

Warping of Radar Data into Camera Image for Cross-Modal Supervision in Automotive Applications

Christopher Grimm, *Graduate Student Member, IEEE*, Tai Fei, *Senior Member, IEEE*, Ernst Warsitz, Ridha Farhoud, Tobias Breddermann, Reinhold Haeb-Umbach, *Fellow, IEEE*

Abstract—We present an approach to automatically generate semantic labels for real recordings of automotive range-Doppler (RD) radar spectra. Such labels are required when training a neural network for object recognition from radar data. The automatic labeling approach rests on the simultaneous recording of camera and lidar data in addition to the radar spectrum. By warping radar spectra into the camera image, state-of-the-art object recognition algorithms can be applied to label relevant objects, such as cars, in the camera image. The warping operation is designed to be fully differentiable, which allows backpropagating the gradient computed on the camera image through the warping operation to the neural network operating on the radar data. As the warping operation relies on accurate scene flow estimation, we further propose a novel scene flow estimation algorithm which exploits information from camera, lidar and radar sensors. The proposed scene flow estimation approach is compared against a state-of-the-art scene flow algorithm, and it outperforms it by approximately 30% w.r.t. mean average error. The feasibility of the overall framework for automatic label generation for RD spectra is verified by evaluating the performance of neural networks trained with the proposed framework for Direction-of-Arrival estimation.

Index Terms—Automotive radar, neural network, lidar, virtual testing, direction-of-arrival, cross-modal supervision

I. INTRODUCTION

In the past decade, deep neural network (DNN) based algorithms have proven excellent performance for a wide range of applications [1]–[3]. Those impressive results rely, however, to a great degree on the availability of large labeled databases. In case of range-doppler (RD) spectrum automotive frequency modulated continuous wave (FMCW) radar signal processing this requirement poses a significant practical problem. Not only are those large-scale databases not available as of today, they would be furthermore very difficult to build, because, unlike, e.g., a camera image, an RD plot is very difficult for a human to interpret and thus to label properly. Another challenge comes from high variations in radar hardware and parameterization, resulting in very specific radar measurement capabilities. Therefore, researchers often require to generate their own dataset, rather than using open public datasets [4].

Some recent works have dealt with automatic RD spectrum labeling by lidar and/or camera [5]–[8]. However, these labels are generally restricted to object level for certain classes as

pedestrian and vehicles, which limits its use to “high-level” radar classification tasks, such as moving object detection. We, on the contrary, are here interested in applying DNNs to low-level radar classification tasks, such as radar target detection, stationary target detection and direction-of-arrivals (DoAs) estimation.

In this paper, we address this problem by automatically generating high-quality and dense labels from a reference system consisting of camera, lidar and state-of-the-art visual object classification algorithms. Neural object classification algorithms are applied to the camera image to label relevant objects. Those labels are then mapped to locations in an RD spectrum by warping the RD spectrum into the camera image, and then serve as labels for a neural network operating on radar spectrum data. The warping task, however, is challenging, as radar and camera measure quite different physical phenomena. Furthermore, the field of view of camera and radar may only partially overlap, and their coordinate systems are not aligned.

We augment the measurement space of the reference camera system to cover the measurement space of the radar system by estimating the 3D-velocity of each pixel in the camera image. To this end, we extend a previously proposed scene flow estimation algorithm to utilize information gained from the radar spectrum. This significantly improves scene flow estimation, and in turn the precision of the warping operation.

As the warping operation is designed to be fully differentiable, the gradient computed on the camera image can be backpropagated through the warping module to the neural network operating on the RD spectrum, allowing for supervised training of the network with real world radar data as input. Fig. 1 illustrates the proposed framework. It can be viewed as a universal framework for cross-modal supervision of neural networks operating on FMCW radar data.

We demonstrate the effectiveness of the proposed framework by applying it to radar based DoA estimation and present both qualitative and quantitative results.

The main contributions of this paper can be summarized as follows

- An improved camera scene flow estimation algorithm by incorporating RD spectrum information
- Dense label generation for RD spectrum of FMCW radar data by utilizing labels from camera vision models
- A method for warping radar data into camera images
- A framework to enable supervised training of DNNs on real-world data for multiple radar applications.

The remainder of this paper is organized as follows. Sec. II summarizes related work in the field of automotive dataset generation. Sec. III introduces the utilized sensors and the

Christopher Grimm, Tai Fei, Ernst Warsitz, Ridha Farhoud and Tobias Breddermann are employees of HELLA GmbH & Co. KGaA, 59552 Lippstadt, Germany e-mail: {christopher.grimm, tai.fei, ernst.warsitz, ridha.farhoud, tobias.breddermann}@forvia.com

Reinhold Haeb-Umbach is with the Department of Communication Engineering, Paderborn University, 33098 Paderborn, Germany e-mail: haeb@nt.uni-paderborn.de

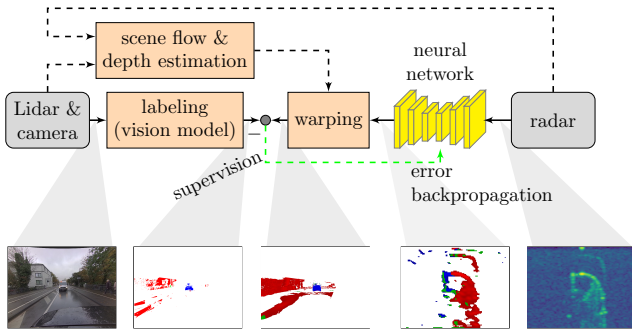


Figure 1: **Overview of the proposed approach:** Lidar and camera provide sensor readings to the vision model (typically a DNN) which automatically generates labels in the camera image. At the same time, the radar’s RD spectrum is processed by a neural network and warped from RD spectrum into the camera image domain. The residual between the generated label and the warped neural network prediction is calculated and then backpropagated through the warping operation to the network operating on the radar data for parameter adjustment (green arrow). The warping operation requires dense scene flow and depth estimation, which is obtained by lidar, camera and radar data. The lower row shows example images of a semantic segmentation application. From left to right: RGB image, reference semantic mask (only relevant pixels shown), predicted semantic mask warped into camera image, predicted semantic mask in RD-spectrum, RD-map.

basic sensor calibration procedure. Sec. IV describes the novel scene flow estimation algorithm which is supported by radar measurements. The scene flow estimates are used in the warping algorithm given in Sec. V. In Sec. VI, camera labels, training objectives for the neural networks (NNs) and network architectures are given as well as techniques on how to make the training more robust against reference label errors. In Sec. VII the NN based DoA estimator’s accuracy is evaluated. The conclusions are drawn in Sec. VIII.

II. RELATED WORK

In 2012, Geiger *et al.* [9] introduced the famous *KITTI* dataset. The development and publication of this dataset was motivated by the intention to foster the development of computer vision for autonomous driving. At that time, algorithms were mostly evaluated under laboratory conditions and performed poorly when applied to field data [10]. Recognizing the need for more realistic datasets, the authors collected camera, lidar and GPS data from approx. 40 km real-world driving scenarios around Karlsruhe, Germany. Starting with benchmarks for stereo vision, optical flow, visual odometry and 3D object detection, in which ground truth was provided either by sensor fusion or, in case of object detection, by human labelers, the range of applications and benchmark systems has been extended since then to depth completion of sparse depth measurements, semantic segmentation and scene flow estimation [11]. However, automotive radar has not been included in the dataset, and, accordingly, *KITTI* is not

suitable for the development of automotive radar applications considered here.

In 2019, Caesar *et al.* [12] introduced the *nuScenes* multi-modal dataset, which incorporates sensor data from multiple cameras, lidar and radars. The ground truth was mostly provided by professional human labelers, for instance in the form of 3D object bounding boxes. The dataset currently contains approx. 242 km of driving data acquired in Boston, USA, and in Singapore. The radar data is given as a point cloud consisting of radar detections, in which each point has an annotated position and velocity. From the perspective of radar signal processing, the radar data in this dataset is already above spectrum level, while raw spectrum level data is missing. Accordingly, it is barely suitable for the development of signal processing algorithms for RD spectrum data.

Similar concerns apply to the dataset provided in [13]. Other automotive datasets can be found e.g., in [14], [15] which, however, are concerned with creating large and densely annotated camera data only.

Recently, more and more synthetic datasets have been developed. The datasets of [16]–[18] contain artificially generated camera images as well as detailed ground truth from simulated environments. Bühren *et al.* [19] synthesized radar targets from traffic simulations. Again, the simulation was carried out on above spectrum level radar data, and thus is inappropriate for the development of signal processing algorithms for RD spectra.

In [20] the authors showcased deep-learning based object detection trained on simulated RD spectra via raytracing. Different driving scenarios were simulated and a trained object detector was able to achieve decent performance on the test data. Note, however, that the test data originated from the same raytracing tool as the training data, and thus performance on real world data remains unknown.

In [21] the authors propose a Differential-GPS (DGPS) based reference system for radar target evaluation. While this provides very precise measurements, the equipment is very expensive and, thus, the acquisition of large and complex datasets is practically impossible as only objects are covered, equipped with such a system. Reference RD spectra can also be generated from lidar data, as proposed in [5]. Vehicles are detected in the lidars point cloud and tracked over time. However, the reference RD spectra are very sparse, since only a reference for vehicles was given, and since vehicles are modeled as single point reflectors.

The idea of cross-modal supervision was followed in a number of works. Zhao *et al.* proposed a framework to leverage camera based label predictions for other sensors [22]. In 2019, [6] and [7] trained DNNs for object detection operating on range-azimuth-Doppler spectrum in real world driving scenarios. The radar data is mapped from polar coordinates to the Cartesian coordinate system of the reference lidar sensor, in which the object labels are presented to the network. However, the Cartesian coordinate transformation relies on accurate azimuth resolution of the radar itself, which can hardly be met by today’s series production radar sensors for automatic driver assistance system (ADAS) applications. Additionally, the framework was designed to only generate

labels for vehicles on highway environments and thus only provides a sparse label subset of the entire scene.

In [23], [24] the authors labeled radar target clusters by fusing camera images, lidar pings and radar targets with the help of state-of-the-art perception models and achieved 72 % classification accuracy compared to human labeling with a significant reduction in labeling time. They achieved 0.48 m average translational error on the *nuScenes* object detection benchmark. But, again, this task is above spectrum level, and, therefore, not applicable here.

One of the very few sets of annotated radar spectrum data is [8], in which object annotations for vehicles, cyclists and pedestrians are given in range-azimuth-Doppler space. The annotations are generated by assigning object classes from objects estimated by camera to matching objects estimated by radar. Although annotations are given for spectrum data, the accuracy of the automatically generated annotations highly depends on the accuracy of the radar based object detections. Furthermore, the matching becomes significantly harder in dense urban environments with multiple objects per scene, potentially resulting in a high number of mismatches and therefore introducing label noise. The framework does not provide labels for stationary environment and thus provides only a subset of all labels in a complex environment.

In 2021, a dataset containing Range-Azimuth-Doppler spectra from radar was presented in [4]. Object level labels are provided for the classes: person, bicycle, car, motorcycle, bus and truck. To auto-generate the labels, radar detection pointclouds are extracted, by applying CFAR and beamforming to the radar spectra. The pointcloud is clustered into radar objects, which are then associated with class labels from instance segmentation on stereo image pairs. This approach is very similar to the approach presented in [8] and therefore we see similar restrictions: Firstly, the label quality highly relies on the quality of the radar pointcloud. The second issue is that the labels provided here are limited to dynamic road users only. Adding labels for stationary objects would likely compromise the association accuracy between radar and camera objects.

Based on our research, we ranked the best suitable frameworks for radar RD spectrum labeling from literature against our proposed framework in Tab. I. We combined the approaches from [6] and [7] as well as [8] and [4] as they present very similar approaches for auto-annotation. As prior publications focus on providing annotations for the application of semantic segmentation in radar spectra, we do see significant difference to our dataset, which is also able to provide dense labels for DoA estimation, as proven in this paper and other applications.

The reader will find information about classical DoA radar applications e.g. phase-monopulse or Bartlett beamforming in e.g. [25], [26].

In 2012 Sit *et al.* [27] proposed a NN based DoA estimator. The NN was trained 859 samples referring two scatters moving in azimuth. The NN was able to achieve better DoA separability than classical MUSIC algorithm whilst offering significant reduced processing time.

More recent NN based approaches were presented in [28]–[30]. The authors installed a radar-under-test on an azimuth

positioner in an anechoic chamber and placed a corner reflector 1.5 m in front of the sensor. Measurement data from the radar sensor was acquired at different azimuth orientations. As the target position was known, it served as ground truth label for multi-layer perceptron (MLP) training. The MLP was trained to estimate DoA for up to two targets, and achieved similar performance as classical algorithms for instance Deterministic Maximum Likelihood, while being significantly faster in inference. As the data collected in anechoic chamber might differ from data obtained in real world traffic scenarios, we will focus on training and evaluating NNs in the latter scenarios.

For the interested reader, a more elaborated overview can be found in [31].

III. SENSOR CALIBRATION

In the following we assume the presence of the following sensors, which reflects our measurement setup described in Section VII-A further below:

- D : A Differential-GPS with inertial navigation system (DGPS-INS) which delivers high-precision position information
- C_1, C_2 : An array of two cameras, one pointing to the rear, and the other to the right.
- L_1, L_2 : Two lidar scanners
- R : the RD spectrum radar unit.

The spatial information they provide is given in their local coordinate systems (COOSs). Since the sensors are not mounted at the same position, the origins of their local COOSs differ. For jointly processing the data, the COOSs of the sensors and of the (ego) vehicle, denoted by E , have to be aligned first.

A transformation of a position \mathbf{x}_j in the COOS of sensor j to the position \mathbf{x}_i in the COOS of sensor i , where $i, j \in \{D, C_1, C_2, L_1, L_2, R, E\}$, is achieved by affine transformation with the rotation matrix ${}^i\mathbf{R}_j \in \mathbb{R}^{3 \times 3}$ and the translation vector ${}^i\mathbf{t}_j \in \mathbb{R}^{3 \times 1}$ as follows:

$$\mathbf{x}_i = {}^i\mathbf{R}_j \mathbf{x}_j + {}^i\mathbf{t}_j. \quad (1)$$

The parameters $({}^E\mathbf{R}_{L_1}, {}^E\mathbf{t}_{L_1})$ and $({}^E\mathbf{R}_{L_2}, {}^E\mathbf{t}_{L_2})$ from the COOS of the lidar sensors L_1 and L_2 to the COOS of the ego vehicle have been estimated via [32]. The tuples $({}^{C_1}\mathbf{R}_{L_1}, {}^{C_1}\mathbf{t}_{L_1})$ and $({}^{C_2}\mathbf{R}_{L_1}, {}^{C_2}\mathbf{t}_{L_1})$ have been extracted from sensor mounting parameters, because the cameras were mounted right next to the lidar sensors, see Fig. 8 in Section VII. The radar position relative to the ego vehicle $({}^R\mathbf{R}_E, {}^R\mathbf{t}_E)$ was measured via [33], including intrinsic calibration. The DGPS-INS position $({}^E\mathbf{R}_D, {}^E\mathbf{t}_D)$ has been measured via measuring tape. Transformation between other COOS pairs are derived from these tuples by subsequent transformation, e.g. the transformation from ego to camera COOS $({}^C\mathbf{R}_E, {}^C\mathbf{t}_E)$ can be derived by transformation from ego to lidar first and then from lidar to camera.

The output of the camera is an image. Angular information is coded in the image pixels \mathbf{p}_i for camera $i \in \{1, 2\}$. This information has to be first converted to camera coordinates \mathbf{x}_{C_1} and \mathbf{x}_{C_2} . This is done by applying a camera-pinhole model [34] from Eq. 2b and 2c with the intrinsic parameters focal lengths

Table I: Comparison of automatic labeling frameworks.

Framework	[6] / [7]	[4] / [8]	ours
Sensors (lidar/camera/radar)	yes/no/yes	no/yes/yes	yes/yes/yes
Sceneflow approach	None	Optical flow based / no	State-of-the-art
Classes (Pedestrian, car, stationary, background)	yes, yes, no, yes	yes, yes, no, yes	yes, yes, yes, yes
Label density	sparse	sparse	dense
Scene complexity in dataset	highway	test track / urban + highway	urban + highway
Applications (Semantic segmentation, DoA, Target detection, RCS estimation from camera)	yes, no, no, no	yes, no, no, no	yes, yes, yes, yes

$\{f_{x1}, f_{y1}, f_{x2}, f_{y2}\}$ and principal points $\{c_{x1}, c_{y1}, c_{x2}, c_{y2}\}$, estimated via [9]

$$\mathbf{p}_i = [u_i, v_i]^T \quad (2a)$$

$$x_{Ci,x} = \frac{u_i - c_{xi}}{f_{xi}} x_{Ci,z} \quad (2b)$$

$$x_{Ci,y} = \frac{v_i - c_{yi}}{f_{yi}} x_{Ci,z}, \quad (2c)$$

where $\mathbf{x}_{Ci} = (x_{Ci,x}, x_{Ci,y}, x_{Ci,z})^T$. As the remainder of this paper applies to both cameras, we will from now on discard the camera index i for the sake of simplicity. The term $x_{Ci,z}$ is known by dense depth estimation, which will be discussed later.

IV. NOVEL SCENE FLOW ESTIMATION APPROACH

In order to warp the radar RD data into camera image, which is part of the overall processing chain shown in Fig. 1, we first have to estimate the 3-D velocity of every pixel in the camera image. This process is commonly referred as scene flow estimation.

In [35], a scene flow estimation algorithm called deep rigid instance scene flow (DRISF) was proposed, which achieves state-of-the-art performance on the *KITTI* data set at the time of the development of this paper. The idea is, that the overall scene flow can be estimated by considering the 3-D motion of each actor individually. DRISF is based on the definition of an objective function consisting of energy terms, that is optimized by a Gauß-Newton (GN) approach.

As we found the accuracy of DRISF insufficient to achieve satisfactory warping performance, we here propose an extension called deep rigid instance scene flow with radar (DRISFwR), which adds an energy term computed from RD spectrum data to the objective function of DRISF. This radar specific adaptation is schematically depicted in Fig. 2.

DRISFwR not only reduces scene flow errors, as will be illustrated in Sec. VII-D, it also integrates automatic alignment of scene flow estimation with radar data. Furthermore, we incorporate precise ego-motion from DGPS-INS into the algorithm to achieve an accurate estimation of static background scene flow.

In the following, we first describe the input features for DRISFwR, then define required pixel subsets, explain the process of static background scene flow estimation and the required adaptation of the energy formulation for foreground objects.

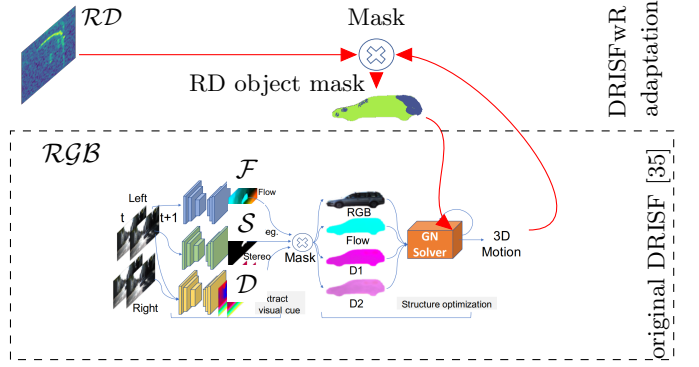


Figure 2: **DRISFwR overview (adapted from [35]):** The original DRISF approach is drawn in the dashed rectangle. Our DRISFwR extension is shown on top. After each iteration of the GN solver, the RD object mask is updated and used for next iteration.

A. Input data to DRISFwR

Four types of knowledge sources are exploited for scene flow estimation: instance segmentation, optical flow, dense depth information and RD spectrum. They are depicted in Fig. 2 as \mathcal{F} , \mathcal{S} , \mathcal{D} and \mathcal{RD} .

Optical Flow mask: For optical flow calculation between two adjacent RGB frames, HD³-Flow [36] has been applied. It learns probabilistic pixel correspondences to provide optical flow confidence at pixel level and achieves state-of-the-art results on the *KITTI* dataset. In the following, let \mathcal{F} denote the computed flow masks, which contain the optical flow values for all the pixels in the camera image \mathcal{RGB} .

Instance Segmentation mask: Segmentation information of relevant objects in the camera's image is obtained by applying a pretrained NN for object segmentation to the camera image. Here, semantic instance segmentation is produced via mask_rcnn_inception_v2_coco from the Tensorflow model zoo [37]. This NN was trained on the COCO dataset and selected here for its good accuracy and fast inference time. During inference, semantic object masks \mathcal{S} are generated as well as object classes. Since the utilized network is not specifically tailored for automotive classes, we retain only the following five relevant classes: pedestrian, car, truck, bicycle and motor bike.

Dense depth mask: In DRISF [35], Ma *et al.* obtained depth estimation based on the stereo camera setup of the *KITTI* dataset. As our sensor setup does not have a stereo camera setup, we utilize sparse depth measurements

$\mathcal{D}_{\text{sparse}} := \{\mathbf{p} \mid \mathbf{p} \in \text{lidar}\}$ of those camera pixels that are in the lidars field of view, and “densify” them with the help of mono-cameras. This process is called depth completion with lots of research done in [9], [38]. We employed the algorithm proposed by [39] which solves depth completion via Markov Random Fields and which requires no adjustment to our sensor setup. The algorithm projects the sparse lidar depth measurements onto the camera image using eqs. (1), (2b) and (2c), and propagates the depth values from lidar occupied pixels into regions with similar brightness values, while enforcing second-order smoothness in depth, see Fig. 3 for an example.




Figure 3: **Depth completion results:** Left-to-right: RGB image, sparse depth information for those pixels that correspond to lidar measurements, dense depth obtained by depth completion.

RD spectrum: As a new cue for scene flow estimation, radar data is employed, which provides valuable velocity measurements from the scene. The non-coherent power spectral density $\mathcal{RD}(\mathbf{p}_s)$ of pixel \mathbf{p}_s is computed from the complex-valued RD spectrum $\mathcal{U}_i \in \mathbb{C}$ of receive antennas $i \in [1, 2, 3]$ of a three-channel radar, as

$$\mathcal{RD}(\mathbf{p}_s) = 10 \log_{10} \sum_{i=1}^3 |\mathcal{U}_i(\mathbf{p}_s)|^2. \quad (3)$$

For further information please refer to [40]. We consider \mathcal{RD} to refer to the totality of all pixels in the spectrum. Alternatively, we could have chosen pointclouds from traditional radar perception instead as a guide for scene flow estimation. However, we selected the RD spectra, because (a) we wanted to be more independent from classical radar perception, and (b) radar pointclouds are obtained from RD spectra and therefore are subsets of those spectra.

B. Pixel sets for scene flow estimation

Utilizing the instance segmentation information, the set of the camera’s foreground pixels is formed by those pixels \mathbf{p} that have been identified to belong to one of the following object classes:

$$\mathcal{P}_{\text{fg}} := \left\{ \mathbf{p} \mid \mathcal{S}(\mathbf{p}) \in \{\text{pedestrian, car, truck, bicycle, motorbike}\} \right\} \quad (4)$$

by the instance segmentation algorithm. Further, let $\mathcal{P}_{\text{radar}}$ be the set of camera pixels covered by the radar’s field of view (FoV), which is defined as follows

$$\mathcal{P}_{\text{radar}} := \left\{ \mathbf{p}(\mathbf{x}_R) \mid |\phi(\mathbf{x}_R)| \leq \frac{135^\circ}{2} \wedge \theta(\mathbf{x}_R) \leq \frac{22^\circ}{2} \right\}, \text{ where} \\ \phi(\mathbf{x}_R) = \tan^{-1}(x_{R,x}, x_{R,y}), \theta(\mathbf{x}_R) = \tan^{-1}(x_{R,x}, x_{R,z}). \quad (5)$$

Here, \mathbf{x}_R is the point in the COOS of the radar system, and $\mathbf{p}(\mathbf{x}_R)$ the pixel in the camera image, corresponding \mathbf{x}_R .

To mitigate the effect of errors in instance segmentation, we perform further object clustering on each pointcloud belonging to an instance in camera image via DBSCAN [41]. Thereby, points which have a spatial distance larger than 0.3m from the closest point in the cluster i are classified as outliers and removed from the main cluster corresponding to the instance, resulting in the set $\mathcal{P}_{\text{DBSCAN}}$ of valid pixels. The threshold of 0.3m was considered appropriate for the used lidar sensor.

For a pixel to be used for scene flow estimation, it must be both in the set of valid pixels after DBSCAN clustering, $\mathcal{P}_{\text{DBSCAN}}$, and in the FoV of the radar sensor:

$$\mathcal{P}_i := \left\{ \mathbf{p} \mid \mathbf{p} \in \mathcal{P}_{\text{DBSCAN}} \wedge \mathbf{p} \in \mathcal{P}_{\text{radar}} \right\}. \quad (6)$$

This set definition automatically handles the sensor parallax between radar and camera. An example of valid pixels is visualized in Fig. 4 for both cameras. Valid pixels are highlighted as yellow pixels. It can be seen, that no pixel of the ego-vehicle was selected as valid. Similarly, only pixels from the radar FoV are marked as valid. Some pixels surrounding objects are automatically marked as invalid, according to $\mathcal{P}_{\text{DBSCAN}}$.




Figure 4: **Masks for selection of valid pixels:** Examples of valid pixels (yellow) for all objects according to \mathcal{P}_i visualized for both cameras (left: camera 1, right: camera 2.).

C. Motion formulation

The overall scene flow $\boldsymbol{\xi} = [\xi_x, \xi_y, \xi_z]^T$ is a superposition of background $\boldsymbol{\xi}_{bg}$ and foreground scene flow $\boldsymbol{\xi}_{fg}$,

$$\boldsymbol{\xi} = \boldsymbol{\xi}_{bg} + \boldsymbol{\xi}_{fg}. \quad (7)$$

The background scene flow encodes the motion of the ego- E over ground and applies to all pixels. Foreground motion is the motion of objects over ground. Similar to DRISF, the foreground motion is computed only for pixels in \mathcal{P}_{fg} . In contrast to DRISF, we decided to encode the foreground motion in 3 translational motion parameters instead of 6 (3 rotational + 3 translational). This simplification was made to improve the robustness and is justified by assuming small object dimensions and low rotational rates for foreground objects for the majority of traffic scenarios.

Background motion formulation: The perceived motion of the background highly depends on the motion of the ego vehicle. Here, rotation cannot be neglected since, even when the ego vehicle's rotation is small, the scene flow can be high for objects that are far away, see Eqs. (8) and (9a) below.

In the following, let k denote the discrete time index. Background motion is understood to be the relative 3-D motion of the environment, which is induced by the movement of the ego vehicle over ground. The position $\mathbf{x}_E(\mathbf{p}, k+1)$ of a stationary point \mathbf{p} at time step $k+1$ can be predicted from its position $\mathbf{x}_E(\mathbf{p}, k)$ at time step k as follows:

$$\mathbf{x}_E(\mathbf{p}, k+1) = \mathbf{R}_{E,bg}(\mathbf{x}_E(\mathbf{p}, k) + \mathbf{x}_{RA}) - \mathbf{x}_{RA} - \mathbf{t}_{E,bg}, \quad (8)$$

in which $\mathbf{R}_{E,bg} \in \mathbb{R}^{3 \times 3}$ and $\mathbf{t}_{E,bg} \in \mathbb{R}^{3 \times 1}$ describe the rotation and translation of the ego vehicle over ground, an information that can be obtained by DGPS-INS. In Eq. (8), before applying the rotation, the point $\mathbf{x}_E(\mathbf{p}, k)$ in the COOS of the ego vehicle is first shifted onto the rear axle of the ego vehicle by adding \mathbf{x}_{RA} , according to the Ackermann vehicle model [42], which assumes that vehicles yaw around their rear axle. Similar as described in [11], $\mathbf{R}_{E,bg}$ and $\mathbf{t}_{E,bg}$ are refined over a sequence of camera images.

By applying the coordinate transformation, Eq. (1), from ego to camera, we arrive at

$$\mathbf{x}_C(\mathbf{p}, k+1) = \mathbf{R}_{C,bg} \mathbf{x}_C(\mathbf{p}, k) + \mathbf{t}_{C,bg} \quad (9a)$$

$$\mathbf{R}_{C,bg}(k) = {}^C\mathbf{R}_E \mathbf{R}_{E,bg} {}^E\mathbf{R}_C \quad (9b)$$

$$\begin{aligned} \mathbf{t}_{C,bg} = & {}^C\mathbf{R}_E \mathbf{R}_{E,bg} \mathbf{x}_{RA} - {}^C\mathbf{R}_E(\mathbf{x}_{RA} - \mathbf{t}_{E,bg}) \\ & + (\mathbf{I} - \mathbf{R}_{C,bg}(k)) {}^C\mathbf{t}_E. \end{aligned} \quad (9c)$$

This eventually gives the background scene flow of point \mathbf{p} :

$$\xi_{bg}(\mathbf{p}) = \mathbf{x}_C(\mathbf{p}, k+1) - \mathbf{x}_C(\mathbf{p}, k), \quad (10)$$

Foreground motion formulation: In DRISFwR, a radar specific energy term is added to the energy formulation of DRISF:

$$\min_{\xi} \left\{ \underbrace{\lambda_{\text{photo}} E_{\text{photo}}(\xi; \mathcal{I}) + \lambda_{\text{rigid}} E_{\text{rigid}}(\xi; \mathcal{I}) + \lambda_{\text{flow}} E_{\text{flow}}(\xi; \mathcal{I})}_{\text{original DRISF}} + \underbrace{\sum_{s_d=1}^3 \lambda_{\text{radar}, s_d} E_{\text{radar}}(\xi; \mathcal{I}, s_d)}_{\text{extension by DRISFwR}} \right\}. \quad (11)$$

For the sake of simplicity, let $\mathcal{I} = \{\mathcal{RGB}^0, \mathcal{RGB}^1, \mathcal{D}^0, \mathcal{D}^1, \mathcal{S}^0, \mathcal{F}, \mathcal{RD}\}$, thus any time an equation requires one or multiple of the arguments, we use \mathcal{I} instead. The λ 's are the weighting factors of the individual energy terms. E_{photo} describes, how well the \mathcal{RGB} pixels from the object agree to those corresponding to \mathcal{RGB} , when predicted according to ξ . E_{rigid} describes how well pixel positions agree between frames assuming a scene flow of ξ . The energy term E_{flow} captures differences between the optical flow mask \mathcal{F} and the obtained scene flow. For more details about these energy terms, we refer the reader to the original DRISF paper [35], as we will focus on the radar specific extension here.

The newly introduced E_{radar} energy captures how well ξ agrees to the relative velocity measured by radar. This term is formulated as

$$E_{\text{radar}}(\xi; \mathcal{I}, s_d) = \sum_{\mathbf{p} \in \mathcal{P}_i} \rho(r_{\text{radar}}(\xi, \mathbf{p}; \mathcal{I}, s_d)) \quad (12a)$$

$$r_{\text{radar}}(\xi, \mathbf{p}; \mathcal{I}, s_d) = \mathcal{RD}_{s_d}(|\mathbf{x}_R(\mathbf{p})|, v_R(\xi, \mathbf{p})) - \mathcal{RD}_{\text{target}}, \quad (12b)$$

where $\rho(r_{\text{radar}}) = (r_{\text{radar}}^2 + 10^{-6})^{0.45}$ (generalized Charbonnier penalty) is the robust fitting function which we adopt from the original DRISF. The power of the RD-map value at range $|\mathbf{x}_R|$ and relative velocity $v_R(\xi)$ seen from radar is $\mathcal{RD}(|\mathbf{x}_R|, v_R(\xi))$, whereas $\mathcal{RD}_{\text{target}}$ is the target value of the power and here set to be the maximum power of the RD-map. This setting enforces the scene flow to maximize the RD-map power of an object in its local RD-map vicinity determined by the other energy terms. We found this behavior similar to what human labelers would expect. The summation index s_d , the scale level, will be discussed later in this section.

The minimization of (11) can be formulated as a weighted Least Squares (wLS) problem, see the supplementary material of [35] for details,

$$\xi^{(m+1)} = \underset{\xi}{\operatorname{argmin}} \left\{ \sum_{\mathbf{p} \in \mathcal{P}_i} \mathbf{r}^T(\mathbf{p}, \xi^{(m)}) \mathbf{W}(\mathbf{p}, \xi^{(m)}) \mathbf{r}(\mathbf{p}, \xi^{(m)}) \right\}. \quad (13)$$

Here, $\mathbf{r} \in \mathbb{R}^K$ is the vector of residuals, in which K is the number of captured residuals from Eq. 11. $\mathbf{W} \in \mathbb{R}^{K \times K}$ is the diagonal weight matrix which is specified by the λ 's and the selected fitting function as shown later in this section.

In each iteration m , the GN solver performs scene flow updates $\xi^{(m+1)} = \xi^{(m)} + \Delta \xi^{(m)}$ by computing

$$\Delta \xi^{(m)} = - \left(\sum_{\mathbf{p} \in \mathcal{P}_i} \mathbf{J}^T \mathbf{W} \mathbf{J} \right)^{-1} \sum_{\mathbf{p} \in \mathcal{P}_i} \mathbf{J}^T \mathbf{W} \mathbf{r}. \quad (14)$$

Here, $\mathbf{J} \in \mathbb{R}^{K \times 3}$ is the Jacobian giving the gradient of the residuals in each scene flow direction. All terms on the right hand side depend on \mathbf{p} and $\xi^{(m)}$, which we omitted, however, for the sake of readability.

The Jacobian $\mathbf{J}_{\text{radar}} \in \mathbb{R}^{1 \times 3}$ corresponding to the residual of the radar term r_{radar} , which defines one row in \mathbf{J} , is computed as follows:

$$\mathbf{J}_{\text{radar}} = \frac{\partial r_{\text{radar}}(\xi; \mathcal{I}, s_d)}{\partial \xi} = \frac{\partial \mathcal{RD}(\xi; \mathcal{I}, s_d)}{\partial v_r(\xi)} \frac{\partial v_r(\xi)}{\partial \xi_{\text{radar}}(\xi)} \frac{\partial \xi_{\text{radar}}(\xi)}{\partial \xi}. \quad (15)$$

The first term on the right hand side represents the dependency of RD-map power on relative radial velocity in form of image gradients. Example gradients are illustrated in Fig. 5, the right-hand figures in the box with the yellow background. At RD-map positions with high power, one can see large negative or positive gradients, the sign depending on the horizontal approaching direction.

Since the image gradients in the RD-map are very steep, a scale-space [43] is created by repeated Gaussian smoothing, max-pooling of \mathcal{RD} and bilinear up-sampling along the Doppler dimension. The smoothed RD-map is referred to as

\mathcal{RD}_{s_d} . A Gaussian kernel is applied to obtain smooth gradients. Max-pooling enhances the receptive field of DRISFwR, which allows object alignment in RD-map over a wider range of velocities. Upscaling is performed to keep the bin resolution constant over scales and to induce further content smoothing. Different scale levels s_d are considered in DRISFwR by accumulating their energy in Eq. (11). To prefer alignment on lower scales, energy weighting λ_{radar} , is bisected at every scale level s_d . Furthermore, to cope with aliasing in the Doppler spectrum, copies of the radar spectrum are concatenated in Doppler direction, see Fig. 5.

The second term on the right hand side of Eq. 15 incorporated the scalar projection of the scene flow onto positional vector \mathbf{x}_R , both observed from radar perspective.

$$v_r(\xi_{\text{radar}}) = \frac{\mathbf{x}_R^T}{|\mathbf{x}_R|} \xi_{\text{radar}}, \quad (16)$$

Its derivative is given by:

$$\frac{\partial v_r(\xi_{\text{radar}})}{\partial \xi_{\text{radar}}} = \frac{\mathbf{x}_R^T}{|\mathbf{x}_R|}. \quad (17)$$

The relation between a position in the radar COOS and in the camera COOS is given by coordinate transformation according to Eq. (1). Thus

$$\begin{aligned} \xi_{\text{radar}} &= \frac{\partial \mathbf{x}_R}{\partial t} = {}^R\mathbf{R}_E {}^E\mathbf{R}_L {}^C\mathbf{R}_L^{-1} \frac{\partial \mathbf{x}_C}{\partial t} \\ &= {}^R\mathbf{R}_E {}^E\mathbf{R}_L {}^C\mathbf{R}_L^{-1} \xi = {}^R\mathbf{R}_C \xi, \end{aligned} \quad (18)$$

which transforms the scene flow from camera coordinates into radar coordinates via the rotational matrix ${}^R\mathbf{R}_C$. We obtain

$$\frac{\partial \xi_{\text{radar}}}{\partial \xi} = {}^R\mathbf{R}_C. \quad (19)$$

The adaptive weighting W_{radar} is given by:

$$W_{\text{radar}} = \lambda_{\text{radar}} \frac{\partial^2 \rho}{\partial r_{\text{radar}}^2} = 0.45 \lambda_{\text{radar}} (r_{\text{radar}}^2 + 10^{-6})^{-0.55}. \quad (20)$$

An update step of the GN algorithm is depicted in Fig. 5. The figure shows two cars highlighted by the red rectangles in the camera image and the RD-map in the blue box. In the yellow box, the position in the RD-map for one car is shown in red for multiple update steps. After 100 iterations, the car's position in RD-map has shifted into the local power maxima in the RD-map. The remaining gray box shows the RD-map power warped into the camera image. At the location of the cars in the camera image, the warped power has significantly increased after 100 iterations. The warping operation will be discussed in the next section.

GN iterations are stopped once $\|\Delta \xi\|_2 < 0.1$ mm. Typically this occurred after approx. 30 iterations but was stopped when convergence was not reached after 100 iterations.

V. WARPING FROM RADAR SPECTRA INTO CAMERA IMAGE

In the previous section, we described the scene flow and the dense depth estimation. As can be seen in Fig. 1, they form the fundamental input for the warping operation which we will describe in the following.

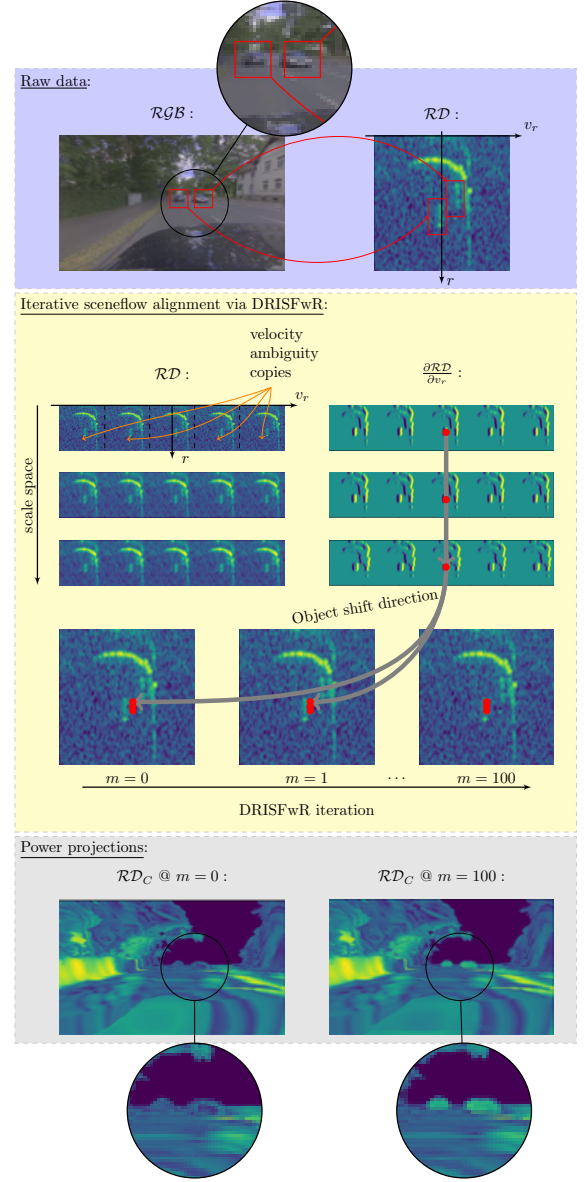


Figure 5: **Automatic scene flow alignment to Radar data via DRISFwR:** Top: RGB image and RD-map with two vehicles. Middle: Scale-space of radar data used in DRISFwR with energy (left) and partial derivative (right). One object is visualized in scalespace as red point. Repositioning via scene flow adjustment in consecutive DRISFwR update steps with gray arrows marking shift direction. Bottom: Power projections. After DRISFwR convergence, energy projection from RD-map is captured correctly for both vehicles (bottom right)

A. Bilinear warping operation

The radar RD grid is given by coordinates of radial relative velocity v_r and range $|\mathbf{x}_r|$ seen from the radar perspective. These can be obtained following Eq. (1) and Eq. (16). Note, that only the radial relative velocity is relevant for the warping operation, as the radar sensor is not able to measure tangential velocities.

In our warping operation $\eta(\cdot, \cdot, \cdot)$, we select for each pixel

in the camera image the corresponding pixel in the RD grid depending on the assigned $v_r(\xi)$ and $|\mathbf{x}_r|$ and project the assigned RD grid pixel value back into the camera image grid. If $v_r(\xi)$ and $|\mathbf{x}_r|$ are not on the RD grid, bi-linear interpolation is applied, since it provides visually appealing projections while being differentiable, which enables the backpropagation in NN training.

Fig. 6 shows an example of warping an RD-map into the camera image. Such a warping helps a human to understand how a radar sees the environment in terms of power. We will refer to the projected RD-map as \mathcal{RD}_C in the following.

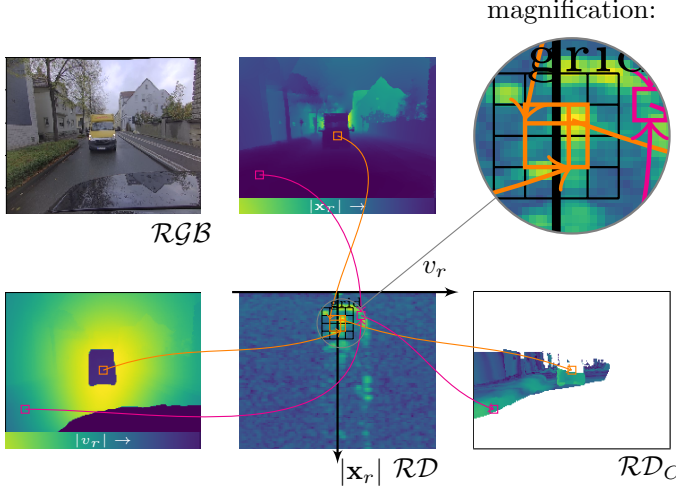


Figure 6: **RD-map warping into camera image:** As an example for the warping operation, here for two pixels in the camera image (upper left) their estimated range (upper center) and radial velocity (lower left) determines their corresponding position in the RD-map (lower center). The RD-maps intensity value is bi-linearly interpolated (upper right) and then projected into the camera image (lower right). Notice that only pixels in the radar’s FoV are warped here (non-white pixels in \mathcal{RD}_C).

\mathcal{RD}_C can be calculated at the pixel position \mathbf{p} , using bilinear interpolation, as:


$$\mathcal{RD}_C(\mathbf{p}) = \eta(\mathcal{RD}, v_r(\xi), |\mathbf{x}_r|, \mathbf{p}). \quad (21)$$

Note, that in Eq. (21) and Fig. 6 the RD-map \mathcal{RD} was depicted as an illustrative example for the warping operation. However, every information in the form of the RD grid can be warped into a camera image in an analog fashion. In fact, we will warp a radar based prediction by NNs in the same manner, as will be discussed in the remainder of this paper.

B. Bilinear vs. trilinear warping

In the bilinear interpolation only velocity and distance from our reference are used for warping. But the proposed framework is flexible and can also employ 3-D data for warping. To demonstrate that, Fig. 7 shows the warping of a 3-D radar spectrum, which consists of velocity, distance and azimuth DoA information, into a camera image. The 3-D radar spectrum was obtained by performing FFT beamforming

[26] on the radars channel returns. For warping, trilinear interpolation was used if the considered camera pixel does not correspond to a point on the 3-D RD grid.



(a) Bilinear interpolation (b) Trilinear interpolation

Figure 7: **Comparison of the interpolation methods:** Left: Projection of the power from RD-map via bilinear interpolation. Right: Projection of the 3D beamforming spectrum via trilinear interpolation. Pixels outside the radars FoV are visualized dark.

We normalized the signal power in the beamforming spectrum according to the number of used antennas, thus highly reflective regions in Fig. 7 do have similar power values. The main difference in the two images can be found in the low power regions, e.g. the street area, in which better noise suppression was achieved by beamforming.

VI. CROSS-MODAL SUPERVISED NN TRAINING ON RADAR SPECTRUM DATA

In this section, we describe an example application for training a “student” neural network operating on radar spectrum data, whose training targets have been provided from a “teacher” neural network, which labeled the camera data.

A. Radar based DoA prediction via Neural Network

To estimate the DoA of the radar reflections, a convolutional neural network (CNN) is used to operate on the radar spectrum \mathcal{U} . To be precise, the radar spectrum of the utilized three-channel radar $\{\mathcal{U}_1, \mathcal{U}_2, \mathcal{U}_3\}$ is preprocessed before being transferred to the network. The first feature map is the RD-map, while the second and third feature maps are the phase spectra $\angle(\mathcal{U}_i, \mathcal{U}_{i-1})$ between adjacent channels:

$$\Gamma = \begin{bmatrix} \mathcal{RD} \\ \angle(\mathcal{U}_2, \mathcal{U}_1) \\ \angle(\mathcal{U}_3, \mathcal{U}_2) \end{bmatrix} \quad (22a)$$

$$\angle(\mathcal{U}_i, \mathcal{U}_{i-1}) = \arg(\mathcal{U}_i \mathcal{U}_{i-1}^*). \quad (22b)$$

The RD-map is provided to allow the neural network to discriminate between targets and noise and the phase spectra are provided as a main clue for DoA inference.

For every pixel in the RD grid, the NN, denoted as $\phi\text{-Net}(\Gamma)$, delivers a DoA estimate $\phi_{\text{RD-grid}}$ in azimuth direction which is warped into the camera image as $\phi_{\text{pred., RD}}$:

$$\phi_{\text{pred., RD}} = \phi\text{-Net}(\Gamma) \quad (23)$$

$$\phi_{\text{pred., cam}}(\mathbf{p}) = \eta\left(\phi_{\text{pred., RD}}, v_r(\xi), |\mathbf{x}_r|, \mathbf{p}\right). \quad (24)$$

For the network training, DoA labels $\phi_{\text{reference}}$ are generated from the dense depth estimation in the camera image, projected into the radar COOS as:

$$\phi_{\text{reference}} = \arctan 2(x_{R,x}, x_{R,y}). \quad (25)$$

Predictions of $\phi\text{-Net}$ are made, so that the objective $l_{\phi\text{-Net}}$ is minimized during network training:

$$l_{\phi\text{-Net}} = \sum_{\mathbf{p} \in \mathcal{P}_{\text{train}}} |\phi_{\text{reference}}(\mathbf{p}) - \phi_{\text{predict, cam}}(\mathbf{p})|, \quad (26)$$

in which the summation is done over all pixels \mathbf{p} in $\mathcal{P}_{\text{train}}$, as we want the network to focus on relevant targets. Examples for $\phi_{\text{reference}}$, $\phi_{\text{predict, cam}}$ and $\phi_{\text{predict, RD}}$ are depicted in Fig. 13. Note, that $\phi_{\text{reference}}(\mathbf{p})$ and $\phi_{\text{predict, cam}}(\mathbf{p})$ provide angles within sensor FoV $< 180^\circ$ by physical properties and therefore 2π wrap must not be treated.

VII. EVALUATION

A. Data Acquisition System

To gather real world sensor data from driving scenarios, we equipped a car with a 77 GHz HELLA radar sensor, with an overview of the sensor specifications given in Tab. II. The sensor captures the raw data, computes RD spectra in the first processing step and forwards them to an in-vehicle recording PC. The radar sensor is mounted at the rear right corner and behind the bumper, as illustrated in Fig. 8.

Table II: Specifications of the tested radar sensor

Operating frequency	77 GHz
Antenna design	static 1Tx3Rx
Modulation scheme	FMCW chirp sequence
Range resolution	0.25 m
Doppler resolution	0.25 m s^{-1}
Maximum range	25 m
Maximum unambiguous Doppler	10 m s^{-1}
FoV (horizontal/vertical)	$140^\circ/20^\circ$

The utilized reference sensor system consisting of 2 Velodyne VLP-32C lidar scanners [44] at 10 FPS and a camera array (2 FirstSensor DC3C-1-E4P-105 [45]) is mounted on a luggage rack on top of the vehicle. In order to sample the surrounding more densely, the system is equipped with two lidar sensors, which are mounted with different orientations on the vehicle. They jointly provide a denser point-cloud in the radar's FoV than a single lidar could provide, see Fig. 3. In order to cover the entire radar FoV (azimuth: 135° , elevation: 22°), two cameras (azimuth: $\approx 104^\circ$) with overlapping FoVs were used. The camera array is able to record a continuous videostream at 30 FPS. The Differential-GPS with inertial navigation system (DGPS-INS) of type GeneSys ADMA-G-Pro+ [46] is used as a reference sensor for precise vehicle

over ground motion estimation. All sensors have been properly synchronised via temporal calibration proposed in [47]. To be specific, the trigger times of the sensors have been recorded and the temporally closest sensor samples were associated. Additionally, motion correction was applied on the data of the lidar sensors.

The driving trajectory for this dataset is given in Fig. 8 and contains mostly urban environment. It has been recorded over a period of approx. 1 h driving, giving rise to 36000 frames at 10 FPS and approx. 7 billion samples on pixel level.




Figure 8: **Data capturing setup:** Top: vehicle sensor configuration; bottom: driving trajectory for recording the database; bottom left: magnified scene flow evaluation area.

B. Scene flow evaluation

The performance of the proposed scene flow algorithm has been benchmarked on a subset of our dataset, with known ground truth for scene flow. A parking lot area (see Fig. 8) with all objects being stationary, was selected. The ground-truth for scene flow estimation was obtained from DGPS-INS, which delivered precise information about the ego motion of the vehicle. Scene flow in stationary scenes has been described in Sec. IV-C, and thus the scene flow ground truth ξ_{gt} is given by $\xi_{gt} = \xi_{bg}$.

The KITTI dataset provides scene flow evaluation metrics mainly in camera image plane direction. In our scenario, scene flow errors both orthogonal and tangential to the image plane are important and thus we decided to use error metrics that

consider both errors equally. The first is the mean absolute error (MAE)

$$\text{MAE}_{sf} = \frac{1}{N} \sum_{\mathbf{p}_{sf} \in \mathcal{P}_{sf}} \|\xi_{gt}(\mathbf{p}_{sf}) - \xi(\mathbf{p}_{sf})\|_2 \quad (27)$$

$$\mathcal{P}_{sf} := \{\mathbf{p} \mid \mathbf{p} \in P_{\text{radar}} \wedge \mathbf{p} \in P_{\text{sparse}} \wedge \mathbf{p} \in P_{\text{fg}}\},$$

between ground truth and estimated scene flow, measured on all pixels \mathbf{p} visible to the radar, have assigned lidar sparse depth measurement and have an instance mask (i.e., are foreground pixels). Here, N is the cardinality of that set \mathcal{P}_{sf} .

The second evaluation metric is the scene flow error rate, which is the number of foreground pixels exceeding an error threshold of 0.25 m/s, which is a typical radar velocity resolution.

As can be seen in Tab. III the DRISFwR algorithm achieves lower MAE and error rates than the current top-ranking scene flow estimation algorithm in *KITTI* benchmarks. Both algorithms assign a strong motion prior, by assuming rigid bodies for each instance from instance segmentation. As an alternative scene flow estimation technique with weak motion constraints, we include an estimation based on optical flow by warping depth images between frames via HD³ [36]. It can be seen, that scene flow performance without strong motion prior assumption is much worse. This is expected, as the nominal accuracy of the reference system is 0.16° tangential and 0.03 m radial giving a theoretical three-dimensional accuracy of $\approx 1.42 \text{ m s}^{-1}$ for a target in 20 m distance. The rigid object assumption often results in hundreds of pixel level samples per object allowing to stay below the nominal accuracy of the reference sensors itself. The runtime of DRISFwR exceeds the one of DRISF slightly, as it includes an additional energy term. The computation is performed offline, allowing us to automatically label our 1 h long dataset within 1 day, which we found acceptable. All algorithms have been implemented in PyTorch [48], and the GPU code has been executed on NVIDIA GeForce GTX TITAN X.

Table III: **Quantitative comparison of scene flow approaches.** DRISFwR reduces scene flow MAE and error rate compared to DRISF and the an estimator without strong motion prior HD³.

Methods	runtime	error rate (%)	MAE _{sf} (m/s)
HD ³	0.12 sec	69.9	4.37
DRISF	0.6 sec	31.2	0.31
DRISFwR	0.8 sec	25.5	0.22

Note, that for our application the radial velocity error is likely to be smaller than the values given in the table, as we only require accurate radial velocity estimates, whereas the metric measures both radial and tangential velocity deviations.

C. Automatic labeling quality

DRISFwR assists the warping of radar spectra into camera images without human supervision, within the scope of the evaluated accuracy, see prior section. However, sometimes scene flow is found to be erroneous. Three error types can be discerned:

- 1) stationary objects are mis-labeled as moving
- 2) significant portions of the RD-map power is not transferred into \mathcal{RD}_C
- 3) erroneous segmentation of point cloud resulting in unreasonable geometric extension in \mathcal{RD} .

Fortunately, these types of errors can easily be spotted by a trained human and labeled as anomalies and thus be excluded from the subsequent comparison between the reference data and the estimates from radar. Thus, with the help of our reference label generation system, the human labeling task can be tremendously simplified from a regression problem (draw DoA reference angles) to a binary classification problem (identify whether proposed label is plausible or not). We think, the first problem cannot even be done by human labelers, while the later one can easily be done. Our test dataset (10800 images) has been labeled by following this process. Examples of anomalous masks are given in Fig. 9. Of all test dataset images, approx. 20.3% received at least one anomaly mask. We consider this value reasonable, as it is close to the warping error rate for DRISFwR, identified in the last section. The average human labeling time per image was approx. 5 s. As the anomaly labeling catches all anomalies induced by erroneous warping, the anomaly masks can be used not only for the DoA application, but also for training other applications, such as target detection and semantic segmentation.

Note that the human anomaly labeling has only been applied to the test dataset in order to produce the most reasonable metrics. The NN training, however, was performed without this human intervention and thus includes even warping errors.

D. Direction-of-arrival evaluation

Typically, the radar based estimation quality highly depends on the signal-to-noise-ratio (SNR). We therefore measured the performance as a function of the SNR. Here, the signal power is estimated as the power of the corresponding pixel in \mathcal{RD} . The noise level estimate is obtained by applying the CFAR target detector [40] to \mathcal{RD} and calculating the 99.5% percentile of all non-target pixels. In stationary scenes with low ego vehicle velocity $< 2 \text{ m s}^{-1}$, the signal power from multiple reflections accumulates in the zero Doppler velocity pixels. These signal powers cannot be recovered, leading to erroneous signal power estimation and therefore, deteriorate the quality metrics unreasonable. Stationary scenes are thus discarded from the test dataset.

The warping operation is non-bijective in the sense that multiple pixels in the camera may correspond to a single pixel in the RD-map and vice versa. In evaluation, this typically means (as the former case is more typical), that one estimation value (one pixel in the RD grid) is compared to an ensemble of reference values \mathcal{P}_s (multiple pixels in the camera image). The set \mathcal{P}_s can be formed by warping the linear pixel indices of RD-map into camera image and selecting all pixels to a certain pixel indice value. Ideally, one would want to assign an electromagnetic contribution to each reference pixel and compute the center of reflection to which the estimated value is then compared to. However, at the current stage of this reference system, the electromagnetic contribution is unknown.




Figure 9: **Examples of anomalous labels:** From left to right: camera images with semantic instance masks from radar FoV only (red: pedestrian, green: vehicle), projected RD-maps, RD-maps with semantic masks and anomaly masks (white rectangles). The orange arrows connect corresponding regions in camera image and RD-map. In the two upper rows, stationary objects were labeled as moving, due to an error in scene flow estimation (error type 1). However, the misalignment in RD-map is relatively small. The bottom row shows an error in instance segmentation, in which background pixels were mapped as pedestrian, resulting in an unreasonably extended object range in RD-map (error type 3).

In practice, we select the best matching pixel from the camera ensemble to the RD grid pixel as the tuple $\{\mathbf{p}_C, \mathbf{p}_{RD}\}$ such that

$$\mathbf{p}_C = \arg \min_{\mathbf{p}_C \in \mathcal{P}_s} \left| \phi_{\text{reference}}(\mathbf{p}_C) - \eta \left(\phi_{\text{predict, RD}}(\mathbf{p}_{RD}); v_r, |\mathbf{x}_r| \right) \right|. \quad (28)$$

The DoA MAE is thus computed as

$$\text{MAE}_{\text{DoA}} = \frac{1}{N} \sum_{\mathbf{p}_C \in \mathcal{P}_{\text{radar}}} \left| \phi_{\text{reference}}(\mathbf{p}_C) - \eta \left(\phi_{\text{predict, RD}}(\mathbf{p}_{RD}); v_r, |\mathbf{x}_r| \right) \right|, \quad (29)$$

where N is the number of RD grid pixels visible to camera and radar, respectively the size of the set $\mathcal{P}_{\text{radar}}$.

We developed a CNN for DoA estimation. It employs ReLU activation on every layer except for the last, where $90 \cdot \tanh()$ is applied to scale the logits' angular values to degrees. The convolutions are operated with unit stride and the layers have the following number of output channels: $[3, 32t, 64t, 128t, 64t, 32t, 32t, 1]$. We experimented with different convolution kernel sizes of 1×1 and 3×3 and equalized the networks parameter count with the layer modifier t . The inference of the 1×1 network will be influenced by one pixel in the RD grid only, while the 3×3 network may use the RD grid pixel vicinity to perform inference. The receptive fields thus compute to 1 ($0.25 \text{ m} \times 0.25 \text{ m s}^{-1}$) and 15 ($3.75 \text{ m} \times 3.75 \text{ m s}^{-1}$), respectively.

We subdivided the pixels in RD map into two groups, the first being pixels from real reflections, thus carrying a noisy

signal, and the second being pixels from noise only. As we want to focus the network training on the first group only, we automatically identify pixels of the second group as those whose SNR is found to be below the threshold of 10dB, giving rise to the following definition of the training set

$$\mathcal{P}_{\text{train}} := \{\mathbf{p} \mid \mathbf{p} \in \mathcal{P}_{\text{radar}} \wedge \mathcal{RD}_C(\mathbf{p}) > 10\text{dB}\}. \quad (30)$$

This threshold has been found by plotting the distribution of \mathcal{RD}_C for the entire test dataset, see Fig. 10, and identifying the SNR value in which the likelihood of the second group samples can be neglected. To this end, a two-component mixture model was fitted to the data, consisting of a Gaussian and a Chi-squared distribution. The parameters of the Gaussian probability density function (PDF), that represents noise, was determined according to central limit theorem [49]. Assuming Swerling type 3 fluctuations for the radar reflections, see e.g. [50], the parameters of the Chi-squared PDF representing the signal were identified.

As can be seen, at +10dB the likelihood of sampling from the identified noise mixture is very low ($< 0.1\%$) and samples above +10dB are mainly provided by signal.

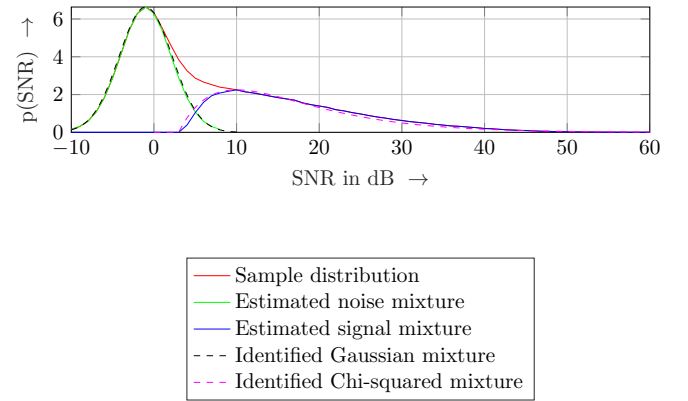


Figure 10: **SNR distribution:** The sample distribution is shown in red, which is decomposed into a noise component (green) and a signal component (blue). Additionally, the sample mixtures are approximated by Gaussian PDF (dashed black) and Chi-squared PDF (dashed magenta).

To explore the influence of the above mentioned threshold selection, we tested each network either with or without the SNR-based sample selection. The parameters of the three tested NN architectures are depicted in Tab. IV.

Table IV: **NN architectures and their configurations.**

Name	kernel size	layer modifier (t)	SNR selection
NN ₀	1×1	3	-
NN ₁	1×1	3	$\mathcal{RD}_C > 10\text{dB}$
NN ₂	3×3	1	-
NN ₃	3×3	1	$\mathcal{RD}_C > 10\text{dB}$

As we want the training and test datasets to be independent and identically distributed (i.i.d.), the 1 h long dataset (Fig. 8) is split into 10 s long sequences of 100 frames each. The sequences are randomly assigned to training, validation and test sets with ratios 70%, 15%, 15% and with no frame overlap.

All networks are optimized via ADAM and a learning rate of 10^{-5} . Early stopping was performed by checking network performance on the validation dataset after each training epoch.

The performance of different DoA estimators in terms of MAE is shown in Fig. 11. In addition to NN-based estimators we include the performance of a phase-comparison monopulse (PM) [25] DoA estimator and Bartlett beamforming (BF) [26] as baseline algorithms which do not require a training phase.

As expected, it can be seen, that BF achieved better MAE than PM over the entire SNR range.

The (3x3) NNs achieved better MAE than their (1x1) counterparts, which suggests, that neighboring RD-map pixels carry valuable information that can be used for a better DoA estimation. A possible explanation for this is (a) the windowing performed before computing the FFT based RD grid, which smears energy over multiple pixels, and (b) the fact that geometrically neighboring pixels in the camera image tend to occupy neighboring locations in the RD-grid

All NNs with enabled SNR selection during training achieved significantly better MAE at higher SNR values.

Finally it can be seen, that the best performing NN (NN₃) achieved equal or better MAE accuracy as the best performing classical DoA estimator (BF).




Figure 11: **Performance of DoA estimation methods:** Mean Absolute Error of azimuth estimation of different DoA estimators and performance of moving target indication (MTI).

Additionally we perform moving target indication (MTI), which heavily relies on DoA predictions to classify stationary versus moving targets. The utilized MTI algorithm is from [51], [52]. If $\|\xi_{fg}\|_2^1 > 0.5 \text{ ms}^{-1}$ we set the reference label to “moving” and otherwise to “stationary”. MTI accuracy is calculated by comparing the prediction to the reference label. It can be observed in the figure that in low SNR regions, the (3x3) NN achieved better accuracy compared to the other estimators, while it is surpassed by beamforming in the high SNR regime.

For a more detailed analysis, Fig. 12 depicts the MAE as a function of the azimuth angle and SNR for selected estimators. Comparing the MAE histogram of NN₀ with NN₁ and NN₂ with NN₃, we see that the NNs with enabled SNR selection (NN₁ and NN₃) achieved better MAE values

especially at the bordering azimuth regions. This suggests, that optimizing the NN on pixels from the signal mixture only, see Fig. 10, in combination with the used loss function, see Eq. 26, is beneficial. However, to get isotropic DoA estimation behaviour over the entire azimuth range, it might be necessary to experiment with other loss functions in the future.

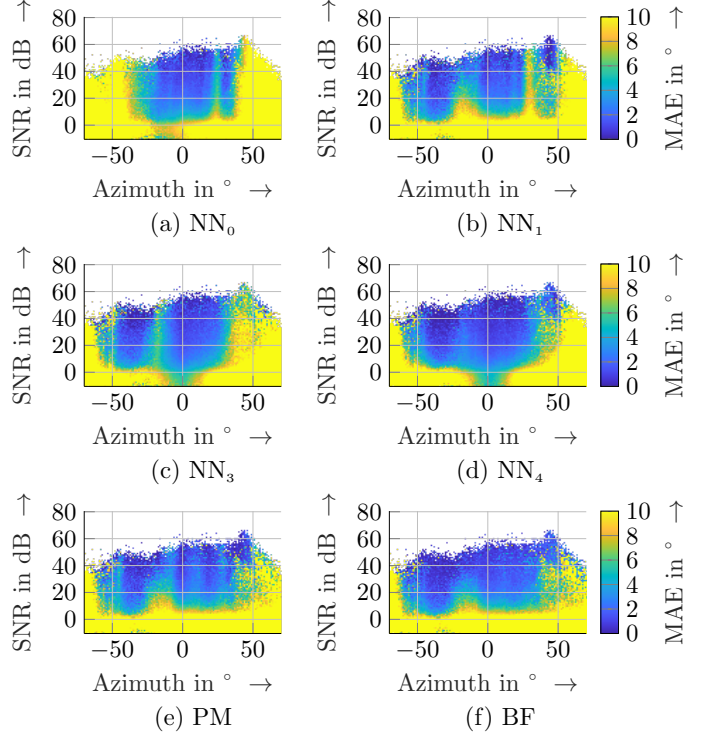


Figure 12: **MAE histograms:** Mean Absolute Error in dependency on azimuth and SNR for different DoA estimators. The NNs with enabled SNR selection (NN₁ & NN₃) achieved better MAE at bordering azimuth regions than their counterparts (NN₀ & NN₂).

Qualitative examples for DoA estimation and MTI are depicted in Fig. 13. It can be seen, that the color hue between DoA label and DoA prediction images mostly match and thus the predictions mostly resemble the DoA labels. The main difference between DoA label and prediction images comes from the selected visualization method, in which the color intensity of the prediction is scaled by the underlying RD-map SNR value. This visualization method helps to focus on regions with high SNR and reduces the distraction by low SNR regions, which mostly correspond to irrelevant noise regions. When comparing the non SNR color coded predictions from NN and PM (right hand side images in the Figure) for instance, it is obvious, that the NN predictions appear to be less noisy in low SNR regions. However, as stated above, DOA predictions in low SNR regions are in general of little interest.

As our framework provides full label density for a camera image with roughly 185.000 pixels per single frame, the DoA NN starts to provide reasonable estimations even after only 10 training steps/images. As the dataset consists of 72000 frames, we found this to be a very efficient usage of data.

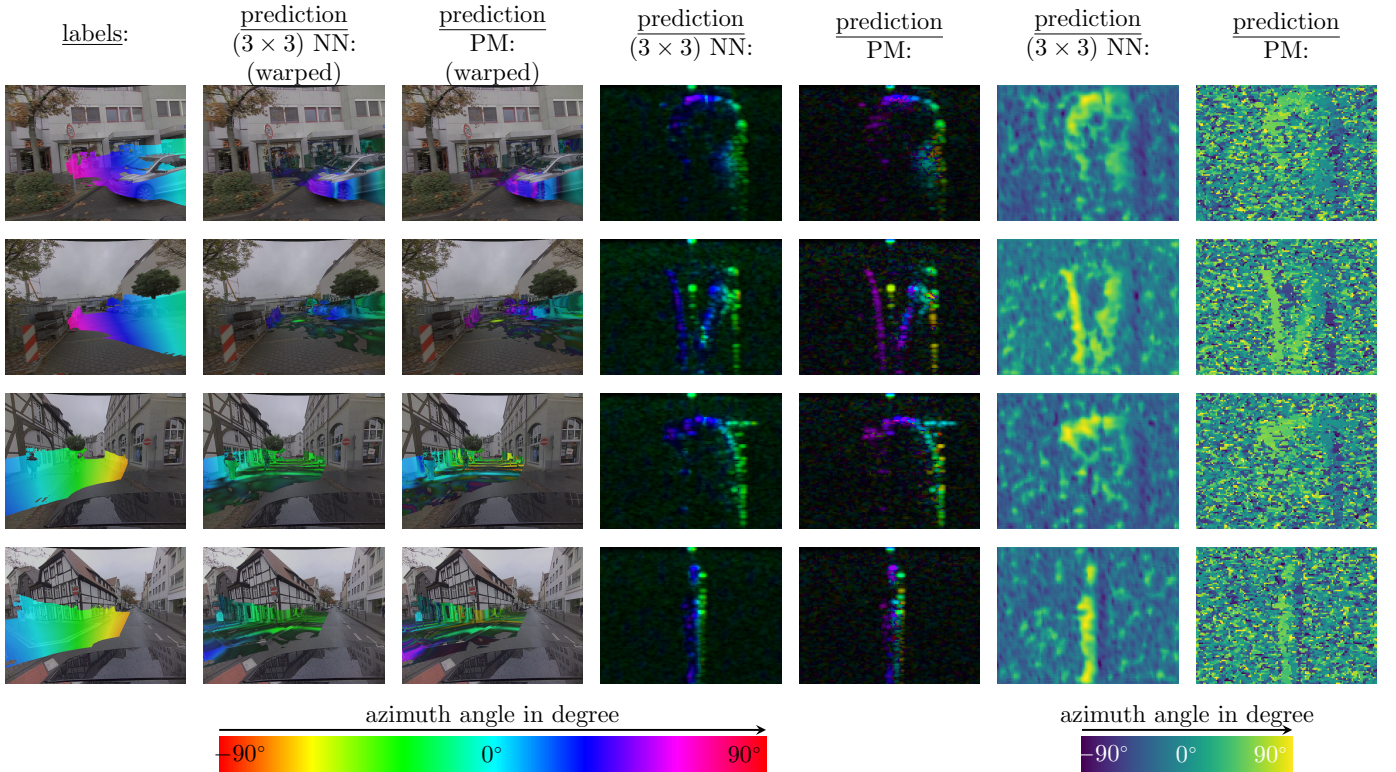


Figure 13: **Qualitative results of direction-of-arrival estimation on test data examples.** From left-to-right: RGB image with color coded azimuth labels, RGB image with color coded azimuth predictions from NN, RGB image with color coded azimuth predictions from PM, RD-map with color coded azimuth prediction from NN, RD-map with color coded azimuth prediction from PM, azimuth predictions from NN in RD-grid, azimuth predictions from PM in RD-grid.

VIII. CONCLUSIONS

In this paper we developed an algorithm for warping FMCW radar data into camera image, enabling us to utilize high-quality labels generated from camera and lidar reference sensors for the supervised training of NNs operating on radar spectra as input. An important factor for a reliable warping operation is the development of a novel scene flow estimation algorithm utilizing radar, camera and lidar data that is proposed in this paper. The automatic labeling method enabled us to densely label high-resolution spectrum level radar data from real world driving scenarios.

The validity of the approach was verified by evaluating the performance of NNs for DoA estimation, which achieved comparable or superior performance to classical DoA estimation methods. As the warping operation projects the spectrum level radar data and NN predictions into camera image, the results are easy to interpret for humans and allow qualitative assessment.

Furthermore, we could show that the use of the pixel vicinity through the NNs receptive field when performing DoA inference was beneficial. Although the general approach of utilizing pixel vicinity for better estimation is known from, e.g., classical image processing, to the best of our knowledge, we are the first to utilize this in radar based DoA estimation. The development of this automatic labeling framework was key contribution for this observation,

REFERENCES

- [1] R. Haeb-Umbach, J. Heymann, L. Drude, S. Watanabe, M. Delcroix, and T. Nakatani, "Far-field automatic speech recognition," *Proceedings of the IEEE*, vol. 109, no. 2, p. 124–148, 2021.
- [2] D. Cireşan, U. Meier, J. Masci, and J. Schmidhuber, "Multi-column deep neural network for traffic sign classification," in *Proceedings of the International Joint Conference on Neural Networks (IJCNN)*, 2011.
- [3] X. Zhu, D. Tuia, L. Mou, G. Xia, L. Zhang, F. Xu, and F. Fraundorfer, "Deep learning in remote sensing: A comprehensive review and list of resources," *IEEE Geoscience and Remote Sensing Magazine*, vol. 5, 2017.
- [4] A. Zhang, F. E. Nowruzi, and R. Laganieri, "Raddet: Range-azimuth-doppler based radar object detection for dynamic road users," 2021.
- [5] S. Schult, "Generierung von synthetischen 77-ghz-radarrohsignalen aus lidar- & kameradaten für die evaluation von objektklassifikatoren in automotiv-radaranwendungen," Master's thesis, Masters Thesis, University of Applied Science Gelsenkirchen Bocholt Recklinghausen, 2018.
- [6] B. Major, D. Fontijne, A. Ansari, R. T. Sukhvasi, R. Gowaikar, M. Hamilton, S. Lee, S. Grzechnik, and S. Subramanian, "Vehicle detection with automotive radar using deep learning on range-azimuth-doppler tensors," in *IEEE/CVF International Conference on Computer Vision Workshop (ICCVW)*, 2019.
- [7] T.-Y. Lim, "Radar and camera early fusion for vehicle detection in advanced driver assistance systems," in *The Thirty-third Annual Conference on Neural Information Processing Systems, NeurIPS 2019, Vancouver, Canada, Dec. 8-14, 2019*, 2019.
- [8] A. Ouaknine, A. Newson, J. Rebut, F. Tupin, and P. Pérez, "Carrada dataset: Camera and automotive radar with range-angle-doppler annotations," 2020.
- [9] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, *Vision meets Robotics: The KITTI Dataset*. International Journal of Robotics Research (IJRR), 2012.

- [10] 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)*, 2012.
- [11] M. Menze and A. Geiger, "Object scene flow for autonomous vehicles," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [12] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liong, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nusScenes: A multimodal dataset for autonomous driving," March 2019, [cs.LG].
- [13] M. Meyer and G. Kuschik, "Automotive radar dataset for deep learning based 3d object detection," *2019 16th European Radar Conference (EuRAD)*, pp. 129–132, 2019.
- [14] M. Cordts, M. Omran, S. Ramos, T. Rehfeld, M. Enzweiler, R. Benenson, U. Franke, S. Roth, and B. Schiele, "The cityscapes dataset for semantic urban scene understanding," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [15] G. Neuhold, T. Ollmann, S. Bulò, and P. Kotschieder, "The mapillary vistas dataset for semantic understanding of street scenes," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [16] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. Lopez, *The SYNTHIA Dataset: A large collection of synthetic images for semantic segmentation of urban scenes*. in CVPR, 2016.
- [17] A. Gaidon, Q. Wang, Y. Cabon, and E. Vig, *Virtual worlds as proxy for multi-object tracking analysis*. in CVPR, 2016.
- [18] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proceedings of Machine Learning Research*, vol. 78. PMLR, 13–15 Nov 2017, pp. 1–16. [Online]. Available: <http://proceedings.mlr.press/v78/dosovitskiy17a.html>
- [19] M. Bühren and B. Yang, "Automotive radar target list simulation based on reflection center representation of objects," in *Proc. Intern. Workshop on Intelligent Transportation (WIT)*, 2006, pp. 161–166.
- [20] A. P. Sligar, *Machine Learning-Based Radar Perception for Autonomous Vehicles Using Full Physics Simulation*. IEEE Access, 2020.
- [21] Z. Shuqing, *Performance Evaluation of Automotive Radars Using Carrier-Phase Differential GPS*. in IEEE Transactions Instrumentation and Measurement, 2010.
- [22] M. Zhao, T. Li, M. Alsheikh, Y. Tian, H. Zhao, A. Torralba, and D. Katabi, "Through-wall human pose estimation using radio signals," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [23] L. Wang, C. Anklam, and F. Baumgaertner, "Labelling of data for classification with automotive radar sensors," in *automotive Forum of European Microwave Integrated Circuit Conference (EuMIC)*, 2019.
- [24] L. Wang, T. Chen, C. Anklam, and B. Goldluecke, "High dimensional frustum pointnet for 3d object detection from camera, lidar, and radar," in *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, 2020.
- [25] V. Stabilito, "Phase-comparison monopulse radar," *United States US US3167765A*, 1961.
- [26] H. Krim and M. Viberg, "Two decades of array signal processing research," *IEEE Signal Processing Magazine*, vol. 13, 1996.
- [27] Y. L. Sit, M. Agatonovic, and T. Zwick, "Neural network based direction of arrival estimation for a mimo ofdm radar," *2012 9th European Radar Conference*, pp. 298–301, 2012.
- [28] M. Gardill, J. Fuchs, C. Frank, and R. Weigel, "A multi-layer perceptron applied to number of target indication for direction-of-arrival estimation in automotive radar sensors," in *Proceedings of the IEEE 28th International Workshop on Machine Learning for Signal Processing (MLSP)*, 2018.
- [29] J. Fuchs, R. Weigel, and M. Gardill, "Single-snapshot direction-of-arrival estimation of multiple targets using a multi-layer perceptron," in *Proceedings of the 2019 IEEE MTT-S International Conference on Microwaves for Intelligent Mobility (ICMIM)*, 2019.
- [30] —, "Model Order Estimation using a Multi-Layer Perceptron for Direction-of-Arrival Estimation in Automotive Radar Sensors," in *Proceedings of the 2019 IEEE Topical Conference on Wireless Sensors and Sensor Networks (WiSNet)*, 2019.
- [31] M.-Y. You, A.-N. Lu, Y.-X. Ye, K. Huang, and B. Jiang, "A review on machine learning-based radio direction finding," *Mathematical Problems in Engineering*, vol. 2020, pp. 1–9, 08 2020.
- [32] J. Levinson and S. Thrun, *Unsupervised Calibration for Multi-beam Lasers*. in Experimental Robotics, 2014.
- [33] U. Kuehnau, "Sensor development for autonomous driving," *Automotive Radar Sensors for Semi-Automatic and Autonomous Driving and Parking Systems*, IWPC Wolfsburg, Tech. Rep., 2017.
- [34] G. Xu and Z. Zhang, *Epipolar Geometry in Stereo, Motion, and Object Recognition: A Unified Approach*. USA: Kluwer Academic Publishers, 1996.
- [35] W. Ma, S. Wang, R. Hu, Y. Xiong, and R. Urtasun, "Deep rigid instance scene flow," in *[cs. CV]*, April 2019.
- [36] Z. Yin, T. Darrell, and F. Yu, "Hierarchical discrete distribution decomposition for match density estimation," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [37] Tensorflow, "Detection model zoo," 2020, gitHub repository. [Online]. Available: https://github.com/tensorflow/models/blob/master/research/object_detection/g3doc/detection_model_zoo.md
- [38] J. Diebel and S. Thrun, "An application of Markov random fields to range sensing," in *proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2006.
- [39] A. Harrison and P. Newman, "Image and sparse laser fusion for dense scene reconstruction," in *Proceedings of the Field and Service Robotics*, 2009.
- [40] C. Schröder, *System Design of an Array Antenna Radar with a Rapid Chirp Waveform*, ser. Berichte aus der Elektrotechnik. Shaker Verlag, 2013. [Online]. Available: <https://books.google.de/books?id=TyKTngEACAAJ>
- [41] M. Ester, H. Kriegel, J. Sander, and X. Xu, "A density-based algorithm for discovering clusters in large spatial databases with noise," in *Proceedings of the International Conference on Knowledge Discovery and Data (KDD)*, 1996.
- [42] M. M. and, and H. Wallentowitz. *Dynamik der Kraftfahrzeuge*: Springer, 2004.
- [43] T. Lindeberg, "Scale-space theory: A basic tool for analysing structures at different scales," *Journal of Applied Statistics*, vol. 21, no. 2, pp. 224–270, 1994.
- [44] Velodyne LiDAR, "Ultra puck," *datasheet*, vol. 27, July 2020. [Online]. Available: http://www.mapix.com/wp-content/uploads/2018/07/63-9378_Rev-D_ULTRA-Puck_VLP-32C_Datasheet_Web.pdf
- [45] FirstSensor, "Blue next network digital camera e4p," *datasheet*. [Online]. Available: https://www.first-sensor.com/cms/upload/datasheets/DS_Blue_Next_DC3C-1-E4P.pdf
- [46] GeneSys Elektronik GmbH, "Adma," 2020, *datasheet*,. [Online]. Available: https://www.genesys-offenburg.de/index.php?eID=tx_securedownloads&p=110&u=0&g=0&t=1595915060&hash=ccb5d292bcfc9f7f7492e7b888d0a8d6ca71736f&file=fileadmin/user_upload/ProdDescr_ADMA_rel_05.2019.pdf
- [47] P. Merriaux, Y. Dupuis, R. Boutteau, P. Vasseur, and X. Savatier, "Lidar point clouds correction acquired from a moving car based on can-bus data," June 2017, [cs.RO].
- [48] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga, A. Desmaison, A. Kopf, E. Yang, Z. DeVito, M. Raison, A. Tejani, S. Chilamkurthy, B. Steiner, L. Fang, J. Bai, and S. Chintala, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems* 32, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Aché-Buc, E. Fox, and R. Garnett, Eds. Curran Associates, Inc., 2019, pp. 8024–8035. [Online]. Available: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>
- [49] R. Durrett, *Probability: Theory and Examples*, 5th ed., ser. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2019.
- [50] L. Mesow, "Multisensorielle datensimulation im fahrzeugumfeld für die bewertung von sensorfusionsalgorithmen," Ph.D. dissertation, Technical University Chemnitz, 2006.
- [51] C. Grimm, T. Breddermann, R. Farhoud, T. Fei, E. Warsitz, and R. Haeb-Umbach, "Hypothesis test for the detection of moving targets in automotive radar," in *2017 IEEE International Conference on Microwaves, Antennas, Communications and Electronic Systems (COMCAS)*, 2017, pp. 1–6.
- [52] C. Grimm, R. Farhoud, T. Fei, E. Warsitz, and R. Haeb-Umbach, "Detection of moving targets in automotive radar with distorted ego-velocity information," in *2017 IEEE Microwaves, Radar and Remote Sensing Symposium (MRRS)*, 2017, pp. 111–116.




Christopher Grimm (Graduate Student Member, IEEE) received the B. Eng. degree in automotive engineering from University of Applied Science and Arts Dortmund, Germany, in 2013, and the M. Sc. in mechanical engineering from Technical University of Braunschweig, Germany, in 2015. Since 2016 he is pursuing the Dr.-Ing. degree in electrical engineering with the Department of Communications Engineering, University of Paderborn, Germany and HELLA GmbH & Co. KGaA, Lippstadt, Germany. Since 2020, he is working as a development engineer

at HELLA GmbH & Co. KGaA where he is responsible for the development of artificial intelligence based radar signal processing algorithms. His research interests are automotive radar signal processing, machine learning and computer vision.



Tobias Breddermann Tobias Breddermann received the Dipl.-Ing. degree in electrical engineering and the Dr.-Ing. degree from RWTH Aachen University, Aachen, Germany, in 2007 and 2013, respectively. From 2007 to 2013, he was with the Institute of Communication Systems and Data Processing at RWTH Aachen University. In 2013, he joined HELLA GmbH & Co. KGaA. His research interests include radar signal processing, advanced DoA estimation and target separation.



Tai Fei (Senior Member, IEEE) received the B.Eng. degree in telecommunication engineering from Shanghai Maritime University, Shanghai, China, in 2005 and the Dipl.-Ing. and Dr.-Ing. degrees in electrical engineering and information technology from Darmstadt University of Technology (TUD), Darmstadt, Germany, in 2009 and 2014, respectively. From 2009 to 2012, he worked as a Research Associate with the Institute of Water-Acoustics, Sonar-Engineering and Signal-Theory at Hochschule Bremen, Bremen, Germany, in collaboration with the

Signal Processing Group at TUD, Darmstadt, Germany, where his research interest was the detection and classification of underwater mines in sonar imagery. From 2013 to 2014, he worked as a Research Associate with the Center for Marine Environmental Sciences at University of Bremen, Bremen, Germany. Since 2014, has been working as a development engineer at HELLA GmbH & Co. KGaA, Lippstadt, Germany, where he is mainly responsible for the development of reliable signal processing algorithms for automotive radar systems.



Reinhold Haeb-Umbach (Fellow, IEEE) is a professor of Communications Engineering at Paderborn University, Germany. He holds a Dr.-Ing. degree from RWTH Aachen University, and has a background both in an industrial and academic research environment. His main research interests are in the fields of statistical signal processing and machine learning, mainly with applications to speech and audio processing, but also to other fields, such as radar signal processing. He has authored more than 200 scientific publications, and he is both a fellow of

the IEEE and of the International Speech Communication Association (ISCA).



Ernst Warsitz received the Dipl.-Ing. and Dr.-Ing. degrees in electrical engineering from Paderborn University, Paderborn, Germany, in 2000 and 2008, respectively. He joined the Department of Communications Engineering of the University of Paderborn in 2001 as a Research Staff Member, where he was involved in several projects relating to single- and multi-channel speech processing and automated speech recognition. From 2007 he worked as a development engineer at HELLA GmbH & Co. KGaA, Lippstadt, Germany, in the field of signal processing

algorithms for automotive radar systems. He is currently the head of the Radar Technology Department at HELLA GmbH & Co. KGaA, Lippstadt, Germany



Ridha Farhoud Ridha Farhoud received the Dipl.-Ing. degree in electrical engineering and the Dr.-Ing. degree from the Leipzig University in Hannover in 1997 and 2009, respectively. From 1998 to 2010, he worked as research engineer at the „Institut für Informationsverarbeitung“ (information processing) in Hannover. He was involved in several projects related to Synthetic Aperture Radar, Image Interpretation and low bit rate video coding. In 2010, he joined Hella KGaA Hueck & Co. He is currently a member of the head of the Radar Signal Processing

Department at HELLA GmbH & Co. KGaA, Lippstadt, Germany