

Sensing and Machine Learning for Automotive Perception

A Review

Pandharipande, Ashish; Cheng, Chih Hong; Dauwels, Justin; Gurbuz, Sevgi Z.; Ibanez-Guzman, Javier; Li, Guofa; Piazzoni, Andrea; Wang, Pu; Santra, Avik

DOI

[10.1109/JSEN.2023.3262134](https://doi.org/10.1109/JSEN.2023.3262134)

Publication date

2023

Document Version

Final published version

Published in

IEEE Sensors Journal

Citation (APA)

Pandharipande, A., Cheng, C. H., Dauwels, J., Gurbuz, S. Z., Ibanez-Guzman, J., Li, G., Piazzoni, A., Wang, P., & Santra, A. (2023). Sensing and Machine Learning for Automotive Perception: A Review. *IEEE Sensors Journal*, 23(11), 11097-11115. <https://doi.org/10.1109/JSEN.2023.3262134>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Green Open Access added to TU Delft Institutional Repository

'You share, we take care!' - Taverne project

<https://www.openaccess.nl/en/you-share-we-take-care>

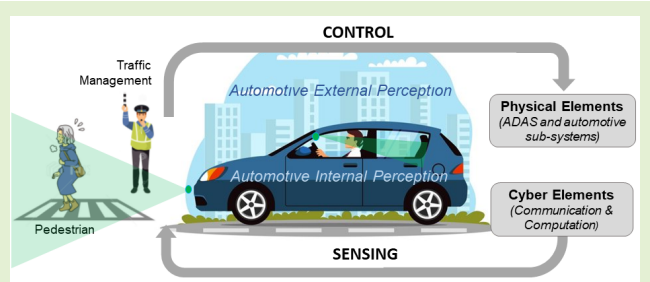
Otherwise as indicated in the copyright section: the publisher is the copyright holder of this work and the author uses the Dutch legislation to make this work public.

Sensing and Machine Learning for Automotive Perception: A Review

Ashish Pandharipande¹, Senior Member, IEEE, Chih-Hong Cheng, Justin Dauwels, Sevgi Z. Gurbuz², Senior Member, IEEE, Javier Ibanez-Guzman³, Member, IEEE, Guofa Li⁴, Member, IEEE, Andrea Piazzoni⁵, Pu Wang⁶, Senior Member, IEEE, and Avik Santra⁷, Senior Member, IEEE

Abstract—Automotive perception involves understanding the external driving environment and the internal state of the vehicle cabin and occupants using sensor data. It is critical to achieving high levels of safety and autonomy in driving. This article provides an overview of different sensor modalities, such as cameras, radars, and light detection and ranging (LiDAR) used commonly for perception, along with the associated data processing techniques. Critical aspects of perception are considered, such as architectures for processing data from single or multiple sensor modalities, sensor data processing algorithms and the role of machine learning techniques, methodologies for validating the performance of perception systems, and safety. The technical challenges for each aspect are analyzed, emphasizing machine learning approaches, given their potential impact on improving perception. Finally, future research opportunities in automotive perception for their wider deployment are outlined.

Index Terms—Advanced driver assistance system (ADAS), automotive perception, autonomous driving, cameras, light detection and ranging (LiDAR), radars, safety, sensor data processing.



I. INTRODUCTION

Different levels of automation are being included in modern vehicles, from an advanced driver assistance system (ADAS)

Manuscript received 11 February 2023; revised 15 March 2023; accepted 15 March 2023. Date of publication 30 March 2023; date of current version 31 May 2023. The work of Chih-Hong Cheng was supported by StMWi Bayern as part of a project to support the thematic development of Fraunhofer IKS. The associate editor coordinating the review of this article and approving it for publication was Prof. Domenico Ciuonzo. (The coauthors contributed equally to this work.) (Corresponding author: Ashish Pandharipande.)

Ashish Pandharipande is with NXP Semiconductors, 5656 AE Eindhoven, The Netherlands (e-mail: ashish.pandharipande@nxp.com).

Chih-Hong Cheng is with Fraunhofer IKS, 80686 Munich, Germany (e-mail: chih-hong.cheng@iks.fraunhofer.de).

Justin Dauwels is with the Department of Microelectronics, Delft University of Technology (TU Delft), 2628 CD Delft, The Netherlands (e-mail: J.H.G.Dauwels@tudelft.nl).

Sevgi Z. Gurbuz is with the Department of Electrical and Computer Engineering, The University of Alabama, Tuscaloosa, AL 35487 USA (e-mail: szgurbuz@ua.edu).

Javier Ibanez-Guzman is with Renault, 78064 Guyancourt, France (e-mail: javier.ibanez-guzman@renault.com).

Guofa Li is with the College of Mechanical and Vehicle Engineering, Chongqing University, Chongqing 400044, China (e-mail: hanshan198@gmail.com).

Andrea Piazzoni is with the Interdisciplinary Graduate Programme-ERI@N, Centre of Excellence for Testing&Research of AVs, Nanyang Technological University, Singapore 639798 (e-mail: andrea006@ntu.edu.sg).

Pu Wang is with the Mitsubishi Electric Research Laboratories, Cambridge, MA 02139 USA (e-mail: pwang@merl.com).

Avik Santra is with Infineon Technologies, Irvine, CA 92618 USA (e-mail: Avik.Santra@infineon.com).

Digital Object Identifier 10.1109/JSEN.2023.3262134

to a fully automated driving system (ADS). These systems use sensor and control technologies to improve driving safety and comfort. Automotive perception is a core module in such systems. Perception information relies on using one or more sensor modalities, such as camera, radar, and light detection and ranging (LiDAR). By suitably processing raw sensor data, information on the environment around the vehicle (external perception) and the state of the vehicle cabin (internal perception) is derived. Each sensor has its strengths and limitations, and its signal response will vary according to the driving environment [1]. The processing of sensor data is, thus, key to deriving reliable environment information for safe vehicle driving. This includes different aspects: 1) processing architectures considering what data to combine from sensors at which level of the processing chain; 2) processing algorithms for external perception capabilities, such as object detection and classification, range, and velocity estimation, or internal perception capabilities, such as occupancy detection and occupant alertness; 3) incorporation of physics-/model-/data-driven methods into data processing; 4) application-level performance metrics and validation approaches; and 5) safety in perception-driven vehicle functions. In this article, our objectives are to identify the key challenges in sensor processing for automotive perception, review the advances toward addressing them, and identify the gaps that exist to attain higher levels of perception.

High-quality, robust automotive perception is needed to reduce the number of traffic accidents and fatalities resulting from human driving errors. Automotive perception has seen

significant progress in the past years due to the emergence of advanced sensors, computing power, and the successful application of machine learning techniques. This has led to the deployment of multiple driving functions with increased levels of autonomy in commercial vehicles. Automotive perception, however, is a challenging problem for several reasons. First, the operational design domain (ODD) is complex, and perception needs to be reliable across different environmental and driving conditions. Second, the interaction between perception and driving controls may lead to propagation errors when the human is no longer part of the control loop. Third, the design and deployment of perception-based autonomous vehicles involve new technological and social challenges.

The remainder of this article is organized as follows. Section II outlines the architecture of ADAS/ADS systems to provide the context within which automotive perception systems are used. Section III describes the different applications of automotive perception systems to infer information about the exterior and interior of the vehicle. Section IV provides details of state-of-the-art methods used to infer information from data acquired from major classes of sensors, namely, radar, camera, and LiDAR. It includes the challenges involved in deriving reliable and robust perception. Section V addresses the emergent validation domain, discussing the norms, safety metrics, and the monitoring of abnormal situations especially considering machine learning. Finally, Section VI concludes this article by presenting future opportunities for automotive perception.

II. GENERIC ADAS/ADS ARCHITECTURE

An ADAS/ADS, as depicted in Fig. 1, has multiple components [2]: a sensor system with sensors, data processing involving sensing and perception, a decision and planning system, and advanced driving controls and automotive services.

The perception system covers awareness information both external and internal to the automotive. External perception covers information on static and dynamic objects on the street, traffic, and street signs, and is obtained using sensors such as cameras, radars, and LiDARs. External perception provides a real-time picture of the dynamic environment around the vehicle, either by advanced processing of data from a sensor modality [3], [4], [5] or by fusing data from multiple modalities [6], [7]. Another aspect of external perception is localization to determine the location of the vehicle. A global positioning system (GPS) is commonly used for localization to provide a global position of the vehicle and a velocity estimate. A limitation of GPS is that the received signals suffer from blockage and heavy multipath conditions typical in dense urban environments. To compensate for resulting localization vulnerabilities in GPS, an inertial measurement unit (IMU) with sensors such as gyroscopes and accelerometers is employed. In such a fusion-based localization system, GPS position errors are corrected using IMU data along with additional constraints on vehicle motion, orientation, and its position on a map. Note that

external perception sensors can also be used to derive relative localization information using cameras [8] and radars [9]. Internal perception provides information on the occupants and objects in an automotive. Using sensor technologies, such as camera and radar, the presence, activity, and attention levels of the driver and passengers may be monitored to support various levels of autonomous driving. Understanding the state of occupants is crucial for effective human-vehicle interaction (HVI) in ADAS/ADS.

The decision and planning system determines the maneuvers for a vehicle. Driving decisions are based on previously acquired knowledge about the environment, such as the drivable area and traffic rules, and also real-time information, such as objects in the vicinity and traffic patterns. Decision and planning can be divided into three stages: global routing, behavior inference, and local motion planning. Global routing determines vehicle routes from source point A to destination point B according to some criteria, such as shortest travel time or least number of traffic signs encountered. This determination is done using graph routing algorithms, a digital map, and a traffic management system. After a global route has been determined, the automotive must be able to navigate the selected route and interact with other traffic participants according to driving conventions and rules. Given a sequence of road segments specifying the selected route, the behavioral inference stage is responsible for selecting an appropriate driving behavior at any point of time based on the perceived behavior of other traffic participants, road conditions, and other available signals from the infrastructure. The local motion planning stage translates the behavioral inference stage decisions into a feasible local path plan. It determines a path that is dynamically feasible for the automotive and comfortable for the passenger, and avoids collisions with obstacles determined by the perception system.

Vehicle driving control executes the reference path defined by the decision and planning system by selecting appropriate actuator inputs to carry out the planned motion path. Controls need to be accurate for safe automotive driving and robust under various driving conditions. As such, the control system should also be able to deal with diverse physical vehicle characteristics and dynamics. The vehicle control system is intimately linked to advanced driving functions, such as adaptive cruise control, emergency braking, and lane-keeping assistance; automotive services, such as assisted parking, traffic alerts, and diagnostics, further enhance the experience and safety of a user.

III. SENSORS FOR AUTOMOTIVE PERCEPTION APPLICATIONS

A. Automotive-External Perception

Different sensor-based vehicle applications, as depicted in Fig. 2, rely on information from the perceived environment for situational understanding and decision-making. These exteroceptive sensors acquire data from the vehicle's environment, which is then transformed into meaningful information, such as the occupancy grid map, the 3-D position of different traffic agents, and road characteristics (e.g., lane markings).

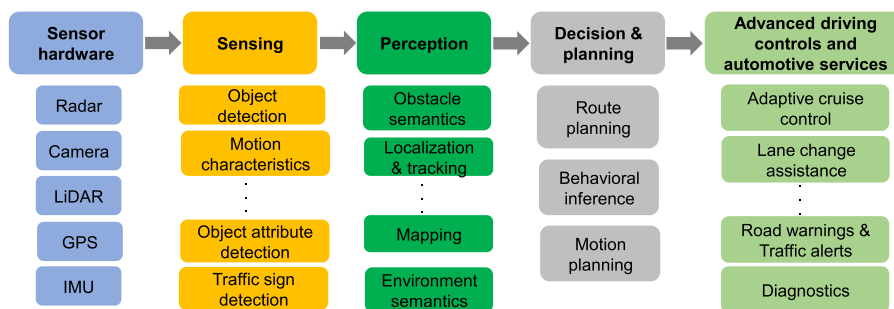


Fig. 1. Architecture of an ADAS/ADS equipped automotive.




Fig. 2. Driving features based on automotive-external perception.

There are two types of sensors: passive sensors, such as video cameras and infrared/thermal imaging sensors, which measure ambient environmental energy entering the sensor; and active sensors that emit energy into the environment and then measure the environmental response. These sensors can manage more controlled interactions with the environment; however, the emitted energy is limited by safety constraints or from interference between its signal and those of other active sensors, which may impact sensing performance (refer to [10] and [11] for automotive radar transmission power limits). Examples of active sensors include ultrasonic, LiDARs, and radar.

Exteroceptive sensors are sensitive to the operating outdoor environment conditions, and their performance can vary considerably according to their deployment location and weather conditions. The performance of perception algorithms depends on the ODD where the automotive operates [12], with various factors coming into play, such as dynamic range (e.g., the ratio of the largest to the smallest measurable signal for a radar), range (e.g., how far or how near a radar or LiDAR can detect), resolution (e.g., the number of pixels in an image), and frame rate (e.g., the rate at which data are acquired or the frames per second of a camera). In addition, the performance of perception sensors is also impacted by their layout. For example, while radars are commonly placed around vehicle bumpers or brand emblems, the same placement may not be suitable for other modalities, such as cameras, due to an impeded field-of-view. Environment conditions have a role in sensor lifetime performance; very low or high temperatures, dust, and humidity are the factors that affect sensor performance, and means to weather-proof sensors or to service them must be accounted for.

B. Automotive-Internal Perception

In automotive-internal perception, also termed in-cabin monitoring, systems are an indispensable feature of vehicle safety systems. They are a part of cyber-physical human systems (CPHSs) capable of responding or taking actions based on the perception of the human condition in the vehicle, as shown in Fig. 3. Four principle tasks have been the focus of development: 1) occupancy detection/characterization; 2) driver monitoring; 3) passenger monitoring; and 4) HVI. Occupancy detection/characterization refers to the ability to automatically detect where people are located in the vehicle and especially the presence of children and infants. Occupancy sensing can provide valuable input on the proper use of child seats and seat belts while also enabling the optimization of airbag function according to the height of the person sitting in the seat. This function is also an important safety feature for preventing the death of children and pets left in vehicles on hot days, as a vehicle perceiving this situation could take preventive measures, such as notifying the owner of the vehicle, turning on the air conditioning, or opening a window. New U.S. federal regulations require all new cars to be equipped with a back-seat alert system [13], while the European New Car Assessment Program also prescribes new child presence detection (CPD) protocols [14] to prevent in-vehicle heat-related child deaths.

Driving monitoring systems are primarily targeted toward ensuring safety by monitoring a driver’s ability to effectively drive the vehicle. This can be indicated by a variety of measurable variables, such as driver vital signs (heart rate and respiration), driver fatigue, drowsiness, and attention. Examples of attention monitoring include tracking the direction that the driver is gazing, head movements, and eye blinking, especially blink duration and frequency. Because there is a correlation between fatigue and heart rate, vital sign monitoring can be used for the detection of critical health events, such as a heart attack, but also for drowsiness detection. Other health-related indicators that have been considered include blood pressure measurement and blood glucose level monitoring. Passenger monitoring includes features of vital signs and health monitoring, as well as general activity within the vehicle. Especially, when children are present, monitoring seat belt usage and whether potentially dangerous passenger activity is occurring (such as children changing seats) can be important for safety.

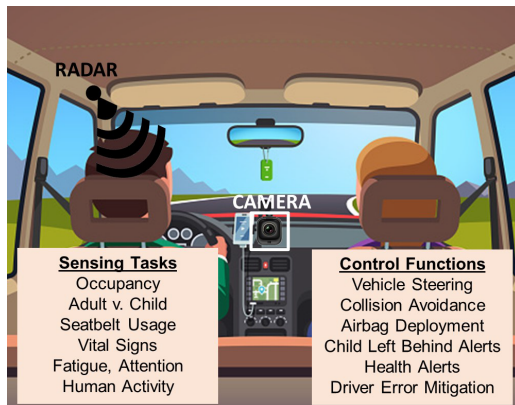


Fig. 3. Driving features based on automotive-internal perception.

HVI has been predominantly considered within the context of noncontact gesture recognition for control of user interfaces and vehicle subsystems, such as the radio, infotainment systems, or air conditioning, among others. However, the prospect of autonomous vehicles raises new dimensions in HVI, whereby two-way human communications with the vehicle must also be considered. Examples motivating such functionality include not just in-cabin HVI, but interactions that might occur if a passengerless autonomous vehicle were pulled over by a police officer. HVI also arises in environmental scene understanding outside the car, as an autonomous vehicle must also be able to navigate based on directions from a police officer directing traffic with the additional consideration that some detours may not be well marked with driving lanes. Moreover, while current collision avoidance systems simply try to detect and steer away from obstacles, HVI systems of the future could also include pedestrian injury mitigation features in the event of a collision.

Currently, automotive-internal perception systems implemented in commercially available vehicles rely on video-based technologies [16] for extracting passenger information, such as occupancy, posture, whether part of the body is outside the vehicle, seat belt usage, use of a car seat, whether a child has been left unattended, whether there are objects or belongings in the vehicle, and detection of pets, animals, or miscellaneous irregular situations. In addition to these functions for safety, video can also be used for security purposes to detect incidents of vandalism, abuse, the presence of dangerous weapons, theft, or other illegal activities. The video has also been used to enable noncontact user interface control via gesture recognition. In this case, a small camera is positioned such that it can monitor the area immediately in front of the touch screen next to the driver.

Although cameras have enabled some in-cabin monitoring features, their performance is adversely impacted by changing and low ambient light conditions. There has been an increased interest in radars recently for in-cabin monitoring, as millimeter-wave radars can be used to remotely measure vital signs, track eye blinking, recognize gestures, and detect vehicle occupancy. Radar offers sensing solutions that are less invasive of privacy in comparison to the video while not being dependent upon ambient lighting conditions. Over the last few years, several studies examining radar-based

occupancy detection have been published. While some studies have considered continuous wave [17], pulsed [18], and impulse radio ultrawideband (IR-UWB) radar [19] systems, most works have focused on the utilization of high-range resolution FMCW radars [20], especially multichannel FMCW [21], [22], [23] due to its greater angular resolution.

Driver attention monitoring studies have focused on radar-based vital sign recognition [24], [25], [26], monitoring of driver head movements [27], eye blinking [28], [29], [30], [31], fatigue, concentration and drowsiness [32], [33], [34], [35], and health indicators, such as blood pressure [36] and blood glucose levels [37]. The presence of body movements and respiration effects both the measurement of heart rate as well as that of eye blinking frequency and duration. Thus, methods for jointly estimating heartbeat and blink rate have also been proposed [38].

Gesture recognition using radar was postulated early in [39] but was made practically possible with the development of integrated, millimeter-wave RF transceivers. Radar-based gesture recognition gained significant attention due to the Google SOLI project [40]. While gesture recognition research has predominantly focused on the design of deep neural networks (DNNs) to improve the classification accuracy of common, ubiquitous hand gestures, such as a virtual knob, slider, and push button, the efficacy of radar has also been demonstrated for sign language recognition [41]. In automotive environments, studies have considered radar-based intelligent driver assistance [42], the utilization of LiDAR, camera, and radar for traffic signaling gesture recognition [43], and radar-enabled HVI for in-car infotainment control [44]. In-vehicle behavior and gesture recognition have also been proposed using Wi-Fi signals [45].

IV. PROCESSING AND LEARNING ALGORITHMS

As discussed in Section III, exteroceptive sensors exhibit considerably varying performance and sensitivity to operating outdoor environment conditions. This is manifested in Fig. 4, where the stereo camera, radar, and LiDAR frames from the RADIATE dataset [15] are plotted under various light/weather conditions (e.g., sun, night, fog, rain, and snow) in different driving scenarios (e.g., urban, suburban, and highway). Specifically, Fig. 4 clearly shows the adverse impact (e.g., blur) of fog, rain, and snow on the stereo camera, while LiDAR frames show dense false point clouds due to snowflakes in snow conditions. On the other hand, radar frame quality is limited by its resolution, particularly at long distances. In the following, we overview the processing and learning algorithms for each sensor modality and the need for sensor fusion.

A. Radar

Automotive radar has been traditionally a part of ADAS for safety features, such as emergency braking, adaptive cruise control, and self-parking systems. These features have been enabled using traditional frequency-modulated continuous-wave (FMCW) chirp signals along with signal processing techniques, such as pulse compression, low-pass filtering, analog-to-digital conversion (ADC), fast Fourier transform

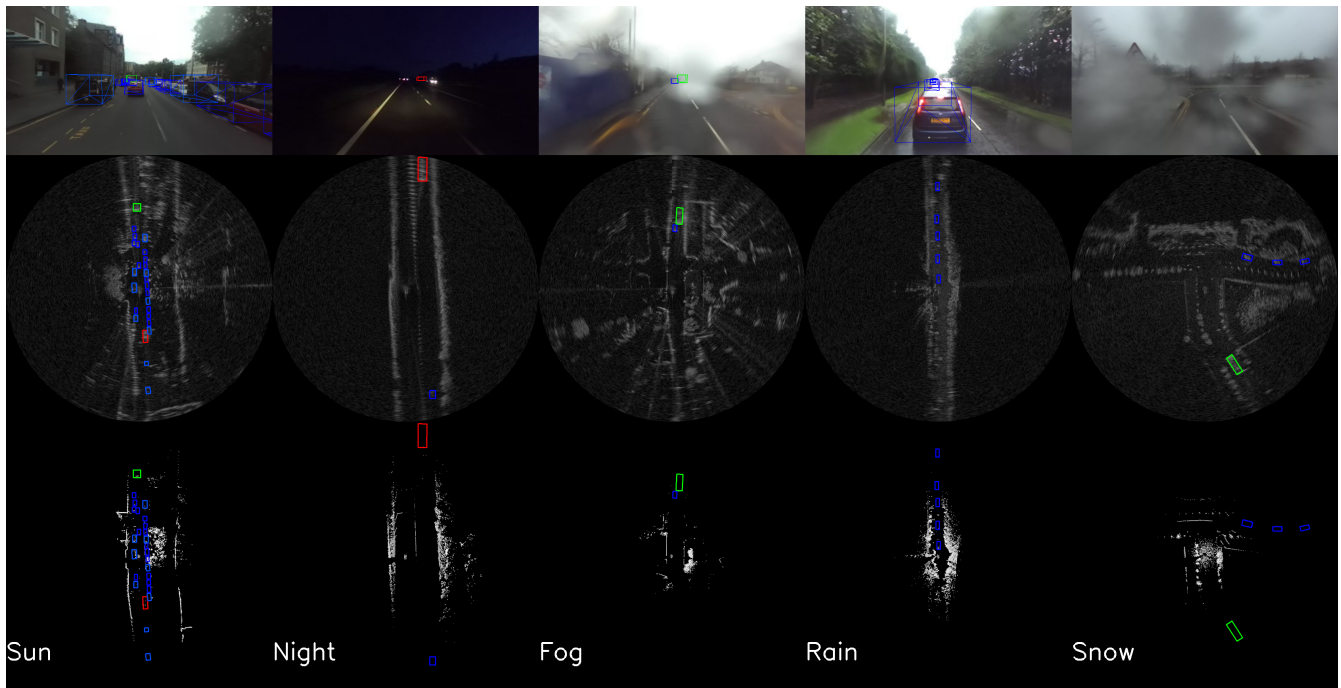


Fig. 4. Automotive external perception frames of the stereo camera (top row), radar (middle row), and LiDAR (bottom row) under different weather and light conditions (sun, fog, rain, snow, and night in five columns). Figures are plotted from data in the open RADIATE dataset [15].

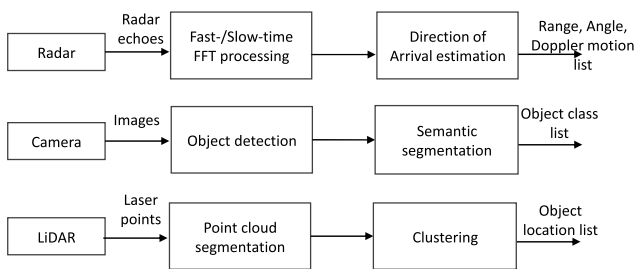


Fig. 5. Illustrative processing for radar, camera, and LiDAR external perception.

(FFT), Doppler processing, clutter removal, and constant false alarm rate (CFAR) detection. For each receiver RF chain, by applying FFTs on the fast-time and slow-time samples of the baseband signal, as depicted in Fig. 5, the range-Doppler (RD) heatmap can be constructed with options to apply a window function to suppress sidelobes. These RD heatmaps from multiple receiver RF chains can be combined to increase the signal-to-interference-and-noise ratio (SINR). Then, radar detection is performed by comparing the output power of a receiving (Rx) filter with a threshold. If the measured power exceeds the threshold, an object is detected, and the associated signals are then processed to estimate object parameters, such as range, radial velocity (or range rate), and angles. In this context, a CFAR detection or its variant is commonly used to maximize detection probability while maintaining a fixed probability of false alarm for a given SINR. One example of CFAR detection is the cell-averaging CFAR, which computes the threshold from the average power of neighboring RD cells that are separated by several guard cells to the cell of interest (COI) to avoid possible object contamination and determine whether the COI contains an object (H_1 hypothesis) or not (H_0 hypothesis) by comparing the detection statistic with

the threshold. A comprehensive overview of basic FMCW-based automotive radar signal processing techniques can be found in [46] and [47]. Conventionally, these automotive radar systems were designed to achieve desired resolution and maximum unambiguous limits in the range and velocity domain by optimizing waveform parameters (e.g., bandwidth and chirp period).

Beyond the safety features, current automotive radar is limited in terms of its pixel information (i.e., limited azimuth and elevation angular resolutions) to deliver high-quality perception information [46]. To meet the challenge, a new wave of chip developments eyes on improving the angular resolution of automotive radar [59] to deliver LiDAR-like radar perception with enriched semantic features. Particularly, high-resolution 3-D (range-velocity-azimuth)/4-D (3-D+elevation) radar images enable the ability to correctly detect and classify objects in the environment. For example, it is crucial to differentiate overhead objects, such as bridges and traffic signs, while being aware of low-lying objects, such as manholes, road debris, and curbs.

Since the angular resolution is determined by the beamwidth that is inversely proportional to the aperture size, chip vendors take various approaches to form the beam and synthesize a large aperture. Mechanically scanned FMCW radars, e.g., Navtech CTS350-X, have been used to collect 360° bird’s-eye view (BEV) radar images in the range-azimuth domain but without the Doppler velocity [15]. Assuming that the ego vehicle’s motion is known, synthetic aperture radar (SAR) techniques can coherently combine returned radar waveforms to create a high-resolution 2-D image of the scene [60]. For instance, a 0.1° azimuth resolution was achieved for imaging static objects and can deliver ~1 000 000 points for a typical scene [61].

TABLE I
OPEN AUTOMOTIVE RADAR DATASETS

Datasets	Year	Size	Annotation	Data Format	Adverse Weather
nuScenes [48]	2020	large	bounding box, Track ID	sparse points	Yes
RadarScenes [49]	2021	large	Point Annotation	sparse points	Yes
RADIATE [15]	2020	medium	bounding box, Track ID	dense points (scanning)	Yes
Oxford robotcar [50]	2020	large	object pose	dense points (scanning)	Yes
MulRan [51]	2020	large	object pose	dense points (scanning)	No
TJ4DRadSet [52]	2022	medium	bounding box, Track ID	dense points	No
VoD [53]	2022	small	bounding box, Track ID	dense points	No
CARRADA [54]	2020	medium	bounding box, Track ID	RAD heatmap	No
CRUW [55]	2021	medium	object class, position	RAD heatmap	No
RaDICAL [56]	2021	large	bounding box	Raw data	No
RADDet [57]	2021	medium	bounding box	RAD heatmap	No
RADial [58]	2021	medium	object, segmentation	RAD heatmap+points	No

To achieve high angular resolution, another popular approach is to use multiple-input–multiple-output (MIMO) radar [62], where multiple N_t transmitting (Tx) and N_r Rx antennas are used to form a virtual array with $N_t N_r$ elements. The combined MIMO-FMCW automotive radar only employs $N_t + N_r$ RF chains to reduce the hardware cost. The shape of the virtual array is determined by the convolution of the transmitter array and the receiver array. To achieve this, one needs to separate the corresponding waveform to each transmitter at each receiver, provided that the Tx waveforms from different Tx antennas can be separable or orthogonal. Several orthogonal MIMO signaling schemes can be realized in time-division multiplexing (TDM), frequency-division multiplexing (FDM), and Doppler-division multiplexing (DDM) (also referred to as slow-time MIMO) modes [63], [64], [65], [66]. Once the waveforms are separated, the received MIMO radar waveforms can be arranged along the fast-time (range/distance), slow-time (Doppler/velocity), azimuth (horizontal array orientation), and elevation (vertical array orientation) dimensions. Depending on the computing resources and power consumption budget, the MIMO-FMCW automotive radar may coherently process the separated MIMO waveforms in one or more dimensions at once or in a cascading fashion (e.g., RD domain first and then angular domains) with standard FFT operations or more advanced super-resolution spectrum estimation methods, such as MUSIC, ESPRIT, and compressed sensing. The TDM-MIMO mode has been commercialized by various chip vendors due to relatively easy implementation and less computational requirements on the waveform separation to achieve more than 100 or even ~ 2000 virtual channels in the azimuth and elevation domains [63], [67], [68].

Open radar-included automotive perception datasets have begun to emerge by using commercial radar chips (see Table I). Earlier efforts focused on the collection of radar detection points for perception. nuScenes [48] is one of the first large-scale (one million annotated frames) automotive perception datasets with commercial automotive radar included. However, due to poor angular resolution, the radar point clouds per vehicle are very sparse (e.g., less than 10). RadarScenes [49] combines detection points from four automotive radar sensors operating at 79 GHz to boost the number of detection points per object. For high angular-resolution automotive radar datasets, the Oxford radar

robotcar [50], MulRan [51], and RADIATE [15] (see the middle row of Fig. 4) datasets used mechanically scanning FMCW radar to get sub- 1^{circ} angular resolution for BEV radar images but without the Doppler velocity information. More recent efforts focus on the radar heatmap in RD, range-angle (RA), and RA-Doppler (RAD) domains [54], [55], [56], [57], [58]. For instance, the CARRADA dataset provides three (RA, RD, and RAD) heatmaps on a scenario of test tracks [54].

Driven by the availability of high-resolution automotive radar hardware platforms and open datasets, advanced model-based signal processing and learning-based pipelines [69] have shown great potential to achieve state-of-the-art performance in radar-assisted object detection, segmentation, multiobject tracking, simultaneous localization and mapping, trajectory and behavior prediction, multimodal sensor fusion, and scene understanding.

1) *Radar Detection Point*: For sparse radar detection points, model-based object detection and tracking algorithms have been considered in the context of extended object tracking (EOT) [70]. One of the key challenges for EOT is to model the spatial distribution of radar detection points over the extent of an automotive vehicle and the subsequent prediction and update of expanded state (e.g., vehicle position, orientation, speed, turn rate, length, and width) using Bayesian filtering [71], [72], [73], [74], [75], [76], [77]. These sparse radar detection points can also be processed by heuristic distance-based clustering algorithms, such as DBSCAN or PointNet-like learning algorithms. For instance, PointNet [78] and PointNet++ [79], [80] were applied to segment radar detection points and estimate 2-D bounding boxes from those segmented radar points by taking into account unique radar features, such as the Doppler velocity and radar cross section (RCS) [81]. A two-branch radar point segmentation network was considered in [82] with a convolution network branch performing semantic segmentation over radar grid maps in the static environment and the other recurrent segmentation network on radar point clouds of moving objects. Then, a merging step takes the output class probabilities of each cell in the grid map from the two classifiers to form point clouds.

Compared with sparse points, dense radar points from high-resolution automotive radar platforms (e.g., scanning-based, SAR, or MIMO-based) may enable radar feature extraction

at a level closer to LiDAR perception networks. Particularly, PointPillars was applied to the 4-D radar data in a new VoD dataset [53], and a reduced performance gap in terms of object detection can be achieved between a 64-line LiDAR sensor and a high-resolution 4-D radar sensor with the utilization of elevation resolution and integration of successive radar frames. Bai et al. [83] proposed a radar transformer that uses both vector and scalar attention mechanisms to construct attention maps over 3-D (spatial, Doppler, and RCS) domains. Exploiting the temporal relation of successive radar frames can further enhance feature extraction of the radar point cloud. Li et al. [84] proposed a cross-attention network that exploits the consistency of objects over successive radar frames in different levels (e.g., the input level by permuting the frame order and the feature level by introducing the cross-attention feature module). These selected temporally enhanced features are then used to regress oriented bounding boxes (OBBs) at each of successive radar frames, similar to the CenterPoint framework [85].

2) *Radar Heatmap*: Radar heatmaps may have more semantic features for low-RCS and static objects than the after-CFAR radar detection points. It might be arguable that the neural network-based feature extraction can be more representative of the heatmap domain for complex-shape objects than the model-based CFAR detection and point-based networks. For instance, the CA-CFAR detection is optimal in the Neyman-Pearson criterion if the noise and interference amplitude are Rayleigh distributed, which may not hold in practice. Moreover, the choice of guard cells is critical to avoid object contamination. Direct heatmap-based approaches, on the other hand, skip the traditional model-based CFAR detection and can directly backpropagate the training loss into the heatmap feature extraction networks. To this end, image-based backbone networks and downstream pipelines have been applied to the radar heatmap [86]. A straightforward way is to treat the 3-D RAD heatmap as an RGB image but with the number of channels the same as the number of Doppler bins. Then, image-domain object detection, segmentation, and tracking pipelines can be applied to the input RGB-like radar heatmap [57], [87], [88]. In [57], a one-stage anchor-based YOLO framework with a dual detection head was considered to generate both 3-D RAD and 2-D Cartesian bounding boxes. Evaluation on their own RADDet dataset shows a 56.3% average precision (AP) at Intersection over Union (IoU) of 0.3 on 3-D bounding box predictions and, respectively, 51.6% at IoU of 0.5 on 2-D bounding box prediction. To reduce the input dimension of a full 4-D heatmap with two angular azimuth and elevation domains, the 4-D RAD heatmaps can be decomposed or projected into multiple 2-D heatmaps. Gao et al. [87] proposed an radar multiple-perspectives (RAMP)-CNN approach to bypass the complexity of 4-D convolutions and fuse extracted features from multiple 2-D heatmaps. Ouaknine et al. [88] take three projected 2-D “views” of the RAD heatmaps as the input, and three feature extraction networks are used for each projected view. These view-dependent features are then concatenated and fed to a decoder for the segmentation task. Evaluation of the CARRADA dataset shows a mean

IoU at 58.7 on the RD heatmap and 41.3% on the RA heatmap over 4 categories of pedestrian, cyclist, car, and background.

The tremendous progress in recent years, driven by the open access of high-resolution automotive radar datasets, provides a promising future for radar-based external perception. For the external perception task, achieving higher angular resolution radar detection points or heatmaps poses a further strain on the cost, computational resources, and power consumption budget. Similar to the ResNet and vision transformer, strong and unified radar-specific feature extraction backbone networks are needed. The need for strong backbone networks also calls for diverse downstream tasks, such as object detection, segmentation, and tracking, or self-supervised learning without any (or with limited accuracy) annotation labels.

B. Camera

As evident from Fig. 5 and the top row of Fig. 4, camera-based images provide distinct features to separate objects of interest (e.g., vehicles and pedestrians) from the background, and such features can be integrated into state-of-the-art deep learning frameworks for downstream tasks, including object detection, classification, depth regression, object association and tracking, pixel, and instance segmentation [89].

Large-scale datasets involving cameras have been collected for the development of autonomous driving technologies. The details of these datasets are listed in Table II. The most widely used ones are the KITTI and BDD100K because of their early release with various sensor signals or a large number of images in various traffic situations to examine the effectiveness of developed methods.

Current learning-based methods can be generally classified into two categories: two-stage detection and one-stage detection. The two-stage detection, also called region-based detection, first scans the complete image to find the potential regions of interest and then focuses on these regions for deeper understanding, which generally imitates the attentional mechanism of the human brain [103]. The two stages for detection are, respectively, responsible for generating a set of proposals and making predictions for these proposals. During the proposal generation phase, a set of proposals is generated. In the prediction phase, the feature vectors of generated proposals are encoded by deep convolutional neural networks (CNNs), and then, classifiers are used to determine the category labels of the proposals [104]. R-CNN is a pioneering two-stage object detector proposed by Girshick et al. [105]. Xie et al. [106] proposed an improved object detection approach based on R-CNN. High-quality oriented proposals were first generated in an almost cost-free way, and then, regression and classification technologies were used for prediction. The testing results on two datasets, including DOTA and HRSC2016, showed that the mean APs (mAPs) were 75.87% and 96.50%, respectively. Despite the advances in learning detectors based on R-CNN networks, proposal generation still relies on traditional methods such as selective search [104]. Studies [107] and [108] show that CNN has a remarkable ability to locate objects in convolution layers. Therefore, the faster CNN method is proposed in [109]

TABLE II
OPEN-SOURCE DATASETS FOR AUTONOMOUS DRIVING

Datasets	Year	Number of images	Viewing Angle	Resolution	Camera Type	Adverse Weather
CamVid [90]	2008	18×10^3	Dashboard	960×720	Monocular	No
KITTI [91]	2012	15×10^3	Vehicle Roof	1392×512	2 Stereo	No
Cityscapes [92]	2016	25×10^3	Windshield	–	Stereo	No
Oxford RobotCar [93]	2017	19×10^3	360 Degree	1280×960 1024×1024	Stereo & Monocular	Yes
Mapillary [94]	2017	25×10^3	–	1920×1080	–	Yes
BDD100K [95]	2017	100×10^6	Windshield	720p	Video	Yes
ApolloScape [96]	2018	144×10^3	360 Degree	3384×2710	Stereo	No
CULane [97]	2018	130×10^3	Windshield	1640×590	–	No
H3D [98]	2019	25×10^3	Vehicle Roof	1920×1200	3 Monocular	–
NuScenes [48]	2019	1.4×10^6	360 Degree	1600×900	6 Monocular	Yes
Foggy [99]	2019	14×10^3	Windshield	960×1280	Stereo	Yes
Sim 10K [100]	2017	10×10^3	Windshield	–	Gaming Engine	Yes
TuSimple [101]	2017	6.5×10^3	Windshield	1280×720	–	No
RDD2020 [102]	2020	26×10^3	Windshield	600×600 720×720	Smartphone	No

by developing a region proposal network based on CNNs. Hnewa and Radha [110] proposed to use faster R-CNN for object detection in rainy weather for autonomous driving. The detection results show that faster R-CNN incorporating image translation and domain adaptation performed the best among the examined methods in rainy weather.

Different from two-stage detection algorithms that divide the detection pipeline into two parts, one-stage detection assumes that each region in the image is with a possible detected object, and each region of interest is categorized into background or target object without a separate stage to generate proposals [104]. Directly mapping from image pixels to bounding boxes with category probabilities generally saves time compared with the two-stage methods [103]. According to the searching methods for areas of interest, the one-stage methods can be further divided into anchor-based methods and anchor-free methods. The main idea of anchor-based methods is to predict searching anchors based on the prior definition. Anchor boxes with different sizes slid over each position of an image, predicting the searching anchor box as background or object based on the ground-truth predefined anchors. YOLO [111], as a typical anchor-based method, considered object detection as a regression problem and spatially divided the whole image into a number of grid cells. Each cell was considered as a proposal to detect the presence of objects [103]. Li et al. [112] proposed three deep learning methods for pedestrian detection in haze weather based on YOLO. The evaluation results showed that the proposed method MNPrioriBoxes-Yolo with separable depthwise convolution and bottlenecks had obvious advantages with fewer parameters for pedestrian detection in haze weather. Improved lightweight detection algorithms based on YOLO (e.g., YOLOv4 and YOLOv5) further enhance the recognition ability on small objects with a limited number of parameters [113].

The main idea of anchor-free methods is to use keypoints to describe the boxes used for detection; hence, the main task is transformed into keypoint detection. The related methods have two branches, namely, corner-based methods and center-based methods [104]. For corner-based methods, also called

multiple keypoints estimation methods [114], the confidence scores of bounding boxes are predicted through joint corner information in the feature map. Compared to RepPoints [115] using nine keypoints, Yang et al. [116] used a large number of adaptive points to model objects, which achieves state-of-the-art performance on instance segmentation tasks. Center-based methods simplify object detection to a central point detection task by estimating the probability of a pixel as the central point. Inspired by the region proposal network in the anchor-based method, Fan et al. [114] proposed the foreground information introduction CenterNet (FII-CenterNet) method for traffic object detection based on CenterNet [117].

For improvement based on these two-stage or one-stage methods, many advanced data augmentation and deep learning technologies have been developed to help learn effective features for better prediction. Data augmentation changes the characteristics of images by cropping, flipping, rotating, scaling, translating, color perturbations, and adding noise to enrich the diversity of data samples for training to learn stable features [108]. The mainly incorporated deep learning modules for performance improvement include attention mechanism, pyramid pooling, linear bottleneck and inverted residuals, depthwise separable convolution, atrous convolution, knowledge distillation, domain adaptation, SEBlock, ResBlock, mask mechanism, network pruning, and quantification [110], [118] [119], [120].

Another advanced technology used to improve the performance of learning-based algorithms in cameras is a transformer proposed in [121]. Inspired by the success of the transformer in neurolinguistic programming, the transformer has also been widely used in computer vision tasks. These methods mainly include a pure transformer, a transformer with convolution, and self-supervised representation learning with a transformer. Transformer with convolution combines transformer modules with convolutional network modules, and self-supervised representation learning uses transformer self-supervised mechanism for training. As a typical pure transformer, vision transformer [122] divides the 2-D image data into image blocks as the input to the standard transformer for supervised training. Liu et al. [123] adopted the transformer

encoder structure and the convolution module for lane detection, and the verification results showed that optimal performance was achieved in both efficiency and accuracy. Here, a novel decoder with dense queries and rectified attention field was proposed, which alleviates the deficiency in pedestrian detection by using the transformer decoder DETection TRansformer (DETR).

Most existing approaches are supervised and rely on large-scale datasets with reliable annotated labels, which is difficult to obtain especially for extreme weather and driving situations. Developing unsupervised and weakly supervised learning algorithms is a promising solution. In [124], an image-level multilabel classifier was integrated into the detection backbone to obtain sparse but critical image regions corresponding to the classification information to bridge the gaps between source and target domains. In [125], a cross-domain adaptive clustering approach was proposed by pulling into distances between peers while simultaneously pulling away distances from different categories, which achieves state-of-the-art performance in semisupervised domain adaptation. Although approaches have been developed for unsupervised solutions [126], [127], [128], more efforts are still needed for further improvement in this research area.

The other remaining challenges in camera-based processing and learning algorithms include the following.

- 1) Deep learning-based methods are end-to-end with insufficient model interpretability. More attention to provide model insights is required in the future.
- 2) Most of the learning-based methods are with complex networks, with high computation requirements on hardware. This makes these ML methods infeasible to implement in the current generation of ADAS. Lightweight ML technologies need to be developed for wide deployment in automotive.
- 3) Detection of small objects is still not satisfactory. Fusing the signals from multiple sources for multimodal fusion may be a solution to this challenging issue.

C. LiDAR

As seen in Fig. 4, LiDAR provides a direct depth profile over the angular (azimuth and elevation) domains with a resolution even finer than the automotive radar. Most LiDAR-type sensors are based on different types of scanning mechanisms that allow for the laser beams to be projected over a large field of view following specific patterns and according to the technology used. Scanning can be through mechanical spinning or a solid state. The former often includes a bulky rotating module, which will spin a mirror around a vertical axis and tilt it along the pitch orientation. The latter refers to a scanning system through micromirrors based on MEMS technology [129]. A LiDAR generates streams of 3-D-points, with intensity data associated with each point (proportional to the reflected signal). Unlike video cameras, LiDARs measure the distance by applying mostly the time of flight method. Photonics principles are part of the technology used to operate with light sources that need to convert signals rapidly and process them to generate the desired measurements while, at the same time, mapping the direction

and position of the beams as they scan to attain the sensor field of view. It must be considered that there is substantial processing prior to the generation of the sensor data; thus, purpose-built processors are often used that need to comply with automotive operating standards (e.g., temperature and vibration).

The data streamed out from the LiDAR are interpreted through perception algorithms into hierarchical object descriptions. This process can be divided into ground segmentation, object detection, tracking, recognition, and motion prediction. Previously, these phases were addressed separately using geometric and early machine learning techniques (e.g., support vector machines, based on statistical learning frameworks). However, the success encountered in the use of deep learning in machine vision is also reflected in its effective use of 3-D-point clouds. Deep learning technologies can automatically extract features from the raw input in a single phase. CNNs and recurrent neural networks (RNNs), such as long short-term memory (LSTM), are the most frequently used models. Ground segmentation can be achieved by applying CNN to LiDAR points represented by multichannel range images [130]. DNN-based solutions achieve object detection by recognition, keeping to the paradigm of supervised learning. For example, vehicles can be detected by CNN-based neural networks on a BEV representation of LiDAR 3-D-points [131]. However due to the low density of 3-D LiDAR points at long distances on the perceived objects, CNN methods on the range image and BEV representation have difficulty in detecting pedestrians [132]. A compact representation of a LiDAR point cloud as a graph was proposed as a point-graph neural network (GNN) method in [133]. A semisupervised using temporal GNNs to leverage the rich spatiotemporal information in 3-D LiDAR point cloud videos for object detection was considered in [134]. A novel approach is to integrate evidential theory into a deep learning architecture for LiDAR-based road segmentation and mapping [135]. Currently, object tracking is implemented mainly using deep learning, replacing the conventional tracking algorithm based on estimation filters [136]. A detection net will process first a sequence of LiDAR 3-D-points and images to generate detection proposals. Then, tracks are estimated by finding the best detection associations. This is achieved by a marching net and a scoring net.

Pointwise semantic segmentation, which was previously difficult to attain using model-based methods, is now possible using deep learning models. One of the most popular networks is PointNet that provides a unified architecture for applications ranging from object classification and part segmentation to scene semantic parsing, directly from LiDAR 3-D-point clouds [78]. As the point cloud density is increased, together with more annotated datasets, LiDAR performance for semantic segmentation should improve providing not only classes of objects but also spatial information. A major constraint for the application of these methods is the need for large datasets. However, different annotated datasets have recently emerged including the SemanticKITTI dataset [137]. It is based on the KITTI dataset and is considered one of the largest pointwise annotated datasets. Synthetic datasets, such

as the PRESIL dataset [138], which provide labeled scenarios for particular situations are also available.

D. Fusion

Table III shows a comparison of different sensor modalities in terms of sensing and operational features. Sensing features depict sensing performance characteristics, while operational features capture robustness and system integration aspects. Typically, radars can provide range, velocity, and angular information with high resolution, in comparison to visible light and infrared cameras. Cameras on their own are unreliable in situations of abrupt change in illumination, such as when entering/exiting a tunnel or extreme weather conditions. LiDARs also suffer from performance degradation in extreme weather conditions, while radars are robust under adverse weather, environmental, and illumination conditions. Compared to radars or LiDARs, cameras can capture contour, texture, and color information of the scene enabling excellent recognition capabilities under nonextreme environments. Although LiDARs are superior to radars in ranging accuracy and denser point cloud, their cost is much higher and has a larger form-factor, making it difficult for integration in a flexible and esthetic way. In summary, visible light cameras are superior in determining object features and, hence, find use in traffic scene/sign understanding; radars have better performance in determining object motion characteristics with high resolution and low cost, while LiDARs have superior ranging performance with a wide detection coverage. However, each sensor modality also has limitations, with no single modality providing the needed sensing and perception functionalities.

The aim of sensor fusion is the collective processing of inputs from various modalities to perceive and derive interpretations with a defined level of certainty about the environment around the vehicle. Based on the discussion in earlier sections and depicted in Fig. 4 and Table III, it is clear that each individual sensor cannot work independently under all scenarios and deliver accurate information with the precision required to operate an autonomous vehicle with the highest degree of safety. Sensor fusion allows information from all sensors to be fused meaningfully to extract the best of all sensors while offsetting the disadvantages of an individual modality.

For ADAS/ADS, it is important that the sensor fusion architecture combines data or processed data at different meaningful stages of the pipeline. To enable perception functions, there are three fundamental sensor fusion approaches to associate and integrate data across modalities to enable an informed decision [139], [140], [141].

- 1) *Late Fusion*: Each sensor is operating individually, and then, the processed data, i.e., likelihood function, get fused at the end to make a collective decision for the system. In [142], a multimodal vehicle detection system employing a late fusion strategy was proposed combining optical image and 3-D LiDAR detections. Individual modalities have their own detection pipeline, and then, the detection information is fused via joint rescaling and

nonmaximum suppression, and demonstrated improved detection performance over single modality detection.

- 2) *Early Fusion*: Sensor fusion happens at the initial data stage with no to minimal data preprocessing to align and normalize the raw data. The fused data are collectively used to improve the detection, classification, segmentation, and monitoring of the objects. In [143], early fusion of LiDAR and camera data into a multidimensional occupation grid representation as input to fully convolutional networks for lane detection was proposed.
- 3) *Mid-Level (or Cross) Fusion*: This fusion approach combines the early and late fusion approaches. Targeted information derived from different sensors is fused at an initial data stage, given that a certain predefined criterion is fulfilled, while other target information is fused at higher levels under other predefined criteria, such as low signal-to-noise ratio conditions. In [144], radar detections were associated with preliminary detection results obtained from a camera image and then generate radar feature maps in addition to image features to estimate 3-D object bounding boxes. Camera and LiDAR features are fused in a shared bird's eye view space in [145] showing improved mAP for 3-D object detection in comparison to [144] and individual sensors in adverse weather and illumination conditions.

V. VALIDATION METHODOLOGIES AND SAFETY CONSIDERATIONS

Evaluating the performance of a perception system is a complex task. It involves regulation concerning software to ensure safety, defining system performance metrics, and designing methodologies for robust ML perception systems.

A. Safety Standards and Guidelines for Perception

Currently, ISO 26262 [146] defines processes and measures for the functional safety of systems including one or more electrical and/or electronic systems. The requirements in ISO 26262 are considered to be sufficient to deal with risks due to random hardware faults or classic systematic software faults (e.g., array-out-of-bounds) for sensors.

While a system involving ML components may (ideally) be free of hardware or systematic software errors as governed by ISO 26262 (functional safety), the performance limitations of ML (functional insufficiencies) within the ODD can still lead to risks. ISO 21448 [147] focuses on processes and measures to ensure the absence of unreasonable risk due to a hazard caused by functional insufficiencies, where the performance limitation of sensors is also explicitly mentioned as one of the sources of functional insufficiencies. The basic safety specification considers the occurrence of an error pattern [148] within the ODD being sufficiently low, commonly reflected as a probability term. While the appendix of the ISO 21448 covers some high-level aspects and some of the process-oriented results aim at offering a general argumentation framework [148], [149], [150], [151], [152], we consider the key technical challenge to be the “implementation aspects” of such a process.

TABLE III
COMPARISON OF SENSORS FOR EXTERNAL PERCEPTION

		Visible Light Camera	Infrared Camera	Radar	LiDAR
Sensing Features	Distance measurement	Medium	Medium	Very High	High
	Velocity measurement	Low	Low	High	Low
	Angle measurement	Low	Low	Medium	High
	Measurement resolution	High	Low	High	High
	Object Features	Color & contour	—	Intensity	Intensity
	Field of View	Medium	Medium	Medium	360°
	Sampling Rate	High	Medium	Medium	Low
Operational Features	Weather (rain, snow, fog)	Vulnerable	Vulnerable	Robust	Vulnerable
	Visibility (dust, smoke)	Vulnerable	Vulnerable	Robust	Vulnerable
	Illumination (low-light, glare)	Vulnerable	Robust	Robust	Robust
	Sensor Interference	None	None	Yes	Yes
	Processing Requirements	Medium	Medium	Medium	Medium
	Sensor Layout & Aesthetics	Good	Medium	Good	Low
	Costs	Low-Medium	High	Low	High

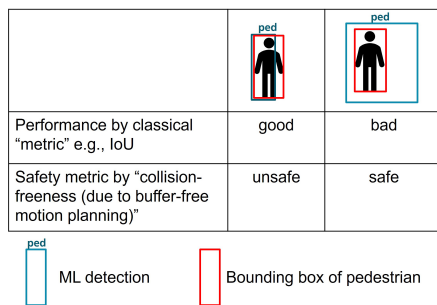


Fig. 6. Illustration where standard performance metrics may not be directly used in safety-critical driving contexts.

ANSI/UL 4600 [153] is a standard to promote a proper consideration of safety issues for generic autonomous systems and specifically adopt autonomous vehicles as a concrete case. Specifically, for perception, the standard describes how an acceptable perception system can only be achieved after a proper ODD definition. The ODD determines a perception ontology (objects and events for perception functionality) and the challenges encountered (e.g., seasonal effects). Thus, a perception module’s performance is defined by its ability to map sensor data to the ontology.

B. Performance Evaluation and Safety-Aware Metrics

State-of-the-art perception algorithms are often compared on benchmark datasets, such as [48], [91], and [154]. Leaderboards in categories such as 3-D camera-only detection and LiDAR segmentation are maintained by developers and researchers employing a set of metrics, often derived from metrics in computer vision literature. Standard metrics are based on IoU, precision, and recall (both at pixel/point level or object level) with some aggregation function applied (e.g., mAP on averaging over all considered prediction classes). Moreover, tracking is a crucial step for perception, so tracking metrics, such as multiple object tracking accuracy/multiple object tracking precision, are also employed [155]. These standard metrics are designed to measure the average difference between specific features of the ground truth and

the perception output. However, the resulting safety is much more relevant than perception accuracy. In fact, we can observe the seemingly subtle fact that optimizing a DNN following standard performance metrics does not necessarily imply that the DNN produces a safe prediction. One intuition can be observed from the fact that, for autonomous driving, pedestrians that are distant should not be equally weighted compared to pedestrians being close by. Another visual example is shown in [156] where the standard IoU metric for bounding box detection will indicate that, within Fig. 6, the prediction in the left image is better than the right. However, when safety is defined at the ML-level to be “completely covering the object” as any region outside the bounding box is considered to be an empty space, the right image, although “worse” in terms of IoU, is safe. To be used in safety-critical ADAS, the developed metrics need to be connected to concrete performance limitations and driving applications, and the acceptance threshold should be justified. Despite recent developments [157], [158], [159], [160] in developing safety-aware metrics by including various factors, such as reaction time or imperfection of the labeling, these metrics are not direct reflectors of safety. To be used in a concrete safety case, one needs to fine-tune these metrics by matching the definition of acceptably safe defined in a concrete application. A similar phenomenon also occurs in associating the degree of robustness related to safety. Within the field of machine learning, researchers formulate robustness using concepts such as L_∞ norms, characterizing the minimum amount of input change that maintains prediction consistency. Nevertheless, deciding the required minimum robustness bound of an ML model in a given application with convincing rationale (so that the noise to the ML model will not be the source of harm) is far from trivial.

Moreover, detection, tracking, and scene segmentation are only subtasks of the overall perception algorithm required in ADAS. This makes the connection between safety and perception evaluation challenging to be tackled with summarized metrics. Vehicles are deployed in uncontrolled traffic, and the same perception error (e.g., misclassification) may be irrelevant in some situations while being crucial in others.

Another reason for this limitation is the decision and planning module (refer to Fig. 1), which is crucial in determining how the ADAS reacts to the perception output. The same perception error can be a safety issue or not depending on the driving style of the ADAS-equipped vehicle. For example, a vehicle driving at a higher speed may demand better perception at longer distances. At the same time, a slower vehicle could achieve adequate safety even with sensors with a shorter operating range. Thus, an end-to-end testing step is necessary to evaluate the perception quality while accounting for both traffic scenarios, and the decision and planning module. On this topic, Piazzoni et al. [161] propose perception error models to test the impact of specific perception errors of safety (e.g., detection accuracy over time and tracking-loss probability) via virtual testing and scripted scenarios. This approach requires accurate modeling of the perception errors that affect the perception algorithm under test [162], [163].

C. Data Considerations and ML Systems

A public dataset for automotive perception is usually a combination of sensor data collected by a vehicle on the road with annotated ground truth, typically 2-D/3-D bounding boxes of obstacles and road users. Each data collection campaign generates a dataset with unique limitations and features, as shown in Table I. Moreover, each dataset may include a different set of labels and detected objects. Perception datasets are expensive, with inaccurate ground truth, and limited in nature. The high cost of data collection results from the need to drive on the road with sensor-equipped vehicles, which is either expensive in terms of time or the number of vehicles. Alongside the cost of collecting data, there is also a constant need for new data collection campaigns that employ more recent sensor hardware or firmware releases.

As part of data collection, data need to be labeled with ground truth. Common solutions are manual labeling that suffers from high cost and possible interpretation errors or the application of offline algorithms that offer higher reliability than online algorithms by exploiting causal information. A data collection campaign can only collect and label a limited amount of data. Besides planned limitations (e.g., type of sensor used, location, and time of the day), features such as weather conditions and traffic events cannot be controlled. Since the traffic environment is not controlled, the accuracy of automated labeling techniques may be limited. Furthermore, data reflecting life-threatening events (e.g., close-to-collision) may not be easily obtained.

D. Virtual Testing and Sensor Simulation

To overcome issues with real-world testing, synthetic data can be obtained in virtual environments [100], [164], [165], [166], e.g., using photorealistic simulation or emulating humans and motion behaviors. This approach has a few key advantages. First, it is much less time-consuming and resource intensive, as no vehicle is driven on the road in uncontrolled traffic. In addition, virtual environments can offer perfect ground truth values, and every simulation aspect can




Fig. 7. Typical architecture of virtual testing involving sensor simulation.

be controlled (e.g., weather conditions or traffic events). Thus, this approach can provide a huge amount of labeled data, which ML algorithms could exploit.

However, the use of synthetic data unavoidably raises the problem of domain gap, where one cannot ensure that the performance demonstrated on synthetic data can be faithfully transferred to the real world. Model training for domain adaptation is under active research (e.g., [167], [168], and [169]). For demonstrating diversity, it is also related to ODD, where one should develop methods to have a systematic understanding of how data are collected. Evidence based on combinatorial testing is based on characterizing the ODD, followed by ensuring that the collected dataset can have a reasonable amount of data for any arbitrary pair of criteria. The idea has been applied in the ML setup [150], [170], [171] for highway and urban autonomous driving.

Moreover, the standard ANSI/UL 4600 [153] recommends using virtual environments for autonomous vehicle testing, for both hardware-in-the-loop and software-in-the-loop modalities. The employment of virtual simulators is common practice [165], [166], [172], [173], [174]. Fig. 7 illustrates a typical architecture for a cosimulation loop. The simulator handles the objects in the scene (e.g., traffic vehicles and pedestrians) and employs sensor models to generate synthetic data. The vehicle stack processes the synthetic data and determines the response, which is then sent to the simulator. Most simulators also include complex vehicle dynamics, road maps, and tools to script traffic scenarios.

Along with individual strengths and weaknesses, most simulators share the advantage of offering safe, scalable, controlled, and scriptable test solutions. However, a major challenge is their fidelity, i.e., their ability to provide results that are representative of real-life situations. This aspect is particularly relevant for perception, as virtual environments have to generate synthetic signals (e.g., images and point clouds) to feed the autonomous vehicle stack. Thus, a high-fidelity sensor model requires accurate modeling of materials, physical properties, and effects of weather conditions [175], [176].

E. Monitoring Against Abnormal Situations

Finally, as the ML model is only expected to function in the ODD, the ML system should have a mechanism to notify if the current input is “outside the ODD.” Practically, the specified operational domain can be ambiguous (e.g., autonomous highway driving in Germany), but the ML function can still produce incorrect outputs due to errors in generalization. This problem is mediated by building up a monitor to “detect unknowns in runtime” by various means. As research in this

field is still under active development, we only discuss some representative ideas.

The simplest method, similar to the work of Hendrycks and Gimpel [177], is to use softmax in the output for proxying the likelihood of being a specific class. Then, a warning is raised when the prediction is not particularly strong for every class. The ODIN approach utilizes the softmax classification but uses the temperature scaling [178] to perform uncertainty calibration. Apart from interpreting the output values, another direction considers the inspection of features within intermediate layers. The work of Lee et al. [179] assumes that feature vectors of intermediate layers produced by the training data are approximately Gaussian-distributed. The authors use the Mahalanobis distance as a confidence score for adversarial or out-of-distribution (OoD) detection. Extending this line of thoughts, recent work also aims at training the DNN to allow directly outputting uncertainty (the deterministic uncertainty quantification (DUQ) network) [180], where, for classification, the network is trained in a way such that, within the feature space, each class has a representative vector. In operation, the training input is translated into the feature vector; the classification and the uncertainty are based on the distance to all class-representative vectors. The application to object detection is reflected in CertainNet [181], where the classification is done similar to DUQ. However, for regression, the estimation is based on computing the variance of all nearby grids having overlapping predictions over the same object. Yet, another possibility is to consider OoD detectors built using abstraction-based approaches [182], [183], [184], where DNN-generated feature vectors from the training dataset are clustered and enclosed using hyperrectangles. Note that input outside the training data distribution may not imply that it is not in the ODD.

Uncertainty can also be measured using redundancy and majority voting. This leads to Bayesian approaches, such as dropout at runtime [185] or ensemble learning. Deep Ensembles [186] achieve state-of-the-art uncertainty estimation but at a large computational overhead (since one needs to train many models by taking different random seeds); thus, recent work attempts to mitigate this with various ideas [187], [188], such as parameter sharing in earlier layers for different models. Beyond direct measurement, the recent work on evidential learning [189] aims to learn parameters in a higher order setup, where the learned parameters are used to uncover the distribution of the uncertainty model.

Apart from technical challenges, a crucial implementation-related challenge for monitors lies in the imbalanced data for “unknowns” or input outside the ODD. Many of the presented techniques here use in-distribution data in a self-supervised fashion and use OoD data for calibrating the boundary. However, to perform such a calibration and test the applicability, one only has “known unknowns” collected before the development.

VI. FUTURE OUTLOOK

The selection of sensor modalities, sensing and perception algorithms, data processing, and multisensor fusion

architectures, to deliver robust, high-quality perception for ADAS/ADS vehicles will remain a topic of interest. Perception systems need to provide high levels of functional performance for safe driving under diverse ODDs during the entire operational lifecycle. It is expected that perception can support high levels of autonomous driving functions in restricted ODDs, e.g., on well-mapped highways in good weather conditions. Validation of sensing and perception methods, especially based on machine learning methods, can, thus, be done in different ODDs to ensure the safety and incremental adoption of autonomous features.

Although sensor fusion strategies leveraging supervised algorithms are able to mitigate the shortcomings in an individual sensing modality, these are still not perfect. Utilizing reinforcement learning paradigms in conjunction with supervised learning algorithms within the sensor fusion context could assist in scenario-based learning. Furthermore, reinforcement learning algorithms can be used to assess the risk of failure of the sensor fusion solution early on and facilitate human intervention.

Besides local fusion discussed in Section IV-D, cooperative perception is another approach to enhance the capabilities of local perception sensors by sharing information among vehicles or communicating with the infrastructure. Vehicle connectivity is a means to enable such information sharing. The role of connectivity in extending automotive perception capabilities is, however, outside the scope of this article; the reader is referred to [190] and [191]. Connectivity also enables updates and, thus, improvements in ML models while offering new perception-driven services [192], [193]. However, it also brings additional challenges, such as transformation errors between the different reference frames, delays, and uncertainty with respect to the shared information, security, and trust, which need to be addressed.

It is expected that ML-driven automotive sensing and perception will lead to “better than human” capabilities like having obstacle information over a 360° field-of-view. However, the response of ML components cannot be guaranteed by traditional system engineering and software validation approaches. Aspects such as explainability and reproducibility of ML processing become critical in ADAS/ADS and, given their safety-critical nature, remain a challenge [194], [195].

One of the challenges in reliable automotive sensing and perception information across diverse ODDs is the limited availability of data in difficult driving and weather conditions, diverse in-cabin conditions, and along the sensor operational lifecycle. The availability of quality sensor datasets is crucial to avoid issues such as a class imbalance in training ML models. As discussed in Sections IV and V, the validation of automotive perception systems is a substantial challenge. Traditional computer vision metrics used in automotive perception are not context-aware; there is a need to consider safety awareness and ADAS/ADS applications. Due to the lack of large real-world datasets, there is a need to use synthetic data. For this purpose, the development of virtual simulators could lead to more effective and efficient ways of end-to-end perception systems and ADAS/ADS testing. Improving their

fidelity via more realistic sensor models can provide better synthetic data for ML training.

With greater driving autonomy, there is also a concern that over-reliance on machine-based decisions may lead to bad driving habits and increased driving distractions. This brings about the need for HVI mechanisms that enable humans to take over driving operations on time by overriding autonomous systems when necessary. New decision-making and control designs that leverage both automotive-external and automotive-internal perception information while taking into account deployment differences are needed.

Research on automotive perception technologies to support the holy grail of fully autonomous driving should address synergy between technology and the fields of ethical, legal, and social sciences. Greater collaboration among diverse disciplines, such as hardware and software reliability, algorithm designs, safety and quality engineering, security and privacy, human-machine designs, and insurance and legal, is required for perception-based ADAS/ADS designs to be successfully deployed at scale to provide enhanced safety and driving comfort.

REFERENCES

- [1] E. Marti, M. A. de Miguel, F. Garcia, and J. Perez, "A review of sensor technologies for perception in automated driving," *IEEE Intell. Transp. Syst. Mag.*, vol. 11, no. 4, pp. 94–108, Sep. 2019.
- [2] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Trans. Intell. Vehicles*, vol. 1, no. 1, pp. 33–55, Jun. 2016.
- [3] Y. Chen, Y. Wang, F. Qu, and W. Li, "A graph-based track-before-detect algorithm for automotive radar target detection," *IEEE Sensors J.*, vol. 21, no. 5, pp. 6587–6599, Mar. 2021.
- [4] O. Bialer, A. Jonas, and T. Tirer, "Super resolution wide aperture automotive radar," *IEEE Sensors J.*, vol. 21, no. 16, pp. 17846–17858, Jun. 2021.
- [5] Y. Sun, T. Fei, and N. Pohl, "A high-resolution framework for range-Doppler frequency estimation in automotive radar systems," *IEEE Sensors J.*, vol. 19, no. 23, pp. 11346–11358, Aug. 2019.
- [6] H. Lian, X. Pei, and X. Guo, "A local environment model based on multi-sensor perception for intelligent vehicles," *IEEE Sensors J.*, vol. 21, no. 14, pp. 15427–15436, Aug. 2021.
- [7] R. Ravindran, M. J. Santora, and M. M. Jamali, "Camera, LiDAR, and radar sensor fusion based on Bayesian neural network (CLR-BNN)," *IEEE Sensors J.*, vol. 22, no. 7, pp. 6964–6974, Feb. 2022.
- [8] D. Kang and D. Kum, "Camera and radar sensor fusion for robust vehicle localization via vehicle part localization," *IEEE Access*, vol. 8, pp. 75223–75236, 2020.
- [9] A. Venon, Y. Dupuis, P. Vasseur, and P. Merriaux, "Millimeter wave FMCW RADARS for perception, recognition and localization in automotive applications: A survey," *IEEE Trans. Intell. Vehicles*, vol. 7, no. 3, pp. 533–555, Sep. 2022.
- [10] *Radar Services in the 76–81 GHz Band; Report and Order—ET Docket No. 15–26*, Federal Communications Commission, document FCCCIRC1707-07, 2017.
- [11] *Systems Characteristics of Automotive Radars Operating in the Frequency Band 76–81 GHz for Intelligent Transport Systems Applications*, International Telecommunications Union, document ITU-R M.2057-1, 2018.
- [12] H. Cho, "Operational design domain (ODD) framework for driver-automation integrated systems," Ph.D. thesis, Massachusetts Inst. Technol., Cambridge, MA, USA, 2020.
- [13] *New Car Assessment Program*, National Highway Traffic Safety Administration, document NHTSA-2021-0002, 2022. [Online]. Available: <https://www.govinfo.gov/content/pkg/FR-2022-03-09/pdf/2022-04894.pdf>
- [14] European New Car Assessment Programme. (2023). *Child Presence Detection Test and Assessment Protocol V1.1*. [Online]. Available: <https://cdn.euroncap.com/media/75474/euro-ncap-cpd-test-and-assessment-protocol-v11.pdf>
- [15] M. Sheeny, E. De Pellegrin, S. Mukherjee, A. Ahrabian, S. Wang, and A. Wallace, "RADIATE: A radar dataset for automotive perception in bad weather," 2020, *arXiv:2010.09076*.
- [16] A. Mishra, S. Lee, D. Kim, and S. Kim, "In-cabin monitoring system for autonomous vehicles," *Sensors*, vol. 22, no. 12, pp. 1–10, 2022. [Online]. Available: <https://www.mdpi.com/1424-8220/22/12/4360>
- [17] E. Yavari, P. Nuti, and O. Boric-Lubecke, "Occupancy detection using radar noise floor," in *Proc. IEEE/ACES Int. Conf. Wireless Inf. Technol. Syst. (ICWITS) Appl. Comput. Electromagn. (ACES)*, Mar. 2016, pp. 1–3.
- [18] A. Lazaro, M. Lazaro, R. Villarino, and D. Girbau, "Seat-occupancy detection system and breathing rate monitoring based on a low-cost mm-wave radar at 60 GHz," *IEEE Access*, vol. 9, pp. 115403–115414, 2021.
- [19] J. H. Huh and S. H. Cho, "Seat belt reminder system in vehicle using IR-UWB radar," in *Proc. Int. Conf. Netw. Infrastruct. Digit. Content (IC-NIDC)*, Aug. 2018, pp. 256–259.
- [20] M. Hoffmann, D. Tatarinov, J. Landwehr, and A. R. Diewald, "A four-channel radar system for rear seat occupancy detection in the 24 GHz ISM band," in *Proc. 11th German Microw. Conf. (GeMiC)*, Mar. 2018, pp. 95–98.
- [21] A. R. Diewald et al., "RF-based child occupation detection in the vehicle interior," in *Proc. Int. Radar Symp.*, 2016, pp. 1–4.
- [22] A. Caddemi and E. Cardillo, "Automotive anti-abandon systems: A millimeter-wave radar sensor for the detection of child presence," in *Proc. 14th Int. Conf. Adv. Technol., Syst. Services Telecommun. (TELSIKS)*, Oct. 2019, pp. 94–97.
- [23] H. Abedi, S. Luo, V. Mazumdar, M. M. Y. R. Riad, and G. Shaker, "AI-powered in-vehicle passenger monitoring using low-cost mm-wave radar," *IEEE Access*, vol. 10, pp. 18998–19012, 2022.
- [24] S. Pisa, E. Pittella, and E. Piuze, "A survey of radar systems for medical applications," *IEEE Aerosp. Electron. Syst. Mag.*, vol. 31, no. 11, pp. 64–81, Nov. 2016.
- [25] S. M. M. Islam, O. Boric-Lubecke, V. M. Lubecke, A.-K. Moadi, and A. E. Fathy, "Contactless radar-based sensors: Recent advances in vital-signs monitoring of multiple subjects," *IEEE Microw. Mag.*, vol. 23, no. 7, pp. 47–60, Jul. 2022.
- [26] G. Paterniani et al., "Radar-based monitoring of vital signs: A tutorial overview," *Proc. IEEE*, vol. 111, no. 3, pp. 277–317, Mar. 2023.
- [27] R. Chae, A. Wang, and C. Li, "FMCW radar driver head motion monitoring based on Doppler spectrogram and range-Doppler evolution," in *Proc. IEEE Topical Conf. Wireless Sensors Sensor Netw. (WiSNet)*, Jan. 2019, pp. 1–4.
- [28] K. Staszek, K. Wincza, and S. Gruszczynski, "Driver's drowsiness monitoring system utilizing microwave Doppler sensor," in *Proc. 19th Int. Conf. Microw., Radar Wireless Commun. (MIKON)*, vol. 2, May 2012, pp. 623–626.
- [29] Y. Kim, "Detection of eye blinking using Doppler sensor with principal component analysis," *IEEE Antennas Wireless Propag. Lett.*, vol. 14, pp. 123–126, 2014.
- [30] K. Yamamoto, K. Toyoda, and T. Ohtsuki, "Doppler sensor-based blink duration estimation by analysis of eyelids closing and opening behavior on spectrogram," *IEEE Access*, vol. 7, pp. 42726–42734, 2019.
- [31] E. Cardillo, G. Sapienza, C. Li, and A. Caddemi, "Head motion and eyes blinking detection: A mm-wave radar for assisting people with neurodegenerative disorders," in *Proc. 50th Eur. Microw. Conf. (EuMC)*, Jan. 2021, pp. 925–928.
- [32] C.-H. Tseng, J.-R. Lin, C.-L. Lin, Y.-C. Wu, and L.-T. Huang, "The prototype of a driver attention level monitoring system: The Sanbao radar," in *Proc. IEEE Int. Conf. Consum. Electron.-Taiwan*, May 2018, pp. 1–2.
- [33] X. Gu, L. Zhang, Y. Xiao, H. Zhang, H. Hong, and X. Zhu, "Non-contact fatigue driving detection using CW Doppler radar," in *IEEE MTT-S Int. Microw. Symp. Dig.*, May 2018, pp. 1–3.
- [34] G. Ciattaglia, S. Spinsante, and E. Gambi, "Slow-time mmWave radar vibrometry for drowsiness detection," in *Proc. IEEE Int. Workshop Metrol. Automot. (MetroAutomotive)*, Jul. 2021, pp. 141–146.
- [35] Z. Dong et al., "A fatigue driving detection method based on frequency modulated continuous wave radar," in *Proc. IEEE Int. Conf. Consum. Electron. Comput. Eng. (ICCECE)*, Jan. 2021, pp. 670–675.
- [36] S. Ishizaka, K. Yamamoto, and T. Ohtsuki, "Non-contact blood pressure measurement using Doppler radar based on waveform analysis by LSTM," in *Proc. IEEE Int. Conf. Commun.*, Jun. 2021, pp. 1–6.

- [37] A. E. Omer, S. Safavi-Naeini, R. Hughson, and G. Shaker, "Blood glucose level monitoring using an FMCW millimeter-wave radar sensor," *Remote Sens.*, vol. 12, no. 3, p. 385, Jan. 2020.
- [38] K. Yamamoto, K. Toyoda, and T. Ohtsuki, "Spectrogram-based simultaneous heartbeat and blink detection using Doppler sensor," in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2019, pp. 1–6.
- [39] J. F. Holzrichter and L. C. Ng, "Speech articulator and user gesture measurements using micropower, interferometric EM-sensors," in *Proc. 18th IEEE Instrum. Meas. Technol. Conf., Rediscovering Meas. Age Inform.*, vol. 3, May 2001, pp. 1942–1946.
- [40] S. Wang, J. Song, J. Lien, I. Poupyrev, and O. Hilliges, "Interacting with soli: Exploring fine-grained dynamic gesture recognition in the radio-frequency spectrum," in *Proc. 29th Annu. Symp. User Interface Softw. Technol.*, Oct. 2016, pp. 851–860.
- [41] M. M. Rahman, E. A. Malaia, A. C. Gurbuz, D. J. Griffin, C. Crawford, and S. Z. Gurbuz, "Effect of kinematics and fluency in adversarial synthetic data generation for ASL recognition with RF sensors," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 58, no. 4, pp. 2732–2745, Aug. 2022.
- [42] P. Molchanov, S. Gupta, K. Kim, and K. Pulli, "Short-range FMCW monopulse radar for hand-gesture sensing," in *Proc. IEEE Radar Conf. (RadarCon)*, May 2015, pp. 1491–1496.
- [43] B. B. James, "Recognizing traffic signalling gestures through automotive sensors," M.S. thesis, Dept. Elect. Comput. Eng., Mississippi State Univ., Mississippi State, MA, USA, 2022.
- [44] K. A. Smith, C. Cseh, D. Murdoch, and G. Shaker, "Gesture recognition using mm-wave sensor for human-car interface," *IEEE Sensors Lett.*, vol. 2, no. 2, pp. 1–4, Jun. 2018.
- [45] M. Raja, V. Ghaderi, and S. Sigg, "WiBot! in-vehicle behaviour and gesture recognition using wireless network edge," in *Proc. IEEE 38th Int. Conf. Distrib. Comput. Syst. (ICDCS)*, Jul. 2018, pp. 376–387.
- [46] G. L. Charvat, *Small and Short-Range Radar Systems*, 1st ed. Boca Raton, FL, USA: CRC Press, 2014.
- [47] H. Rohling and M.-M. Meinecke, "Waveform design principles for automotive radar systems," in *Proc. CIE Int. Conf. Radar*, Oct. 2001, pp. 1–4.
- [48] H. Caesar et al., "NuScenes: A multimodal dataset for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 11621–11631.
- [49] O. Schumann et al., "RadarScenes: A real-world radar point cloud data set for automotive applications," 2021, *arXiv:2104.02493*.
- [50] D. Barnes, M. Gadd, P. Murcutt, P. Newman, and I. Posner, "The Oxford radar RobotCar dataset: A radar extension to the Oxford RobotCar dataset," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Aug. 2020, pp. 6433–6438.
- [51] G. Kim, Y. S. Park, Y. Cho, J. Jeong, and A. Kim, "MulRan: Multimodal range dataset for urban place recognition," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Aug. 2020, pp. 6246–6253.
- [52] L. Zheng et al., "TJ4DRadSet: A 4D radar dataset for autonomous driving," 2022, *arXiv:2204.13483*.
- [53] A. Palffy, E. Pool, S. Baratam, J. F. P. Kooij, and D. M. Gavrilu, "Multi-class road user detection with 3+1D radar in the view-of-delft dataset," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 4961–4968, Apr. 2022.
- [54] A. Ouaknine, A. Newson, J. Rebut, F. Tupin, and P. Pérez, "CARRADA dataset: Camera and automotive radar with range- angle- Doppler annotations," in *Proc. 25th Int. Conf. Pattern Recognit. (ICPR)*, Jan. 2021, pp. 5068–5075.
- [55] Y. Wang, Z. Jiang, Y. Li, J.-N. Hwang, G. Xing, and H. Liu, "RODNet: A real-time radar object detection network cross-supervised by camera-radar fused object 3D localization," *IEEE J. Sel. Topics Signal Process.*, vol. 15, no. 4, pp. 954–967, Jun. 2021.
- [56] T.-Y. Lim, S. A. Markowitz, and M. N. Do, "RaDiCaL: A synchronized FMCW radar, depth, IMU and RGB camera data dataset with low-level FMCW radar signals," *IEEE J. Sel. Topics Signal Process.*, vol. 15, no. 4, pp. 941–953, Jun. 2021.
- [57] A. Zhang, F. E. Nowruzi, and R. Laganieri, "RADDet: Range-Azimuth-Doppler based radar object detection for dynamic road users," in *Proc. 18th Conf. Robots Vis. (CRV)*, May 2021, pp. 95–102.
- [58] J. Rebut, A. Ouaknine, W. Malik, and P. Pérez, "Raw high-definition radar for multi-task learning," 2021, *arXiv:2112.10646*.
- [59] K. Doris, A. Filippi, and F. Jansen, "Reframing fast-chirp FMCW transceivers for future automotive radar: The pathway to higher resolution," *IEEE Solid-State Circuits Mag.*, vol. 14, no. 2, pp. 44–55, Spring 2022.
- [60] A. Laribi, M. Hahn, J. Dickmann, and C. Waldschmidt, "Performance investigation of automotive SAR imaging," in *IEEE MTT-S Int. Microw. Symp. Dig.*, Apr. 2018, pp. 1–4.
- [61] M. Mostajabi, C. M. Wang, D. Ranjan, and G. Hsyu, "High resolution radar dataset for semi-supervised learning of dynamic objects," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 450–457.
- [62] J. Li and P. Stoica, *MIMO Radar Signal Processing*. Hoboken, NJ, USA: Wiley, 2008.
- [63] S. Rao, "White paper: MIMO radar," Texas Instruments, Dallas, TX, USA, Tech. Rep., SWRA554A, 2017.
- [64] P. Wang, P. Boufounos, H. Mansour, and P. V. Orlik, "Slow-time MIMO-FMCW automotive radar detection with imperfect waveform separation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process. (ICASSP)*, May 2020, pp. 8634–8638.
- [65] S. Sun, A. P. Petropulu, and H. V. Poor, "MIMO radar for advanced driver-assistance systems and autonomous driving: Advantages and challenges," *IEEE Signal Process. Mag.*, vol. 37, no. 4, pp. 98–117, Jul. 2020.
- [66] F. G. Jansen, "Automotive radar Doppler division MIMO with velocity ambiguity resolving capabilities," in *Proc. 16th Eur. Radar Conf. (EuRAD)*, Oct. 2019, pp. 245–248.
- [67] A. Och, C. Pfeffer, J. Schrattecker, S. Schuster, and R. Weigel, "A scalable 77 GHz massive MIMO FMCW radar by cascading fully-integrated transceivers," in *Proc. Asia-Pacific Microw. Conf. (APMC)*, Nov. 2018, pp. 1235–1237.
- [68] (2021). *4D Imaging Radar: The World's First 2K Ultra-High Resolution Radar Platform*. Accessed: Sep. 27, 2022. [Online]. Available: <https://arberobotics.com/wp-content/uploads/2021/05/4D-Imaging-radar-product-overview.pdf>
- [69] A. Santra and S. Hazra, *Deep Learning Applications of Short-Range Radars*. Norwood, MA, USA: Artech House, 2020.
- [70] K. Granström, M. Baum, and S. Reuter, "Extended object tracking: Introduction, overview, and applications," *J. Adv. Inf. Fusion*, vol. 12, no. 2, pp. 1–10, 2017.
- [71] J. W. Koch, "Bayesian approach to extended object and cluster tracking using random matrices," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 44, no. 3, pp. 1042–1059, Jul. 2008.
- [72] U. Orguner, "A variational measurement update for extended target tracking with random matrices," *IEEE Trans. Signal Process.*, vol. 60, no. 7, pp. 3827–3834, Jul. 2012.
- [73] M. Baum and U. D. Hanebeck, "Extended object tracking with random hypersurface models," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 50, no. 1, pp. 149–159, Jan. 2014.
- [74] N. Wahlström and E. Özkan, "Extended target tracking using Gaussian processes," *IEEE Trans. Signal Process.*, vol. 63, no. 16, pp. 4165–4178, Aug. 2015.
- [75] P. Broßheit, B. Duraisamy, and J. Dickmann, "The volcanormal density for radar-based extended target tracking," in *Proc. IEEE 20th Int. Conf. Intell. Transp. Syst. (ITSC)*, Oct. 2017, pp. 1–6.
- [76] K. Granström, M. Fatemi, and L. Svensson, "Poisson multi-Bernoulli mixture conjugate prior for multiple extended target filtering," *IEEE Trans. Aerosp. Electron. Syst.*, vol. 56, no. 1, pp. 208–225, Feb. 2020.
- [77] Y. Xia et al., "Learning-based extended object tracking using hierarchical truncation measurement model with automotive radar," *IEEE J. Sel. Topics Signal Process.*, vol. 15, no. 4, pp. 1013–1029, Jun. 2021.
- [78] R. Q. Charles, H. Su, M. Kaichun, and L. J. Guibas, "PointNet: Deep learning on point sets for 3D classification and segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 652–660.
- [79] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, "PointNet++: Deep hierarchical feature learning on point sets in a metric space," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.*, 2017, pp. 5105–5114.
- [80] J. Liu et al., "Deep instance segmentation with automotive radar detection points," *IEEE Trans. Intell. Vehicles*, vol. 8, no. 1, pp. 84–94, Jan. 2023.
- [81] O. Schumann, M. Hahn, J. Dickmann, and C. Wöhler, "Semantic segmentation on radar point clouds," in *Proc. 21st Int. Conf. Inf. Fusion (FUSION)*, Jul. 2018, pp. 2179–2186.
- [82] O. Schumann, J. Lombacher, M. Hahn, C. Wöhler, and J. Dickmann, "Scene understanding with automotive radar," *IEEE Trans. Intell. Vehicles*, vol. 5, no. 2, pp. 188–203, Jun. 2020.
- [83] J. Bai, L. Zheng, S. Li, B. Tan, S. Chen, and L. Huang, "Radar transformer: An object classification network based on 4D MMW imaging radar," *Sensors*, vol. 21, no. 11, p. 3854, Jun. 2021.

- [84] P. Li, P. Wang, K. Berntorp, and H. Liu, "Exploiting temporal relations on radar perception for autonomous driving," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2022, pp. 17071–17080.
- [85] X. Zhou, D. Wang, and P. Krähenbühl, "Objects as points," 2019, *arXiv:1904.07850*.
- [86] S. Gurbuz, *Deep Neural Network Design for Radar Applications*. Rijeka, Croatia: SciTech, 2020.
- [87] X. Gao, G. Xing, S. Roy, and H. Liu, "RAMP-CNN: A novel neural network for enhanced automotive radar object recognition," *IEEE Sensors J.*, vol. 21, no. 4, pp. 5119–5132, 2021.
- [88] A. Ouaknine, A. Newson, P. Perez, F. Tupin, and J. Rebut, "Multi-view radar semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 15671–15680.
- [89] R. Qian, X. Lai, and X. Li, "3D object detection for autonomous driving: A survey," *Pattern Recognit.*, vol. 130, Oct. 2022, Art. no. 108796.
- [90] G. J. Brostow, J. Shotton, J. Fauqueur, and R. Cipolla, "Segmentation and recognition using structure from motion point clouds," in *Proc. Eur. Conf. Comput. Vis. Cham, Switzerland: Springer*, 2008, pp. 44–57.
- [91] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? The KITTI vision benchmark suite," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2012, pp. 3354–3361.
- [92] M. Cordts et al., "The cityscapes dataset for semantic urban scene understanding," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 3213–3223.
- [93] W. Maddern, G. Pascoe, C. Linegar, and P. Newman, "1 year, 1000 km: The Oxford RobotCar dataset," *IJ Robot. Res.*, vol. 36, n. 1, pp. 3–15, 2016.
- [94] G. Neuhold, T. Ollmann, S. R. Bulo, and P. Kotschieder, "The Mapillary vistas dataset for semantic understanding of street scenes," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 4990–4999.
- [95] F. Yu et al., "BDD100 K: A diverse driving dataset for heterogeneous multitask learning," 2018, *arXiv:1805.04687*.
- [96] X. Huang, P. Wang, X. Cheng, D. Zhou, Q. Geng, and R. Yang, "The apolloscape open dataset for autonomous driving and its application," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 10, pp. 2702–2719, Oct. 2020.
- [97] X. Pan, J. Shi, P. Luo, X. Wang, and X. Tang, "Spatial as deep: Spatial CNN for traffic scene understanding," in *Proc. AAAI Conf. Artif. Intell.*, 2018, vol. 32, no. 1, pp. 1–4.
- [98] A. Patil, S. Malla, H. Gang, and Y.-T. Chen, "The H3D dataset for full-surround 3D multi-object detection and tracking in crowded urban scenes," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2019, pp. 9552–9557.
- [99] C. Sakaridis, D. Dai, and L. Van Gool, "Semantic foggy scene understanding with synthetic data," *Int. J. Comput. Vis.*, vol. 126, no. 9, pp. 973–992, Sep. 2018.
- [100] M. Johnson-Roberson, C. Barto, R. Mehta, S. N. Sridhar, K. Rosaen, and R. Vasudevan, "Driving in the matrix: Can virtual worlds replace human-generated annotations for real world tasks?" in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, Jun. 2017, pp. 746–753.
- [101] TuSimple. (2017). *TuSimple lane detection benchmark*. [Online]. Available: <https://github.com/TuSimple/tusimple-benchmark>
- [102] D. Arya, H. Maeda, S. K. Ghosh, D. Toshniwal, and Y. Sekimoto, "RDD2020: An annotated image dataset for automatic road damage detection using deep learning," *Data Brief*, vol. 36, Jun. 2021, Art. no. 107133.
- [103] Z.-Q. Zhao, P. Zheng, S.-T. Xu, and X. Wu, "Object detection with deep learning: A review," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 11, pp. 3212–3232, Nov. 2019.
- [104] X. Wu, D. Sahoo, and S. C. Hoi, "Recent advances in deep learning for object detection," *Neurocomputing*, vol. 396, pp. 39–64, Jul. 2020.
- [105] R. Girshick, J. Donahue, T. Darrell, and J. Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 580–587.
- [106] X. Xie, G. Cheng, J. Wang, X. Yao, and J. Han, "Oriented R-CNN for object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 3520–3529.
- [107] B. Zhou, A. Khosla, A. Lapedriza, A. Oliva, and A. Torralba, "Learning deep features for discriminative localization," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2921–2929.
- [108] L. Liu et al., "Deep learning for generic object detection: A survey," *Int. J. Comput. Vis.*, vol. 128, no. 2, pp. 261–318, Oct. 2020.
- [109] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 28, 2015, pp. 1–15.
- [110] M. Hniewa and H. Radha, "Object detection under rainy conditions for autonomous vehicles: A review of state-of-the-art and emerging techniques," *IEEE Signal Process. Mag.*, vol. 38, no. 1, pp. 53–67, Dec. 2020.
- [111] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi, "You only look once: Unified, real-time object detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 779–788.
- [112] G. Li, Y. Yang, and X. Qu, "Deep learning approaches on pedestrian detection in hazy weather," *IEEE Trans. Ind. Electron.*, vol. 67, no. 10, pp. 8889–8899, Oct. 2020.
- [113] Y. Cai et al., "YOLOv4–5D: An effective and efficient object detector for autonomous driving," *IEEE Trans. Instrum. Meas.*, vol. 70, pp. 1–13, 2021.
- [114] S. Fan et al., "FII-CenterNet: An anchor-free detector with foreground attention for traffic object detection," *IEEE Trans. Veh. Technol.*, vol. 70, no. 1, pp. 121–132, Jan. 2021.
- [115] Z. Yang, S. Liu, H. Hu, L. Wang, and S. Lin, "RepPoints: Point set representation for object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9657–9666.
- [116] Z. Yang et al., "Dense reppoints: Representing visual objects with dense point sets," in *Proc. Eur. Conf. Comput. Vis. Cham, Switzerland: Springer*, 2020, pp. 227–244.
- [117] K. Duan, S. Bai, L. Xie, H. Qi, Q. Huang, and Q. Tian, "CenterNet: Keypoint triplets for object detection," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 6569–6578.
- [118] Z. Zou, K. Chen, Z. Shi, Y. Guo, and J. Ye, "Object detection in 20 years: A survey," *Proc. IEEE*, vol. 111, no. 3, pp. 1–20, Mar. 2023.
- [119] G. Li, Z. Ji, and X. Qu, "Stepwise domain adaptation (SDA) for object detection in autonomous vehicles using an adaptive CenterNet," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 10, pp. 17729–17743, Oct. 2022.
- [120] G. Li, Z. Ji, X. Qu, R. Zhou, and D. Cao, "Cross-domain object detection for autonomous driving: A stepwise domain adaptive YOLO approach," *IEEE Trans. Intell. Vehicles*, vol. 7, no. 3, pp. 603–615, Sep. 2022.
- [121] A. Vaswani et al., "Attention is all you need," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 30, 2017, pp. 1–20.
- [122] A. Dosovitskiy et al., "An image is worth 16×16 words: Transformers for image recognition at scale," 2020, *arXiv:2010.11929*.
- [123] L. Liu, X. Chen, S. Zhu, and P. Tan, "CondLaneNet: A top-to-down lane detection framework based on conditional convolution," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2021, pp. 3773–3782.
- [124] C.-D. Xu, X.-R. Zhao, X. Jin, and X.-S. Wei, "Exploring categorical regularization for domain adaptive object detection," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 11724–11733.
- [125] J. Li, G. Li, Y. Shi, and Y. Yu, "Cross-domain adaptive clustering for semi-supervised domain adaptation," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2021, pp. 2505–2514.
- [126] G. Wang, C. Zhang, H. Wang, J. Wang, Y. Wang, and X. Wang, "Unsupervised learning of depth, optical flow and pose with occlusion from 3D geometry," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 1, pp. 308–320, Jan. 2022.
- [127] W. Kim, A. Kanazaki, and M. Tanaka, "Unsupervised learning of image segmentation based on differentiable feature clustering," *IEEE Trans. Image Process.*, vol. 29, pp. 8055–8068, 2020.
- [128] H. Fan, P. Liu, M. Xu, and Y. Yang, "Unsupervised visual representation learning via dual-level progressive similar instance selection," *IEEE Trans. Cybern.*, vol. 52, no. 9, pp. 8851–8861, Sep. 2022.
- [129] Y. Li and J. Ibanez-Guzman, "LiDAR for autonomous driving: The principles, challenges, and trends for automotive LiDAR and perception systems," *IEEE Signal Process. Mag.*, vol. 37, no. 4, pp. 50–61, Jul. 2020.
- [130] M. Velas, M. Spanel, M. Hradis, and A. Herout, "CNN for very fast ground segmentation in Velodyne LiDAR data," in *Proc. IEEE Int. Conf. Auto. Robot. Syst. Competitions (ICARSC)*, Apr. 2018, pp. 97–103.
- [131] N. Deo and M. M. Trivedi, "Convolutional social pooling for vehicle trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2018, pp. 1468–1476.

- [132] X. Chen, H. Ma, J. Wan, B. Li, and T. Xia, "Multi-view 3D object detection network for autonomous driving," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1907–1915.
- [133] W. Shi and R. Rajkumar, "Point-GNN: Graph neural network for 3D object detection in a point cloud," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2020, pp. 1708–1716.
- [134] J. Wang, H. Gang, S. Ancha, Y.-T. Chen, and D. Held, "Semi-supervised 3D object detection via temporal graph neural networks," in *Proc. Int. Conf. 3D Vis. (3DV)*, Dec. 2021, pp. 413–422.
- [135] E. Capellier, F. Davoine, V. Cherfaoui, and Y. Li, "Transformation-adversarial network for road detection in LiDAR rings, and model-free evidential road grid mapping," in *Proc. 11th Workshop Planning, Perception, Navigat. Intell. Vehicles*, 2019, pp. 47–52.
- [136] D. Frossard and R. Urtasun, "End-to-end learning of multi-sensor 3D tracking by detection," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, May 2018, pp. 635–642.
- [137] J. Behley et al., "SemanticKITTI: A dataset for semantic scene understanding of LiDAR sequences," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9297–9307.
- [138] B. Hurl, K. Czarniecki, and S. Waslander, "Precise synthetic image and LiDAR (PreSIL) dataset for autonomous vehicle perception," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2019, pp. 2522–2529.
- [139] M. Haberjahn and K. Kozempel, "Multi level fusion of competitive sensors for automotive environment perception," in *Proc. 16th Int. Conf. Inf. Fusion*, Jul. 2013, pp. 397–403.
- [140] M. Aeberhard and T. Bertram, "Object classification in a high-level sensor data fusion architecture for advanced driver assistance systems," in *Proc. IEEE 18th Int. Conf. Intell. Transp. Syst.*, Sep. 2015, pp. 416–422.
- [141] D. Feng et al., "Deep multi-modal object detection and semantic segmentation for autonomous driving: Datasets, methods, and challenges," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 3, pp. 1341–1360, Mar. 2021.
- [142] A. Asvadi, L. Garrote, C. Premebida, P. Peixoto, and U. J. Nunes, "Multimodal vehicle detection: Fusing 3D-LiDAR and color camera data," *Pattern Recognit. Lett.*, vol. 115, pp. 20–29, Nov. 2017.
- [143] F. Wulff, B. Schäufele, O. Sawade, D. Becker, B. Henke, and I. Radusch, "Early fusion of camera and LiDAR for robust road detection based on U-Net FCN," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2018, pp. 1426–1431.
- [144] R. Nabati and H. Qi, "CenterFusion: Center-based radar and camera fusion for 3D object detection," in *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis.*, Jan. 2021, pp. 1527–1536.
- [145] Z. Liu et al., "BEVFusion: Multi-task multi-sensor fusion with unified bird's-eye view representation," in *Proc. IEEE Int. Conf. Robot. Autom. (ICRA)*, 2023, pp. 1–12.
- [146] *Road Vehicles—Functional Safety*, Standard ISO 26262:2018, 2018. [Online]. Available: <https://www.iso.org/standard/68383.html>
- [147] *Road Vehicles—Safety of the Intended Functionality*, Standard ISO 21448:2022, 2022. [Online]. Available: <https://www.iso.org/standard/77490.html>
- [148] R. Salay et al., "The missing link: Developing a safety case for perception components in automated driving," 2021, *arXiv:2108.13294*.
- [149] S. Burton, L. Gauerhof, and C. Heinzemann, "Making the case for safety of machine learning in highly automated driving," in *Proc. Int. Workshop Assurance Cases Softw.-Intensive Syst.* Cham, Switzerland: Springer, 2017, pp. 5–16.
- [150] C.-H. Cheng, C.-H. Huang, and G. Nührenberg, "nn-dependability-kit: Engineering neural networks for safety-critical autonomous driving systems," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Nov. 2019, pp. 1–6.
- [151] X. Zhao et al., "A safety framework for critical systems utilising deep neural networks," in *Proc. Int. Conf. Comput. Saf., Rel., Secur.* Cham, Switzerland: Springer, 2020, pp. 244–259.
- [152] Y. Jia, T. Lawton, J. McDermid, E. Rojas, and I. Habli, "A framework for assurance of medication safety using machine learning," 2021, *arXiv:2101.05620*.
- [153] *ANSI/UL 4600, Standard for Safety for Evaluation of Autonomous Products*. Accessed: Oct. 2022. [Online]. Available: <https://ulse.org/UL4600>
- [154] *Waymo Open Dataset: An Autonomous Driving Dataset*. Accessed: Oct. 2022. [Online]. Available: <https://www.waymo.com/open>
- [155] K. Bernardin and R. Stiefelwagen, "Evaluating multiple object tracking performance: The CLEAR MOT metrics," *J. Image Video Process.*, vol. 2008, pp. 1–10, Feb. 2007.
- [156] C.-H. Cheng, T. Schuster, and S. Burton, "Logically sound arguments for the effectiveness of ML safety measures," in *Proc. Int. Conf. Comput. Saf., Rel., Secur.* Cham, Switzerland: Springer, 2022, pp. 343–350.
- [157] C.-H. Cheng, C.-H. Huang, H. Ruess, and H. Yasuoka, "Towards dependability metrics for neural networks," in *Proc. 16th ACM/IEEE Int. Conf. Formal Methods Models Syst. Design (MEMOCODE)*, Oct. 2018, pp. 1–4.
- [158] M. Lyssenko, C. Gladisch, C. Heinzemann, M. Woehrle, and R. Triebel, "From evaluation to verification: Towards task-oriented relevance metrics for pedestrian detection in safety-critical domains," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2021, pp. 38–45.
- [159] G. Volk, J. Gamedinger, A. von Bernuth, and O. Bringmann, "A comprehensive safety metric to evaluate perception in autonomous systems," in *Proc. IEEE 23rd Int. Conf. Intell. Transp. Syst. (ITSC)*, Sep. 2020, pp. 1–8.
- [160] C.-H. Cheng, A. Knoll, and H.-C. Liao, "Safety metrics for semantic segmentation in autonomous driving," in *Proc. IEEE Int. Conf. Artif. Intell. Test. (AITest)*, Aug. 2021, pp. 57–64.
- [161] A. Piazzoni, J. Cherian, M. Slavik, and J. Dauwels, "Modeling perception errors towards robust decision making in autonomous vehicles," in *Proc. 29th Int. Joint Conf. Artif. Intell.*, Jul. 2020, pp. 3494–3500.
- [162] M. Hoss, M. Scholtes, and L. Eckstein, "A review of testing object-based environment perception for safe automated driving," *Automot. Innov.*, vol. 5, pp. 1–28, Feb. 2022.
- [163] J. Sadeghi et al., "A step towards efficient evaluation of complex perception tasks in simulation," 2021, *arXiv:2110.02739*.
- [164] S. R. Richter, V. Vineet, S. Roth, and V. Koltun, "Playing for data: Ground truth from computer games," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2016, pp. 102–118.
- [165] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun, "CARLA: An open urban driving simulator," in *Proc. 1st Annu. Conf. Robot Learn.*, 2017, pp. 1–16.
- [166] G. Rong et al., "LGSVL simulator: A high fidelity simulator for autonomous driving," in *Proc. IEEE 23rd Int. Conf. Intell. Transp. Syst. (ITSC)*, Sep. 2020, pp. 1–6.
- [167] L. Yang, X. Liang, T. Wang, and E. Xing, "Real-to-virtual domain unification for end-to-end autonomous driving," in *Proc. Eur. Conf. Comput. Vis.* Cham, Switzerland: Springer, 2018, pp. 530–545.
- [168] T.-H. Vu, H. Jain, M. Bucher, M. Cord, and P. P. Perez, "DADA: Depth-aware domain adaptation in semantic segmentation," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 7364–7373.
- [169] A. Robey, G. J. Pappas, and H. Hassani, "Model-based domain generalization," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 34, 2021, pp. 20210–20229.
- [170] C.-H. Cheng, C.-H. Huang, and H. Yasuoka, "Quantitative projection coverage for testing ML-enabled autonomous systems," in *Proc. Int. Symp. Automated Technol. Verification Anal.* Cham, Switzerland: Springer, 2018, pp. 126–142.
- [171] C. Gladisch, C. Heinzemann, M. Herrmann, and M. Woehrle, "Leveraging combinatorial testing for safety-critical computer vision datasets," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2020, pp. 324–325.
- [172] MSC. (2022). *Vires VTD*. [Online]. Available: <https://vires.msccsf.com>
- [173] (2022). *IPG CarMaker*. [Online]. Available: <https://ipg-automotive.com/en/products-solutions/software/carmaker/>
- [174] Siemens. (2022). *Siemens Simcenter Prescan*. [Online]. Available: <https://plm.automation.siemens.com/global/en/products/simcenter/prescan.html>
- [175] J. Zhao, Y. Li, B. Zhu, W. Deng, and B. Sun, "Method and applications of LiDAR modeling for virtual testing of intelligent vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 22, no. 5, pp. 2990–3000, May 2021.
- [176] C. Linnhoff, P. Rosenberger, and H. Winner, "Refining object-based LiDAR sensor modeling—Challenging ray tracing as the magic bullet," *IEEE Sensors J.*, vol. 21, no. 21, pp. 24238–24245, Nov. 2021.
- [177] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," in *Proc. Int. Conf. Learn. Represent.*, 2017, pp. 1–10. [Online]. Available: <https://openreview.net/forum?id=Hkg4TI9xl>

- [178] S. Liang, Y. Li, and R. Srikant, "Enhancing the reliability of out-of-distribution image detection in neural networks," in *Proc. Int. Conf. Learn. Represent.*, 2018, pp. 1–15. [Online]. Available: <https://openreview.net/forum?id=H1VGkIXRZ>
- [179] K. Lee, K. Lee, H. Lee, and J. Shin, "A simple unified framework for detecting out-of-distribution samples and adversarial attacks," in *Proc. Conf. Neural Inf. Process. Syst.*, vol. 31, 2018, pp. 1–12.
- [180] J. Van Amersfoort, L. Smith, Y. W. Teh, and Y. Gal, "Uncertainty estimation using a single deep deterministic neural network," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 9690–9700.
- [181] S. Gasperini et al., "CertainNet: Sampling-free uncertainty estimation for object detection," *IEEE Robot. Autom. Lett.*, vol. 7, no. 2, pp. 698–705, Apr. 2022.
- [182] T. A. Henzinger, A. Lukina, and C. Schilling, "Outside the box: Abstraction-based monitoring of neural networks," in *Proc. 24th Eur. Conf. Artif. Intell.*, in *Frontiers in Artificial Intelligence and Applications*, vol. 325, G. D. Giacomo, A. Catalá, B. Dilkina, M. Milano, S. Barro, A. Bugarín, and J. Lang, Eds. Santiago de Compostela, Spain: IOS Press, 2020, pp. 2433–2440.
- [183] C.-H. Cheng, C.-H. Huang, T. Brunner, and V. Hashemi, "Towards safety verification of direct perception neural networks," in *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, Mar. 2020, pp. 1640–1643.
- [184] C. Wu, Y. Falcone, and S. Bensalem, "Customizable reference runtime monitoring of neural networks using resolution boxes," 2021, [arXiv:2104.14435](https://arxiv.org/abs/2104.14435).
- [185] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. Int. Conf. Mach. Learn.*, 2016, pp. 1050–1059.
- [186] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," in *Proc. 31st Int. Conf. Neural Inf. Process. Syst.* Red Hook, NY, USA: Curran Associates Inc., 2017, pp. 6405–6416.
- [187] M. Dusenberry et al., "Efficient and scalable Bayesian neural nets with rank-1 factors," in *Proc. Int. Conf. Mach. Learn.*, 2020, pp. 2782–2792.
- [188] M. Havasi et al., "Training independent subnetworks for robust prediction," in *Proc. Int. Conf. Learn. Represent.*, 2021, pp. 1–13. [Online]. Available: <https://openreview.net/forum?id=OGg9XnKxFAH>
- [189] M. Sensoy, L. Kaplan, and M. Kandemir, "Evidential deep learning to quantify classification uncertainty," in *Proc. Adv. Neural Inf. Process. Syst.*, vol. 31, 2018, pp. 1–11.
- [190] A. Caillot, S. Ouerghi, P. Vasseur, R. Bouteau, and Y. Dupuis, "Survey on cooperative perception in an automotive context," *IEEE Trans. Intell. Transp. Syst.*, vol. 23, no. 9, pp. 14204–14223, Sep. 2022.
- [191] B. Hafner, V. Bajpai, J. Ott, and G. A. Schmitt, "A survey on cooperative architectures and maneuvers for connected and automated vehicles," *IEEE Commun. Surveys Tuts.*, vol. 24, no. 1, pp. 380–403, 1st Quart., 2022.
- [192] S. Wang et al., "Federated deep learning meets autonomous vehicle perception: Design and verification," 2022, [arXiv:2206.01748](https://arxiv.org/abs/2206.01748).
- [193] A. C. Marosi, R. Lovas, A. Kisari, and E. Simonyi, "A novel IoT platform for the era of connected cars," in *Proc. IEEE Int. Conf. Future IoT Technol.*, Jan. 2018, pp. 1–11.
- [194] A.-M. Leventi-Peetz and T. Östreich, "Deep learning reproducibility and explainable AI (XAI)," 2022, [arXiv:2202.11452](https://arxiv.org/abs/2202.11452).
- [195] B. Li et al., "Trustworthy AI: From principles to practices," *ACM Comput. Surv.*, vol. 55, no. 9, pp. 1–46, Jan. 2023.



Ashish Pandharipande (Senior Member, IEEE) received the M.S. degree in electrical and computer engineering, the M.S. degree in mathematics, and the Ph.D. degree in electrical and computer engineering from the University of Iowa, Iowa City, IA, USA, in 2000, 2001, and 2002, respectively.

Subsequently, he was a Postdoctoral Researcher with the University of Florida, Gainesville, FL, USA, a Senior Researcher with the Samsung Advanced Institute of Technology,

Suwon, South Korea, a Senior Scientist with Philips Research, Eindhoven, The Netherlands, and a Lead R&D Engineer with Signify, Eindhoven. He has held visiting positions at AT&T Laboratories, Middletown, NJ, USA, and the Department of Electrical Communication Engineering, Indian Institute of Science, Bengaluru, India. He is currently the Innovation Director of NXP Semiconductors, Eindhoven. His research interests are in sensing, networking and controls, data analytics, and their applications in domains such as autonomous mobility, smart lighting systems, energy monitoring and control, and cognitive wireless systems. He has around 200 international conference and journal publications, and more than 100 patent grants/applications in these fields.

Dr. Pandharipande is also a Senior Editor of IEEE SIGNAL PROCESSING LETTERS, a Topical Area Editor of IEEE SENSORS JOURNAL, and an Associate Editor of IEEE JOURNAL OF BIOMEDICAL AND HEALTH INFORMATICS.



Chih-Hong Cheng received the M.Sc. degree in electrical engineering from the National Taiwan University, Taipei, Taiwan, in 2008, and the Ph.D. in informatics from the Technical University of Munich, Munich, Germany, in 2012.

After research stints at the fortiss Research Institute, Munich, ABB Corporate Research, Ladenburg, and DENSO, Eching, he is currently Department Head with Fraunhofer IKS, Munich, and an Adjunct Faculty with Research Center for Information Technology Innovation (CITI), Academia Sinica, Taipei. His main research interests include formal methods, AI safety, and software engineering.



Justin Dauwels received the Ph.D. degree in electrical engineering from the Swiss Polytechnical Institute of Technology (ETH), Zürich, Switzerland, in December 2005.

He was a Postdoctoral Fellow with the RIKEN Brain Science Institute, Wako, Japan, from 2006 to 2007, and a Research Scientist with the Massachusetts Institute of Technology, Cambridge, MA, USA, from 2008 to 2010. He was an Associate Professor with the School of Electrical and Electronic Engineering, Nanyang Technological University (NTU), Singapore, till the end of 2020. He was the Deputy Director of the ST Engineering–NTU Corporate Lab, Singapore, which comprises more than 100 Ph.D. students, research staff, and engineers, developing novel autonomous systems for airport operations and transportation. He is an Associate Professor with the Department of Microelectronics, Delft University of Technology (TU Delft), Delft, The Netherlands. His research interests are in data analytics with applications to intelligent transportation systems, autonomous systems, and analysis of human behavior and physiology.

Dr. Dauwels was a JSPS Invited Fellow from 2010 to 2011. He has been a JSPS Postdoctoral Fellow since 2007, a BAEF Fellow in 2008, and a Henri-Benedictus Fellow of the King Baudouin Foundation in 2008. He has been an elected member of the IEEE Signal Processing Theory and Methods Technical Committee and the IEEE Biomedical Signal Processing Technical Committee since 2018. He has been serving as a member of the Editorial Advisory Board of the *International Journal of Neural Systems* and an organizer of IEEE conferences and special sessions. His research team has won several best paper awards at international conferences and journals. He has served as the Chairperson of the IEEE CIS Chapter in Singapore from 2018 to 2020. He has been serving as an Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING since 2018 and the journal *Signal Processing* (Elsevier) since 2021.



Sevgi Z. Gurbuz (Senior Member, IEEE) received the B.S. degree in electrical engineering with a minor in mechanical engineering and the M.Eng. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1998 and 2000, respectively, and the Ph.D. degree in electrical and computer engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2009.

From February 2000 to January 2004, she worked as a Radar Signal Processing Research Engineer at the U.S. Air Force Research Laboratory, Sensors Directorate, Rome, NY, USA. She was an Assistant Professor with the Department of Electrical and Electronics Engineering, TOBB University, Ankara, Turkey, and a Senior Research Scientist with the TUBITAK Space Technologies Research Institute, Ankara. She is currently an Assistant Professor with the Department of Electrical and Computer Engineering, University of Alabama, Tuscaloosa, AL, USA. She has recently received a patent in April 2022 relating to radar-based American Sign Language (ASL) recognition. Her current research interests include RF sensor-enabled cyber-physical systems, radar signal processing, physics-aware machine learning, sensor networks, human motion recognition for biomedical, automotive autonomy, and human-computer interaction (HCI) applications.

Dr. Gurbuz is a member of the SPIE and Association for Computing Machinery (ACM). She also serves as a member of the IEEE Radar Systems Panel. She was a recipient of the 2023 NSF CAREER Award, the 2022 American Association of University Women Research Publication Grant in Engineering, Medicine and Science, the IEEE Harry Rowe Mimno Award for the Best IEEE AES Magazine Paper of 2019, the 2020 SPIE Rising Researcher Award, an EU Marie Curie Research Fellowship, and the 2010 IEEE Radar Conference Best Student Paper Award. She is an Associate Editor of the IEEE TRANSACTIONS ON AEROSPACE AND ELECTRONIC SYSTEMS (T-AES), the IEEE TRANSACTIONS ON RADAR SYSTEMS (T-RS), and the *IET Radar, Sonar & Navigation* (RSN) journal.



Javier Ibanez-Guzman (Member, IEEE) received the M.S.E.E. degree from the University of Pennsylvania, Philadelphia, PA, USA, in 1988, and the Ph.D. degree from the University of Reading, Reading, U.K., in 1993, on a U.K. SERC fellowship.

He was a Visiting Scholar with the University of California at Berkeley, Berkeley, CA, USA. He was a Senior Scientist with the Singapore Institute of Manufacturing Technology (SimTech), A-Star Research Institute, Singapore, where he spearheaded work on autonomous ground vehicles. He is currently a Corporate Expert on Autonomous Systems with Renault S.A., Guyancourt, France, and the Co-Director of the SIVALab Common Laboratory between the CNRS, UTC Compiègne, and Renault, working on intelligent vehicle technologies. He is a Fulbright Scholar. He is an expert in the EU and Eureka research programs.

Dr. Ibanez-Guzman is a Chartered Engineer (C.Eng.) and a Fellow of the Institute of Engineering Technology, U.K. He is a Senior Editor and an Associate Editor for related IEEE TRANSACTIONS and a representative of ISO groups associated with autonomous vehicles and AI.



Guofa Li (Member, IEEE) received the Ph.D. degree in mechanical engineering from Tsinghua University, Beijing, China, in 2016.

He is currently a Professor with the College of Mechanical and Vehicle Engineering, Chongqing University, Chongqing, China. His research interests include environment perception, driver behavior analysis, and human-like decision-making based on artificial intelligence technologies in autonomous vehicles and intelligent transportation systems.

He has published more than 70 papers in his research areas.

Dr. Li was a recipient of the Young Elite Scientists Sponsorship Program in China. He received the Best Paper Awards from the China Association for Science and Technology (CAST) and the *Automotive Innovation* journal. He also serves as an Associate Editor for IEEE SENSORS JOURNAL and a lead Guest Editor for *IEEE Intelligent Transportation Systems Magazine* and *Automotive Innovation*.



Andrea Piazzoni received the M.Sc. degree in computer science from the University of Milano-Bicocca, Milan, Italy, in 2016. He is currently pursuing the Ph.D. degree with the Interdisciplinary Graduate Programme, Nanyang Technological University (NTU), Singapore.

He has worked as a researcher in academia, focusing on the areas of robotics and autonomous vehicles. His research interests are virtual simulation, robotic perception and

decision-making, and probabilistic graphical models.



Pu (Perry) Wang (Senior Member, IEEE) received the Ph.D. degree in electrical engineering from the Stevens Institute of Technology, Hoboken, NJ, USA, in 2011.

He is currently a Senior Principal Research Scientist with the Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA, USA, where he was an Intern in summer 2010. His current research interests include signal processing, Bayesian inference, deep learning, and their applications to (mmWave and THz)

sensing, wireless communications, networks, and automotive perception.

Dr. Wang is an Associate Member of the IEEE SPS Sensor Array and Multichannel (SAM) Technical Committee and a member of the IEEE ComSoc Integrated Sensing and Communication Emerging Technology Initiative (ISAC-ETI) and the IEEE SPS Signal Processing Theory and Methods (SPTM) Technical Committee. He was a recipient of the IEEE Jack Neubauer Memorial Award from the IEEE Vehicular Technology Society in 2013. He was selected as a Distinguished Speaker of the Society of Petrophysicists and Well Log Analysts (SPWLA) in 2017. He was an Associate Editor of the IEEE SIGNAL PROCESSING LETTERS and a Guest Editor of *IEEE Signal Processing Magazine*, IEEE SENSORS JOURNAL, and *Journal of Advances in Information Fusion* (JAIF). He is also a Senior Area Editor of the IEEE SIGNAL PROCESSING LETTERS.



Avik Santra (Senior Member, IEEE) received the M.S. (Hons.) degree in signal processing from the Indian Institute of Science, Bengaluru, India, in 2010, and the Ph.D. degree in electrical, electronics, and informatics from FAU University, Erlangen, Germany, in 2022.

Earlier in his career, he worked as a System Engineer at Broadcom, India, and a Research Engineer at Airbus, Germany. He is currently a Principal Machine Learning Engineer responsible for developing system solutions for Wi-Fi, image sensors, and radar sensors at Infineon, Irvine, CA, USA. He is a co-author of two books on AI published at Artech and Wiley-IEEE, has filed over 70 patents, and published over 55 research papers.

Dr. Santra is an Associate Editor of IEEE SENSORS JOURNAL and *Machine Learning With Applications* (Elsevier).