion Architecture for Object Detection. In *2019 Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, pages 1–7, Bonn, Germany, 2019.
- [11] Rodrigo Pérez, Falk Schubert, R H Rasshofer, and Erwin Biebl. Deep Learning Radar Object Detection and Classification for Urban Automotive Scenarios. 2019.
- [12] Joseph Redmon and Ali Farhadi. YOLOv3: An Incremental Improvement. apr 2018.
- [13] Matthias Rottmann, Pascal Colling, Thomas-Paul Hack, Fabian Hüger, Peter Schlicht, and Hanno Gottschalk. Prediction Error Meta Classification in Semantic Segmentation: Detection via Aggregated Dispersion Measures of Softmax Probabilities. *CoRR*, abs/1811.0, 2018.
- [14] Kira Maag, Matthias Rottmann, and Hanno Gottschalk. Time-Dynamic Estimates of the Reliability of Deep Semantic Segmentation Networks. *CoRR*, abs/1911.0, 2019.
- [15] Matthias Rottmann and Marius Schubert. Uncertainty measures and prediction quality rating for the semantic segmentation of nested multi resolution street scene images. *CoRR*, abs/1904.04516, 2019.
- [16] Marius Schubert, Karsten Kahl, and Matthias Rottmann. Metadetect: Uncertainty quantification and prediction quality estimates for object detection, 2020.
- [17] Robin Chan, Matthias Rottmann, Fabian Hüger, Peter Schlicht, and Hanno Gottschalk. MetaFusion: Controlled False-Negative Reduction of Minority Classes in Semantic Segmentation. *CoRR*, abs/1912.0, 2019.
- [18] Jeff Phillips. Achieving Safe Autonomous Driving. [Online]. Available: <https://www.ni.com/de-de/perspectives/imminent-trade-offs-for-achieving-safe-autonomous-driving.html>, 2020.

- [19] Zhangjing Wang, Yu Wu, and Qingqing Niu. Multi-Sensor Fusion in Automated Driving: A Survey, 2020.
- [20] Martin Schneider. Automotive Radar – Status and Trends. 2005.
- [21] Fabian de Ponte Müller. Survey on Ranging Sensors and Cooperative Techniques for Relative Positioning of Vehicles. *Sensors*, 17, 2017.
- [22] Ryan Aldrich and Thanuka Wickramaratne. Low-cost radar for object tracking in autonomous driving: A data-fusion approach. In *2018 IEEE 87th Vehicular Technology Conference (VTC Spring)*, pages 1–5, 2018.
- [23] Paul Fritsche, Simon Kueppers, Gunnar Briese, and Bernardo Wagner. Radar and LiDAR Sensorfusion in Low Visibility Environments. pages 30–36, 2016.
- [24] Ashish Pidurkar, Ranjit Sadakale, and A. K. Prakash. Monocular camera based computer vision system for cost effective autonomous vehicle. In *2019 10th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*, pages 1–5, July 2019.
- [25] Jonathan Long, Evan Shelhamer, and Trevor Darrell. Fully Convolutional Networks for Semantic Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [26] Bilel Benjdira, Taha Khursheed, Anis Koubaa, Adel Ammar, and Kais Ouni. Car Detection using Unmanned Aerial Vehicles: Comparison between Faster R-CNN and YOLOv3. In *2019 1st International Conference on Unmanned Vehicle Systems-Oman (UVS)*, pages 1–6, feb 2019.
- [27] Paulius Tumas, Adam Nowosielski, and Arturas Serackis. Pedestrian Detection in Severe Weather Conditions. *IEEE Access*, 8:62775–62784, 2020.
- [28] Guofa Li, Yifan Yang, and Xingda Qu. Deep Learning Approaches on Pedestrian Detection in Hazy Weather. *IEEE Transactions on Industrial Electronics*, 67(10):8889–8899, 2020.
- [29] Yuxuan Xiao, Aiwen Jiang, Jihua Ye, and Mingwen Wang. Making of Night Vision: Object Detection Under Low-Illumination. *IEEE Access*, 8:123075–123086, 2020.
- [30] Jerome H Friedman. Stochastic gradient boosting. *Computational Statistics & Data Analysis*, 38(4):367–378, 2002.
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft COCO: Common Objects in Context. may 2014.
- [32] François Chollet. Keras - Deep learning library [Online]. Available: <https://github.com/keras-team/keras>, 2015.
- [33] Martin Abadi, Ashish Agarwal, Paul Barham, et al. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, 2015.