

# Tackling Clutter in Radar Data – Label Generation and Detection Using PointNet++

Johannes Kopp<sup>1</sup>, Dominik Kellner<sup>2</sup>, Aldi Piroli<sup>1</sup> and Klaus Dietmayer<sup>1</sup>

**Abstract**—Radar sensors employed for environment perception, e.g. in autonomous vehicles, output a lot of unwanted clutter. These points, for which no corresponding real objects exist, are a major source of errors in following processing steps like object detection or tracking. We therefore present two novel neural network setups for identifying clutter. The input data, network architectures and training configuration are adjusted specifically for this task. Special attention is paid to the downsampling of point clouds composed of multiple sensor scans. In an extensive evaluation, the new setups display substantially better performance than existing approaches. Because there is no suitable public data set in which clutter is annotated, we design a method to automatically generate the respective labels. By applying it to existing data with object annotations and releasing its code, we effectively create the first freely available radar clutter data set representing real-world driving scenarios. Code and instructions are accessible at [www.github.com/kopp-j/clutter-ds](http://www.github.com/kopp-j/clutter-ds).

## I. INTRODUCTION

Radar sensors play an important role in the environment perception of modern autonomous robot systems. This is especially true for outdoor applications that require a large field of view of 100 m or more. Here, radar stands out against other sensor technologies like camera and lidar due to its long range and robustness to adverse weather conditions such as rain or fog [1]. Moreover, it allows for a direct measurement of the velocities of objects. For these reasons, research is driven particularly by the automotive sector, where radar sensors are vital to the development of new driver assistance systems and a continuously rising degree of automation.

Automotive radar sensors perceive their surroundings by emitting electromagnetic waves and analyzing the reflections from obstacles in the propagation path. This way, the position of objects and their velocity in radial direction relative to the sensor are estimated. At the end of each “scan”, a list of such detection points, also called targets, is output. What is problematic, however, is that this includes many detections that do not actually match any real object in the scene. This is what is called clutter<sup>3</sup>. Clutter is not limited to (seemingly) random noise. Rather, a large part of it is linked to objects located elsewhere according to certain rules. Causes for this are e.g. indirect propagation paths or erroneous ambiguity resolution by the sensor’s signal processing. Fig. 1 visualizes some of the corresponding effects.

<sup>1</sup>Institute of Measurement, Control and Microtechnology, Ulm University, Albert-Einstein-Allee 41, 89081 Ulm, Germany {firstname}.{lastname}@uni-ulm.de

<sup>2</sup>BMW AG, Petuelring 130, 80809 Munich, Germany dominik.m.kellner@bmw.de

<sup>3</sup>Some authors refer to the same detections or a subset of them as ghosts, anomalies, artifacts or noise.

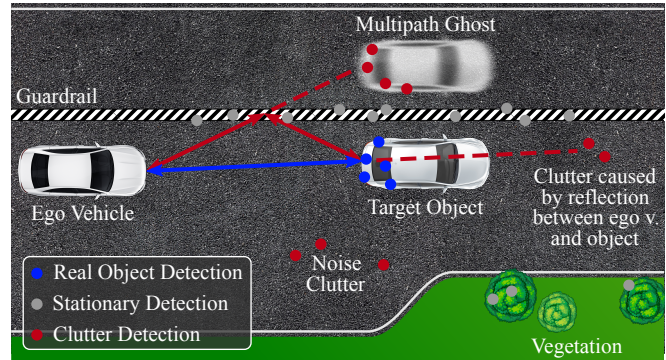


Fig. 1: Visualization of exemplary causes of clutter. Arrows indicate propagation paths. Multipath ghosts are clutter resulting from specular reflection at objects like guardrails or concrete walls, which act like mirrors for radar waves.

Clutter linked to real objects behaves very similar to detections that correctly represent road users. This poses a substantial challenge to environment perception methods like object detection or tracking. Thus, it is desirable to identify and filter the problematic points previous to such steps. In this work, we present neural network setups realizing this functionality. Our main contributions are the following:

- We describe a method for automatic label generation for an existing data set whose annotations do not yet incorporate clutter. The method’s code is published.
- Building on a setup for the semantic segmentation of radar data, we design two novel network variants specifically tailored to the task of detecting clutter.
- We show how it is possible to generate a prediction for each relevant detection in a point cloud while keeping the number of processed points constant. To this end, we introduce accumulation-aware downsampling.

## II. RELATED WORK

In recent years, clutter in automotive radar data has been recognized as an omnipresent and increasingly important issue [2], [3]. Causes for its occurrence and how they are reflected in a sensor’s output are being studied e.g. in [4]–[7]. [4] describes and models a wide range of different effects and locates their origin in the processing chain. The most frequent and problematic propagation paths leading to clutter are analyzed in even more detail in [5]. [6] and [7] focus exclusively on specular reflection. Both contain a geometric model of the phenomenon. In addition to that, [6] describes

the effect on the level of electromagnetic signals, while the authors of [7] conduct practical experiments for verification.

Regarding the detection of clutter in automotive radar data, several different approaches exist. It must be noted, though, that many authors simplify the task by considering only a subset of possible clutter or restricting the variety of scenarios in their data. For example, in both [8] and [9], model-based algorithms are used to detect multipath ghosts only. The authors of [10] and [11] employ PointNet++ to distinguish between detections stemming from real objects, different types of multipath ghosts and other points (among further configurations). Their data is recorded with experimental radar sensors in a controlled environment. The ego vehicle is always stationary and, for the most part, there is only a single moving object. Other works limited to certain types of clutter are [12], where PointNet++ is extended by a new grouping mechanism, and [13]. A wider scope is covered e.g. in [14]. Here, clutter of any kind is searched using an adapted PointNet architecture. The authors of [15] compare a random forest, a convolutional neural network and PointNet++ regarding their effectiveness for clutter detection. But they restrict the evaluation to the area directly in front of the sensor. Lastly, we present a rule-based approach for identifying clutter in [5].

The varying scopes and different definitions of erroneous detections in literature can mostly be explained by the fact that up until recently, there existed no public radar clutter data set. All authors used their own proprietary data. Even now, the only available data set [11] is limited to multipath ghosts and does not adequately represent road traffic, as outlined above. These problems are solved by the label generation method we present in the next section. Comparable approaches have previously been realized in [14] and [13]. However, both of them suffer from major drawbacks. The former requires a 3D radar sensor, something still rather uncommon, and the availability of additional lidar data for label generation. The second approach relies not only on lidar points but also existing annotations regarding the positions and dimensions of objects. Neither of the two takes velocity information into account. On top of that, there is no public code or a description sufficiently detailed for reimplementing.

### III. LABEL GENERATION AND DATA SET

As mentioned, there is no publicly available radar data set in which all types of clutter detections are annotated. For this reason, we devise a method to automatically generate the necessary labels for the public RadarScenes data set [16]. The approach is based solely on the positions and velocities of detections and the already present semantic annotations for moving objects. Additional modalities like lidar data are not needed (in contrast to [13], [14]). The method can thus be easily applied to any other data set with similar annotations. Its code is available at [www.github.com/kopp-j/clutter-ds](http://www.github.com/kopp-j/clutter-ds). By providing a ready-to-run script, we effectively create a public radar clutter data set that may be adopted by other authors.

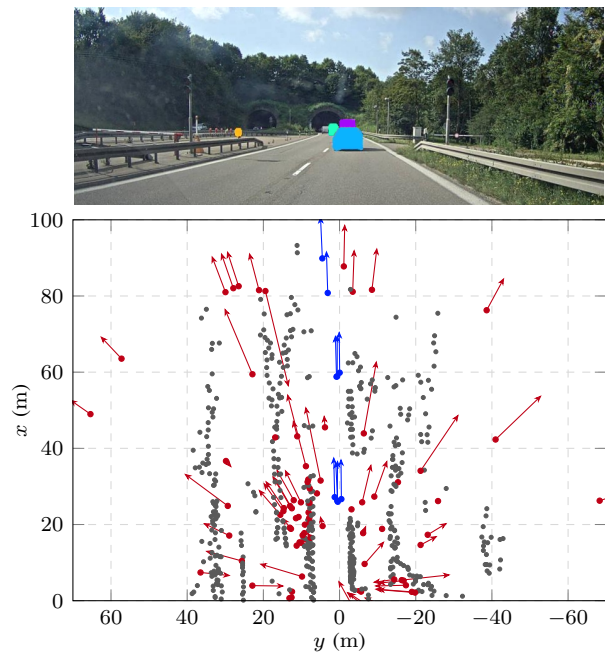


Fig. 2: Example of radar scans (one of each sensor) containing clutter. Points mark the positions of detections relative to the ego vehicle, arrows visualize their velocity over ground. The generated labels are indicated by color: detections annotated as *moving object* are blue ●, *clutter* red ● and *stationary* detections gray ●. Colored patches in the camera image cover nearby objects for legal reasons and have no meaning [17].

RadarScenes is a large automotive radar data set. It contains about 4.3 h of manually annotated recordings. These include urban scenarios, rural routes and sequences where the ego vehicle is stationary. Four 2D 77 GHz radar sensors are used, producing a total of 119 M detections. Each sensor covers an area of  $\pm 60^\circ$  horizontal angle and 100 m range and faces either to the front or side of the test vehicle. The data set’s annotations focus on the positions and types of moving objects. Each detection whose position matches one of those objects is labeled with the respective class, e.g. *car* or *pedestrian*. All other detections are marked as background. This includes those with slight measurement errors, those stemming from stationary objects and all types of clutter.

For label generation, we define three new classes to be distinguished: *moving object*, *stationary* and *clutter*. Radar sensors are rather inaccurate (compared, e.g., to lidar), especially regarding the azimuth angle. We thus tolerate ordinary measurement errors within reasonable bounds when deciding which detections correspond to real *moving objects*. This is different from the original annotations of RadarScenes, where a detection must lie perfectly inside an object’s bounding box to be assigned to the respective class. The label *stationary* identifies detections stemming from nonmoving objects. The remaining detections are *clutter*.

Automatic label generation works as follows. All detections originally labeled with an object class are marked as *moving object*. To introduce the appropriate tolerances for

TABLE I: Class distribution of the relabeled data set

Class	Moving Object	Clutter	Stationary
Ratio	3.35 %	5.57 %	91.08 %

measurement errors, background detections in close proximity to one of those points are searched and assigned to the class as well. This avoids the need for knowing explicit object bounding boxes. The size of each tolerance area depends on the angle of view. The maximum difference from the object detection’s position is set to 0.3m in range and between 2° and 4° in azimuth, linearly increasing for higher view angles. This is because the accuracy of azimuth angles measured by the sensor degrades as the angle increases.

The remaining background detections must be split into *stationary* and *clutter*. But for those for which low absolute velocities were measured, it is impossible to determine if they resulted from a clutter effect, or from actual reflections at stationary objects or the ground. However, detections with velocity below a certain threshold are usually processed by specialized methods like occupancy grid mapping [1], [18]–[20]. In these, clutter is much less of an issue than e.g. for object detection. It is therefore sufficient to label only background detections above such a threshold as *clutter*, and all others as *stationary*. We set the limit to  $|v_{\text{abs}}| \geq 0.5 \text{ m/s}$ . For RadarScenes, that is the velocity at which a detection is certain not to represent a stationary object. The value is composed of three times the standard deviation of the sensor’s velocity measurement  $3\sigma_{v_{\text{rel}}} = 0.3 \text{ m/s}$  and an estimate of the error introduced during ego motion compensation.

We verify the validity of the described approach both by visually inspecting the resulting labels and experimentally, see Section V-D. The class distribution of the final data set is given in Tab. I. Exemplary sensor scans and the respective generated labels are shown in Fig. 2.

#### IV. METHODS

##### A. Accumulation-aware Downsampling

When employing neural networks that process point clouds directly, like PointNet++, it is beneficial to ensure a constant number of input points. This is required e.g. for batch processing, which greatly speeds up training and enables the use of batch normalization [21]. However, sensors like lidar or radar typically produce point clouds of varying size. Fluctuations are especially pronounced for radar sensors because the number of measured detections heavily depends on the quantity and nature of objects in the environment and the sensor’s velocity. For example, in the RadarScenes data set a single scan usually contains anything between 20 and 330 detections [22]. Moreover, it is common practice to increase the density of radar point clouds by accumulating scans over a sliding time window of fixed length [23]. Since consecutive scans include similar numbers of detections, this amplifies fluctuations of the point clouds’ size even more. To provide the network with a set number of points regardless, resampling must be applied. When the sensors output too

few detections, an upsampling scheme must artificially add data, e.g. by duplication. In the opposite case, downsampling selects only the necessary number of points. This can be done, for example, by random sampling or via a rule like removing the detections with the lowest radar cross-section (RCS) [10], [14].

The problem when applying downsampling is that the network will not generate a prediction for some of the measured detections. Therefore, it is often omitted during inference, where a constant number of input points is not strictly necessary. This comes with a disadvantage for practical application, however: Outliers in the input point cloud’s size are reflected in an unusually high processing time. Furthermore, the different configuration for training and inference can potentially reduce performance for such exceptionally large point clouds.

We notice that the accumulation of multiple radar scans, which is a main reason for the heavy fluctuation of point cloud sizes, also opens the possibility for new downsampling methods resolving these issues. For use in a car or robot, the only predictions that are actually relevant are those for the most recent scan in the point cloud. Older detections have already been processed before and are included primarily to combat the sparsity of radar data. These considerations lead to the solution of the above problems: Downsampling must be restricted to remove only detections that belong to one of the old scans in the point cloud. That way, a prediction is produced for all relevant detections even if downsampling is applied. The number of processed points can thus be kept constant also during inference.

A second, even more efficient method that achieves the same goals is the following. Instead of employing an active downsampling step, detections can be stored in a fixed-size queue, i.e. a data structure where pushing new points in at the front results in the same number of oldest points getting dropped from the back. The size of this queue should be set to the desired number of points. Then, its use is equivalent to downsampling by removing as many of the oldest scans in the accumulation time window as needed. The oldest scan remaining after that usually does not fit into memory in full. Which of its detections are kept is determined by the order in which they were initially pushed into the queue. By applying this method, the amount of memory required for buffering sensor data is reduced to a minimum and the time needed for active downsampling is saved.

Despite the simplicity of the underlying ideas, we are, to the best of our knowledge, the first to explore accumulation-aware downsampling of radar point clouds and to propose approaches like this. The only other work in which the time of measurement is even considered during downsampling is [11]. There, discarding old detections is preferred, but some of the latest ones are likely to be removed as well.

##### B. Network Architectures and Setup

A basic network architecture that is well suited for processing point clouds is PointNet++ [24]. Fundamentally, it

recreates the encoder-decoder structure known from convolutional neural networks. In multiple so-called set abstraction (SA) levels, the number of considered points is progressively reduced by sampling. At the same time, the degree of abstraction and the area represented by each point are increased. Following this, the matching number of feature propagation (FP) levels interpolate features (and finally a classification) for all points in the original input point cloud.

The starting point of our work is the PointNet++ variant and setup described in [22], [25]. There, the original architecture is modified for the semantic segmentation of object classes in 2D radar data. Three SA and FP levels are used, the former of which employ multi-scale grouping (MSG). The input to the network are point clouds of 3072 detections. These are created by accumulating scans from four radar sensors over a time window of 500 ms. All details can be found in [22]. While it is possible to use this setup for clutter detection without major adjustments, as is done in [10] and [11], we decide to adapt it to the task.

We propose two improved variants with modified input data, network parameters and training setup. The first, in the following referred to as variant A, is designed for processing point clouds of only 1280 detections. This reduces the inference time and the required amount of memory. Variant B pushes that approach to the limit by working on single sensor scans, completely forgoing accumulation. An overview of the two network architectures is shown in Fig. 3.

Analyzing the effects that cause clutter reveals that for many of them the clutter detections and the objects involved in their formation may lie far apart (cf. Fig. 1). For this reason, we test different ways to incorporate a wider spatial context into the calculation of features, e.g. by utilizing an attention mechanism similar to [26]. We find that increasing the radii of neighborhood regions considered during MSG works best. For variant B, we make further modifications to compensate for the much lower density of points when processing single scans (only 144 detections on average). In particular, we prepend the feature preprocessing module proposed in [27], reduce the number of points sampled in the SA levels and add a third scale (i.e. query region) to grouping<sup>1</sup>.

Regarding the input to the networks, more changes beyond the reduction of point cloud sizes are made. In the original setup, each detection is described by its position, ego motion compensated velocity, RCS and timestamp. We choose to use the position in the Cartesian vehicle coordinate system (CS) at the time of recording the most recent scan in the point cloud. The timestamp is relative to that moment. Moreover, we add the detection’s position in the polar sensor CS, which is more useful in some cases, and the measuring sensor’s ID. When upsampling of the input point cloud is required, that is done by duplicating randomly selected detections. Replicas

<sup>1</sup>Details: In variant A, 1024, 512 and 256 points are sampled in the three SA levels, respectively. MSG uses two scales of radii (1, 3), (2, 5) and (4, 10) m. Layer sizes are the same as in [22]. Variant B samples only 256, 128 and 64 points. Grouping radii are (1, 3, 6), (2, 4, 8) and (3, 6, 12) m. The unit PointNet for preprocessing employs (64, 64, 32) channels.

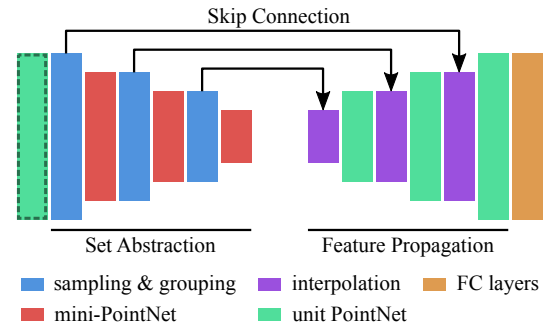


Fig. 3: Fundamental network architectures employed by the two proposed variants. The preprocessing module marked by the dashed line is used only in variant B.

created this way are excluded from loss calculation to avoid skewing gradients. For downsampling, either of the methods presented in Section IV-A can be used.

Finally, we also adjust the training setup. Input data is normalized via standardization. To prevent a distortion of following distance calculations, the variances of the Cartesian coordinates are averaged before scaling. The learning rate is varied between  $10^{-9}$  and  $10^{-3}$  according to a cyclical learning rate policy. And we employ focal loss [28] to increase weighting of detections that are especially hard to classify correctly and to better address the severe class imbalance. The remaining training parameters are as described in [22].

## V. EXPERIMENTS AND RESULTS

We evaluate the described network variants for clutter detection regarding their performance and inference time. For this, the RadarScenes data set together with the newly generated labels is used (see Section III). When a setup requires the accumulation of multiple radar scans to form a larger point cloud, only predictions for the most recent scan are analyzed, as motivated in Section IV-A. Points that are added by upsampling are ignored. To minimize the influence of chance on results, we choose five different initial seeds for random number generators and repeat training of each network five times, once per seed. The mean performance of those runs on the validation set is then reported. Stated inference times are the average time a model requires for processing one point cloud on an Nvidia RTX 2080 Ti GPU.

### A. Comparison with Other Approaches

To assess the effectiveness of our setups, we compare them with other approaches in literature. As baseline, we use the setup for the semantic segmentation of objects presented in [22]. It was the starting point for our work and is attuned specifically to RadarScenes. Except for adjustments to the input, the same network is employed for identifying multipath ghosts in [10]. We also recreate that setup and test it on our data. Finally, we evaluate the setup and best performing architecture from [12] (PointNet++ with multi-form grouping). Regarding our own networks, variant A, which processes accumulated point clouds, is tested in combination with both proposed accumulation-aware downsampling methods. For

TABLE II: Quantitative comparison of different network setups for clutter detection. Class-specific and mean precision (Prec.), recall and F1 score (F1) are given in %.

Setup	Accumulation-aware Downsampling	Moving Object			Prec.	Clutter		Stationary F1	Mean F1	Inference Time
		Prec.	Recall	F1		Recall	F1			
Baseline [22]	–	80.16	<b>94.67</b>	86.81	66.08	86.47	74.91	98.43	86.72	7.78 ms
[10]	–	67.96	92.55	78.36	56.91	81.81	67.12	97.30	80.93	6.56 ms
[12]	–	83.12	89.94	86.39	91.69	88.48	90.05	99.55	92.00	<b>4.11 ms</b>
Ours, variant A	Restriction to old points	<b>88.98</b>	93.84	<b>91.34</b>	93.11	<b>94.89</b>	<b>93.99</b>	<b>99.75</b>	<b>95.03</b>	7.19 ms
Ours, variant A	Fixed-size queue	88.36	94.13	91.15	92.72	94.76	93.73	99.73	94.87	7.19 ms
Ours, variant B	–	85.21	91.44	88.22	<b>93.25</b>	92.15	92.69	99.74	93.55	5.90 ms

the one that restricts downsampling to the removal of old detections, data is accumulated over 300 ms and removed points are selected by random sampling. When instead a fixed-size queue is employed, the detections of a scan are stored in such order that the slowest are dropped first.

We follow the original training setups as closely as possible. Changes are made only when they are necessary to adapt to the new data. Accordingly, the model from [22] is trained for 25 epochs, others for 20. Class weights for loss calculation are determined the same way for all setups except [10], which specifies its own formula to be applied. The weight of stationary detections is empirically set to  $w_s = 0.6$ . The other two classes are weighted proportionately to their frequencies  $f_o, f_c$  for the respective configuration, i.e. such that  $w_o \cdot f_o = w_c \cdot f_c$  and  $\sum_i w_i \cdot f_i = 1$ .

The results of the experiments are listed in Tab. II. As can be seen, employing any of our new network variants leads to significant improvements over the existing approaches. At the same time, the reduced number of input points results in a speed-up compared to the baseline setup. Restricting downsampling to remove only old detections performs slightly better than utilizing a fixed-size queue. However, the differences are small enough that the superior memory efficiency of the queue could still justify the tradeoff. Nonetheless, we always use the former method with variant A in the following sections. The setups that process sensor scans individually, i.e. [12] and our variant B, work remarkably well. Despite the extreme sparsity of the point clouds and the complete lack of temporal information, they perform much better than the baseline. Our setup has the higher accuracy, and still lowers the inference time by about 24% compared to [22].

Some exemplary data samples and the corresponding outputs of our models are visualized in Fig. 4.

### B. Accumulation-aware Downsampling

Next, we test our new accumulation-aware downsampling methods in isolation. We compare them with other ways to produce a prediction for every detection from the most recent scan in a point cloud. To this end, we repeat the evaluation of the baseline models trained in the previous section but exchange the employed method. The original setup uses a post-processing step: Points that were removed by downsampling receive a prediction by copying it from their nearest neighbor. Another alternative is to omit downsampling entirely.

The comparison in Tab. III shows that the post-processing step from [22] is not suited for application in clutter de-

tection. This is because, unlike detections stemming from real objects, clutter occurs not only in local groups but also as single detections scattered between other points. Copied predictions are therefore oftentimes incorrect. The omission of downsampling and the two accumulation-aware methods result in similar, vastly improved performance. The advantage of the latter, however, is that processed point clouds always contain the same number of detections. This reduces both the average inference time and its fluctuations. In the most extreme cases, up to about 10.5k detections are accumulated over the time window of 500 ms. For this number of points, inference takes 10.94 ms on average when downsampling is omitted. The accumulation-aware methods prevent such spikes. Furthermore, they can be applied already during training to adapt the model to the used scheme.

### C. Sensor-specific Models

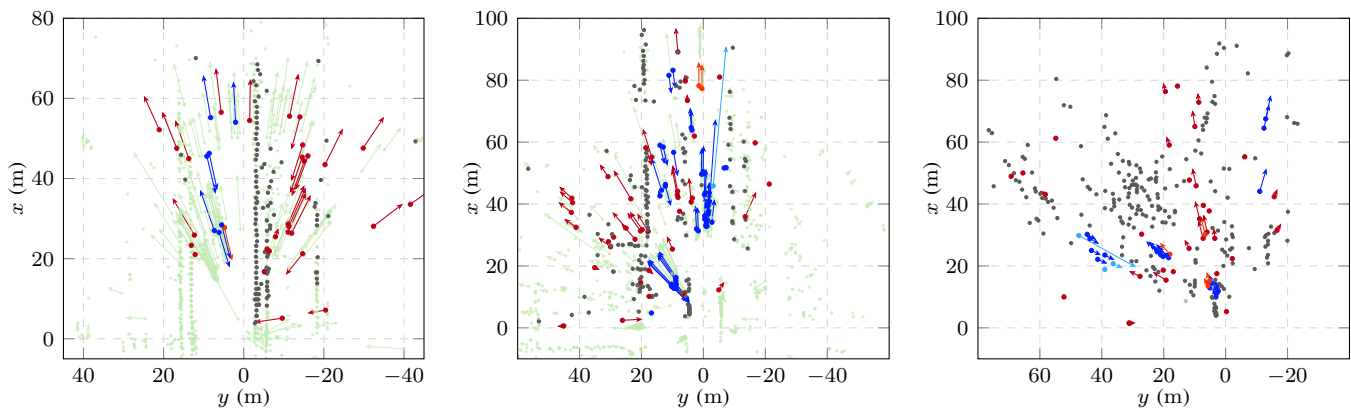
In the previous sections, data from all four available radar sensors is mixed. Thus, models must learn to generalize between the different mounting positions. We now investigate the effect of using a model that operates exclusively on data of a specific sensor. Experiments are conducted for the sensor facing to the front left since the view onto oncoming traffic leads to a lot of difficult situations with multiple moving objects. For the sensor-specific version of setup variant A, we use an extended accumulation time window of 1.1 s so that point clouds of similar size as before are obtained.

TABLE III: Comparison of ways to obtain predictions for all points in the latest scan. F1 scores of baseline models are given in %, inference time and its variance in ms and (ms)<sup>2</sup>.

Method	F1 Score		Inference Time	
	Clutter	Mean	Mean	Variance
Post-processing [22]	74.91	86.72	7.78	0.665
Omit downsampling	90.37	92.48	7.79	0.704
Remove only old points	90.62	<b>92.50</b>	<b>7.40</b>	<b>0.582</b>
Fixed-size queue	<b>90.69</b>	92.04	<b>7.40</b>	<b>0.582</b>

TABLE IV: F1 scores for data of the front left sensor (in %) achieved by general and sensor-specific models

Variant	Sensor-specific	Mov. Obj.	Clutter	Statio.	Mean
A		<b>92.50</b>	<b>92.58</b>	99.70	<b>94.93</b>
A	✓	91.81	91.68	<b>99.71</b>	94.40
B		89.63	90.50	99.67	93.27
B	✓	89.90	90.40	99.66	93.32



(a) Drive over a four-lane bridge. Specular reflections of cars at the guardrails cause a lot of clutter. Network variant A.

(b) Highly complex scenario in urban traffic. Buildings on both sides, trees on the median strip. Network variant A.

(c) Left turn at an intersection with refuge island. Variant B. Even the scattered object detections on the right are identified.

Fig. 4: Examples of point clouds and the respective network outputs. Colors indicate the ground truth and the correctness of predictions. Detections annotated as *moving object*, *clutter* or *stationary* are drawn in light or dark blue ●, red ● and gray ●, respectively. The darker version of a color is used when the prediction matches the label, while the lighter one marks falsely classified detections. Pale green points ● belong to old sensor scans, for which the network’s output is irrelevant.

The performance for data of the front left sensor of general and sensor-specific versions of our setups is reported in Tab. IV. Contrary to what one might expect, performance of variant B is barely increased by training specifically for the selected sensor. Using the sensor-specific version of variant A even leads to a loss of accuracy. This suggests that a short accumulation of data from multiple sensors with overlapping fields of view is preferable to treating sensors separately.

#### D. Transfer of Models To Other Data and Comparison with Rule-based Approach

In an earlier work, we recorded a small radar clutter data set with our own test vehicle. Details can be found in [5]. Since here the correctness of all annotations was manually checked, this data can be utilized to validate the proposed label generation method.

The best performing model of setup A from Section V-A, which was trained only on the relabeled RadarScenes data set, achieves a mean F1 score of 93.72% for the new data. This is especially remarkable considering that just a single sensor with a different mounting position and orientation than before is used. The only adjustment to the network setup that is needed is an extension of the accumulation time to 500 ms. Interestingly, performance can hardly be

further improved by fine-tuning the model on the new data, cf. Tab. V. Together with the impressive results right from the start, this shows that the method we devised for label generation is highly accurate. A model trained with the produced annotations learns to identify the characteristics of classes sufficiently well that generalization to data of a different sensor configuration is easily possible.

The evaluation on the second data set has an additional benefit. It enables a direct comparison of our new setups with the rule-based approach described in [5], see again Tab. V. Even though that algorithm has to make a prediction only for points with significant velocity, the neural networks exhibit much better performance. Their drawback, of course, is their comparatively high computational cost and innate lack of explainability.

## VI. CONCLUSION

In this work, we present a method to automatically generate ground-truth for a radar clutter data set. It requires only the basic sensor outputs and the object annotations of freely available data. The method’s code is published. This results in the first data set of realistic driving scenarios in which clutter is labeled that is accessible to all researchers. We then modify an existing setup of PointNet++ for the detection of clutter in automotive radar point clouds. The number and form of input points, the architecture and the training setup are all tailored specifically to the task. By introducing accumulation-aware downsampling, we further increase efficiency. Our best performing setup detects almost 95% of clutter while only rarely misclassifying other points. This is a significant improvement over existing machine-learning approaches and a comparable rule-based method. An alternative setup variant reduces the inference time considerably and still manages to achieve similar performance.

TABLE V: Comparison of approaches on the data set from [5]. Models are pretrained on RadarScenes and then fine-tuned on the new data. Reported values are F1 scores measured using 4-fold cross-validation in %.

Approach	Nonclutter & Ambiguous	Clutter	Stationary	Mean
Ours, variant A	<b>85.66</b>	<b>96.18</b>	<b>99.63</b>	<b>93.82</b>
Ours, variant B	80.51	94.00	99.40	91.30
Rule-based [5]	44.36	87.78	–	66.07

## REFERENCES

- [1] F. Engels, P. Heidenreich, M. Wintermantel, L. Stacker, M. Al Kadi, and A. M. Zoubir, "Automotive Radar Signal Processing: Research Directions and Practical Challenges," *IEEE Journal of Selected Topics in Signal Processing*, 2021.
- [2] S. Sun, A. P. Petropulu, and H. V. Poor, "MIMO Radar for Advanced Driver-Assistance Systems and Autonomous Driving: Advantages and Challenges," *IEEE Signal Processing Magazine*, 2020.
- [3] J. Dickmann, J. Klappstein, M. Hahn, N. Appenrodt, H.-L. Bloecher, K. Werber, and A. Sailer, "Automotive Radar the Key Technology for Autonomous Driving: From Detection and Ranging to Environmental Understanding," in *2016 IEEE Radar Conference (RadarConf)*. IEEE, 2016.
- [4] M. F. Holder, C. Linnhoff, P. Rosenberger, C. Popp, and H. Winner, "Modeling and Simulation of Radar Sensor Artifacts for Virtual Testing of Autonomous Driving," *9. Tagung Automatisiertes Fahren*, 2019.
- [5] J. Kopp, D. Kellner, A. Piroli, and K. Dietmayer, "Fast Rule-Based Clutter Detection in Automotive Radar Data," in *2021 IEEE International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2021.
- [6] N. Scheiner, F. Kraus, F. Wei, B. Phan, F. Mannan, N. Appenrodt, W. Ritter, J. Dickmann, K. Dietmayer, B. Sick, and F. Heide, "Seeing Around Street Corners: Non-Line-of-Sight Detection and Tracking In-the-Wild Using Doppler Radar," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2020.
- [7] A. Kamann, P. Held, F. Perras, P. Zaumseil, T. Brandmeier, and U. T. Schwarz, "Automotive Radar Multipath Propagation in Uncertain Environments," in *2018 IEEE International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2018.
- [8] F. Roos, M. Sadeghi, J. Bechter, N. Appenrodt, J. Dickmann, and C. Waldschmidt, "Ghost Target Identification by Analysis of the Doppler Distribution in Automotive Scenarios," in *2017 International Radar Symposium (IRS)*. IEEE, 2017.
- [9] C. Liu, S. Liu, C. Zhang, Y. Huang, and H. Wang, "Multipath Propagation Analysis and Ghost Target Removal for FMCW Automotive Radars," in *2020 IET International Radar Conference (IRC)*. Institution of Engineering and Technology, 2020.
- [10] F. Kraus, N. Scheiner, W. Ritter, and K. Dietmayer, "Using Machine Learning to Detect Ghost Images in Automotive Radar," in *2020 IEEE International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2020.
- [11] —, "The Radar Ghost Dataset – An Evaluation of Ghost Objects in Automotive Radar Data," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021.
- [12] T. Griebel, D. Authaler, M. Horn, M. Henning, M. Buchholz, and K. Dietmayer, "Anomaly Detection in Radar Data Using PointNets," in *2021 IEEE International Conference on Intelligent Transportation Systems (ITSC)*. IEEE, 2021.
- [13] L. Wang, S. Giebenhain, C. Anklam, and B. Goldluecke, "Radar Ghost Target Detection via Multimodal Transformers," *IEEE Robotics and Automation Letters*, 2021.
- [14] M. Chamseddine, J. Rambach, D. Stricker, and O. Wasenmüller, "Ghost Target Detection in 3D Radar Data using Point Cloud based Deep Neural Network," in *2020 IARP International Conference on Pattern Recognition (ICPR)*. IEEE, 2020.
- [15] Y. Jin, R. Prophet, A. Deligiannis, I. Weber, J.-C. Fuentes-Michel, and M. Vossiek, "Comparison of Different Approaches for Identification of Radar Ghost Detections in Automotive Scenarios," in *2021 IEEE Radar Conference (RadarConf)*. IEEE, 2021.
- [16] O. Schumann, M. Hahn, N. Scheiner, F. Weishaupt, J. F. Tilly, J. Dickmann, and C. Wöhler. (2021) RadarScenes: A Real-World Radar Point Cloud Data Set for Automotive Applications. [Online]. Available: <https://doi.org/10.5281/zenodo.4559821>
- [17] —, "RadarScenes: A Real-World Radar Point Cloud Data Set for Automotive Applications," *arXiv*, 2021. [Online]. Available: <http://arxiv.org/abs/2104.02493>
- [18] O. Schumann, J. Lombacher, M. Hahn, C. Wöhler, and J. Dickmann, "Scene Understanding With Automotive Radar," *IEEE Transactions on Intelligent Vehicles*, 2020.
- [19] M. Li, Z. Feng, M. Stolz, M. Kunert, R. Henze, and F. Küçükay, "High Resolution Radar-based Occupancy Grid Mapping and Free Space Detection," in *2018 International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS)*. SCITEPRESS, 2018.
- [20] R. Prophet, A. Deligiannis, J.-C. Fuentes-Michel, I. Weber, and M. Vossiek, "Semantic Segmentation on 3D Occupancy Grids for Automotive Radar," *IEEE Access*, 2020.
- [21] S. Ioffe and C. Szegedy, "Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift," *arXiv*, 2015. [Online]. Available: <http://arxiv.org/abs/1502.03167>
- [22] O. Schumann, "Machine Learning Applied to Radar Data: Classification and Semantic Instance Segmentation of Moving Road Users," Doctoral Thesis, TU Dortmund, 2021.
- [23] N. Scheiner, F. Kraus, N. Appenrodt, J. Dickmann, and B. Sick, "Object Detection for Automotive Radar Point Clouds – A Comparison," *AI Perspectives*, 2021.
- [24] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space," *Advances in Neural Information Processing Systems*, 2017.
- [25] O. Schumann, M. Hahn, J. Dickmann, and C. Wöhler, "Semantic Segmentation on Radar Point Clouds," in *2018 ISIF International Conference on Information Fusion (FUSION)*. IEEE, 2018.
- [26] X. Yan, C. Zheng, Z. Li, S. Wang, and S. Cui, "PointASNL: Robust Point Clouds Processing Using Nonlocal Neural Networks With Adaptive Sampling," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2020.
- [27] A. Cennamo, F. Kaestner, and A. Kummert, "Leveraging Radar Features to Improve Point Clouds Segmentation with Neural Networks," in *2020 International Conference on Engineering Applications of Neural Networks (EANN)*. Springer, Cham, 2020.
- [28] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollár, "Focal Loss for Dense Object Detection," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2017.