Making Radar Detections Safe for Autonomous Driving: A Review

Tim Brühl[1,2], Lukas Ewecker[1], Robin Schwager[1,2], Tin Stribor Sohn[1] and Sören Hohmann[2]
[1]*Dept. of Highly Automated Driving Systems, Dr. Ing. h.c. F. Porsche AG, Weissach, Germany*
[2]*Institute of Control Systems, Karlsruhe Institute of Technology, Germany*

Keywords: SOTIF, Radar Perception, Radar Signal Processing, Radar-Camera Fusion, Functional Safety.

Abstract: Radar sensors rank among the most common sensors used for highly automated driving functions due to their solid distance and velocity measurement capabilities and their robustness against adversarial environmental conditions. However, radar point clouds are noisy and must therefore be filtered. This work reviews current research with the aim to make radar detections usable for safe perception functions which require a guarantee for correctness of the measured environmental representation. The impact on radar errors on the distinct downstream tasks is explained. Besides, the term of safety for automated driving functions is illuminated under consideration of the current standards and state-of-the-art research interpreting these standards is presented. Furthermore, this work discusses safe radar signal processing and filtering, approaches to enrich radar data points by information fusion, e.g. from cameras and other radars, and development tools for safe radar-based perception functions. Finally, next steps on the way towards safety guarantees for radar sensors are identified.

1 INTRODUCTION

Autonomous driving and parking are two of the major emerging fields in the current development of the automotive industry. In contrast to assisting functions, the driver is not involved during autonomous functions which exceed SAE level 2 (SAE International., 2021). With this, the safety demands on elements of the function rise significantly. This applies for both planning and acting, but these two items will only work properly if the sensing element performs accurately. Critical situations must not be missed as falling back on the driver is not an option anymore. The system needs to handle a variety of tasks in the operational design domain, which includes unforeseen, probably even unimaginable situations. During development, it needs to be ensured that these situations are covered. Typical sensors which are used for environment perception are radar sensors, lidar sensors and cameras (Yeong et al., 2021), which are combined to sensor sets in case of autonomous vehicles.

Radar is an acronym for Radio Detection and Ranging, which provides an indication on its working principle. Main parts of a radar sensor are the Voltage-Controlled Oscillator (VCO) which generates an electromagnetic radio wave, which is transmitted by a set of antennas. These radio waves reflect in the world. The reflected waves are received by a second set of antennas and processed to determine attributes like position, signal power and relative velocity of the reflected obstacle. Initially, radar sensors were primarily applied in military, aviation and nautical fields. Meanwhile, millimeter-wave radar sensors also gained a considerable proportion in automotive sensor systems, starting as a sensor for assistance systems like adaptive cruise control and safety functions such as autonomous emergency braking, and now being an enabler for autonomous driving functions (Waldschmidt et al., 2021). Advantages of radar sensors are the compact form factor, due to high integration density of components, as well as the strong performance when it comes to measuring even high distances as well as velocities in a single scan (Steinbaeck et al., 2017). Common modern sensor systems achieve a range of over 200 m with an accuracy and resolution of 0.1 m, and a velocity accuracy of $0.05\,\frac{\mathrm{m}}{\mathrm{s}}$ (Aptiv PLC, 2023) (Robert Bosch GmbH, 2023). The capability of radar waves to pervade plastics as well as even thin coating layers provides the opportunity of integrating the sensor invisible behind the vehicle's bumpers. Additionally, the radar is the only sensor showing exceptional environmental and atmospheric robustness (Marti et al., 2019) Considering the complementary of sensors, a combination of radar sensors and camera sensors seems particularly promising (Zhou et al., 2022). The major drawbacks of radar

Brühl, T., Ewecker, L., Schwager, R., Sohn, T. and Hohmann, S.
Making Radar Detections Safe for Autonomous Driving: A Review.
DOI: 10.5220/0012630400003702

In *Proceedings of the 10th International Conference on Vehicle Technology and Intelligent Transport Systems (VEHITS 2024)*, pages 299-310
ISBN: 978-989-758-703-0; ISSN: 2184-495X

sensor systems like moderate orientation measuring capability or weak classification performance are addressed by the camera.

As lidar is an acronym for Light Detection and Ranging, these sensors share their working principle of transmission and reception of electromagnetic wave with the radar sensors. However, the wavelength of lidar sensors is close to the visible part of the spectrum and hence shorter compared to radar sensors. Elementary disparities in the physical properties of the measures come with this difference in wavelength. Lidar sensors exceed radar sensors in spatio-temporal consistency, meaning that data points are very congruent in two consecutive measures of the same scene. This is not given for radar sensors, as the points spread and scatter along an object (Bilik, 2023). Another asset is the high lateral and, thanks to multiple lidar scan layers, even elevational resolution of the lidar point cloud. Hence, lidar sensors are superior to radars in object detection and classification tasks. However, except for new approaches which try to integrate the Frequency-Modulated Continuous Wave (FMCW) technique to lidars (Sayyah et al., 2022), common lidar sensors are not capable of directly measuring the relative velocity per point. The achievable range of lidar sensors is slightly less compared to radars. In addition, lidar sensors are heavier and more cumbersome, which is another aspect for automotive use cases.

Most important for the considerations of this work is that, due to its working principle, lidar suffers from impairments in adversarial weather conditions like rain, snow and especially fog (Zang et al., 2019). Looking from the perspective of safe environment perception, that is why lidar does not qualify that well for this task in compared to the radar sensor. Contrary to radar sensors, lidars need a cleaning solution which often includes washing fluid. This work aims to show that radars have the potential to serve as sensors for safe environmental perception even in bad lighting and weather, if attention is paid to certain singularities during the signal processing.

2 STRUCTURE

The paper is structured as follows. At first, an overview of current radar tasks in the field of automated driving functions is given. Next, we show the current radar processing approaches and point out that safety considerations are a special recess in the bunch of radar processing methods. The term of safe radar detections is elucidated and current literature in this area is introduced. In this context, we highlight the role of the two most important safety standards for automotive applications. We look at fusion techniques with the potential to add information and reduce uncertainty of radar points, and we describe data sets for advancement of radar point treatment. Eventually, we outline the results of our review, draw conclusions and demonstrate further potential research activities.

3 LITERATURE REVIEW

The literature section review comprises the radar capabilities, a survey of radar signal processing, an analysis of error sources in radar sensors, an introduction into the concept of safe detection, a definition of critical radar points and development approaches to alleviate the criticality of radar points.

3.1 Radar Capabilities

This section gives an introduction into the different tasks in automated driving functions which a radar can handle by today. The four presented fields of application are depicted in figure 1.

3.1.1 Radar Localization

Currently, measurement data acquired by radar sensors is mainly used for object detection (Schumann et al., 2020)(Scheiner et al., 2021). However, new application fields for radar offer to the research community. The property that radar measurements have a point cloud as an output makes them suitable for localization in already seen environments via point cloud registration. In the past, several works dealt with the challenge to apply Simultaneous Localization and Mapping (SLAM) algorithms to radar data. (Hong et al., 2021) describe a SLAM algorithm to deal with bad weather and illumination conditions. A blob detector finds key points in a cartesian radar image and detections on moving vehicles are deleted by a graph-based outlier detection algorithm. The key points are tracked, and motion is compensated using a pose graph. After every measure, the minimally required number of features is completed by adding new features in case that some were lost during the last tracking cycle. The loop closure is done by comparing candidates in a Principal Component Analysis (PCA) and rejecting the least likely ones. Comparable performance to vision-based and lidar-based approaches was achieved. Similar work was done by (Schuster et al., 2016), who used an occupancy grid based key point extractor and constructed a graph, adding poses as well as odometry measures

(a) Matching radar point clouds (red dots) to map data (black structure) for localization purposes

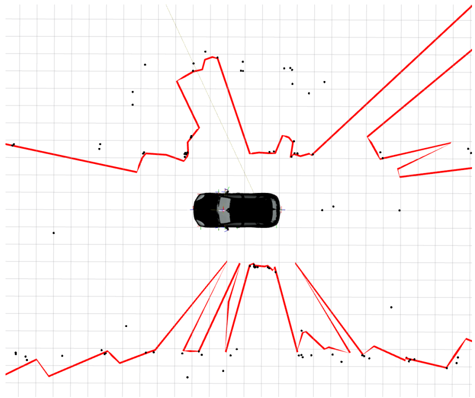

(b) Radar free space detection (red line) based on radar points (black dots)

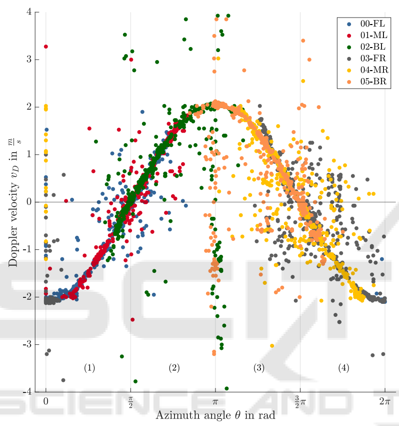

(c) Radar odometry, doppler velocity of radar points at constant vehicle speed of $v = 2.0\,\frac{\mathrm{m}}{\mathrm{s}}$

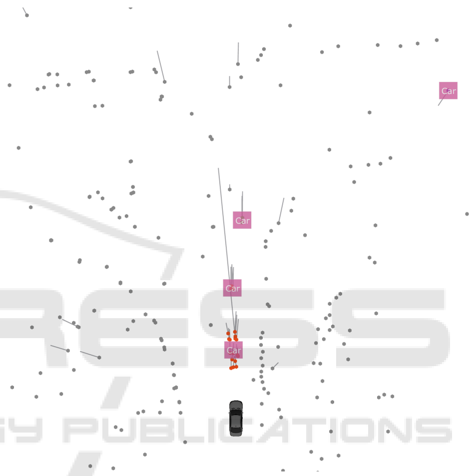

(d) Radar object detection and tracking (Schumann et al., 2021)

Figure 1: Different application fields for radar sensors.

of the vehicle. The outliers are detected with the Random Sample Consensus (RANSAC) algorithm in this work. The graph is optimized using Levenberg-Marquardt-optimization.

3.1.2 Radar Free Space Detection

Another field of frequent research activity is radar free space detection. This can also be considered as the inverse problem of the detection of static obstacles in the surrounding of the sensor. The Radar Cross Section (RCS) is proportional to the backscatter signal power density (Knott, 2012), meaning that a higher RCS increases the peak of the detection in the noise floor. Hence, given that the backscatter power is not constant, but normally distributed, the probability that the object is detected increases as well. The probability to receive a backscatter signal and the average signal power depends on the RCS of the object which is measured (Knott, 2012). A typical approach to find drivable space is to discretize the environment to occupancy grids, filling cells with reflections which feature reflections, and leaving out the ones without reflections (Xu et al., 2020)(Li et al., 2018). Due to the noisiness of the radar point cloud, filtering algorithms are applied to reduce the false positive rate of the boundary detection. (Popov et al., 2023) tackle the problem of free space detection by a deep neural network, whose input is a radar point cloud which is accumulated over a time of 0.5 s, ego-motion compensated, and projected in a 2D bird's-eye-view format. Human-annotated bounding boxes based on lidar detections are used as the training data. The network features not only a free space segmentation head, but also a bounding box detection and a classification.

3.1.3 Radar Odometry

Recent approaches showed that radar sensors can also be used for motion estimation due to their ability to precisely determine the Doppler velocity. Assuming that most of the perceived targets are stationary, the reflections from these points can be used, together with the angular information, to estimate the longitudinal and lateral velocity as well as the yaw rate of a vehicle. (Kellner et al., 2014) showed the feasibility of this approach and investigated the precision improvement when using multiple radars. The geometric coherence between the measurement is derived and based on that, the vehicle's motion state is estimated, using the Ackermann conditions. Radar odometry can be tackled on a second way, which is closer to the SLAM approach described before. Point clouds can be matched and the vehicle's motion is the first derivative of the pose transform, the transformation between two successive point clouds. The work of Adolfsson et al. proposes a conservative filtering approach which keeps only a set of the strongest radar detections with constant size. Additionally, it is assumed that true positives represent objects with a surface. Thus, a surface vector is estimated for every radar point. The points with a surface vector correlating only weakly with the other points are deleted. The thinned-out point cloud is registered to the previous one using the Broyden-Fletcher-Goldfarb-Shanno (BFGS) line search method (Adolfsson et al., 2021).

3.2 Radar Signal Processing

Radar errors may also happen in the digital signal processing steps and should be prevented here as best as possible. Modern radars send FMCW waveforms and determine the range of a target point and its relative velocity by the frequency shift of the received signal (Patole et al., 2017). The whole radar signal processing is a many-layered process, starting with interference mitigation, fast-fourier transformation and beamforming in the preprocessing step, followed by the creation of the target list (Engels et al., 2021). In this step, a Constant False Alarm Rate (CFAR) mechanism is often applied to determine the threshold for a signal peak to be considered as a target point. The peak's position allows for calculation of the other point attributes, the so-called parameter estimation (Engels et al., 2021). The resulting point cloud is the basis for all the applications which were described before, such as object detection or localization. However, the techniques described in section 3.1 all try to achieve a regularly working result, filtering the point cloud and keeping only strong reflections while accepting to ignore weak reflections of objects in favor of a low false positive rate. These filtering methods are not sufficient when undercutting low failure in time rates must be guaranteed. In the following, an overview of techniques in the state-of-the-art for guaranteeing safe perception is given. For this section, we state the term of the Object of Interest (OoI), which can be both an object which should be detected during an object detection or free space estimation task or an arbitrary object with backscattering, generating at least one true positive point, that is considered for measures of radar odometry or localization.

3.3 Radar Hardware Development

Starting in front of the chain, it needs to be ensured that the OoI generates at least one point in the point cloud. This requires a sufficiently strong backscatter signal which exceeds the required threshold. The radar sensor itself is accountable for parts of the overall perception performance. (Gerstmair et al., 2019) show that phase noise in the transmitted signal caused by the voltage controlled oscillator plays a significant role for the signal-to-noise ratio. If phase noise can be kept on a minor level, Vulnerable Road Users (VRUs) such as pedestrians elicit reflections which exceed the noise floor level, while a poor phase noise can lead to a pedestrian masked by the noise floor. Furthermore, they outline methods to estimate the power spectrum density of the phase noise and monitoring approaches using a cascaded Monolithic Microwave Integrated Circuit (MMIC) system. Since the interference phenomenon may also lead to a decreased signal-to-noise level, other works consider its mitigation. (Aydogdu et al., 2020) discuss the effects of interference and propose proactive strategies, proactive meaning to avoid or reduce interference. These works are undoubtedly relevant to reduce the occurrence of false negative points, existing objects of the environment which are not represented by a radar target point. However, they do also not focus on the safety-aware selection of radar points. Hence, in the following, the methods to treat the radar points which are returned by the sensor for consecutive purposes shall be discussed.

3.4 Origin and Classification of False Radar Points

Radar point clouds are noisy, meaning that they contain points which do not represent an OoI. (Bilik et al., 2019) names clutter as a reason for noisiness and a major challenge for radar signal processing. Clutter points are road echoes in front of the vehicle which are barely distinct from signal returns of OoIs. Multipath clutter is described as another effect which can obscure true positive targets due to the doppler spread effect (Yu and Krolik, 2012). On the other hand, this may cause false positive radar points as well. The impact of these false points has been described in various works in the past. (Barnes and Posner, 2020) explains a method to identify those radar points which are robust enough to serve as key points for tasks such as motion estimation or localization. A U-Net like architecture is used and trained based on ground-truth transformation between consecutive point clouds. The output of this network is the location and quality of the key point as well as a uniqueness attribute. The drawback of this approach is that the selection of key points might be suitable for the scan matching task, whereas the selection decisions are not explainable.

In general, to make radar points safe and discover potential unsafe points, it needs to be expounded what safety is and which mechanisms in perception impair the safety of the function.

3.5 What Is Safety?

One established safety concept is Functional Safety (FuSa), which is described by the ISO 26262 standard (International Organization for Standardization, 2018). FuSa means that electrical systems need to fulfill safety guarantees depending on their potential to injure their user. The safety goals to avoid malfunction in a situation are harder, the more likely the user is exposed to a situation, the worse the situation is controllable and the more severe resulting injuries at a malfunction are. In automated driving functions, it turned out that FuSa, originally thought as a concept for hardware flawlessness, cannot be applied well for malfunctions which are caused by poor perception capabilities or misbehavior of an algorithm. Hence, the Safety Of The Intended Functionality (SOTIF) standard was defined, which amplifies safety definitions and targets to provide a description for these cases. According to (International Organization for Standardization, 2022), SOTIF is defined as the "absence of unreasonable risk due to hazards resulting from insufficiencies of the intended functionality or by reasonably foreseeable misuse by persons", which must be guaranteed for autonomous driving functions.

(Peng et al., 2023) name two classes for risks of autonomous functions, external risk and internal risk. Safe radar perception focuses mainly on external risks, meaning that e.g., velocities and positions of adversarial road participants are not determined correctly. In contrast to this, internal risk describes the probability that an algorithm performs erroneously, and external measurements are appropriate. Furthermore, they establish the term "SOTIF entropy" based on the Shannon entropy, which describes the uncertainty of a label prediction in an object classification task. This entropy can be determined for perception as well as for prediction and planning.

Ensuring that autonomous vehicles make safe decisions, the term "safety" can be interpreted in different ways. (Bila et al., 2017) (Muhammad et al., 2021) name different tasks such as collision avoidance, detection and tracking of vehicles or pedestrian detection. Mastering each of these tasks will lead to a safe autonomous driving function. Additionally, they highlight the existence of the three stages measurement, analysis, and execution. Each stage of this so-called cognitive control cycle can impair the safety of the function.

(Khatun et al., 2020) combines the two concepts of SOTIF and FuSa by a scenario-based hazard analysis and risk assessment (HARA). Hazards can emerge as vehicle-based or functional malfunction, covering typical FuSa parts as well as SOTIF parts of the safety contemplation. Scenarios can result from both branches and are preselected. To rate the risk, the three criteria of FuSa are employed. Hazard and Operability (HAZOP) keywords are collected and transferred to SOTIF functional imperfections. Various works (Chu et al., 2023)(Peng et al., 2021) emphasize that dealing with uncertainties is an inevitable part when it comes to making automated functions safer. (Dietmayer, 2016) categorizes uncertainties in determining the state (i.e., position, orientation, motion state etc.), the existence (presence detection) and the classification of an OoI.

3.6 Safety-Aware Function Evaluation

The review of (Hoss et al., 2022) discovers testing methods to ensure methodically that functions meet the safety requirements. They describe metrics to measure relevance of certain scenarios, introduce ways to specify scenarios as well as the Operational Design Domain (ODD) and propose to create a test catalog on either a knowledge-driven, a data-driven or a hybrid way. Furthermore, it is outlined how to acquire data for safety evaluation through different ways, either by the vehicle-under-test itself or by using data which is received e.g., via vehicle-to-X.

Solving tasks such as object detection with machine learning methods is problematic due to their black-box characteristics and lacking measure of uncertainty of the decision made. However, one work shows how to apply SOTIF on these methods using the example of a lane keeping function (Abdulazim et al., 2021). They use a baseline model for the function which consists of a filter block, a memory unit, and a predictor. Trigger events for a safety-critical scenario are defined and tested during the verification process. The verification and validation process happened on both a synthesized data set and a data set recorded by a real vehicle during a road test. Summarizing, the work offers a guideline on how to apply SOTIF concepts on machine learning methods, and describes the procedure in a very simplified example.

As mentioned before, evaluation of the automated function should already take place on the perception stage, since failures in the perception affect all downstream algorithms. One way to identify the threats of perception functions is to dissect the event chain during perception, consisting of the information reception and information processing (Philipp et al., 2020). Errors, leading to failure, can occur in the raw scan as well as in the feature level. The work focuses on feature level errors and thus it lacks a detailed analysis of the radar sensor disturbance and failure mechanisms.

Schönemann outlines a procedure to ensure safe behavior of an autonomous function, in this case the Automated Valet Parking (AVP) function, already in the function design. He splits his observations in three aspects: the minimum safety requirements define the situations which an automated function should be able to handle, the minimum required perception zone defines the free space needed to monitor to set the vehicle in a safe state, i.e., from driving to standstill, and the minimum functional requirements, deviating elements such as perception or planner for the automated function to work (Schönemann, 2019).

Chu et al. made valuable work in the field of connecting SOTIF with perception challenges. Same to Schönemann, they also use the safety distance concept to model a minimum required perception area. They derive the velocity of objects near to the vehicle to become a potential collision candidate. Additionally, they connect these findings to a sensor model with uncertainties to identify the required capabilities of the perception system for collision avoidance. Weaknesses in the coverage of the considered sensor set were identified during a case study, and they also suggest including degrading conditions regarding lighting and weather or to apply the ideas to more complex tasks (Chu et al., 2023).

3.7 Definition of Critical Radar Points

The previously introduced works define the safety term and enable applying it to the perception of radar systems. We emphasize that the ignorance of a radar point can lead to a safety hazard as well as misinterpretation or wrong assessment of a situation. However, as false positive radar points occur abundantly, filtering needs to be done while being aware of the hazards this brings along. Commonly known and widespread are the terms to define the fraction of points which are relevant from the given data set, precision, and the fraction of points which are covered from the total relevant object set, recall (Powers, 2011)(Sun et al., 2023). Precision and recall are defined as

$$Precision = \frac{TP}{TP+FP} \tag{1}$$

$$Recall = \frac{TP}{TP+FN} \tag{2}$$

See figure 2 as an explanation of how the terms are used in this context. The grey, transparent area shows the field of view of a radar sensor. The backscatter of the car on the left hand side generate strong detection points, same for the posterior walker. These two objects would be considered as true positives. The anterior walker ("On-edge object") generates only one weak detection. Assuming that weak detections are filtered, this object would be a false negative. Whereas some clutter from the ground generates strong detections and is likely to be interpreted as an obstacle, illustrated by the red-dashed box. Lowering the threshold for detections to count as "strong" would include the on-edge object, but also clutter from the ground. In succession, recall increases and precision declines. The reverse effect happens if the threshold is elevated. The on-edge object is not included in the object list anymore and also the false positive clutter is sparser.

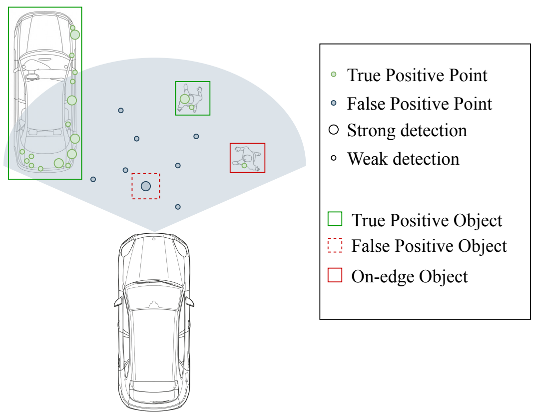

Figure 2: Tradeoff between recall and precision.

This tradeoff has been described elaborately in the literature (Buckland and Gey, 1994)(Yang et al., 2019). With reference to radar points, in urban scenarios with increased multi-path propagation, much interference by other vehicles and many objects, the noise level is high compared to free field scenarios. Additionally, pedestrians may have a RCS of $0.01\,\mathrm{m}^2$ (Chen et al., 2014), which is a thousand or less of the RCS of the truck next to them (Waldschmidt et al., 2021). As VRUs with low RCS do hardly backscatter and may vanish in the noise floor, a recall of 1.0 will not be reached, even if every point delivered by the radar sensor is considered. Hence, increasing the maximum recall rate is a main working field for radar specialists. Still, optimizing the recall rate will involve more false positives, which is to the disadvantage for the precision rate. We assume that a certain precision rate needs to be reached to ensure that the function is robust enough to master most of the situation without undergoing a deadlock, i.e., too many false positives are detected in crowded situations and the vehicle does not move anymore. Enriching radar points which may be a threat for SOTIF conformity with information may be one of the keys. This information can descend from other sensors such as cameras or radar sensors.

3.8 Camera-Radar Fusion

Fusion is the combination of information from more than one source, utilizing the redundancy and the congruence of data. According to (Khaleghi et al., 2013), it can be distinguished between four different fusion problems: imperfection, correlation, inconsistency, and disparateness. Speaking about adding information to radar point clouds via fusion, especially two aspects are relevant. The radar data may be deficient (imperfection) or miss information (inconsistency). Correlation could become problematic if the data from two sensors of a kind or the data of one sensor at different time stamps shall be incorporated in the fusion. Attributes that are usually fused, such as position and velocity, don't exhibit a disparity between different perception sensors.

(Tang et al., 2022) highlight that the use of camera information is beneficial to use for fusion with radar sensors due to its complementarity. Moreover, they propose different fusion schemes, e.g., object fusion before and after tracking and region-of-interest based fusion. However, their work does not concern about early fusion stages such as fusion at the level of radar detection points or camera pixels, but aim to solve the task of object detection for radar and camera data separately before fusing the object lists.

The fusion in these levels is explained more precisely in (Zhou et al., 2022). Data-level fusion extracts the features and generates the object list from the fusion of radar and camera data. Target-level fusion extracts the features separately and combines these features together. The structure called decision-level fusion in their work is an object fusion which is similar to a fusion before tracking. Furthermore, the transformation of radar points to camera pixels, the camera calibration required to that end, and the synchronization process is discussed. Feature level fusion is recognized as commonly used due to the erroneousness of radar point clouds. Thinking about safe radar detection points, this may pertain for most non-critical areas, while a fusion level at an earlier stage might be beneficial to keep the precision rate high.

(Velasco-Hernandez et al., 2020) mention that deep learning plays a major role for multi sensor object detection and fusion. According to (Zhou et al., 2020), deep learning architectures can be split into a two-stage detection network, where the radar assumes the task of a region proposal network in a deep learning image-based object detection algorithm such as (Girshick, 2015)(Brazil and Liu, 2019), and a one-stage detection network which solves the object detection task inside of the image separately and projects radar points in the 2D front camera image, generating a sparse radar image (John and Mita, 2019). (Zhou et al., 2020) also name special attention fusion (SAF) which can be combined with the already known Fully Convolutional One-Stage Object Detection (FCOS) framework in a one-stage detection network manner for the purpose of radar-camera-fusion. The SAF block is a modified ResNet-50 block which is applied on a radar image and results are multiplied with those of the vision branch. By this means, the training can be performed end-to-end (Chang et al., 2020).

While the attention mechanism is applied to control how deep learning image recognition works, the attention focuses purely on the presence of radar points, not evaluating if these radar points can be hazardous or relevant for the downstream automated driving function. Prior aspiration for the method is to reach a high average precision on data sets, leaving out safety aspects for the fusion.

3.9 Multi-Radar Fusion

The work on the fusion of multiple radar point clouds which are generated by different sensors of a radar belt configuration at a vehicle is not as widespread, compared to camera fusion. The reason might be that the fusion can be trivial, aggregating points of all sensors to a common point cloud, as it was also done in data sets like (Schumann et al., 2021). However, the distribution of detections can be interesting for safety aspects. Speaking about the identification of false positive points, the probability of a point being false positive reduces significantly if a second radar raises a detection one cycle later in the same area. (Diehl et al., 2020) handle multi-radar fusion by transferring all measurements to a grid representation, modeling the sensor uncertainties by 2D gaussian distributions. The tracking can also be done in this grid representation using multiple particle filters. (Li et al., 2022) shed light on the temporal relations between consecutive radar frames. Two frames, called current frame and previous frame, are both passed through a backbone neural network to extract features. A filter layer selects object features, and the feature matrices of multiple object candidates are passed to the temporal relational layer, which encodes object spatial positions with the features of the two consecutive frames. Eventually, a decoding step leads to the object list representation. The achieved results exceeding the state-of-the-art show the importance of temporal information in radar data.

Overlapping fields-of-view in a radar belt offer the opportunity to attain more information about a specific region of interest. Imagine a point that has a low signal strength and is hence considered to be deleted. If it is identified as being a critical point, multiple occurrences of backscatter in the same region at previous time steps can help to amplify evidence on the existence of the point.

3.10 Data Sets

(Yurtsever et al., 2020) review a couple of data sets, listing lighting conditions, weather conditions and recorded data, but do not show which of the data sets include radar data. For this paper, their work is taken as a basis for an own small review about current data sets which promote the development of radar data processing algorithms. The results can be found in table 1, and the most important data sets are explained in the following passage.

The Oxford Radar RobotCar data set provides radar point clouds in urban environment together with a lidar point cloud and six cameras, indeed they used a Navtech CTS350-X which is a 360-degree scanning radar without doppler information and holds properties different from a belt consisting of multiple radars, as it is conventionally used in today's automotive industry (Barnes et al., 2020).

(Schumann et al., 2021) encounter this issue and serve RadarScenes, which features a point cloud captured by four radar sensors and doppler velocities. Additionally, all points are manually labeled, which is to our knowledge unique for a radar data set. Unfortunately, the radar sensors cover only the front and sides of the vehicle, but no rear area. RadarScenes works well when multi radar fusion shall be investigated, but is not useful for radar-camera fusion. NuScenes is one of the most popular radar data sets, featuring drives in Boston and Singapore, six camera perspectives and point clouds from five radars. Additionally, the 3D ground truth bounding boxes are available (Caesar et al., 2020). Although this data set has a comparatively large size, it does not feature a lot of hazardous situations. RADIATE is a radar data set which puts special effort in including bad weather and lighting conditions such as night, snow, and rain, but unfortunately lacking 3D annotations (Sheeny et al., 2021).

We ascertained that there are various data sets which are made for the development of autonomous driving functions. However, we could not find a data set which addresses safety critical scenes only. This is a substantial lack in research, since improving fusion algorithms on the existing data sets will not necessarily improve the performance of the system handling safety-critical situations. We suggest evaluating the methods on data of the own system and putting effort in a general description of how to make sensor setups safe by identifying their individual weaknesses.

Using own data requires to record ground truth data in parallel as a reference. Generally, annotated data must be highly precise, and labels need to be complete (Xiao et al., 2021). The ground truth data must hold significantly more details than the system which is under development. For tasks like object or occupancy detection, 3D bounding boxes represent the ground truth. As 3D bounding box labeling by hand is tedious, Lee et al. propose to use 3D object detectors applied on a lidar point cloud to generate labeling proposals while the annotator's only task is to select the individual instances (Lee et al., 2018). The lidar setup for own ground truth generation depends on the task, but the setups of the vehicles used to generate data sets (Caesar et al., 2020) (Geyer et al., 2020) can serve as a source of inspiration. Looking to other radar tasks such as odometry or SLAM, the ground truth is a motion measurement respective a pose. (Maddern et al., 2020) describe the procedure to post-process raw Global Positioning System (GPS) data and data from an inertial measurement unit to receive a centimeter-accurate representation of the vehicle's movements and positions during the data recording. Techniques like these allow for comparison between distinct algorithm modifications.

Table 1: Recent data sets of autonomous driving including radar data.

Data Set	Author and Year	Radar Sensors	Advantages	Drawbacks
NuScenes	(Caesar et al., 2020)	5x Continental ARS 408-21 (2Tx/6Rx) ea. near/far	• calibrated camera data provided • lidar data provided • variance of situations	• radar points sparse in near range • vacant spots at side of the vehicle
Oxford Radar RobotCar	(Barnes et al., 2020)	1x Navtech CTS350-X	• lidar data included • odometry data included • stereo-camera images provided	• no conventional multi-radar setup • no radar point cloud provided • no annotations provided
RadarScenes	(Schumann et al., 2021)	4x 77 GHz corner radar	• semantic class annotation per radar point	• no 360° radar • documentary camera only
RADIATE	(Sheeny et al., 2021)	1x Navtech CTS350-X	• focus on adversarial weather conditions • lidar data provided	• no conventional multi-radar setup • no radar point cloud provided • no 360° camera coverage • no annotations provided
CARRADA	(Ouaknine et al., 2021)	1x 77 GHz radar (2Tx/4Rx)	• good camera images provided • range/azimuth annotations provided • numerous VRUs in object selection	• no 360° radar/camera view provided • only static scenes • 4 GHz sweep used
RADDet	(Zhang et al., 2021)	1x AWR1843-BOOST (3Tx/4Rx)	• annotations per radar point available • stereo-camera images provided	• no adversarial weather conditions • only static scenes • only one radar
CRUW	(Wang et al., 2021)	1x TI AWR1843 (3Tx/4Rx)	• camera-only annotations provided • camera-radar-fused annotations provided • good distribution of various road types	• no 360° radar/camera view provided • no radar point annotations provided • no adversarial weather conditions
K-Radar	(Paek et al., 2022)	1x 4D radar 77GHz	• 4D tensor with elevation information provided • lidar data provided • 360° camera data provided • good road type and weather distribution	• only one radar • no radar point annotations provided

4 DISCUSSION

As this work points out, the potential of radar sensors for automated driving and parking functions is notable. When it comes to environmental perception in poor weather and illumination conditions or measuring the relative velocity of an object in a direct manner, there is no way around this technique.

However, while a large volume of works addresses the overall performance and accuracy of e.g., SLAM or object detection, there is minor research focused on how radar sensors could provide error guarantees. It can be stated that rethinking needs to happen in the signal processing, where the works comprised in this survey rely on simple probabilistic filtering methods. such as CFAR. At this, sensors stay behind their potential in radar target detection, since a criticality-based threshold adaption would outperform targeting constant false alarm rates. Failure can already occur by mistakes in the hardware design. Thus, the possible error sources must be investigated during the design process, as this paves the way for a general evaluation of failure rates.

The safety term got extended recently by the SOTIF definition, and various work is done to describe hazard analyses. However, the perception elements are hardly included in these considerations. Concepts like the safety distance are a meaningful foundation, but propagation through the whole perception task, respecting the radar specifics, still needs to be done. It is essential to process radar points differently and to make the difference not with the probability of the situation, but with its effect, especially regarding the harm of VRUs. A raw point in front of the vehicle in its passage route needs to be treated completely different than a raw point behind the vehicle, which will not be hit in the near future. Not many works respond to this new point of view, as the survey showed.

Handling the still enormous amount of potentially critical raw points will only be possible by considering data of more than one sensor. The fact that research is still at its beginning is reflected in the small number of works which were presented for such sensor fusion algorithms which aim at providing better safety. While a lot of works focus on improving the overall detection performance and stability, none of the works try to confirm or disprove the object existence at locations of radar points, which could advantageously be approached by radar-camera fusion or multi-camera fusion.

The sparseness of work on this area continues when it comes to suitable data sets for the development of such algorithms. We did not find one radar data set containing critical road situations only. It is desirable to have a data set containing different corner case scenarios including critical scenarios with VRUs for radar sensors, in diverse weather and illumination conditions, with the most important perturbing effects included. Data should contain radar data as a point cloud, calibrated camera data, annotated 3D bounding boxes, and vehicle odometry to infer the position from the starting point.

The new standards which are resumed in this paper show that level 3 or level 4 automated functions will not find their way to the road without an approval with a pervasive failure-in-time guarantee as the most important component. This guarantee requires a new way to process data of the perception sensors, such as radar, lidar or cameras. If no special action for such systems is taken at the beginning of the design process, long time before verification and validation of these systems become imminent, the approval will be almost impossible.

5 CONCLUSION AND OUTLOOK

In this paper, we pointed out the notable potential which radar sensors show to promote automated driving functions. The diversity of tasks which can be tackled by radar sensors is shown, making their usage attractive. Next, the current progress in the area of radar hardware development and signal processing is shown. These works serve as a basis for the consecutive inspection of safety-aware radar data processing. To explain the line of reasoning, we scrutinized interpretations of the safety term in context of automated driving functions. The role of critical false positive radar points was highlighted. Selected, safety-focused works in the field of function evaluation were investigated to work out how the perception data can be connected to the criticality of scenarios. As enrichment of data points with information is indispensable to declare their significance, the area of sensor data fusion is illuminated. Radar-camera fusion and multi-radar-fusion were especially emphasized in this survey. We ended with giving an overview about current radar data sets and rate them with respect to the usability for safety-aware function development.

It is recognized that future work should be done investigating the impact of erroneously filtered positive radar points on the individual tasks described in this work. A method to identify critical situations is important to control enriching the right regions with additional information, originating e.g. from camera setups. Pursuing these aspects will bring the research community one step closer to perception which is safe under a guarantee.

REFERENCES

Abdulazim, A., Elbahaey, M., and Mohamed, A. (2021). Putting Safety of Intended Functionality SOTIF into Practice. Technical report, Warrendale, PA.

Adolfsson, D., Magnusson, M., Alhashimi, A., Lilienthal, A. J., and Andreasson, H. (2021). CFEAR Radarodometry - Conservative Filtering for Efficient and Accurate Radar Odometry. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 5462–5469. ISSN: 2153-0866.

Aptiv PLC (2023). Gen7 radar family.

Aydogdu, C., Keskin, M. F., Carvajal, G. K., Eriksson, O., Hellsten, H., Herbertsson, H., Nilsson, E., Rydstrom, M., Vanas, K., and Wymeersch, H. (2020). Radar Interference Mitigation for Automated Driving: Exploring Proactive Strategies. *IEEE Signal Processing Magazine*, 37(4):72–84.

Barnes, D., Gadd, M., Murcutt, P., Newman, P., and Posner, I. (2020). The Oxford Radar RobotCar Dataset: A Radar Extension to the Oxford RobotCar Dataset. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6433–6438. ISSN: 2577-087X.

Barnes, D. and Posner, I. (2020). Under the Radar: Learning to Predict Robust Keypoints for Odometry Estimation and Metric Localisation in Radar. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9484–9490. ISSN: 2577-087X.

Bila, C., Sivrikaya, F., Khan, M. A., and Albayrak, S. (2017). Vehicles of the Future: A Survey of Research on Safety Issues. *IEEE Transactions on Intelligent Transportation Systems*, 18(5):1046–1065.

Bilik, I. (2023). Comparative Analysis of Radar and Lidar Technologies for Automotive Applications. *IEEE Intelligent Transportation Systems Magazine*, 15(1):244–269.

Bilik, I., Longman, O., Villeval, S., and Tabrikian, J. (2019). The Rise of Radar for Autonomous Vehicles: Signal Processing Solutions and Future Research Directions. *IEEE Signal Processing Magazine*, 36(5):20–31.

Brazil, G. and Liu, X. (2019). M3d-rpn: Monocular 3d region proposal network for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*.

Buckland, M. and Gey, F. (1994). The relationship between Recall and Precision. *Journal of the American Society for Information Science*, 45(1):12–19.

Caesar, H., Bankiti, V., Lang, A. H., Vora, S., Liong, V. E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., and Beijbom, O. (2020). nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.

Chang, S., Zhang, Y., Zhang, F., Zhao, X., Huang, S., Feng, Z., and Wei, Z. (2020). Spatial Attention Fusion for Obstacle Detection Using MmWave Radar and Vision Sensor. *Sensors*, 20(4):956.

Chen, M., Belgiovane, D., and Chen, C.-C. (2014). Radar characteristics of pedestrians at 77 GHz. In *2014 IEEE Antennas and Propagation Society International Symposium (APSURSI)*, pages 2232–2233. ISSN: 1947-1491.

Chu, J., Zhao, T., Jiao, J., Yuan, Y., and Jing, Y. (2023). SOTIF-Oriented Perception Evaluation Method for Forward Obstacle Detection of Autonomous Vehicles. *IEEE Systems Journal*, 17(2):2319–2330.

Diehl, C., Feicho, E., Schwambach, A., Dammeier, T., Mares, E., and Bertram, T. (2020). Radar-based Dynamic Occupancy Grid Mapping and Object Detection. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–6.

Dietmayer, K. (2016). Predicting of Machine Perception for Automated Driving. pages 407–424. Springer, Berlin, Heidelberg.

Engels, F., Heidenreich, P., Wintermantel, M., Stäcker, L., Al Kadi, M., and Zoubir, A. M. (2021). Automotive Radar Signal Processing: Research Directions and Practical Challenges. *IEEE Journal of Selected Topics in Signal Processing*, 15(4):865–878.

Gerstmair, M., Melzer, A., Onic, A., and Huemer, M. (2019). On the Safe Road Toward Autonomous Driving: Phase Noise Monitoring in Radar Sensors for Functional Safety Compliance. *IEEE Signal Processing Magazine*, 36(5):60–70.

Geyer, J., Kassahun, Y., Mahmudi, M., Ricou, X., Durgesh, R., Chung, A. S., Hauswald, L., Hoang Pham, V., Mühlegg, M., Dorn, S., Fernandez, T., Jänicke, M., Mirashi, S., Savani, C., Sturm, M., Vorobiov, O., Oelker, M., Garreis, S., and Schuberth, P. (2020). A2D2: Audi Autonomous Driving Dataset. Technical report. ADS Bibcode: 2020arXiv200406320G Type: article.

Girshick, R. (2015). Fast r-cnn. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*.

Hong, Z., Petillot, Y., Wallace, A., and Wang, S. (2021). Radar SLAM: A Robust SLAM System for All Weather Conditions. Technical report. arXiv:2104.05347 [cs] type: article.

Hoss, M., Scholtes, M., and Eckstein, L. (2022). A Review of Testing Object-Based Environment Perception for Safe Automated Driving. *Automotive Innovation*, 5(3):223–250.

International Organization for Standardization (2018). ISO 26262:2018 - Road vehicles - Functional safety. Technical report.

International Organization for Standardization (2022). ISO 21448:2022 - Road vehicles - Safety of the intended functionality. Technical report.

John, V. and Mita, S. (2019). RVNet: Deep Sensor Fusion of Monocular Camera and Radar for Image-Based Obstacle Detection in Challenging Environments. In Lee, C., Su, Z., and Sugimoto, A., editors, *Image and Video Technology*, Lecture Notes in Computer Science, pages 351–364, Cham. Springer International Publishing.

Kellner, D., Barjenbruch, M., Klappstein, J., Dickmann, J., and Dietmayer, K. (2014). Instantaneous ego-motion estimation using multiple Doppler radars. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1592–1597. ISSN: 1050-4729.

Khaleghi, B., Khamis, A., Karray, F. O., and Razavi, S. N. (2013). Multisensor data fusion: A review of the state-of-the-art. *Information Fusion*, 14(1):28–44.

Khatun, M., Glaß, M., and Jung, R. (2020). Scenario-based extended hara incorporating functional safety & sotif for autonomous driving. In Baraldi, P., editor, *E-proceedings of the 30th European Safety and Reliability Conference and 15th Probabilistic Safety Assessment and Management Conference*, volume 2020 of *(ESREL2020 PSAM15) ; 01 - 05 November 2020 Venice, Italy*, pages 53 – 59.

Knott, E. F. (2012). *Radar cross section measurements*. Springer Science & Business Media.

Lee, J., Walsh, S., Harakeh, A., and Waslander, S. L. (2018). Leveraging Pre-Trained 3D Object Detection Models for Fast Ground Truth Generation. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2504–2510. ISSN: 2153-0017.

Li, M., Feng, Z., Stolz, M., Kunert, M., Henze, R., and Küçükay, F. (2018). High resolution radar-based occupancy grid mapping and free space detection.

Li, P., Wang, P., Berntorp, K., and Liu, H. (2022). Exploiting temporal relations on radar perception for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17071–17080.

Maddern, W., Pascoe, G., Gadd, M., Barnes, D., Yeomans, B., and Newman, P. (2020). Real-time Kinematic Ground Truth for the Oxford RobotCar Dataset. Technical report. ADS Bibcode: 2020arXiv200210152M Type: article.

Marti, E., de Miguel, M. A., Garcia, F., and Perez, J. (2019). A Review of Sensor Technologies for Perception in Automated Driving. *IEEE Intelligent Transportation Systems Magazine*, 11(4):94–108.

Muhammad, K., Ullah, A., Lloret, J., Ser, J. D., and de Albuquerque, V. H. C. (2021). Deep Learning for Safe Autonomous Driving: Current Challenges and Future Directions. *IEEE Transactions on Intelligent Transportation Systems*, 22(7):4316–4336.

Ouaknine, A., Newson, A., Rebut, J., Tupin, F., and Pérez, P. (2021). CARRADA Dataset: Camera and Automotive Radar with Range- Angle- Doppler Annotations. In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 5068–5075. ISSN: 1051-4651.

Paek, D.-H., KONG, S.-H., and Wijaya, K. T. (2022). K-radar: 4d radar object detection for autonomous driving in various weather conditions. In Koyejo, S., Mohamed, S., Agarwal, A., Belgrave, D., Cho, K., and Oh, A., editors, *Advances in Neural Information Processing Systems*, volume 35, pages 3819–3829. Curran Associates, Inc.

Patole, S. M., Torlak, M., Wang, D., and Ali, M. (2017). Automotive radars: A review of signal processing techniques. *IEEE Signal Processing Magazine*, 34(2):22–35.

Peng, L., Li, B., Yu, W., Yang, K., Shao, W., and Wang, H. (2023). SOTIF Entropy: Online SOTIF Risk Quantification and Mitigation for Autonomous Driving. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–17.

Peng, L., Wang, H., and Li, J. (2021). Uncertainty Evaluation of Object Detection Algorithms for Autonomous Vehicles. *Automotive Innovation*, 4(3):241–252.

Philipp, R., Schuldt, F., and Howar, F. (2020). Functional decomposition of automated driving systems for the classification and evaluation of perceptual threats. In *13. Uni-DAS eV Workshop Fahrerassistenz und automatisiertes Fahren 2020*. Walting.

Popov, A., Gebhardt, P., Chen, K., and Oldja, R. (2023). NVRadarNet: Real-Time Radar Obstacle and Free Space Detection for Autonomous Driving. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6958–6964.

Powers, D. (2011). Evaluation: From precision, recall and f-measure to roc, informedness, markedness & correlation. *Journal of Machine Learning Technologies*, 2(1):37–63.

Robert Bosch GmbH (2023). Front-radar premium.

SAE International. (2021). J3016 - Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. Technical report.

Sayyah, K., Sarkissian, R., Patterson, P., Huang, B., Efimov, O., Kim, D., Elliott, K., Yang, L., and Hammon, D. (2022). Fully Integrated FMCW LiDAR Optical Engine on a Single Silicon Chip. *Journal of Lightwave Technology*, 40(9):2763–2772.

Scheiner, N., Kraus, F., Appenrodt, N., Dickmann, J., and Sick, B. (2021). Object detection for automotive radar point clouds – a comparison. *AI Perspectives*, 3(1):6.

Schönemann, V. (2019). Safety requirements and distribution of functions for automated valet parking.

Schumann, O., Hahn, M., Scheiner, N., Weishaupt, F., Tilly, J. F., Dickmann, J., and Wöhler, C. (2021). RadarScenes: A Real-World Radar Point Cloud Data Set for Automotive Applications. In *2021 IEEE 24th International Conference on Information Fusion (FUSION)*, pages 1–8.

Schumann, O., Lombacher, J., Hahn, M., Wöhler, C., and Dickmann, J. (2020). Scene Understanding With Automotive Radar. *IEEE Transactions on Intelligent Vehicles*, 5(2):188–203.

Schuster, F., Keller, C. G., Rapp, M., Haueis, M., and Curio, C. (2016). Landmark based radar SLAM using graph optimization. In *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, pages 2559–2564. ISSN: 2153-0017.

Sheeny, M., De Pellegrin, E., Mukherjee, S., Ahrabian, A., Wang, S., and Wallace, A. (2021). RADIATE: A Radar Dataset for Automotive Perception in Bad Weather. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1–7. ISSN: 2577-087X.

Steinbaeck, J., Steger, C., Holweg, G., and Druml, N. (2017). Next generation radar sensors in automotive sensor fusion systems. In *2017 Sensor Data Fusion: Trends, Solutions, Applications (SDF)*, pages 1–6.

Sun, C., Zhang, R., Lu, Y., Cui, Y., Deng, Z., Cao, D., and Khajepour, A. (2023). Toward Ensuring Safety for Autonomous Driving Perception: Standardization Progress, Research Advances, and Perspectives. *IEEE Transactions on Intelligent Transportation Systems*, pages 1–19.

Tang, X., Zhang, Z., and Qin, Y. (2022). On-Road Object Detection and Tracking Based on Radar and Vision Fusion: A Review. *IEEE Intelligent Transportation Systems Magazine*, 14(5):103–128.

Velasco-Hernandez, G., Yeong, D. J., Barry, J., and Walsh, J. (2020). Autonomous Driving Architectures, Perception and Data Fusion: A Review. In *2020 IEEE 16th International Conference on Intelligent Computer Communication and Processing (ICCP)*, pages 315–321.

Waldschmidt, C., Hasch, J., and Menzel, W. (2021). Automotive Radar — From First Efforts to Future Systems. *IEEE Journal of Microwaves*, 1(1):135–148.

Wang, Y., Jiang, Z., Gao, X., Hwang, J.-N., Xing, G., and Liu, H. (2021). Rodnet: Radar object detection using cross-modal supervision. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 504–513.

Xiao, P., Shao, Z., Hao, S., Zhang, Z., Chai, X., Jiao, J., Li, Z., Wu, J., Sun, K., Jiang, K., Wang, Y., and Yang, D. (2021). PandaSet: Advanced Sensor Suite Dataset for Autonomous Driving. In *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, pages 3095–3101.

Xu, F., Wang, H., Hu, B., and Ren, M. (2020). Road Boundaries Detection based on Modified Occupancy Grid Map Using Millimeter-wave Radar. *Mobile Networks and Applications*, 25(4):1496–1503.

Yang, M., Wang, S., Bakita, J., Vu, T., Smith, F. D., Anderson, J. H., and Frahm, J.-M. (2019). Re-thinking cnn frameworks for time-sensitive autonomous-driving applications: Addressing an industrial challenge. In *2019 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 305–317. IEEE.

Yeong, D. J., Velasco-Hernandez, G., Barry, J., and Walsh, J. (2021). Sensor and Sensor Fusion Technology in Autonomous Vehicles: A Review. *Sensors*, 21(6):2140.

Yu, J. and Krolik, J. (2012). MIMO multipath clutter mitigation for GMTI automotive radar in urban environments. In *IET International Conference on Radar Systems (Radar 2012)*, pages 1–5.

Yurtsever, E., Lambert, J., Carballo, A., and Takeda, K. (2020). A Survey of Autonomous Driving: Common Practices and Emerging Technologies. *IEEE Access*, 8:58443–58469.

Zang, S., Ding, M., Smith, D., Tyler, P., Rakotoarivelo, T., and Kaafar, M. A. (2019). The Impact of Adverse Weather Conditions on Autonomous Vehicles: How Rain, Snow, Fog, and Hail Affect the Performance of a Self-Driving Car. *IEEE Vehicular Technology Magazine*, 14(2):103–111.

Zhang, A., Nowruzi, F. E., and Laganiere, R. (2021). RADDet: Range-Azimuth-Doppler based Radar Object Detection for Dynamic Road Users. In *2021 18th Conference on Robots and Vision (CRV)*, pages 95–102.

Zhou, T., Yang, M., Jiang, K., Wong, H., and Yang, D. (2020). MMW Radar-Based Technologies in Autonomous Driving: A Review. *Sensors*, 20(24):7283.

Zhou, Y., Dong, Y., Hou, F., and Wu, J. (2022). Review on Millimeter-Wave Radar and Camera Fusion Technology. *Sustainability*, 14(9):5114.