A Review of Testing Object-Based Environment Perception for Safe Automated Driving

Michael Hoss¹ · Maike Scholtes¹ · Lutz Eckstein¹

Received: 30 June 2021 / Accepted: 9 December 2021 / Published online: 24 February 2022
© The Author(s) 2022

Abstract

Safety assurance of automated driving systems must consider uncertain environment perception. This paper reviews literature addressing how perception testing is realized as part of safety assurance. The paper focuses on testing for verification and validation purposes at the interface between perception and planning, and structures the analysis along the three axes (1) test criteria and metrics, (2) test scenarios, and (3) reference data. Furthermore, the analyzed literature includes related safety standards, safety-independent perception algorithm benchmarking, and sensor modeling. It is found that the realization of safety-oriented perception testing remains an open issue since challenges concerning the three testing axes and their interdependencies currently do not appear to be sufficiently solved.

Keywords Automated driving · Environment perception · Safety assurance · Testing

Abbreviations

AV	Automated vehicle
AD	Automated driving
ADAS	Advanced driver assistance system
ALKS	Automated lane keeping system
FOV	Field-of-view
GNSS	Global navigation satellite system
IMU	Inertial measurement unit
MOT	Multiple object tracking
ODD	Operational design domain
OEDR	Object and event detection and response
OSPA	Optimal sub-pattern assignment
RTK	Real-time kinematic correction

1 Introduction

To date, the safety assurance of automated vehicles of Levels 4 and 5 [1] (AVs) has been an open issue, which is partially because the environment perception of AVs is subject to uncertainties. AVs perceive their environment through a sensor setup that typically consists of at least camera, radar, and lidar sensors. Perception algorithms, which typically use machine learning, then process and fuse the raw sensor data to a world model that is handed over to the subsequent functional modules. Typical world models contain a list of objects, where each object corresponds to a road user or other movable entity surrounding the AV [2].

Uncertainties in this world model can be caused by environmental influences, sensing hardware, perception software, and other factors. These uncertainties can be divided into state, existence, and classification uncertainties [3]. Respective examples are that a road user's velocity is estimated too slow, the road user is not detected at all, or a pedestrian is classified as a cyclist.

For AV subsystems that are mostly defined by software, the real-world testing effort can be substantially reduced through simulated tests. However, as the sensor hardware interacts with the real environment in complex ways that are hard to replicate in simulations, real-world tests remain pivotal.

In the current literature, there does not yet exist a common methodology for the execution of real-world perception

Academic editor: Weichao Zhuang

Michael Hoss^{*} and Maike Scholtes^{*} work was done while affiliated with Institute for Automotive Engineering, RWTH Aachen University.

✉ Michael Hoss
michael.hoss@rwth-aachen.de

¹ Institute for Automotive Engineering, RWTH Aachen University, Steinbachstr. 7, 52074 Aachen, Germany

Fig. 1 Taxonomy of Stellet et al. [4] applied to perception testing

tests. Under which criteria is the environment perceived “well enough”? In which scenarios should the environment perception be tested? How can one obtain the true environment state during these scenarios? Summarizing these questions, the primary research question of this review is:

How can perception tests be realized to assess that an uncertain perception subsystem allows safe AV behavior?

To structure the diverse literature contributions to this general question, the taxonomy by Stellet et al. [4] is used. It states that testing is an evaluation of

“A statement on the system-under-test (*test criteria*) that is expressed quantitatively (*metric*) under a set of specified conditions (*test scenario*) with the use of knowledge of an ideal result (*reference*)”.

Applying the axes of the taxonomy to perception testing (Fig. 1) means that a real-world scenario is simultaneously observed by a reference perception system and by a perception subsystem under test. Their results are compared using metrics, which allows to evaluate test criteria.

1.1 Motivation for a Review

The current literature from the areas of perception development, automotive safety assurance, and safety in artificial intelligence contains valuable approaches to all mentioned testing axes. For a successful safety proof, approaches to individual axes have to seamlessly fit into an overall methodology. Prior to writing, it was found difficult to identify how the available approaches complement each other, which is why this paper is aiming at providing a structured overview of the state of the art. The goal of this review is to enable the development of holistic concepts for testing the safety-relevant uncertainties of environment perception.

1.2 Structure and Contributions

Related Reviews, Surveys, and Overviews (Sect. 2) reviews previous secondary literature. Methods (Sect. 3) defines the

scope and the literature search process. The main contributions are overviews on:

- Perception Testing in Safety Standards (Sect. 4)
- Established perception testing that does not directly target the primary research question (Sect. 5)
 - Perception Algorithm Benchmarking (Sect. 5.1)
 - Object-Level Data-Driven Sensor Modeling (Sect. 5.2)
- Approaches to the research question, structured by the three individual testing axes
 - Test Criteria and Metrics (Sect. 6)
 - Test Scenarios (Sect. 7)
 - Reference Data (Sect. 8)

Finally, the Research Gaps and Challenges (Sect. 9) highlights the largest issues for answering the research question, points out intersection topics between the testing axes, and puts the topics of this review in a larger context.

1.3 Term Definitions

The following term definitions are used throughout the rest of this paper. Figure 2 illustrates the relations between different terms by using the layered decomposition of an automated driving (AD) system in Ref. [5].

Perception Subsystem: The AV’s subsystem for environment perception is defined by at least sensor mounting positions (part of Layer 0), the sensors themselves that generate raw data (Layer 1), and perception software (Layer 2) (Fig. 2).

System-Under-Test (SUT): Perception subsystem of a subject vehicle/ego vehicle.

Environment Perception: The process of environment perception consists of the mentioned perception subsystem, which is being affected by environmental influences that are outside the control of it (rest of Layer 0; Fig. 2). These environmental influences determine if and how the individual parts of the ground truth information are accessible [5].

Perception Algorithm: Software that computes an object-based world model from the raw sensor data (adapted from Ref. [6]). Generally includes both traditional code and components that are trained on data. Corresponds to Layer 2 (Fig. 2).

World Model: Representation of the ego vehicle’s environment and own state that is computed by a perception algorithm (adapted from Ref. [6]). Equal to a subjective scene at the interface between Layers 2 and 3 (Fig. 2). The

Fig. 2 Functional decomposition layers and interfaces of Amersbach and Winner [5], focused on perception. Right of the boxed layers are corresponding terms used in this paper. HW=hardware

most essential world model component in this paper is the current state of dynamic object tracks.

Safety: “Absence of accidents, where an accident is an event involving an unplanned and unacceptable loss” [7]. Unless stated otherwise, safety in this paper means *safety of the intended functionality* (SOTIF) of the AV (see Ref. [8]) rather than functional safety according to ISO 26262 [9].

Reliability: “Probability that something satisfies its specified behavioral requirements over time and under given conditions – this is, it does not fail” [7].

Uncertainty: “Doubt about the validity of the result of a measurement,” or a quantitative measure of this doubt [10]. Uncertainty in object-based world models can be state, existence, or classification uncertainty [3], as mentioned before.

Scenario: “A scenario describes the temporal development between several scenes in a sequence of scenes.” [11]. For open-loop perception testing, it is assumed that the temporal development of all ground truth scenes is given, including environmental influences and all road user trajectories. What is to be determined in a test scenario is how the SUT subjectively observes the objective scenes in its world model (more in Sect. 7.1).

Metrics: Functions that compute properties of interest about the SUT. The microscopic/macrosopic terminology for metrics of Ref. [12] is used, where *microscopic* refers to a single scene for a single ego vehicle within a single scenario, and *macroscopic* refers to the average over an entire fleet of identical vehicles that encounters various scenarios and scenes.

2 Related Reviews, Surveys, and Overviews

Previous reviews partially cover the targeted topic of this paper. For example, the surveys by Stellet et al. [4, 13], Riedmaier et al. [14], Nalic et al. [15], and the PEGASUS project overview [16] all deal with testing and safety assurance of

AVs in general, but without specifically analyzing the interface between perception and planning.

A systematic literature review of the Swedish SMILE project targets verification and validation for machine-learning-based systems in AD [17], however, without focusing specifically on environment perception.

Literature reviews about AV sensing and perception like Refs. [18, 19], and [20] provide the technological state of the art, but do not explicitly focus on testing or safety assurance.

Furthermore, there are publications that provide a valuable overview over current safety challenges for AV environment perception, but which are not explicitly literature reviews [21, 22]. For example, Ref. [21] identifies safety concerns and mitigation approaches for safety-critical perception tasks that rely on deep learning.

In summary, this paper aims to fill the existing gap in the literature reviews by analyzing literature related to the main research question in depth throughout the rest of this paper.

3 Literature Review Methods

The methods of this literature review are loosely inspired by the guidelines for systematic literature reviews [23] and snowballing search [24], and respective example applications of both in Refs. [25] and [17]. The suggested guidelines are purposely not followed exactly, because they were designed for reviewing quantitative and empirical research rather than qualitative or position papers, which are, however, common in the current automotive safety assurance domain.

3.1 Thematic Scope

Table 1 summarizes the focus of this literature review. The literature that matches more focused aspects is more likely to be included and discussed in detail, whereas the literature about related aspects outside the main focus might still be

Table 1 Thematic focus of this review paper

Category	Focused aspect	Related aspects outside of main focus
Location in V-model	Right-hand side (offline testing for verification and validation purposes)	Left-hand side; developer testing; runtime monitoring; data gathering for requirements identification ([26]); shadow mode
Test method and target	Black-box testing of interface between perception and planning; description of perception output	White-box testing inside the perception or planning sub-systems; description of planning sensitivity
Safety concept	SOTIF (safety of the intended functionality, see [8]), specifically object and event detection and response (OEDR) for collision mitigation	Functional safety according to ISO 26262; Extension of ISO 26262 w.r.t. machine learning; cybersecurity; physical harm caused by sensors
Open/closed-loop testing	Open-loop testing (enables offline testing of various perception algorithms on recorded raw sensor data)	Closed-loop testing
Environmental entity	Moving and non-moving other road users	Ego vehicle localization, static obstacles, road infrastructure, traffic signs, traffic lights, drivable space, auditory information
Environment representation format	Object list	Sensor raw data, object-independent environmental features, occupancy grid map, parametric drivable space

referenced to provide context where the actually targeted literature was sparse.

3.2 Literature Search Process

Three search processes were performed to accumulate the references of this paper. First, undocumented searches provided an initial knowledge base. Second, a keyword-based search complemented the starting set for a final iterative snowballing search. While the undocumented search provided most references about standards and established testing approaches, the inclusion criteria of the documented search targeted specifically the primary research question and its test axes.

3.2.1 Undocumented Search

All references that are not part of Table 2 were obtained through undocumented search processes before and during the writing process.

3.2.2 Keyword Search

The following strings were searched in Google Scholar, allowing results in the time frame between 01/01/2010 and 20/08/2020:

- (“automated driving” OR automotive) AND (perception OR sensing) AND (verification OR validation OR safety OR reliability OR sotif OR standard OR requirements OR specification OR testing OR metric)
- (“automated driving” OR automotive) AND (perception OR sensing) AND (reference OR data OR test OR criteria)

Table 2 Iterations of snowballing literature search. *Considered leaf search results without further search because of thematic distance to primary research question

Snowballing iteration	Found sources for further forward and backward snowballing search (58 publications total)
Start set	From undocumented search: Salay et al. [27, 28], Salay and Czarnecki [29], Czarnecki and Salay [30], Berk et al. [31, 32], Shalev-Shwartz et al. [33], Amersbach and Winnner [5], Stellet et al. [13], Brahmi et al. [34], Cassel et al. [35]; From undocumented and keyword search: Berk [36], Stellet et al. [4], Salay et al. [37], Ulbrich et al. [11]; From keyword search: Borg et al. [17], Willers et al. [21], Burton et al. [22], Taş et al. [38], Kaprocki et al. [39], Martin et al. [40], Bai et al. [41] (22 publications).
Iteration 1	Ulbrich et al. [42]*, Althoff [43]*, Bagschik et al. [44]*, Burton et al. [45], Koopman and Fratrick [26, 46], Kubertschak et al. [47], Strigel et al. [48], Berk et al. [49, 50]*, [51]*, Johansson and Nilsson [52]*, Johansson et al. [53]*, Bock et al. [54]*, Rivero et al. [55]*, Raschofer et al. [56]*, Fawcett [57]*, Cao and Huang [58], Czarnecki et al. [59]*, Meyer [60]*, Guo [61]*, Salay and Czarnecki [62], Schwaiger et al. [63], Arnez et al. [64], Piazzoni et al. [65], Weast [66], Klamann et al. [67]*, Aravatinos and Schlicht [68]*, Breitenstein et al. [69]* (29 publications).
Iteration 2	Henne et al. [70], Fleck et al. [71], Spanfelner et al. [72]*, Rahimi et al. [73], Czarnecki [74]*, Feng et al. [75]*, [76] (7 publications).

For both search strings, the top 250 results were analyzed. 20 publications were identified as candidates for the start set of the subsequent snowballing search, out of which 11 were actually included (see Table 2).

3.2.3 Snowballing Search

The results from the other search processes formed a start set, based on which a forward and backward snowballing search [24] was performed in fall 2020 using Google Scholar to fill possible gaps. Forward snowballing means searching through the citations that a publication has received, while backward snowballing means searching through the references it itself has included. Both directions were searched within the same iteration, respectively (Table 2).

Included snowballing search results that are not directly related to the primary research question or to at least one of the three testing axes were excluded from the further search (marked with * in Table 2) to keep the search process efficient and the overall effort feasible. Eventually, some sources were removed due to irrelevance for the final version, particularly a third iteration, which consisted of two sources, and some sources from iterations 1 and 2.

4 Perception Testing in Safety Standards

Safety standards, guidelines, and frameworks can prescribe certain aspects around perception testing in safety assurance. Thus, the explicit role of perception testing in standardizing literature is reviewed before moving on to more concrete realizations of perception tests in the subsequent sections.

4.1 ISO 26262

ISO 26262 [9] addresses the assurance of functional safety of electric and electronic systems in the automotive domain. It has been the most essential safety standard for advanced driver assistance systems (ADAS), but has not been developed to assure the safety of AD systems that rely on environment perception or use machine learning [37]. Consequently, the literature investigates how ISO 26262 could be adapted or extended to also cover machine-learning-based components such as modern environment perception systems, see Ref. [17].

One major issue that prevents the concepts of ISO 26262 from being applied to AV environment perception is that the perception subsystem is said to be not fully specifiable (except for defining a training set), while ISO 26262 implicitly assumes that all functionality must be specified [37, 62]. This means that the environment perception is also not fully verifiable and perception testing is in general incomplete [37].

Nevertheless, ISO 26262's workflow for Safety Elements out of Context (SEooC) [9] might be a hint on how the development of a perception subsystem (safety element) could at least be generally organized while the rest of the driving function (context) is not known.

4.2 ISO/PAS 21448

A vehicle that is perfectly functionally safe according to ISO 26262 can still cause accidents if its behavioral logic is wrong [33]. Addressing this issue, ISO/PAS 21448 [8] introduces the concept of SOTIF. The standard describes how SOTIF should be assured for ADAS up to Level 2 and acknowledges that it is likely not sufficient for higher levels of automation.

According to Ref. [8], so-called *triggering events* “that can trigger potentially hazardous behavior shall be identified,” and their impact on the SOTIF shall be assessed. Triggering events related to sensor disturbances can be caused for example by poor weather conditions or poor-quality reflections. They can cause errors in the world model if corresponding *functional insufficiencies* in the perception subsystem exist [21]. Once triggering events for the perception subsystem have been identified (see Refs. [8, 40]) as part of the so-called sensor verification strategy [8], they can also serve as test scenarios for it. Furthermore, Ref. [8] lists non-comprehensive example aspects to consider in verification testing of the perception subsystem, but leaves open the concrete realization of those tests.

4.3 NHTSA Vision and Framework

In “Automated Driving Systems 2.0: A Vision for Safety” [77], the National Highway Traffic Safety Administration (NHTSA) and the Department of Transportation (DOT) follow a “nonregulatory approach to automated vehicle technology safety.” AV manufacturers are allowed and required to follow and document their own individual safety assurance processes, which intends to support innovation rather than hindering it. Twelve so-called safety design elements are mentioned, where the one most relevant to this paper is *Object and Event Detection and Response* (OEDR)—the capability to “detect any circumstance that is relevant to the immediate driving task” and respond to it.

A more recent and more detailed publication of NHTSA defines a framework for testing AD systems [78]. This framework also addresses OEDR and further defines a taxonomy that can help describing an operational design domain (ODD) for an AV. Such a clear ODD description is necessary to express the requirements of how to also cope with unusual situations such as emergency vehicles.

According to Ref. [78], testing the perception outputs can significantly facilitate the assessment of OEDR capabilities

as compared to testing only the resulting trajectories. It can answer questions such as at which range obstacles are detected, if obstacles are correctly classified, and if their location and size are correctly estimated. However, this NHTSA document also leaves open many details of technical realization of such tests.

4.4 UL4600

Another approach to achieve a safety standard is UL4600 by Underwriters Laboratories and Edge Case Research [79] (voting version). Its full name is *Standard for Safety for the Evaluation of Autonomous Products*. UL4600 is a voluntary industry standard for autonomous products and thus also applies to AD functions. The creation of this standard is described in Ref. [80].

Similar to the previously addressed standards, UL4600 does not prescribe how to conduct concrete perception tests. Instead, it specifies claims within a safety case for which manufacturers should deliver arguments and evidence. These claims also go into details about sensing, perception, and machine learning [79]. Manufacturers are given a certain freedom in how to provide evidence for their claims. For example, the claim that the environment perception can provide acceptable functional performance [79] must be supported at least by the metrics false negative and false positive rate, and can be supported by “other relevant metrics.” In terms of existence uncertainty, false negatives (FNs) are objects missed by the SUT, whereas false positives (FPs) are wrongly detected ghost objects. For completeness, true positives (TPs) are existing objects which are also perceived by the SUT.

4.5 UNECE R157

The UNECE regulation [81] provides a regulatory framework for automated lane keeping systems (ALKS) of Level 3 on highways. It describes the role of a technical service to whom the manufacturer has to demonstrate its compliance with the framework. Stated requirements related to OEDR are a minimum field-of-view (FOV) of the perception subsystem and the ability to detect conditions that impair the FOV range [81]. A minimum set of scenarios for testing the FOV range is specified verbally in Ref. [81]. For instance, the perception of small targets like pedestrians or powered two-wheelers should be tested near the edges of the minimum FOV. The technical service is supposed to select and vary the concrete parameters of these test scenarios. Furthermore, the scenario parameters for system-level collision avoidance tests include also the perception-related conditions of the roadway and of lighting and weather [81]. Even though this regulation provides more specific test scenarios than other standardization approaches, they are not explicitly

targeted to perceptual uncertainty, and also lack detailed specifications of environmental conditions.

4.6 Safety First for Automated Driving

This white paper by a consortium of startups, OEMs, and suppliers intends to contribute to an industry-wide safety standard for AD by expanding the considerations of ISO/PAS 21448 to Levels 3 and 4 [82].

Related to this review, Ref. [82] highlights that testing is essential, but not sufficient to assure the safety of AVs. It suggests a test methodology that includes decomposing the AD system into elements such as the perception subsystem, and testing these elements separately. For the statistical validation of the perception subsystem in real-world tests, the usage of a reference perception system and the so-called scenario-based approach (see Ref. [14]) are emphasized. Recorded raw data from the field should be re-processed offline upon perception algorithm updates. The scenario-based approach can argue about a sufficient coverage of relevant traffic scenarios by grouping recorded data that include certain influencing factors into equivalence classes of scenarios (more about interpreting recorded data as logical scenarios in Sect. 7). Four constraints regarding perception testing are outlined [82]:

- The re-processing environment of the field recordings must be validated in terms of hardware and software.
- The test scenario catalog must be both statistically significant and covering the ODD sufficiently (see also Sects. 7.2, 7.3).
- The reference data quality must be appropriate for the validation objective (see also Sect. 8.5).
- Concrete test scenarios or data must be separate from those used in development (see also Sect. 7.5).

Additionally, Ref. [82] explicitly addresses safety challenges regarding machine learning, which is omnipresent in an AV’s perception, but has not yet been sufficiently addressed in ISO 26262 and ISO/PAS 21448.

4.7 ISO/TR 4804

The technical report “Road vehicles - Safety and cybersecurity for automated driving systems - Design, verification and validation” [83] aims at supplementing existing standards and documents on a more technical level. Since its content is closely related to the previously published “Safety First for Automated Driving” white paper [82], a repetition of the key aspects is omitted (Sect. 4.6).

4.8 Summary of Safety Standards

Contemporary automotive safety standards require arguments about successful OEDR for safety assurance, but mostly leave the realization of perception tests open to the AV manufacturer. This motivates the subsequent review of how perception testing is currently realized, and how it could be performed in future to assess if the perception allows safe vehicle behavior.

5 Established Activities of Perception Testing

This section expresses perception algorithm benchmarking outside the safety domain (Sect. 5.1) and data-driven sensor modeling on object level (Sect. 5.2) in terms of the testing taxonomy [4]. Both established activities are included into this separate section because they only partially target this paper's research question. Thereby, this paper aims at providing a complete context for AV perception testing while distinguishing between directly safety-relevant and less safety-relevant testing activities.

5.1 Perception Algorithm Benchmarking

Developers of computer vision and moving object tracking algorithms across multiple industries are already using established benchmark datasets and test metrics for their algorithm's results. The idea behind these testing activities is typically to provide a quantitative ranking among different algorithms rather than evaluating whether safety-relevant pass–fail criteria are met. The hardware and software that provide the raw sensor data are usually not analyzed.

In public research, the providers of public benchmark datasets typically determine all of the three testing axes test scenarios, reference data, and metrics (Fig. 3). Test criteria usually mean that novel algorithms should rank higher than previous algorithms.

5.1.1 Test Scenarios

Test scenarios are captured when the dataset provider records the raw sensor data for the benchmark. The aim for such raw data is typically not to capture safety-relevant scenarios, but rather to provide a diverse set of road user types that the algorithms under test will have to detect and distinguish. Therefore, for the purposes of perception algorithm benchmarking so far, it is sufficient that the recorded scenarios are only characterized in the data format of the recordings, and not in a dedicated scenario description language (Sect. 7.2). Examples for such datasets are KITTI [84], nuScenes [85], the Waymo Open Dataset [86], or Argoverse [87].

Fig. 3 Perception algorithm benchmarking expressed in taxonomy (Fig. 1)

5.1.2 Reference Data

Besides raw sensor data, algorithm benchmarking datasets typically also provide reference data on object level in the form of bounding boxes that are labeled offline with the help of humans [84–86]. For testing the world model output of the perception subsystem, reference data in a metric coordinate system over the ground plane are relevant, which can be labeled based on dense lidar point clouds. In contrast, pixel-level class labels in camera images do not provide a reference for the resulting world model.

Other developer testing activities provide the reference information of target vehicles by means of high-precision global navigation satellite systems (GNSS) [88].

5.1.3 Test Criteria and Metrics

After a perception algorithm has produced an object-based world model from the provided raw sensor data, its tracks can be compared to the provided reference tracks using metrics. Metrics that assess untracked object lists are also widely used, but are not discussed here because tracked object lists seem to be more common at the interface between the perception and planning subsystems. Since a common goal of object tracking metrics is to enable a ranking of algorithms, they often need to summarize various aspects of the perception performance into a small number of scalar values that are used for comparison. This concept of aggregating various detailed low-level metrics into fewer summarizing high-level metrics has already been proposed in 2005 for the performance evaluation of automotive data fusion [88].

Computing lower-level metrics such as the rates of FPs and FNs requires to *associate* (or *match*, often used synonymously) the estimated tracks of the SUT to their corresponding reference tracks, which is realized in the following way. The pairwise distances between all estimated and all reference objects or tracks are computed by means of an object distance function. Examples for such distance functions are the Euclidean distance between object centroids in the ground plane (applied in Refs. [85–87]) or the Intersection

over Union (IoU) of bounding box areas or volumes [84]. Mathematically speaking, however, IoU is not a proper distance function because it is bounded between 0 and 1 and only its complementary value ($1 - \text{IoU}$) expresses distance.

Using these pairwise distances, a multi-object association algorithm can compute an optimal association such that the sum of distances of all associated object pairs is minimized. Often, a threshold distance is applied to only allow reasonable object associations. For example, two objects with a Euclidean distance larger than 2m [85], or with an IoU smaller than 50% [84], could be prevented from becoming associated object pairs. Examples for association algorithms include the Hungarian/Munkres algorithm [89] and the auction algorithm [90].

Once this association is obtained, different higher-level metrics can be computed to describe the overall perception performance. Since the last step of a perception algorithm is usually to track the frame-individual objects over time while potentially fusing them into a single list of tracks, the following metrics are often called multiple object tracking (MOT) metrics.

The most common MOT metrics are the CLEAR MOT metrics MOTA and MOTP, which stand for MOT accuracy and MOT precision, respectively [91]. MOTA measures existence uncertainty by penalizing FPs, FNs, and ID switches, where ID switches are times when the ID of the estimated object changes while the ID of its associated reference object stays the same. MOTP measures state uncertainty by penalizing TPs that are not estimated precisely.

Another state uncertainty metric that additionally measures whether an estimated object state is consistent with the tracker's estimate of its own Gaussian state uncertainty is the Normalized Estimation Error Squared (NEES, see Ref. [92]).

The various forms and derivations of the optimal sub-pattern assignment (OSPA) metric [93–99] penalize existence and state uncertainties at the same time by weighing them to obtain a single score. Unlike most other mentioned metrics, the OSPA metric is a metric in the mathematical sense, meaning it satisfies the three properties (1) identity of indiscernibles, (2) symmetry, and (3) the triangle inequality. This property makes it interesting for research on object tracking on a detailed mathematical level [100, 101].

Besides existence and state uncertainties, the nuScenes detection score (NDS) [85] also considers classification uncertainties, as its goal is to represent the entire object tracking performance by only one scalar. Further metrics that are commonly used for either classification or existence uncertainty are precision, recall, and the area under the so-called receiver-operating characteristic (ROC) curve [57].

One evaluation concept that explicitly considers temporal aspects of the object tracking performance is the Mostly tracked/Partially tracked/Mostly lost approach, which counts

the number of reference trajectories that are tracked during more than 80% or less than 20% of their lifetimes [102].

The so-called Higher-Order Tracking Accuracy (HOTA) metric has been proposed to balance various sub-metrics in a single higher-order metric [103]. While the sub-metrics can describe individual performance aspects in detail, the higher-order metric can balance those sub-metrics without over-emphasizing one aspect over another [103]. The same publication [103] contains a more general introduction to MOT algorithm benchmarking metrics and an in-depth analysis of popular metrics, which further include the mean average precision for tracking (Track-mAP) and the so-called IDF1-score.

Recently, the “Planning Kullback–Leibler divergence” (PKL) [104] has emerged from the field of perception algorithm benchmarking. It is the only metric from this field that authors are aware of that explicitly considers the actual influence of perceptual uncertainties on the downstream motion planner. Various influences on the metric have been analyzed on the submissions of the nuScenes object detection challenge [105]. Interestingly, the submission rankings would be significantly different if PKL was used as the main benchmark instead of mAP [105]. Moreover, the PKL was found to be more consistent with human intuition than the NDS about which perceptual uncertainties are actually dangerous [104] and has been included into the official submission evaluation. Due to its direct relevance for this paper's research question, technical details on the PKL are discussed later along with other safety-oriented microscopic metrics (Sect. 6.2.3).

5.1.4 Difficulties with Association Uncertainty

There seems to be a fuzzy border between *state* and *existence* uncertainties, because a fixed threshold on the distance function that distinguishes TPs from FPs/FNs is not likely to produce intuitive associations under all circumstances. Therefore, Refs. [106] and [107] tune their offline object association in a way that it reproduces human annotations of object association. These approaches seem to lead to subjectively better associations, but have the cost of less human understanding of how the trained associator and subsequently also the test metric works.

5.1.5 Relevance for Vehicle Safety

Except for PKL [104], the mentioned state-of-the-art metrics for object perception and tracking do not consider safety, but rather provide information about the average similarity to a reference dataset [21]. This issue is described in more detail in Ref. [65]. According to Ref. [29], it is possible

to completely specify pass–fail criteria for such safety-independent metrics. Also ISO/PAS 21448 suggests to use safety-independent perception metrics [8]. However, those criteria would only represent non-functional requirements, while functional requirements needed for safety assurance remain an open issue [29].

Similarly, the authors of Ref. [68] point out that metrics like a FP rate cannot provide safety-relevant information because some FPs might be highly safety-relevant while others are not. Therefore, they suggest to formulate realistic fault models, which already exist in ISO 26262 for hardware faults also for perception algorithms. Such fault models should depend on the safety requirements of the overall vehicle. Given examples include the differentiation of a FP for a pedestrian from a FP for a bicycle, or to discretize continuous state uncertainty to obtain Boolean faults. Nevertheless, formulating such fault models in accordance with the safety requirements seems non-trivial, which motivates for a further analysis of metrics and test criteria in Sect. 6.

5.2 Object-Level Data-Driven Sensor Modeling

The idea of data-driven sensor modeling approaches is usually to treat the complex behavior of a sensor as a black box and replace it by a model which can generate artificial sensor data in simulations. To do so, the output of a sensor under test is compared to reference data in order to train or parametrize a model that describes the sensor's perception performance. Hence, sensor modeling can also be interpreted as a testing activity according to the taxonomy of this paper (Fig. 4).

5.2.1 System-Under-Test/System to be Modeled

In the context of sensor modeling, this paper uses the terms *system-under-test* and *system to be modeled* synonymously. Similar to the research interest of this paper, sensor modeling approaches like in Refs. [108–111] model object lists, which have been generated through both sensor hardware and perception software. This is common if the sensor's built-in perception algorithm is inaccessible to the modeling engineer due to sensor supplier intellectual property. Otherwise, the virtually tested AD function could contain parts of the perception algorithm, reducing the system to be modeled to mostly the sensor hardware [112]. However, modeling the already processed perception data has the advantages of less data volume and more comparable data formats.

5.2.2 Scenarios

Like in Perception Algorithm Benchmarking (Sect. 5.1), test scenarios for the SUT are captured during data recording. Already seemingly simple test scenarios with dry and

Fig. 4 Object-level data-driven sensor modeling expressed in taxonomy (Fig. 1)

sunny weather conditions and with only few other road users often provide enough difficulty for data-driven modeling of the sensor hardware behavior. For example, modeling radar reflections on guardrails on a motorway is a topic of current research [112]. Unlike for perception software development, the research field of sensor modeling seems to be lacking common benchmarking datasets that would allow a straightforward comparison of different sensor models on the same data.

5.2.3 Reference Data

Modeling the behavior of sensor hardware in a data-driven approach requires an independent reference to the sensor hardware. This makes RTK-GNSS-IMUs (explanation in Sect. 8.2.1) a suitable and typically used source of reference data [108, 113].

5.2.4 Metrics

Metrics in this paper are quantitative statements on the SUT. Since parametrized or trained sensor models can describe how similar the SUT's perception is to the reference perception, they are interpreted as test metrics here. Note that sensor models are usually not single-score metrics, but rather for example probability distributions of the sensor's errors [108]. Note also that the term *metric* in sensor modeling literature is used differently—there, metrics make a statement on the modeling approach rather than on the SUT [114].

Actual statements on the SUT can be for example the mean values and standard deviations of Gaussian distributions that describe the SUT objects' state errors in position and velocity [110]. A nonparametric distribution for such errors that is based on a Gaussian mixture model is used in Ref. [108]. There are various further ways of expressing the SUT's errors by sensor models, which are however mostly outside the focus of this paper.

5.2.5 Relevance for Safety-Oriented Perception Testing

Typically, sensor models aim at describing the specific phenomena of a sensor modality as detailed as feasible, no matter how safety-relevant these phenomena are [109, 114]. However, some sensor modeling activities explicitly address this paper's research question. For example, Ref. [65] models perception errors while considering the effect they have on robust decision making. In the context of validation of sensor models, Ref. [115] argues that a key property of modeled sensor data should be that they induce the same behavior in the downstream driving function like the real sensor data would do. Such topics are further elaborated in Sect. 6.2.2 about the perception-control linkage.

6 Test Criteria and Metrics

This section and the following two sections each cover one testing axis of the used taxonomy from Ref. [4] (Fig. 1) and are dedicated to the primary research question of this paper.

According to the taxonomy, test criteria and metrics are “a statement on the system-under-test (test criteria) that is expressed quantitatively (metric)” [4]. For example, a criterion that is qualitative at first could be quantified by means of specifying intervals on a related metric that determine passing or failing a test.

After examining how to specify perception requirements and test criteria (Sect. 6.1), this section deals with safety-oriented microscopic metrics and criteria (Sect. 6.2). Furthermore, metrics on the self-reporting and confidence estimation capabilities of the SUT are discussed (Sect. 6.3), as well as macroscopic metrics toward approval (Sect. 6.4).

6.1 Specification of Requirements and Criteria

A perception subsystem shall enable the AV to reach its destination safely, comfortably, and in reasonable time. Therefore, it shall provide information in sufficient quality about all road users that are relevant for fulfilling the driving task. However, such requirements are not specific enough to be tested [116]. Thus, how can one specify “relevant for fulfilling the driving task,” or “sufficient quality” for usage as binary test criteria?

In the following, it may be useful to differentiate between concept specification and performance specification of environment perception [30]. In the mentioned source, concept specification refers to defining the properties to be perceived such as a pedestrian's pose, extent, and dynamic state, given an ODD. In contrast, the performance specification defines how well these properties should be perceived, for example in terms of detection range, confidence, and timing [30].

6.1.1 The Difficulty of Specifying Perception

AVs are expected to operate in unstructured, public, real-world environments, which are called *open context* in Ref. [117]. According to Ref. [72], a complete concept specification of the environment perception may not be possible because a model about the environment generally cannot cover all necessary relations and properties in such an open context. This issue has also been called ontological uncertainty [118]. For example, to specify the perception of pedestrians, one would have to specify what a pedestrian is, which is, however, only partially possible using rules such as necessary or sufficient conditions [29]. Providing examples of pedestrians in a training set is how machine learning engineers specify the concept of a pedestrian. On the one hand, this can enable driving automation of Levels 3 and above, but on the other hand, it prevents a traditional specification according to ISO 26262 [29].

Besides specifying the perception concept and performance for discrete environmental aspects like the classification of a road user, another key challenge is to identify when uncertainty in continuous and dynamic environmental aspects, like a car's velocity, leads to safety-relevant failures [116].

6.1.2 Concrete Approaches of Specifying Perception

Contributions from the field of machine learning investigate how to specify the perception subsystem by using pedestrian detection as a benchmark example [29, 73]. On a higher level, Ref. [29] emphasizes the importance of an adequate language for specification and the potential of deliberate *partial* specifications. The paper proposes and evaluates several methods for incorporating partial specifications into the development process. Further literature with concrete specification approaches includes creating an ontology of the exemplary “pedestrian” domain (concept specification) [73], and taking human perception performance as a reference (performance specification) [119]. A formal language for specifying requirements on the performance of object detection in the absence of reference data is proposed in Refs. [120, 121].

Instead of using an environmental concept like a pedestrian as the center of investigation, the methodology of Ref. [67] outlines how particular test criteria for AV subsystems can be defined in a top-down way, starting from overall safety goals. The methodology is, however, not yet applied specifically to the perception subsystem.

As mentioned earlier, the complexity of the open context can cause gaps in the specified requirements. To fill these gaps, many organizations collect large numbers of real-world mileage to discover so far unknown scenarios [26].

These recorded scenarios can also serve as test scenarios if they are identified as test-worthy (Sect. 7.3).

Besides the approaches mentioned so far, there are also considerations on perception requirements that aim at complying with the traditional ISO 26262 functional safety standard. The sources [35, 52, 53] propose dynamically associating an Automotive Safety Integrity Level (ASIL) with a given driving situation such that for example, the perception is required to comply with the stricter ASIL D in high-risk situations and with the less strict ASIL A in low-risk situations. An example of how functional safety requirements for the perception subsystem can be derived based on a fault tree analysis (FTA) is given in Ref. [122] in the context of automated valet parking.

The following subsections deal with metrics for measuring quantitatively whether specified test criteria are met.

6.2 Microscopic Test Criteria and Metrics

The currently most established perception performance metrics, which are typically used in machine learning, do not represent whether the perception output is sufficient for safe vehicle operation [21] (Sect. 5.1.5). This section therefore reviews literature about metrics and criteria that explicitly distinguish safety-relevant from safety-irrelevant perception errors.

6.2.1 Heuristic for the Safety-Relevance

The authors of Ref. [49] provide a simple approach toward the safety-criticality of perception errors in the existence uncertainty domain. Certain fractions of the binary error types FP and FN are assumed to be safety-critical, where the fractions depend on the error's position of occurrence within the ego vehicle's FOV. For example, perception errors directly in front of the ego vehicle can heuristically be estimated to be more likely to be safety-critical than perception errors occurring farther away or with a lateral offset [49]. The benefit of this approach is that once its numerical values are set, it does not need to consider any downstream driving function for computing safety-critical failure rates.

6.2.2 Modeling the Perception-Control Linkage

However, whether a perception error turns out to be safety-critical or not generally does depend on the downstream driving function. Thus, this section describes different ways of modeling the interface between perception and planning/control, which is also called the perception-control linkage [28].

The above-mentioned fractions of safety-critical perception errors could be determined by means of closed-loop

fault injection simulations [49]. However, if the downstream driving function receives a minor update, the perception metrics would have to be recomputed, which would render such a direct approach unpractical for iterative development.

Alternatively, one could abstract the perception and planning subsystems such that test results of the perception subsystem can be re-used for varying planners. For this purpose, modular functional system architectures [5, 38, 42, 116] could be implemented with contracts, assumptions, and guarantees at the interfaces between the perception and planning subsystems [22, 35, 45, 60].

Such a modular approach would enable different safety assurance methodologies for the different subsystems. For example, the following section discusses data-driven perception testing and formal safety assurance of the planner. The idea behind such formal methods is to always assure a certain safe planning behavior, given that the environment is perceived well enough. Popular formal models for safe planning include Responsibility-Sensitive Safety (RSS) [33] and reachability analysis [43, 123].

6.2.3 Downstream Comparison of SUT and Reference Data

Traditionally, perception metrics compare the SUT with the reference system by means of their object lists that are handed over from Layer 2 to Layer 3 (Fig. 2). With an additionally given planner or assumptions about it, safety-oriented perception metrics can be computed further downstream, namely after Layers 3 and 4 (Fig. 5). The planner, or a formal abstraction, computes behavior outputs for world model inputs from the SUT and from the reference system. Safety-oriented metrics then penalize if the behavior induced by the SUT is different from the one induced by the reference system, especially in terms of danger and worst-case severity.

The previously mentioned PKL metric [104, 105] (Sect. 5.1.3) works in this fashion. It uses an end-to-end machine-learned planner, which is trained to imitate the human ego vehicle driving behavior of the recording of the dataset on which it will be applied. During inference, this planner is fed with the reference object list and with the object list of the SUT. For both object lists in a given scene, it computes the likelihoods of the recorded ego vehicle trajectory during a short time interval after the given scene. A "Planning Kullback–Leibler divergence" between these likelihoods then represents the degree to which the SUT's uncertainties influence the planner.

While PKL uses a concrete planner, the RSS framework [33] defines safety-relevant perception failures using a formal abstraction of the planner and its formal definition of *dangerous* situations. According to the publication, *safety-critical ghosts* are situations in which the environment perceived by the SUT is formally considered *dangerous* for the

Fig. 5 Safety-oriented perception testing of Refs. [28, 33, 104], expressed in taxonomy (Fig. 1) and layers (Fig. 2)

ego vehicle even though according to the reference data, it is not. Likewise, *safety-critical misses* are situations in which the environment perceived by the SUT is formally not considered *dangerous* even though it actually is. For example, a safety-critical ghost can be a false-positive pedestrian detection in front of the ego vehicle that could cause a false and dangerous braking maneuver. A safety-critical miss could be a pedestrian at the same location as a false-negative detection, which could lead to the ego vehicle hitting the pedestrian. Safety-critical ghosts and misses can not only be caused by existence uncertainties, but also by state and classification uncertainties [33] because their definition is agnostic of the type of uncertainty. This avoids a potentially ambiguous differentiation between state and existence uncertainties (see Sect. 5.1.4).

However, this concept of safety-critical perception failures [33] so far seems to have been treated mostly theoretically. For example, its extension to consider multiple time steps or its practical demonstration on real data are not found.

Salay et al. [28] expand this concept of binary safety-critical perception failures to a potentially continuous description of the *risk* that a perception failure can cause. The work formalizes a concept to analyze the so-called *incurred severity* that an uncertain environment perception can cause. It also requires knowledge about the planning subsystem, for example, the planner's policy for computing actions based on a world model. Furthermore, for actions that can cause harm, there must be a way to assess the worst-case severity of this harm. With these assumptions, the microscopic and safety-oriented perception metric *incurred severity* is defined as the difference between the SUT's and the reference system's induced worst-case severities of control actions [28] (Fig. 5).

The same work [28] applies this concept in a case study dealing with classification uncertainty of road users. For example, if a pedestrian is correctly perceived as a pedestrian, then the worst-case *incurred severity* is zero. If a pedestrian is however classified as a cyclist, and if the control action for cyclists is less cautious than for pedestrians,

then the *incurred severity* of this misclassification is likely positive. Future work not yet covered in Ref. [28] is to generalize the computation of *incurred severity* to world models that are also subject to existence and state uncertainty. Moreover, representative real-world exposures (probabilities of occurrence) of scenes would be needed for valid risk computations.

A practical challenge in the approach of Ref. [28] might be the availability of the planner's behavior policy. This challenge is addressed in the subsequent positional paper by Salay et al. [27] on perceptual uncertainty-aware RSS (PURSS). It proposes to let the RSS framework handle world model uncertainties in a more advanced and more practical way than originally published. A behavior policy of the concrete planner is not needed; instead, it is assumed that the planner complies with the formal RSS rules for guaranteed safe control actions. The RSS model, which is explained in detail in Ref. [33], can provide a set of guaranteed safe control actions for any given world model input. Using this set, safety-oriented test criteria for the perception subsystem can be defined. The set of safe control actions is computed for both the world model from the SUT and for the reference world model. The safe control actions S_{SUT} for the SUT's world model is what a real planner would execute, whereas the safe control actions S_{ref} for the reference system's world model is what is actually safe. If there are control actions in S_{SUT} that are not part of S_{ref} , the SUT potentially causes a safety risk [27] (Fig. 6).

Furthermore, the authors of Ref. [27] propose how RSS could be extended to handle imprecise world models, which are sets of SUT world models that contain the true world model with a certain confidence. For example, the position of a pedestrian could be estimated by an interval of [29m, 31m] with a confidence of 95%. In practice, contemporary perception algorithms do output confidences about their world models, for example Gaussians for state uncertainties, which could be sampled at the mentioned confidence intervals. A corresponding test criterion for an imprecise world model is that its RSS-defined set of safe control actions must be safe for each precise world model contained

Fig. 6 Testing whether the SUT's world model can induce control actions that would not be safe for the reference world model [27]. Figure follows taxonomy of Fig. 1

in it [27]. The authors show that the mentioned perception confidence can be used as a lower bound for the probability of guaranteed safe control actions [27].

6.3 Metrics for Uncertainty/Confidence Calibration

As described in the previous section and also in Ref. [26], the explicit consideration of world model uncertainty (or inversely, confidence) can facilitate safety assurance, for example through runtime monitoring (more in Sect. 9.3.3). In this context, it is crucial that the self-reported uncertainty of the SUT correctly reflects its true uncertainty [21]. If this holds, then the SUT is *calibrated* [61], which means that it is neither over-confident nor too doubtful of itself.

6.3.1 Types of Uncertainty

The literature on uncertainty in AD (e.g., [64]) describes that *epistemic* uncertainty, or model uncertainty, constitutes how uncertain the SUT's model of deep neural networks (DNNs) is in correctly describing the environment. In contrast, *alea-toric* uncertainty is caused by physical sensor properties like finite fields-of-views, resolutions, and sensor noise. Both types are relevant for correct calibration. Additionally, Ref. [118] proposes the concept of *ontological* uncertainty to describe the complete unawareness of certain aspects of the environment, even in the reference data.

6.3.2 Representations of Uncertainty

In object perception, state estimation is a regression task (continuous true value), while existence and classification estimation are classification tasks (discrete true values) [64]. State uncertainties are typically quantified by continuous probability distributions or confidence intervals, while existence and classification uncertainties are usually expressed by scalar probabilities between 0 and 1.

6.3.3 Calibration Metrics

This section focuses on metrics that only describe how well-calibrated the world model uncertainties are. In contrast, uncertainty-aware classification or regression metrics [64] are a topic of perception algorithm benchmarking (Sect. 5.1).

Multiple literature sources set up so-called calibration curves [61, 124, 125] for visualization and numerical analysis. These curves plot the accuracy (or empirical frequency) of a prediction over the prediction's confidence, often using bins. For example, in classification tasks, one specific bin could contain all events where the detector reports a confidence/existence probability in between 80% and 90%. If there is a pedestrian present for only 70% of all events in that bin, the calibration curve would be tilted off from the ideal diagonal line. In regression tasks, a similar calibration curve can be set up by defining the accuracy as the empirical frequency of when the estimated confidence interval contained the true value.

Single-score calibration metrics can be computed from such calibration curves [64]. In classification tasks, the *Expected Calibration Error* (ECE) and the *Maximum Calibration Error* (MCE) [124] represent the expected and the maximum difference between confidence and accuracy in such calibration curves, respectively. Similar metrics for regression tasks are the *calibration error* [125] or the related *Area Under the Calibration Error Curve* (AUCE) [126].

Concrete applications of uncertainty calibration and evaluation in 3D object detection from lidar point clouds are given for example in Ref. [75] and the follow-up publication [76]. The latter publication explicitly uses the mentioned calibration curves and the ECE metric.

A further uncertainty evaluation approach that explicitly addresses perception systems in safety-relevant domains is provided in Ref. [70] and subsequently Ref. [63]. The publications, which focus on out-of-distribution detection for classification tasks, distinguish four different cases in classification results. The cases are combinations of the properties certain/uncertain (decided by a threshold), and correct/incorrect. Metrics are defined based on the fractions of these individual cases among all classification results. The fraction of certain, but incorrect results is defined as the *Remaining Error Rate* (RER), and the fraction of certain and correct results is defined as the *Remaining Accuracy Rate* (RAR).

All mentioned metrics so far assume that a precise world model is available as a reference. However, also reference data are generally uncertain and imprecise (Sect. 8.4), which makes test metrics more complex.

6.4 Macroscopic Metrics Toward Approval

So far in this paper, the used metrics and test criteria have been statements on the SUT in either individual scenes

(microscopic metrics) or in relatively small amounts of test data with a research purpose. While microscopic metrics and criteria are necessary for a detailed analysis of the SUT, they alone are not sufficient for the overall safety assurance of AVs because their results need to be extrapolated to estimate *macroscopic* [127] metrics (definition in Sect. 1.3).

6.4.1 Terminology: Safety vs. Reliability

Since the term *reliability* is often used in the context of macroscopic safety metrics [49], it is first briefly differentiated from the term *safety*. The perception subsystem of an AV must enable the overall vehicle to drive safely. Safety, however, is a property of the overall system and not of the perception subsystem. The perception subsystem can at best be *reliable*, which is generally different from being *safe* [7] (Sect. 1.3). Note that a property that includes both safety and reliability is *dependability* [128].

6.4.2 Mean Time Between Failures

The RSS publication [33] argues that for the acceptance of AVs, the overall probability of occurrence of safety-critical failures must be in the order of magnitude of 10^{-9} per hour. This would imply a test criterion of no such failures during about 10^9 hours of driving, where safety-critical failures in their publication are the previously introduced safety-critical ghosts and misses. The related paper [66], which focuses on perception within RSS, states another large required mean time between failures of 10^7 h for the perception subsystem. Only perception failures that can possibly lead to unsafe vehicle behavior should be counted in that number.

Major contributions to the macroscopic reliability assessment of the perception subsystem have come from Berk et al. [36, 49]. The macroscopic metrics they use are the rate of all perception failures and the rate of safety-critical perception failures, both measured in failures per hour. Safety-critical perception failures are computed by filtering all perception failures with a field-of-view-dependent safety-criticality factor between 0 and 1, as explained in Sect. 6.2.1. An overall approval criterion according to Ref. [36] is that the sum of the failure rates of the perception, planning, and actuation subsystems is smaller than a given threshold rate, which is called the target level of safety. A corresponding testing methodology to compute these macroscopic perception metrics is discussed in Sect. 9.3.4.

6.5 Summary of Test Criteria and Metrics

Different metrics describe different properties of the SUT like its safety-relevant performance in a single scene, the calibration of its self-reported uncertainties, or its statistical safety impact. Specifically, the representation of

safety-relevance in microscopic metrics seems challenging. These mentioned different classes of metrics are typically coupled. For example, the SUT's statistical safety impact depends on its safety impact in individual scenes, which in turn may depend on whether the SUT has correctly reported its uncertainty in that scene. Thus, one challenge is to harmonize individually dedicated metrics into an overall framework. Moreover, the specification of test criteria is non-trivial due to the complexity of the open-world context.

7 Test Scenarios

According to the taxonomy of Ref. [4], a test scenario is “a set of specified conditions” under which a test is executed. The following subsections deal with adapting the term *scenario* to the perception context (Sect. 7.1), describing test scenarios and the ODD (Sect. 7.2), generating a test scenario catalog (Sect. 7.3), executing scenarios as test cases (Sect. 7.4), and splitting scenarios into training and test sets (Sect. 7.5). Much of the literature on scenario-based testing does not yet focus on perceptual uncertainty, which is why this section also cites related literature with a broader focus.

7.1 Adapting the Term Scenario to the Perception Context

The definition of the term *scenario* has been discussed extensively in the context of scenario-based safety assurance [11]. According to the mentioned source, a scenario contains actors (mostly road users) whose actions influence the temporal development of scenes within the scenario. For testing purposes, one of such actors within the scenario is the subject vehicle. So far, the focus of scenario-based testing of AVs has been on the plan and act subsystems instead of on the perception subsystem of the subject vehicle [14]. However, if only the perception subsystem is tested in an open-loop fashion, it typically only has the role of a passive observer of the scenario instead of an influencing actor. At this point, closed-loop testing is excluded because the interest is in the offline test of potentially multiple perception algorithms on the same recorded raw sensor data from the real world.

To the best of the authors knowledge, we are not aware of the literature that explicitly considers the influence that an open-loop perception subsystem may have on the temporal development of a scenario. Such an influence is in fact possible through electromagnetic waves emitted by active sensors, which could disturb the sensors of other AVs and thus theoretically influence their behavior. However, in this review, for simplicity it is assumed that perception-specific scenarios are scenarios whose temporal development of objective scenes is pre-determined, but

where the subjective observation of these scenes in world models is left open. This corresponds to *static scenarios* with an additionally pre-determined ego vehicle behavior according to Ref. [129].

Furthermore, it seems that the different abstraction levels of scenarios (functional, logical, concrete) of Ref. [130] can also be used for perception scenarios. A functional scenario is described verbally (e.g., it rains, among other aspects), whereas a logical scenario has parameters (e.g., uniform precipitation in millimeters per hour), and a concrete scenario gives each parameter a numerical value (e.g., 2mm/h). Generally, the logical scenario parameters of interest differ between perception tests and planning tests [131].

While the safety assurance domain already uses the previously mentioned terms and concepts, the authors are not aware of similar in-depth definitions of the term *scenario* originating from the perception domain. However, the perception domain partially uses different terms for similar concepts (see also Sec. 5.1.1). For example, a *scene* as used by the nuScenes dataset [85] (perception domain) mostly corresponds to the recording of a *scenario* (safety assurance domain).

7.2 Description of Scenarios and ODD

Arguing whether test scenarios cover the entire ODD requires ways to describe scenarios and the ODD.

Current ODD description approaches are the taxonomy of BSI PAS 1883 [132], the NHTSA framework [78], and the ongoing standardization of a machine-readable ODD format in ASAM OpenODD [133]. Furthermore, ontologies are used to describe ODDs [59] and scenes within scenarios [44]. Ontologies are intended to be standardized in ASAM OpenXOntology [134] and have also been proposed to define a schema for world model data [27]. For ontology-generated scenarios, the notion of *abstract scenarios* was introduced to describe functional scenarios that are formal and machine-readable, but do not yet have a logical parameter space [131]. The ontology of Ref. [44], which is used to generate traffic scenes for scenario-based testing, organizes all scenario entities by means of a 5-layer-model. This model was first defined using four layers in Ref. [135], and has been adapted to six layers in Ref. [136] and most recently in Ref. [137]. The model's individual layers according to the most recent and also summarizing Ref. [137] are

- (1) Road network and traffic guidance objects
- (2) Roadside structures
- (3) Temporary modifications of Layers 1 and 2
- (4) Dynamic objects
- (5) Environmental conditions
- (6) Digital information.

These scenario description layers should not be confused with the functional decomposition layers in Ref. [5] (Fig. 2). Even though the present review only focuses on perceiving dynamic objects in Layer 4, the properties of all other layers can influence the perception of dynamic objects and are therefore also relevant.

So far, the layer models for scenario description only provide a way for an objective description of the environment within a traffic scenario. Aspects related to a subjective perception of such scenarios through machine perception are not yet explicitly covered, such as surface materials that affect radar reflections [137]. Such shortcomings of scenario descriptions for perceptual aspects have also already been identified in the context of sensor modeling [112].

The previously mentioned scenario description methods are suited to describe functional, abstract, and perhaps logical scenarios. In contrast, describing concrete scenarios for the actual test execution requires a certain data format. The ASAM OSI [138] classes GroundTruth and SensorView appear to be capable of describing objective scenes and the circumstances for their subjective perception, respectively. However, they are designed for virtual simulations instead of real-world test scenarios. Similarly, the established format OpenSCENARIO [139] is also designed for virtual closed-loop testing of the planner instead of open-loop testing of the perception subsystem.

7.3 Generating a Test Scenario Catalog

The ability to describe and structure test scenarios in an ideally formal way is a prerequisite to subsequently generate a representative test scenario catalog. Such a test scenario catalog might be part of the specification of the perception subsystem (Sect. 6.1).

The number of test scenarios has to stay feasible, which renders naive sampling of a high-dimensional scenario space unsuitable and instead motivates for the following approaches. This paper distinguishes between knowledge-driven and data-driven scenario generation, as already done in Refs. [4, 15, 131]. As pointed out by Ref. [15], approaches from both categories usually need to complement each other for a holistic testing strategy. Likewise, also ISO/PAS 21448 mentions dedicated expert scenario generation (knowledge-driven), as well as large-scale random data recording (data-driven) for the perception verification strategy [8].

In either way, the goal is usually to obtain certain *triggering events*, or similarly, *triggering conditions*, which are defined in ISO/PAS 21448 [8] as “specific conditions of a driving scenario that serve as an initiator for a subsequent system reaction possibly leading to a hazardous event.” Such a system reaction can occur either in the sensor hardware, in the perception software, or the downstream driving function. Similar to triggering events, the terms *external influencing*

factor [41] and *criticality phenomenon* [131] are also used to describe safety-relevant scenario influences related to perception.

7.3.1 Knowledge-Driven Scenario Generation

Since human knowledge of traffic scenarios is often qualitative, knowledge-driven scenario generation usually first yields functional or abstract scenarios. The literature on driving scenarios describes the further steps to logical scenarios [140] and the discretization to concrete scenarios [141].

ISO/PAS 21448 enumerates a list of influencing factors, which shall be used to construct scenarios to identify and evaluate triggering conditions for the perception subsystem [8]. This list is included into the standard as an informative example and includes factors like climate, which can be for example fine, cloudy, or rainy.

A specific identification of triggering events according to ISO/PAS 21448 is provided in Ref. [40] by means of a hazard and operability study (HAZOP). The approach is applied to the used sensor modalities camera and lidar as part of an extended automated emergency braking system. Sensor modality-specific influencing factors can be analyzed in even more detail by experts using dedicated tests or simulations. Example studies analyze the influences of weather, road dirt, or rainfall on lidar ([51, 55, 56], respectively), or corner cases specific to visual perception [69].

Such influencing factors or triggering events could also be queried from a domain ontology for driving scenarios [44] or from one that also includes perceptual aspects [131]. The former source argues that without the use of an ontology, human experts are not likely to generate all possible scenarios based on their knowledge. In contrast, if the knowledge is first encoded in an ontology, the generation of scenarios or scenes from that ontology can cover all possibilities based on the encoded knowledge.

A further approach toward completeness of a scenario catalog is to crowd-source a list of concrete OEDR aspects that are relevant for testing [46]. The mentioned publication provides a detailed list that is supposed to serve as a starting point for further extension by the community.

While so far, the mentioned influencing factors were selected mostly according to their influence on the perception subsystem, one can also construct perception test scenarios with the downstream driving function in mind. This particularly includes scenarios that would be challenging to the driving function in case a perception failure occurs. Such a concept is applied under the name of application-oriented test case generation in Refs. [41] and [58] with adaptive cruise control (ACC) systems in mind. For an ACC function, it appears possible to manually investigate the perception-based influencing factors that negatively affect its behavior.

In contrast, if the targeted application is a general-purpose urban Level 4 system, it will be harder to enumerate all related influencing factors.

7.3.2 Data-Driven Scenario Generation

If test scenarios are obtained by analyzing large amounts of randomly recorded vehicle fleet data, then the abstraction level of available data is closer to concrete scenarios than to functional scenarios. However, the recorded data is initially only available in the data schema of the recording format, which usually does not share the same parameters like a logical scenario space. Thus, a challenge in data-driven approaches is to generalize the snippets from the recorded data in a way that they correspond to points in a logical scenario space. This is necessary in order to vary individual parameters such that potentially relevant, but not yet observed scenarios can also be generated. The authors are not aware of scientific publications that present data-driven scenario generation in the safety context specifically for perception testing, and thus will refer to related literature about general testing, which has already been reviewed in more detail in Ref. [15].

One example approach for a data-driven test case generation is described in Ref. [142]. In the mentioned publication, a database system of relevant scenarios is described. Its input data from various sources such as field operational testing, naturalistic driving studies, accident data, or others, must first be converted to a common input data format of traffic recordings. Afterward, individual snippets of the input data are affiliated to pre-determined logical scenario types, such that each snippet corresponds to one point in a logical scenario space. Logical scenario parameters can also be obtained through unsupervised learning on the input snippets, as applied to vehicle trajectories in Ref. [143]. If the input data were obtained from large amounts of randomly collected driving data, one can further compute statistical measures such as the scenario's real-world exposure and its parameters' real-world probability distributions [142]. By additionally computing criticality metrics on the given snippets, one could estimate the potential severity of a section of the logical scenario space. The exposure and severity as meta-information on the scenario space then help to turn the obtained logical scenarios into relevant test cases. For example, one could put more testing efforts on scenarios with a high risk, where the risk of a scenario increases with its exposure and severity, and decreases with its controllability [144].

7.3.3 Combined Scenario Generation

As pointed out by Ref. [15], knowledge-driven and data-driven approaches for scenario generation often complement

each other in practice. For example, the previously mentioned database approach from Ref. [142] assumes that a concept for logical scenario spaces is given, potentially from expert knowledge. In turn, knowledge-driven scenario generation can also be backed up by measurements and data [15].

One example of a combined test case generation for the perception subsystem is given in Ref. [5]. Besides the functional decomposition layers (Fig. 2), the publication also proposes a method for defining particular test cases for each of those decomposition layers. The method takes driving scenarios from a database such as Ref. [142], assumes an accident if there was none, and then uses a fault tree analysis (FTA) to define pass/fail criteria for the individual subsystems in a given scenario. This combined approach allows the derivation of perception-specific test scenarios from general driving scenarios.

7.3.4 Test Scenarios Specific to the SUT

In the end, it could also be that the test scenarios that result in the highest risk for a given SUT are *not* found by only considering external influencing factors from the space of all possible scenarios. Instead, the unexpected insufficiencies that the DNNs of the SUT have learned might lead to more critical behavior than what a usually suspected external influence could trigger [21]. For example, Ref. [145] identifies perception subsystem inputs that lead to highly uncertain world models for a given SUT by propagating uncertainties through the individual perception components.

7.3.5 Difficulties in Covering the ODD with Scenarios

Any attempt to formally describe the ODD and to cover it sufficiently with test scenarios is generally difficult because of the open context problem of the real world [117]. Related to this, Ref. [21] expresses that the data or scenarios used for the development and training of the perception subsystem are usually not a good approximation of the system's ODD in the real open world. Furthermore, the real world changes over time, which makes it necessary to iteratively update the ODD description and its coverage with test scenarios [21].

Moreover, in practice only a finite number of concrete test scenarios can be sampled from the infinite number of concrete scenarios that could be theoretically generated from a logical scenario. This sampling makes it generally unclear whether the results of two similar concrete test scenarios also hold in between these discrete points. In fact, DNNs used for object detection can show a potentially extreme nonlinear behavior with respect to only slight input variations for example in weather conditions. This brittleness, which has been pointed out in Ref. [21], makes it difficult to argue about a sufficient ODD coverage with sampled test scenarios.

Nevertheless, combinatorial testing [146], which is suggested by ISO/PAS 21448 and already applied to computer vision for AVs [147], can potentially keep the number of discretized concrete tests feasible while providing certain coverage guarantees. Furthermore, combinatorial testing that explicitly exploits the functional decomposition approach for AVs [5] is performed in Ref. [148] for simulated AV testing.

7.4 Executing Scenarios as Test Cases

This paper assumes the execution of a test scenario to be an offline execution of a perception algorithm on previously recorded raw sensor data. In this way, only the sensor mounting positions and sensor hardware are fixed by the recording, whereas multiple perception algorithms can be applied.

For scenarios that correspond to snippets from existing real-world vehicle recordings, the corresponding raw data can directly be used for the test without having to be re-recorded. Otherwise, the raw sensor data for the test scenarios must be captured. Public roads offer less influence on the scenario details, but allow encountering many different scenarios over time with little effort per scenario. In contrast, proving grounds allow to set up specific scenarios, which however requires high efforts per scenario and may lack realism, e.g., when reproducing precipitation.

7.5 Training and Test Sets of Scenarios

If the DNNs of the SUT are trained on data and scenarios that are also used for testing, overfitting likely takes place during training and an unbiased testing for verification and validation purposes is not possible. Thus, it is important to properly separate recorded perception data into individual datasets for training and for testing [21, 82]. Furthermore, if the developers evaluate the SUT on a test set multiple times, an unintentional optimization with respect to the test set result might take place over multiple development iterations [21].

7.6 Summary of Test Scenarios

Scenario-based approaches to minimize the overall testing effort of AD functions also seem applicable to the perception subsystem [82]. An agreed-upon description language and data format of perception scenarios would be beneficial for further research. Various approaches have been proposed to obtain test scenarios, where their combination for a sufficient ODD coverage is an active research topic. Furthermore, it is challenging to obtain real-world recordings of all identified test scenarios.

8 Reference Data

The third and final testing axis that this paper covers is the “knowledge of an ideal result (*reference*)” [4]. For object-based environment perception, such an ideal result is an object list containing all road users that the SUT is supposed to perceive. The terms *ground truth* and *reference* are distinguished in a way that a ground truth perfectly and objectively describes the world. A reference is a subjective approximation of the ground truth in terms of data that are generated by a perception system that is superior to the SUT.

Requirements for reference systems for environment perception have been analyzed in depth in the context of ADAS [34]. Apart from the qualitative requirements mobility/portability and reporting of its own uncertainty, the paper states the quantitative requirements reliability, sufficient field-of-view, accuracy, and proper timing of the measurements. The following sections review different ways of generating reference data, where each way has its own advantages and disadvantages in the mentioned categories. This article structures the literature according to the position of the reference sensors, which can be mounted only on the ego vehicle (Sect. 8.1), also on other road users (Sect. 8.2), or externally and not part of any road user (Sect. 8.3). Uncertainty in reference data and an appropriate choice of the reference data source are discussed in Sect. 8.4 and Sect. 8.5, respectively.

8.1 Reference Data From Ego Vehicle Sensors

The biggest advantage of reference data from sensors of the ego vehicle is that no external measurement equipment is needed. Without an external perspective, however, road users that are occluded for the sensors under test can also hardly be present in the reference data.

8.1.1 Reference from Sensors Under Test

Perhaps the most common approach of generating reference data is to use the raw sensor data from the SUT such that humans can label the reference objects and their properties. Labels in a sensor-specific coordinate system, such as in a camera’s pixel space, need to be transformed to the world model’s coordinate system (often in meters and defined over the ground plane) to potentially serve as a reference for it. In practice, the human labeling process is often at least semi-automated by offline perception algorithms, which can go back and forth in time. Such reference data is typically used by datasets for perception algorithm benchmarking such as from Refs. [84–86].

Using only offline perception algorithms without human intervention (as also done in Ref. [149]) dramatically reduces the overall effort and could still result in superior data to the SUT, but generally impairs the data quality. A standardization attempt for human-labeled reference data is ASAM OpenLABEL [150].

In either case, no additional sensor hardware is required, which is valid and a big advantage if the perception algorithm is the only component of interest. At the same time, it means that the SUT’s sensor mounting positions and sensor hardware do not have an independent reference, which might impede their rigorous assessment.

8.1.2 Separate Reference Sensors on Ego Vehicle

Additional sensors on the ego vehicle that are not part of the SUT can also potentially provide a reference for validating the SUT’s sensor mounting positions and hardware performance, see Ref. [149]. This approach seems to be more popular in ADAS developments than for AVs because AVs often already use the best available sensors for their regular operation, and hence there are no other sensors of higher quality left as reference.

8.1.3 k-out-of-n-vote of High-Level Fusion Inputs

The present review paper focuses on testing the world model that is generated from fused inputs of all environment sensors. In contrast, Berk et al. [32] proposed testing the object-level outputs of individual sensor systems *before* those are fused (more in Sect. 9.3.4). Such testing is needed for the authors’ proposed safety assessment strategy for the perception subsystem, where the corresponding reference data could be generated without additional sensors and without manual human effort.

If a majority of fusion input systems detects a certain road user while one input system does not, it can be argued that the majority of agreeing input systems provides the reference data for identifying a perception failure in the single disagreeing system. This idea is generalized to a so-called k-out-of-n vote of sensor systems. It could potentially scale to billions of kilometers of data collection using large vehicle fleets with feasible effort, but its practical implementation on real vehicle data seems to have not yet been described in any available literature.

8.2 Reference Data from Other Road Users

Other road users can be either equipped with specific test equipment for the estimation of their position and state, or they could communicate their states to the ego vehicle

as part of a potential vehicle-to-everything (V2X)-based regular operation in future.

8.2.1 RTK-GNSS-IMUs

Global navigation satellite systems (GNSS) with real-time kinematic correction (RTK) and an attached inertial measurement unit (IMU) have been popular for generating reference data of individual road users. Paired systems of those can be installed both in the ego vehicle and at other road users to determine relative positions, velocities, and orientation angles. RTK-GNSS can measure the absolute position up to a few centimeters, but face difficulties when the signal is affected by nearby structures such as buildings, trees, or a tunnel. IMUs work independently of the surroundings, but only provide relative and incremental positioning information. Data fusion algorithms can combine the strengths of both components to estimate a consolidated state of the respective road user.

Despite the size of the measurement hardware, vulnerable road users (VRUs) can also wear such hardware as a backpack [151], which, however, alters their appearance to other sensors. In general, only specifically equipped road users can provide reference data using RTK-GNSS-IMUs, which excludes the majority of public road users.

The accuracy of RTK corrections of GNSS signals in automotive applications has been analyzed in more detail in Ref. [152]. Another work has specifically analyzed the accuracy of GNSS systems that are used for relative positioning of local groups of vehicles [153]. For more information on GNSS in general, please refer to the textbook of Ref. [154].

8.2.2 Collaborative World Model Through V2X

If target vehicles have the capability to accurately localize themselves without dedicated testing hardware, they could communicate their state to a cloud-based or edge-based collective environment model [155]. The mentioned publication hypothesizes how this could enable a mutual verification of the environment perception of individual road users. The key difference to the previously described RTK-GNSS-IMU approaches is that Ref. [155] aims at the regular operation of future production vehicles rather than dedicated testing activities.

The targeted crowd-sourced reference data generation follows a similar idea like the previously described k-out-of-n-vote of high-level fusion inputs (Sect. 8.1.3). The difference here is that the inputs to the vote would originate from vehicles instead of from sensors. However, also this approach does not seem to have been demonstrated with real-world vehicles yet.

8.3 Reference Data from Non-Road Users

The analyzed literature on reference data generation with sensors external to the road users includes stationary infrastructure sensors, unmanned aerial vehicles (UAVs, or drones), and helicopters. Those data are often intended as either stand-alone naturalistic traffic data or as online enhancements of the AV perception. They can still serve as a reference for the SUT if their quality is superior to the SUT according to context-specific reference data requirements. A data quality superior to the SUT on the ground can be achieved by a more advantageous bird's eye perspective, by intentional overfitting of the perception algorithms to the sensor's location, and by human error-checking. Humans can detect errors in reference object lists easier when comparing them with a raw video from a bird's eye perspective than by comparing them to a potentially disturbed GNSS signal or to potentially incomplete SUT raw data.

8.3.1 Reference Data from Stationary Infrastructure Sensors

Stationary infrastructure sensors are sensors that are mounted on buildings, streetlamps, gantries, bridges, or other nearby structures. They are used in urban locations for example in the Ko-PER dataset [48], in the test area autonomous driving Baden-Württemberg [71], and at the AIM test site [156]. Infrastructure sensors at highways or motorways are used in Ref. [157] and at the Providentia sensor system [158], where the latter source assesses its performance using a helicopter-based system, which is covered in Sect. 8.3.3. The projects HDVMess and ACCorD aim at installing mobile and fixed infrastructure sensors at various sites [159] and evaluate their data generation with a UAV-based reference system [160].

8.3.2 Reference Data from UAVs

Camera-based reference data generation from a bird's eye perspective has been investigated at least since 2014 [47]. The necessary processing steps to track ground vehicles from UAVs were later described in more detail in Refs. [161] and [162]. UAVs are used as reference sensors for sensor modeling in Ref. [110] and for assessing infrastructure sensors in Ref. [160]. A qualitative assessment of UAV-based reference systems in contrast to other reference systems is given in Ref. [163]. Reference data from a UAV's perspective tend to suffer from less occlusions than from perspectives closer to the ground, however, new sorts of occlusions from treetops or bridges occur.

8.3.3 Reference Data from Helicopters

Additionally, helicopters have been used to record the ground traffic for verification and validation purposes of AD systems. The DLRAD dataset [164] was partially recorded using a helicopter that follows a target vehicle with active environment perception. A previous publication from the same project [165] describes how such a helicopter-based system can be used to validate the behavior of ADAS. Furthermore, a helicopter-based reference data generation is also described and used in Ref. [158] to assess a static infrastructure sensor system.

8.4 Uncertainty in Reference Data

Reference systems do not provide an absolute ground truth, but only have to come closer to it than the SUT does. This has been pointed out and quantified with respect to labeling uncertainty in Ref. [166]. Thus, the quality of the reference data (or labels) should be taken into account such that test results are not misleading [21].

For automatically generated reference data from external sensors, there might not be labeling uncertainty, but uncertainty in temporal synchronization of reference and SUT data, as well as uncertainty in their spatial alignment.

Moreover, any automatically generated reference data without human error-checking can suffer from unexpected inaccuracies in the reference measurement system. Those inaccuracies might be rare enough to not matter in small-scale R&D testing, but over large amounts of test scenarios or kilometers, as needed for safety assurance, the probability of encountering relevant reference system errors increases.

8.5 Choice of Reference Data Source

Different ways of generating reference data are used for different aspects of perception testing, where each way has its unique advantages. For example, human labels might be best for testing what a perception algorithm can infer from given raw data in the absence of additional data sources. Separate high-quality sensors on the ego vehicle could be the most useful choice for benchmarking individually inferior sensor systems. RTK-GNSS-IMUs might be best for testing the perception of small numbers of vehicles in the absence of high buildings. Data from an aerial perspective might be most suitable for testing the perception performance under static and dynamic occlusions in complex naturalistic traffic. Infrastructure test fields that spread over a large area can be best suited for investigations that require a large field-of-view. Finally, reference data from the k-out-of-n-vote might have the best potential for keeping the testing effort feasible for large amounts of kilometers or scenarios.

8.6 Summary of Reference Data

The most appropriate choice of reference data depends on the detailed objectives of the respective testing activity within the safety case [82]. In any case, the reference data's limitations should be considered.

9 Research Gaps and Challenges

Section 9.1 summarizes the particular conclusions of each testing axis in terms of the largest issues regarding the primary research question. Intersection topics between the axes are covered in Sect. 9.2. Finally, testing as reviewed in this paper is likely necessary, but not sufficient for safety assurance of AVs. Therefore, Sect. 9.3 places the reviewed topics into a broader context by drawing connections to related activities around perceptual aspects in safety assurance.

9.1 Open Issues per Testing Axis

In terms of metrics and test criteria (Sect. 6), current challenges are the incompleteness of criteria due to the open-world context and also the consideration of a potential safety impact in the metrics.

Open issues about perception-specific test scenarios (Sect. 7) are a common description language and format, a combination of scenario generation approaches for sufficient ODD coverage, and the difficulty of recording the identified test scenarios in the real world.

The most appropriate source of reference data (Sect. 8) depends on the investigated aspect of environment perception. A combination of multiple such investigations into a harmonized methodology targeting the safety proof appears to be not yet demonstrated in the literature.

The mentioned open issues from the literature might be the reason why as of now, safety standards (Sect. 4) leave many details around perception testing open.

9.2 Open Issues Between the Testing Axes

Solutions to the individual testing axes must be able to fit into a consolidated methodology that also considers intersection topics between the axes. For example, at the intersection of test criteria and scenarios, the criteria on the SUT's FOV might be relaxed if the ego vehicle drives slowly due to rain in the respective scenario. In between the axes scenarios and reference data, the fact that rain is present in the scenario might affect some ways of reference data generation more than others. The intersection between reference data and metrics/criteria includes the metrics' sensitivity to details in the reference data. For example, a metric that penalizes false negatives behind occlusions produces different results

depending on whether the reference data contain these occluded objects or not.

While an all-encompassing methodology that covers all of such details seems to be not yet demonstrated, approaches such as Refs. [30] or [49] already propose general ways of how the dependability of the perception subsystem could be managed and tested.

9.3 Further Safety Assurance Activities Regarding Perception

Testing as described in this review (see Table 1) is only one activity out of many within an all-encompassing safety argumentation. The following subsections intend to reveal this larger context by drawing connections to related activities in perception dependability. The approaches are structured according to the means for attaining dependability by Ref. [128], which have recently been adopted to uncertainties in AD [118]. They consist of prevention, removal, tolerance, and forecasting of faults or uncertainty.

9.3.1 Uncertainty Prevention

Formal verification of the perception subsystem could ideally prevent the occurrence of certain world model uncertainties. So far, the authors are not aware of formal verification methods that cover the entire perception subsystem. Also, the review paper Ref. [14] only mentions formal verification methods for the downstream driving function. However, formal analysis and verification of DNNs for perception tasks is an active topic of research [167, 168]. Moreover, the way in which uncertain world models affect formally verified downstream driving functions is actively being researched, for example in the context of RSS [27].

Another means of uncertainty prevention is the restriction to certain ODDs [118].

9.3.2 Uncertainty Removal

Perceptual uncertainties can be removed during development and during its use after the initial release. Uncertainty removal during development includes safety analysis methods such as FTA [118], or insightful developer testing with the goal of reducing found uncertainties afterward. For example, *white-box-testing of the perception subsystem* can reveal uncertainties related to the physically complex measurement principles (see Ref. [112]). Human-understandable insights into the DNNs for object detection are, however, difficult [21].

Testing on simulated raw data allows the perception algorithm to be tested without the need for physical sensor hardware (e.g., [169] for grid mapping). For testing object perception algorithms, a simulation environment such as

Ref. [170] could produce both the raw sensor data as an input to the algorithm under test and a ground truth object list as a reference. A hybrid approach is to test the perception algorithm on recorded real data with simulated injected faults [171].

After the initial release, *incremental safety assurance* through system updates is a necessary means to remove uncertainties because the ODD of the AV changes continuously [21]. Ideally, the safety argumentation should allow minor system updates without having to repeat major testing efforts [26]. Agile safety cases [172] could enable the safe deployment of such incremental updates through a DevOps-like workflow [74].

9.3.3 Uncertainty Tolerance

Runtime monitoring of the uncertainty of the produced world model can help mitigating the safety consequences of perception errors [173]. An RSS-compliant planner would become more cautious for uncertain world models due to a reduced set of guaranteed safe control actions [27] (Sect. 6.2.3). Similarly, in a service-oriented architecture, sensor data that is monitored as too uncertain or faulty could lead to a bypass of the usual information processing pipeline in favor of a safe behavior for degraded modes [174]. An overview of self-representation and monitoring of subsystems is given in Ref. [175]. Similarly, in the ISO 26262 context, it was proposed that the perception subsystem dynamically outputs one world model for each ASIL level such that the planner can dynamically take just that world model that it needs to comply with the ASIL requirement of the current driving situation [35, 52, 53]. In any of such applications, it is crucial to assure that the reported uncertainties are correctly calibrated (Sect. 6.3).

Furthermore, the mentioned *collaborative world model through V2X* (Sect. 8.2.2) can also serve as a means for uncertainty tolerance if participating cars can reduce their world model uncertainties through ways like a cloud-based fusion [176].

9.3.4 Uncertainty Forecasting

The failure rates in a fused world model must be extremely low, which means that the direct test effort for proving or forecasting it is extremely high.

One approach to overcome this so-called approval trap [177] is to test the more failure-prone world models of individual sensor systems *before* they are fused [33, 49]. The latter source investigates how this could be realized without separately generated reference data, allowing the tests to be potentially executed in shadow mode by production vehicles. In this approach (Fig. 7), the reference data take the form of a k-out-of-n-vote of object-level fusion inputs

Fig. 7 Perception testing for macroscopic safety assurance without separately generated reference data [32, 49], expressed in taxonomy (Fig. 1)

(Sect. 8.1.3). Such a mutual cross-referencing mechanism is also suggested in *Safety First for Automated Driving* [82].

The mentioned testing approach exploits the redundancy of fusion inputs in the following way [31]. If a mean time between failures of 10^4 h can be demonstrated for two statistically independent inputs to a data fusion system, the fusion output has a mean time between failures of about 10^8 h. Such a redundancy exploitation has also been stated in the RSS papers [33, 66] and in the literature on perception subsystem architectures [39, 178] and perception requirements [35]. The actual number of necessary test kilometers is then computed in a statistically sound way based on how statistically dependent the inputs to the object-level fusion system actually are [50, 54].

Whether or not perceptual uncertainties lead to safety-critical failures on the vehicle level might require a *sensitivity analysis of the driving function*. Reachability analysis [43, 123] could be used to analyze in detail whether an uncertain world model can lead to unsafe future states, as outlined in the context of the *incurred severity* metric [28]. An impact analysis of perceptual errors on the downstream driving function has also already been described in Ref. [65].

Furthermore, uncertainty forecasting for the release argumentation could make use of the tests mentioned under *Uncertainty Removal* (Sect. 9.3.2) if they are validated and executed on separate test datasets.

10 Conclusion

This paper analyzes literature from multiple neighboring fields, attempting to provide an overview of safety-relevant testing activities for the environment perception of AVs. These neighboring fields included, but are not limited to, AV safety assurance, perception algorithm development, and safeguarding of artificial intelligence. The reviewed information is structured according to the focus topics test criteria and metrics, test scenarios, and reference data. A

combined literature search process that consists of undocumented search, keyword-based search, and snowballing search required large efforts and appears to be suitable for covering the relevant literature.

The analyzed test methods from the available literature seem to be not yet capable of demonstrating the dependability of the perception subsystem for series-production of Level 4 vehicles. Research gaps concerning the individual focus topics, as well as their intersection topics, are summarized. To overcome this issue, the provided overview can serve as a common basis for a further harmonization of primary literature contributions into all-encompassing novel testing methodologies.

Acknowledgements The research leading to these results is funded by the Federal Ministry for Economic Affairs and Energy within the project “VVM - Verification and Validation Methods for Automated Vehicles Level 4 and 5.” The authors would like to thank the consortium for the successful cooperation.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflicts of interest Michael Hoss, Maike Scholtes, and Lutz Eckstein declare no conflict of interest.

Despite all efforts, we might have unintentionally overlooked individual literature contributions to the research question. In that case, we apologize to the corresponding authors and are welcoming such feedback from the community.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Taxonomy and definitions for terms related to driving automation (J3016) (2018)
2. Dietmayer, K.C.J., Reuter, S., Nuss, D.: Representation of fused environment data, pp. 1–30. Springer International Publishing, Cham (2014). https://doi.org/10.1007/978-3-319-09840-1_25-1
3. Dietmayer, K.: Predicting of machine perception for automated driving. In: *Autonomous Driving*, pp. 407–424. Springer, New York (2016)
4. Stellet, J.E., Zofka, M.R., Schumacher, J., Schamm, T., Niewels, F., Zöllner, J.M.: Testing of advanced driver assistance

- towards automated driving: A survey and taxonomy on existing approaches and open questions. In: IEEE 18th International Conference on Intelligent Transportation Systems, pp. 1455–1462. IEEE (2015)
5. Amersbach, C., Winner, H.: Functional decomposition: An approach to reduce the approval effort for highly automated driving. In: 8. Tagung Fahrerassistenz. Technical University of Munich, Munich (2017)
 6. Glossary of research project VVM - Verification and Validation Methods for Automated Vehicles Level 4 and 5. Internal (2020)
 7. Leveson, N.G.: Engineering a safer world: Systems thinking applied to safety. The MIT Press, United States (2016)
 8. Technical Committee ISO/TC 22, Road vehicles, Subcommittee SC 32, Electrical and electronic components and general system aspects: ISO/PAS 21448:2019 road vehicles — safety of the intended functionality (2019). <https://www.iso.org/standard/70939.html>
 9. Technical Committee ISO/TC 22, Road vehicles Subcommittee, SC 32, Electrical and electronic components and general system aspects: ISO 26262 road vehicles - functional safety (2018)
 10. Working Group 1 of the Joint Committee for Guides in Metrology: Evaluation of measurement data — guide to the expression of uncertainty in measurement (2008). GUM 1995 with minor corrections
 11. Ulbrich, S., Menzel, T., Reschka, A., Schuldt, F., Maurer, M.: Defining and substantiating the terms scene, situation, and scenario for automated driving. In: IEEE 18th International conference on Intelligent Transportation Systems. IEEE (2015). <https://doi.org/10.1109/itsc.2015.164>
 12. Junietz, P.M.: Microscopic and macroscopic risk metrics for the safety validation of automated driving. Ph.D. Dissertation, TU Darmstadt, Darmstadt (2019). <http://tuprints.ulb-tu-darmstadt.de/9282/>
 13. Stellet, J.E., Woehle, M., Brade, T., Poddey, A., Branz, W.: Validation of automated driving—a structured analysis and survey of approaches. In: 13. Uni-DAS e.V. Workshop Fahrerassistenz und automatisiertes Fahren (2020)
 14. Riedmaier, S., Ponn, T., Ludwig, D., Schick, B., Diermeyer, F.: Survey on scenario-based safety assessment of automated vehicles. IEEE Access **8**, 87456–87477 (2020). <https://doi.org/10.1109/access.2020.2993730>
 15. Nalic, D., Mihalj, T., Baeumler, M., Lehmann, M., Eichberger, A., Bernsteiner, S.: Scenario based testing of automated driving systems: A literature survey. In: FISITA Web Contress (2020)
 16. PEGASUS method: An overview. Online (2019). <https://www.pegasusprojekt.de/files/tmpl/Pegasus-Abschlussveranstaltung/PEGASUS-Gesamtmethode.pdf>
 17. Borg, M., Englund, C., Wnuk, K., Duran, B., Levandowski, C., Gao, S., Tan, Y., Kaijser, H., Lönn, H., Törnqvist, J.: Safely entering the deep: A review of verification and validation for machine learning and a challenge elicitation in the automotive industry. [arXiv:1812.05389](https://arxiv.org/abs/1812.05389) (2018)
 18. Marti, E., de Miguel, M.A., Garcia, F., Perez, J.: A review of sensor technologies for perception in automated driving. IEEE Intell. Trans. Syst. Mag. **11**(4), 94–108 (2019)
 19. Rosique, F., Navarro, P.J., Fernández, C., Padilla, A.: A systematic review of perception system and simulators for autonomous vehicles research. Sensors **19**(3), 648 (2019). <https://doi.org/10.3390/s19030648>
 20. Zhu, H., Yuen, K.V., Mihaylova, L., Leung, H.: Overview of environment perception for intelligent vehicles. IEEE Trans. Intell. Trans. Syst. **18**(10), 2584–2601 (2017). <https://doi.org/10.1109/tits.2017.2658662>
 21. Willers, O., Sudholt, S., Raafatnia, S., Abrecht, S.: Safety concerns and mitigation approaches regarding the use of deep learning in safety-critical perception tasks. In: Casimiro, A., Ortmeier, F., Schoitsch, E., Bitsch, F., Ferreira, P. (eds.) Computer Safety, Reliability, and Security. SAFECOMP Workshops, pp. 336–350. Springer International Publishing, Cham (2020)
 22. Burton, S., Gauerhof, L., Heinzemann, C.: Making the case for safety of machine learning in highly automated driving. In: Tonetta, S., Schoitsch, E., Bitsch, F. (eds.) Computer Safety, Reliability, and Security, pp. 5–16. Springer International Publishing, Cham (2017)
 23. Kitchenham, B., Charters, S.: Guidelines for performing systematic literature reviews in software engineering. Tech. Rep. v2.3, Evidence-Based Software Engineering (EBSE) Project (2007)
 24. Wohlin, C.: Guidelines for snowballing in systematic literature studies and a replication in software engineering. In: 18th Int. Conf. on Evaluation and Assessment in Software Engineering, pp. 1–10 (2014)
 25. Schalling, F., Ljungberg, S., Mohan, N.: Benchmarking lidar sensors for development and evaluation of automotive perception. In: 4th Int. Conf. and Workshops on Recent Advances and Innovations in Engineering (ICRAIE), 1–6 (2019)
 26. Koopman, P., Wagner, M.: Toward a framework for highly automated vehicle safety validation. In: SAE Technical Paper Series. SAE International (2018). <https://doi.org/10.4271/2018-01-1071>
 27. Salay, R., Czarnecki, K., Alvarez, I., Elli, M.S., Sedwards, S., Weast, J.: PURSS: Towards perceptual uncertainty aware responsibility sensitive safety with ML. In: Workshop on Artificial Intelligence Safety (SafeAI). New York, USA (2020)
 28. Salay, R., Angus, M., Czarnecki, K.: A safety analysis method for perceptual components in automated driving. IEEE 30th Int. Symposium on Software Reliability Engineering (ISSRE) 24–34 (2019)
 29. Salay, R., Czarnecki, K.: Improving ML safety with partial specifications. In: Romanovsky, A., Troubitsyna, E., Gashi, I., Schoitsch, E., Bitsch, F. (eds.) Computer Safety, Reliability, and Security, pp. 288–300. Springer International Publishing, Cham (2019)
 30. Czarnecki, K., Salay, R.: Towards a framework to manage perceptual uncertainty for safe automated driving. Lecture Notes in Computer Science pp. 439–445 (2018). https://doi.org/10.1007/978-3-319-99229-7_37
 31. Berk, M., Schubert, O., Kroll, H.M., Buschardt, B., Straub, D.: Exploiting redundancy for reliability analysis of sensor perception in automated driving vehicles. IEEE Trans. Intell. Trans. Syst. **21**(12), 5073–5085 (2019)
 32. Berk, M., Schubert, O., Kroll, H.M., Buschardt, B., Straub, D.: Reliability assessment of safety-critical sensor information: Does one need a reference truth? IEEE Trans. Reliab. **68**(4), 1227–1241 (2019). <https://doi.org/10.1109/tr.2019.2923735>
 33. Shalev-Shwartz, S., Shammah, S., Shashua, A.: On a formal model of safe and scalable self-driving cars. [arXiv:1708.06374](https://arxiv.org/abs/1708.06374) (2017). V6
 34. Brahmi, M., Siedersberger, K.H., Siegel, A., Maurer, M.: Reference systems for environmental perception: Requirements, validation and metric-based evaluation. In: 6. Tagung Fahrerassistenzsysteme (2013)
 35. Cassel, A., Bergenhem, C., Christensen, O.M., Heyn, H.M., Leadarsson-Olsson, S., Majdandzic, M., Sun, P., Thorsén, A., Trygvesson, J.: On perception safety requirements and multi sensor systems for automated driving systems. In: SAE Technical Paper. SAE International (2020). <https://doi.org/10.4271/2020-01-0101>
 36. Berk, M.: Safety assessment of environment perception in automated driving vehicles. Dissertation, Technical University Munich, Munich (2019)

37. Salay, R., Queiroz, R., Czarnecki, K.: An analysis of ISO 26262: Using machine learning safely in automotive software. [arXiv:1709.02435](#) (2017)
38. Taş, Ö.Ş., Kuhnt, F., Zöllner, J.M., Stiller, C.: Functional system architectures towards fully automated driving. In: IEEE Intelligent Vehicles Symposium (IV), pp. 304–309. IEEE (2016)
39. Kaprocki, N., Velicki, G., Teslic, N., Krunić, M.: Multiunit automotive perception framework: Synergy between AI and deterministic processing. In: IEEE 9th Int. Conf. on Consumer Electronics (ICCE-Berlin). IEEE (2019). <https://doi.org/10.1109/icce-berlin47944.2019.8966168>
40. Martin, H., Winkler, B., Grubmüller, S., Watzenig, D.: Identification of performance limitations of sensing technologies for automated driving. In: IEEE Int. Conf. on Connected Vehicles and Expo (ICCVE), pp. 1–6 (2019)
41. Bai, J., Zhan, Y., Cao, P., Huang, L., Xu, Y., Bi, X.: Application oriented identification of external influence factors on sensing for validation of automated driving systems. In: IEEE MTT-S International Wireless Symposium (IWS), pp. 1–3 (2019)
42. Ulbrich, S., Reschka, A., Rieken, J., Ernst, S., Bagschik, G., Dierkes, F., Nolte, M., Maurer, M.: Towards a functional system architecture for automated vehicles. [arXiv:1703.08557](#) (2017)
43. Althoff, M.: Reachability analysis and its application to the safety assessment of autonomous cars. Ph.D. Dissertation, Technical University Munich (2010)
44. Bagschik, G., Menzel, T., Maurer, M.: Ontology based scene creation for the development of automated vehicles. In: IEEE Intelligent Vehicles Symposium (IV). IEEE (2018). <https://doi.org/10.1109/ivs.2018.8500632>
45. Burton, S., Gauerhof, L., Sathy, B.B., Habli, I., Hawkins, R.: Confidence arguments for evidence of performance in machine learning for highly automated driving functions. In: Romanovsky, A., Troubitsyna, E., Gashi, I., Schoitsch, E., Bitsch, F. (eds.) Computer Safety, Reliability, and Security, pp. 365–377. Springer International Publishing, Cham (2019)
46. Koopman, P., Fratrik, F.: How many operational design domains, objects, and events? In: SafeAI@AAAI (2019)
47. Kubertschak, T., Wittenzellner, M., Maehlich, M.: Kamera-basiertes Referenzsystem für Fahrerassistenzsysteme. In: Forum Bildverarbeitung, vol. 82, p. 261. KIT Scientific Publishing (2014)
48. Strigel, E., Meissner, D., Seeliger, F., Wilking, B., Dietmayer, K.: The Ko-PER intersection laserscanner and video dataset. In: 17th International IEEE Conf. on Intelligent Transportation Systems (ITSC). IEEE (2014). <https://doi.org/10.1109/itsc.2014.6957976>
49. Berk, M., Schubert, O., Kroll, H.M., Buschardt, B., Straub, D.: Assessing the safety of environment perception in automated driving vehicles. SAE Int. J. Trans. Saf. 8(1), 49–74 (2020). <https://doi.org/10.4271/09-08-01-0004>
50. Berk, M., Kroll, H.M., Schubert, O., Buschardt, B., Straub, D.: Bayesian test design for reliability assessments of safety-relevant environment sensors considering dependent failures. In: SAE Technical Paper Series. SAE Technical Paper 2017-01-0050, 2017. <https://doi.org/10.4271/2017-01-0050>
51. Berk, M., Dura, M., Rivero, J.V., Schubert, O., Kroll, H.M., Buschardt, B., Straub, D.: A stochastic physical simulation framework to quantify the effect of rainfall on automotive lidar. In: SAE Technical Paper Series. SAE Int. J. Adv. & Curr. Prac. in Mobility 1(2), 531–538 (2019). <https://doi.org/10.4271/2019-01-0134>
52. Johansson, R., Nilsson, J.: The need for an environment perception block to address all ASIL levels simultaneously. In: IEEE Intelligent Vehicles Symposium (IV). IEEE (2016). <https://doi.org/10.1109/ivs.2016.7535354>
53. Johansson, R., Alissa, S., Bengtsson, S., Bergenhem, C., Bridal, O., Cassel, A., Chen, D.J., Gassilewski, M., Nilsson, J., Sandberg, A., Ursing, S., Warg, F., Werneman, A.: A strategy for assessing safe use of sensors in autonomous road vehicles. In: Tonetta, S., Schoitsch, E., Bitsch, F. (eds.) Computer Safety, Reliability, and Security, pp. 149–161. Springer International Publishing, Cham (2017)
54. Bock, F., Siegl, S., German, R.: Mathematical test effort estimation for dependability assessment of sensor-based driver assistance systems. In: 42th Euromicro Conf. on Software Engineering and Advanced Applications (SEAA). IEEE (2016). <https://doi.org/10.1109/seaa.2016.49>
55. Rivero, J.R.V., Tahiraj, I., Schubert, O., Glassl, C., Buschardt, B., Berk, M., Chen, J.: Characterization and simulation of the effect of road dirt on the performance of a laser scanner. In: IEEE 20th Int. Conf. on Intelligent Transportation Systems (ITSC). IEEE (2017). <https://doi.org/10.1109/itsc.2017.8317784>
56. Rasshofer, R.H., Spies, M., Spies, H.: Influences of weather phenomena on automotive laser radar systems. Adv. Radio Sci.: ARS 9, 49–60 (2011)
57. Fawcett, T.: An introduction to ROC analysis. Pat. Recognit. Lett. 27(8), 861–874 (2006). <https://doi.org/10.1016/j.patrec.2005.10.010>
58. Cao, P., Huang, L.: Application oriented testcase generation for validation of environment perception sensor in automated driving systems. SAE Technical Paper 2018-01-1614, 2018, <https://doi.org/10.4271/2018-01-1614>.
59. Czarnecki, K.: Operational world model ontology for automated driving systems - Part 2: Road users, animals, other obstacles, and environmental conditions. Tech. rep., Waterloo Intelligent Systems Engineering Lab (WISE) (2018). <https://doi.org/10.13140/RG.2.2.11327.00165>
60. Meyer, B.: Applying design by contract. Computer 25(10), 40–51 (1992). <https://doi.org/10.1109/2.161279>
61. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. [arXiv:1706.04599](#) (2017)
62. Salay, R., Czarnecki, K.: Using machine learning safely in automotive software: An assessment and adaption of software process requirements in ISO 26262. [arXiv:1808.01614](#) (2018)
63. Schwaiger, A., Sinhamahapatra, P., Gansloser, J., Roscher, K.: Is uncertainty quantification in deep learning sufficient for out-of-distribution detection? In: Workshop on Artificial Intelligence Safety (AISafety). Yokohama, Japan (2020)
64. Arnez, F., Espinoza, H., Radermacher, A., Terrier, F.: A comparison of uncertainty estimation approaches in deep learning components for autonomous vehicle applications. [arXiv:2006.15172](#) (2020)
65. Piazzoni, A., Cherian, J., Slavik, M., Dauwels, J.: Modeling sensing and perception errors towards robust decision making in autonomous vehicles. In: 29th International Joint Conf. on Artificial Intelligence (IJCAI), 3494–3500 (2020). <https://doi.org/10.24963/ijcai.2020/483>
66. Weast, J.: Sensors, safety models and a system-level approach to safe and scalable automated vehicles. [arXiv:2009.03301](#) (2020)
67. Klamann, B., Lippert, M., Amersbach, C., Winner, H.: Defining pass-/fail-criteria for particular tests of automated driving functions. In: IEEE Intelligent Transportation Systems Conf. (ITSC). IEEE (2019). <https://doi.org/10.1109/itsc.2019.8917483>
68. Aravantinos, V., Schlicht, P.: Making the relationship between uncertainty estimation and safety less uncertain. In: Design, Automation & Test in Europe Conf. & Exhibition (DATE). IEEE (2020). <https://doi.org/10.23919/date48585.2020.9116541>
69. Breitenstein, J., Termöhlen, J.A., Lipinski, D., Fingscheidt, T.: Systematization of corner cases for visual perception in automated driving. 2020 IEEE Intelligent Vehicles Symposium (IV),

- 2020, pp. 1257–1264, <https://doi.org/10.1109/IV47402.2020.9304789>.
70. Henne, M., Schwaiger, A., Roscher, K., Weiss, G.: Benchmarking uncertainty estimation methods for deep learning with safety-related metrics. In: *SafeAI@AAAI*, 83–90 (2020)
 71. Fleck, T., Daaboul, K., Weber, M., Schörner, P., Wehmer, M., Doll, J., Orf, S., Sußmann, N., Hubschneider, C., Zofka, M.R., Kuhnt, F., Kohlhaas, R., Baumgart, I., Zöllner, R., Zöllner, J.M.: Towards large scale urban traffic reference data: Smart infrastructure in the test area autonomous driving Baden-Württemberg. In: Strand, M., Dillmann, R., Menegatti, E., Ghidoni, S. (eds.) *Intelligent autonomous systems*, 15th edn., pp. 964–982. Springer International Publishing, Cham (2019)
 72. Spanfelner, B., Richter, D., Ebel, S., Wilhelm, U., Patz, C.: Challenges in applying the ISO 26262 for driver assistance systems. In: *Tagung Fahrerassistenz*. Munich (2012)
 73. Rahimi, M., Guo, J.L., Kokaly, S., Chechik, M.: Toward requirements specification for machine-learned components. In: *IEEE 27th International Requirements Engineering Conf. Workshops (REW)*. IEEE (2019). <https://doi.org/10.1109/rew.2019.00049>
 74. Czarnecki, K.: Software engineering for automated vehicles: Addressing the needs of cars that run on software and data. In: *IEEE/ACM 41st Int. Conf. on Software Engineering: Companion Proc. (ICSE-Companion)*. IEEE (2019). <https://doi.org/10.1109/icse-companion.2019.00024>
 75. Feng, D., Rosenbaum, L., Dietmayer, K.: Towards safe autonomous driving: Capture uncertainty in the deep neural network for lidar 3D vehicle detection. In: *21st Int. Conf. on Intelligent Transportation Systems (ITSC)*. IEEE (2018). <https://doi.org/10.1109/itsc.2018.8569814>
 76. Feng, D., Rosenbaum, L., Glaeser, C., Timm, F., Dietmayer, K.: Can we trust you? On calibration of a probabilistic object detector for autonomous driving. *arXiv:1909.12358* (2019)
 77. *Automated driving systems 2.0: A vision for safety* (2017). USA DOT NHTSA
 78. Thorn, E., Kimmel, S.C., Chaka, M.: A framework for automated driving system testable cases and scenarios. Tech. rep, USA DOT NHTSA (2018)
 79. UL4600 - standard for safety for the evaluation of autonomous products (2019)
 80. Koopman, P., Ferrell, U., Fratrick, F., Wagner, M.: A Saf. Stand. Approach Fully Auto. Veh. (2019). https://doi.org/10.1007/978-3-030-26250-1_26
 81. Uniform provisions concerning the approval of vehicles with regard to automated lane keeping systems (2020). ECE/TRANS/WP.29/2020/81
 82. Wood, M., Robbel, P., Maass, M., Tebbens, R., Meijs, M., Harb, M., Schlicht, P., et al.: Safety first for automated driving. Online (2019). <https://newsroom.intel.com/wp-content/uploads/sites/11/2019/07/Intel-Safety-First-for-Automated-Driving.pdf>
 83. ISO/TR 4804 — road vehicles — safety and cybersecurity for automated driving systems — design, verification and validation (2020)
 84. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? The KITTI vision benchmark suite. In: *Conference on Computer Vision and Pattern Recognition (CVPR)* (2012). http://www.cvlibs.net/datasets/kitti/eval_tracking.php
 85. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuScenes: A multimodal dataset for autonomous driving. *arXiv:1903.11027* (2019)
 86. Sun, P., Kretzschmar, H., Dotiwalla, X., Chouard, A., Patnaik, V., Tsui, P., Guo, J., Zhou, Y., Chai, Y., Caine, B., Vasudevan, V., Han, W., Ngiam, J., Zhao, H., Timofeev, A., Ettinger, S., Krivokon, M., Gao, A., Joshi, A., Zhang, Y., Shlens, J., Chen, Z., Anguelov, D.: Scalability in perception for autonomous driving: Waymo open dataset. *arXiv:1912.04838* (2019)
 87. Chang, M.F., Lambert, J., Sangkloy, P., Singh, J., Bak, S., Hartnett, A.T., Wang, D., Carr, P., Lucey, S., Ramanan, D., Hays, J.: Argoverse: 3D tracking and forecasting with rich maps. *IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR)* 8740–8749 (2019)
 88. Stampfle, M., Holz, D., Becker, J.C.: Performance evaluation of automotive sensor data fusion. In: *IEEE Intelligent Transportation Systems*, 50–55. IEEE (2005)
 89. Kuhn, H.W.: The hungarian method for the assignment problem. *Naval Res. Log. Q.* **2**(1–2), 83–97 (1955). <https://doi.org/10.1002/nav.3800020109>
 90. Bertsekas, D.P., Castanon, D.A.: The auction algorithm for the transportation problem. *Annals Op. Res.* **20**(1), 67–96 (1989). <https://doi.org/10.1007/bf02216923>
 91. Bernardin, K., Stiefelhagen, R.: Evaluating multiple object tracking performance: the clear MOT metrics. *J. Image Video Process.* **2008**, 1 (2008)
 92. Chen, Z., Heckman, C., Julier, S., Ahmed, N.: Weak in the NEES?: Auto-tuning kalman filters with bayesian optimization. In: *21st International Conference on Information Fusion (FUSION)* (2018). <https://doi.org/10.23919/ICIF.2018.8454982>
 93. Schuhmacher, D., Vo, B.T., Vo, B.N.: A consistent metric for performance evaluation of multi-object filters. *IEEE Trans. Signal Process.* **56**(8), 3447–3457 (2008)
 94. Ristic, B., Vo, B.N., Clark, D.: Performance evaluation of multi-target tracking using the OSPA metric. In: *13th International conference on Information Fusion*, 1–7. IEEE (2010)
 95. He, X., Tharmarasa, R., Kirubarajan, T., Thayaparan, T.: A track quality based metric for evaluating performance of multitarget filters. *IEEE Trans. Aerosp. Electron. Syst.* **49**(1), 610–616 (2013)
 96. Vu, T., Evans, R.: A new performance metric for multiple target tracking based on optimal subpattern assignment. In: *17th International conference on Information Fusion (FUSION)*, 1–8. IEEE (2014)
 97. Beard, M., Vo, B.T., Vo, B.N.: OSPA (2): Using the OSPA metric to evaluate multi-target tracking performance. In: *International Conference on Control, Automation and Information Sciences (ICCAIS)*, 86–91. IEEE (2017)
 98. Rahmathullah, A.S., García-Fernández, Á.F., Svensson, L.: Generalized optimal sub-pattern assignment metric. In: *20th International Conference on Information Fusion (Fusion)*, 1–8. IEEE (2017)
 99. Shi, X., Yang, F., Tong, F., Lian, H.: A comprehensive performance metric for evaluation of multi-target tracking algorithms. In: *3rd International Conference on Information Management (ICIM)*, 373–377. IEEE (2017)
 100. Reuter, S., Vo, B., Vo, B., Dietmayer, K.: The labeled multi-bernoulli filter. *IEEE Trans. Signal Process.* **62**(12), 3246–3260 (2014). <https://doi.org/10.1109/TSP.2014.2323064>
 101. Granstrom, K., Baum, M., Reuter, S.: Extended object tracking: Introduction, overview and applications. *arXiv:1604.00970* (2016)
 102. Li, Y., Huang, C., Nevatia, R.: Learning to associate: Hybrid-boosted multi-target tracker for crowded scene. In: *IEEE Conference on Computer Vision and Pattern Recognition*, 2953–2960. IEEE (2009)
 103. Luiten, J., Osep, A., Dendorfer, P., Torr, P., Geiger, A., Leal-Taixe, L., Leibe, B.: HOTA: A higher order metric for evaluating multi-object tracking. *Int. J. Comput. V.* (2020). <https://doi.org/10.1007/s11263-020-01375-2>

104. Philion, J., Kar, A., Fidler, S.: Learning to evaluate perception models using planner-centric metrics. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (2020). <https://doi.org/10.1109/cvpr42600.2020.01407>
105. Guo, Y., Caesar, H., Beijbom, O., Philion, J., Fidler, S.: The efficacy of neural planning metrics: A meta-analysis of pkl on nusenes. [arXiv:2010.09350](https://arxiv.org/abs/2010.09350) (2020)
106. Florbäck, J., Tornberg, L., Mohammadiha, N.: Offline object matching and evaluation process for verification of autonomous driving. In: IEEE 19th International Conference on Intelligent Transportation Systems (ITSC), 107–112 (2016). <https://doi.org/10.1109/ITSC.2016.7795539>
107. Sondell, D., Svensson, K.: Matching vehicle sensors to reference sensors using machine learning methods. Master's Thesis, Chalmers University of Technology (2018)
108. Hirsenkorn, N., Hanke, T., Rauch, A., Dehlink, B., Rasshofer, R., Biebl, E.: A non-parametric approach for modeling sensor behavior. In: 16th International Radar Symposium (IRS), 131–136 (2015). <https://doi.org/10.1109/IRS.2015.7226346>
109. Hanke, T., Hirsenkorn, N., Dehlink, B., Rauch, A., Rasshofer, R., Biebl, E.: Classification of sensor errors for the statistical simulation of environmental perception in automated driving systems. In: IEEE 19th International Conference of Intelligent Transportation Systems (ITSC), 643–648. IEEE (2016)
110. Krajewski, R., Hoss, M., Meister, A., Thomsen, F., Bock, J., Eckstein, L.: Using drones as reference sensors for neural-networks-based modeling of automotive perception errors. In: IEEE Intelligent Vehicles Symposium (IV) (2020)
111. Zec, E.L., Mohammadiha, N., Schliep, A.: Statistical sensor modelling for autonomous driving using autoregressive input-output HMMs. In: 21st International Conference on Intelligent Transportation Systems (ITSC). IEEE (2018). <https://doi.org/10.1109/itsc.2018.8569592>
112. Holder, M.F., Linnhoff, C., Rosenberger, P., Popp, C., Winner, H.: Modeling and simulation of radar sensor artifacts for virtual testing of autonomous driving. In: 9. Tagung Automatisiertes Fahren (2019)
113. Schaermann, A., Rauch, A., Hirsenkorn, N., Hanke, T., Rasshofer, R., Biebl, E.: Validation of vehicle environment sensor models. In: IEEE Intelligent Vehicles Symposium (IV), 405–411. IEEE (2017)
114. Rosenberger, P., Wendler, J.T., Holder, M.F., Linnhoff, C., Berghöfer, M., Winner, H., Maurer, M.: Towards a generally accepted validation methodology for sensor models - challenges, metrics, and first results. In: Graz Symposium Virtual Vehicle (GSVF) (2019)
115. Holder, M.F., Thielmann, J.R., Rosenberger, P., Linnhoff, C., Winner, H.: How to evaluate synthetic radar data? Lessons learned from finding driveable space in virtual environments. In: 13. Uni-DAS e.V. Workshop Fahrerassistenz und automatisiertes Fahren (2020)
116. Philipp, R., Schuldt, F., Howar, F.: Functional decomposition of automated driving systems for the classification and evaluation of perceptual threats. In: 13. Uni-DAS e.V. Workshop Fahrerassistenz und automatisiertes Fahren (2020)
117. Poddey, A., Brade, T., Stellet, J.E., Branz, W.: On the validation of complex systems operating in open contexts. [arXiv:1902.10517](https://arxiv.org/abs/1902.10517) (2019)
118. Gansch, R., Adey, A.: System theoretic view on uncertainties. In: 2020 Design, Automation & Test in Europe Conference & Exhibition (DATE). IEEE (2020). <https://doi.org/10.23919/date48585.2020.9116472>
119. Hu, B.C., Salay, R., Czarnecki, K., Rahimi, M., Selim, G., Chechik, M.: Towards requirements specification for machine-learned perception based on human performance. In: IEEE 7th International Workshop on Artificial Intelligence for Requirements Engineering (AIRE). IEEE (2020). <https://doi.org/10.1109/aire51212.2020.00014>
120. Dokhanchi, A., Amor, H.B., Deshmukh, J.V., Fainekos, G.: Evaluating perception systems for autonomous vehicles using quality temporal logic. In: Colombo, C., Leucker, M. (eds.) Runtime Verification, pp. 409–416. Springer International Publishing, Cham (2018)
121. Balakrishnan, A., Puranic, A.G., Qin, X., Dokhanchi, A., Deshmukh, J.V., Amor, H.B., Fainekos, G.: Specifying and evaluating quality metrics for vision-based perception systems. In: 2019 Design, Automation & Test in Europe Conference & Exhibition (DATE). IEEE (2019). <https://doi.org/10.23919/date.2019.8715114>
122. Schönmeyer, V., Winner, H., Verhaeg, G., Tronci, F., Padilla, G.A.G.: Fault tree-based derivation of safety requirements for automated driving on the example of cooperative valet parking (2019)
123. Pek, C., Manzinger, S., Koschi, M., Althoff, M.: Using online verification to prevent autonomous vehicles from causing accidents. *Nat. Mach. Intell.* 2(9), 518–528 (2020). <https://doi.org/10.1038/s42256-020-0225-y>
124. Naeini, M.P., Cooper, G.F., Hauskrecht, M.: Obtaining well calibrated probabilities using bayesian binning. In: AAAI Conference on Artificial Intelligence, vol. 2015, 2901–2907. NIH Public Access (2015)
125. Kuleshov, V., Fenner, N., Ermon, S.: Accurate uncertainties for deep learning using calibrated regression. [arXiv:1807.00263](https://arxiv.org/abs/1807.00263) (2018)
126. Gustafsson, F.K., Danelljan, M., Schon, T.B.: Evaluating scalable bayesian deep learning methods for robust computer vision. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2020)
127. Junietz, P., Bonakdar, F., Klamann, B., Winner, H.: Criticality metric for the safety validation of automated driving using model predictive trajectory optimization. In: 21st International conference on Intelligent Transportation Systems (ITSC). IEEE (2018). <https://doi.org/10.1109/itsc.2018.8569326>
128. Avizienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. *IEEE Trans. Depend. Secure Comput.* 1(1), 11–33 (2004). <https://doi.org/10.1109/tdsc.2004.2>
129. Rocklage, E., Kraft, H., Karatas, A., Seewig, J.: Automated scenario generation for regression testing of autonomous vehicles. In: IEEE 20th International conference on Intelligent Transportation Systems (ITSC). IEEE (2017). <https://doi.org/10.1109/itsc.2017.8317919>
130. Menzel, T., Bagschik, G., Maurer, M.: Scenarios for development, test and validation of automated vehicles. In: IEEE Intelligent Vehicles Symposium (IV). IEEE (2018). <https://doi.org/10.1109/ivs.2018.8500406>
131. Neurohr, C., Westhofen, L., Butz, M., Bollmann, M., Eberle, U., Galbas, R.: Criticality analysis for the verification and validation of automated vehicles. *IEEE Access* 9, 18016–18041 (2021). <https://doi.org/10.1109/access.2021.3053159>
132. Khastgir, S.: Operational design domain (odd) taxonomy for an automated driving system (ads) — specification (2020)
133. ASAM OpenODD (2020). <https://www.asam.net/project-detail/asam-openodd/>. Accessed on 09/12/2020
134. ASAM OpenXOntology (2020). <https://www.asam.net/project-detail/asam-openxontology/>. Accessed on 09/12/2020
135. Schuldt, F.: Ein Beitrag für den methodischen Test von automatisierten Fahrfunktionen mit Hilfe von virtuellen Umgebungen. Ph.D. Dissertation, TU Braunschweig (2017). <https://doi.org/10.23919/date48585.2020.9116472>

- 10.24355/dbbs.084-201704241210. https://publikationsserver.tu-braunschweig.de/receive/dbbs_mods_00064747
136. Bock, J., Krajewski, R., Eckstein, L., Klimke, J., Sauerbier, J., Zlocki, A.: Data basis for scenario-based validation of had on highways. In: 27th Aachen Colloquium Automobile and Engine Technology (2018)
 137. Scholtes, M., Westhofen, L., Turner, L.R., Lotto, K., Schuldes, M., Weber, H., Wagener, N., Neurohr, C., Bollmann, M.H., Körtke, F., Hiller, J., Hoss, M., Bock, J., Eckstein, L.: 6-layer model for a structured description and categorization of urban traffic and environment. *IEEE Access* **9**, 59131–59147 (2021). <https://doi.org/10.1109/ACCESS.2021.3072739>
 138. van Driesten, C., Schaller, T.: Overall approach to standardize AD sensor interfaces: Simulation and real vehicle. In: Bertram, T. (ed.) *Fahrerassistenzsysteme 2018*, pp. 47–55. Springer Fachmedien Wiesbaden, Wiesbaden (2019)
 139. ASAM OpenSCENARIO (2021). <https://www.asam.net/standards/detail/openscenario/>. Accessed on 09/02/2021
 140. Menzel, T., Bagschik, G., Isensee, L., Schomburg, A., Maurer, M.: From functional to logical scenarios: Detailing a keyword-based scenario description for execution in a simulation environment. In: *IEEE Intelligent Vehicles Symposium (IV)*. IEEE (2019). <https://doi.org/10.1109/ivs.2019.8814099>
 141. Neurohr, C., Westhofen, L., Henning, T., de Graaff, T., Mohlmann, E., Bode, E.: Fundamental considerations around scenario-based testing for automated driving. In: *IEEE Intelligent Vehicles Symposium (IV)*. IEEE (2020). <https://doi.org/10.1109/iv47402.2020.9304823>
 142. Pütz, A., Zlocki, A., Bock, J., Eckstein, L.: System validation of highly automated vehicles with a database of relevant traffic scenarios. In: 12th ITS European Congr., vol. 1, 19–22 (2017)
 143. Krajewski, R., Moers, T., Neger, D., Eckstein, L.: Data-driven maneuver modeling using generative adversarial networks and variational autoencoders for safety validation of highly automated vehicles. In: 21st International Conference on Intelligent Transportation Systems (ITSC). IEEE (2018). <https://doi.org/10.1109/itsc.2018.8569971>
 144. Hans-Peter Schoener: Challenging highway scenarios beyond collision avoidance for autonomous vehicle certification. ResearchGate preprint (2020). <https://doi.org/10.13140/RG.2.2.29355.05926>
 145. Kurzidem, I., Saad, A., Schleiss, P.: A systematic approach to analyzing perception architectures in autonomous vehicles. In: Zeller, M., Höfig, K. (eds.) *Model-based safety and assessment*, pp. 149–162. Springer International Publishing, Cham (2020)
 146. Kuhn, R., Kacker, R., Lei, Y., Hunter, J.: Combinatorial software testing. *Computer* **42**(7), 94–96 (2009). <https://doi.org/10.1109/mc.2009.253>
 147. Gladisch, C., Heinzemann, C., Herrmann, M., Woehrle, M.: Leveraging combinatorial testing for safety-critical computer vision datasets. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
 148. Weber, N., Frerichs, D., Eberle, U.: A simulation-based, statistical approach for the derivation of concrete scenarios for the release of highly automated driving functions. In: *AmE - Automotive meets Electronics; 11th GMM-Symposium*, 1–6. Dortmund, Germany (2020). <https://doi.org/10.13140/RG.2.2.15306.31683/1>
 149. RS-Reference - efficiently access massive and reliable ground truth data. Online (2020). <http://www.robosense.ai/en/rslidar/RS-Reference>. Robosense
 150. ASAM OpenLABEL (2020). <https://www.asam.net/project-detail/scenario-storage-and-labelling/>. Accessed on 09/12/2020
 151. Scheiner, N., Appenrodt, N., Dickmann, J., Sick, B.: Automated ground truth estimation of vulnerable road users in automotive radar data using GNSS. In: *IEEE MTT-S International conference on Microwaves for Intelligent Mobility (ICMIM)*. IEEE (2019). <https://doi.org/10.1109/icmim.2019.8726801>
 152. Ho, V., Rauf, K., Passchier, I., Rijks, F., Witsenboer, T.: Accuracy assessment of RTK GNSS based positioning systems for automated driving. In: 15th Workshop on Positioning, Navigation and Communications (WPNC). IEEE (2018). <https://doi.org/10.1109/wpnc.2018.8555805>
 153. Tahir, M., Afzal, S.S., Chughtai, M.S., Ali, K.: On the accuracy of inter-vehicular range measurements using GNSS observables in a cooperative framework. *IEEE Trans. Intell. Trans. Syst.* **20**(2), 682–691 (2019). <https://doi.org/10.1109/TITS.2018.2833438>
 154. Subirana, J.S., Zornoza, J.J., Hernández-Pajares, M.: GNSS data processing. volume 1: Fundamentals and algorithms. ESA Communications ESTEC, PO Box **299**, 2200 (2013)
 155. Lampe, B., Woopen, T., Eckstein, L.: Collective Driving - Cloud Services for Automated Vehicles in UNICARagil. In: 28th Aachen Colloquium. Aachen, Germany (2019)
 156. Knake-Langhorst, S., Gimm, K., Frankiewicz, T., Köster, F.: Test site AIM - toolbox and enabler for applied research and development in traffic and mobility. *Trans. Res. Procedia* **14**, 2197–2206 (2016). <https://doi.org/10.1016/j.trpro.2016.05.235>
 157. Notz, D., Becker, F., Kühbeck, T., Watzenig, D.: Extraction and assessment of naturalistic human driving trajectories from infrastructure camera and radar sensors. [arXiv:2004.01288](https://arxiv.org/abs/2004.01288) (2020)
 158. Krämmmer, A., Schöller, C., Gulati, D., Lakshminarasimhan, V., Kurz, F., Rosenbaum, D., Lenz, C., Knoll, A.: Providentia - a large-scale sensor system for the assistance of autonomous vehicles and its evaluation. [arXiv:1906.06789](https://arxiv.org/abs/1906.06789) (2019)
 159. Kloeker, L., Kloeker, A., Thomsen, F., Erraji, A., Eckstein, L.: Traffic detection using modular infrastructure sensors as a data basis for highly automated and connected driving. In: 29th Aachen Colloquium - Sustainable Mobility, 1835–1844 (2020)
 160. Kloeker, L., Geller, C., Kloeker, A., Eckstein, L.: High-precision digital traffic recording with multi-lidar infrastructure sensor setups. [arXiv:2006.12140](https://arxiv.org/abs/2006.12140) (2020)
 161. Guido, G., Gallelli, V., Rogano, D., Vitale, A.: Evaluating the accuracy of vehicle tracking data obtained from unmanned aerial vehicles. *Int. J. Trans. Sci. Technol.* **5**(3), 136–151 (2016). <https://doi.org/10.1016/j.ijst.2016.12.001>. <http://www.sciencedirect.com/science/article/pii/S2046043016300272>
 162. Kruber, F., Morales, E.S., Chakraborty, S., Botsch, M.: Vehicle position estimation with aerial imagery from unmanned aerial vehicles. [arXiv:2004.08206](https://arxiv.org/abs/2004.08206) (2020)
 163. Krajewski, R., Bock, J., Eckstein, L.: Drones as a tool for the development and safety validation of highly automated driving. In: *Aachener Kolloquium*, 1449–1462 (2019)
 164. Kurz, F., Waigand, D., Pekezou-Fouopi, P., Vig, E., Henry, C., Merkle, N., Rosenbaum, D., Gstaiger, V., Azimi, S., Auer, S., et al.: DLRAD - a first look on the new vision and mapping benchmark dataset for autonomous driving. *Int. Arch. Photogramm., Remote Sens. Spatial Inf. Sci.* **42**(1) (2018)
 165. Kurz, F., Rosenbaum, D., Runge, H., Reinartz, P.: Validation of advanced driver assistance systems by airborne optical imagery. *Proc. of Mobil.TUM 1–11* (2015)
 166. Wang, Z., Feng, D., Zhou, Y., Zhan, W., Rosenbaum, L., Timm, F., Dietmayer, K., Tomizuka, M.: Inferring spatial uncertainty in object detection. [arXiv:2003.03644](https://arxiv.org/abs/2003.03644) (2020)
 167. Wang, S., Pei, K., Whitehouse, J., Yang, J., Jana, S.: Efficient formal safety analysis of neural networks. [arXiv:1809.08098](https://arxiv.org/abs/1809.08098) (2018)
 168. Seshia, S.A., Sadigh, D., Sastry, S.S.: Towards verified artificial intelligence. [arXiv:1606.08514](https://arxiv.org/abs/1606.08514) (2020). V4 (after first version from 2016)
 169. Barbier, M., Renzaglia, A., Quilbeuf, J., Rummelhard, L., Paigwar, A., Laugier, C., Legay, A., Ibañez-Guzmán, J., Simonin, O.: Validation of perception and decision-making systems for

- autonomous driving via statistical model checking. In: IEEE Intelligent Vehicles Symposium (IV), 252–259 (2019)
170. von Neumann-Cosel, K.: Virtual test drive: Simulation umfeld-basierter Fahrzeugfunktionen. Dissertation, Technical University Munich, Munich (2014)
 171. Rao, D., Pathrose, P., Huening, F., Sid, J.: An approach for validating safety of perception software in autonomous driving systems. In: Papadopoulos, Y., Aslansefat, K., Katsaros, P., Bozzano, M. (eds.) Model-based safety and assessment, pp. 303–316. Springer International Publishing, Cham (2019)
 172. Myklebust, T., Stålhan, T., Hanssen, G.: Agile safety case and DevOps for the automotive industry. In: 30th European Safety and Reliability Conference and the 15th Probabilistic Safety Assessment and Management Conference (2020)
 173. Antonante, P., Spivak, D.I., Carlone, L.: Monitoring and diagnosability of perception systems. [arXiv:2011.07010](https://arxiv.org/abs/2011.07010) (2020). V3
 174. Woopen, T., Lampe, B., Böddker, T., Eckstein, L., Kampmann, A., Alrifae, B., et al.: UNICARagil - Disruptive Modular Architectures for Agile, Automated Vehicle Concepts. In: 27th Aachen Colloquium. Aachen, Germany (2018)
 175. Nolte, M., Jatzkowski, I., Ernst, S., Maurer, M.: Supporting safe decision making through holistic system-level representations & monitoring – a summary and taxonomy of self-representation concepts for automated vehicles. [arXiv:2007.13807](https://arxiv.org/abs/2007.13807) (2020)
 176. Lampe, B., van Kempen, R., Woopen, T., Kampmann, A., Alrifae, B., Eckstein, L.: Reducing uncertainty by fusing dynamic occupancy grid maps in a cloud-based collective environment model. In: IEEE Intelligent Vehicles Symposium (IV). IEEE (2020). <https://doi.org/10.1109/iv47402.2020.9304816>
 177. Winner, H.: ADAS, Quo Vadis?. In: Winner H., Hakuli S., Lotz F., Singer C. (eds) Handbook of Driver Assistance Systems. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-12352-3_62
 178. Grubmüller, S., Stettinger, G., Sotelo, M.A., Watzenig, D.: Fault-tolerant environmental perception architecture for robust automated driving. In: IEEE International conference on Connected Vehicles and Expo (ICCVE). IEEE (2019). <https://doi.org/10.1109/iccve45908.2019.8965112>

Michael Hoss received his B.Sc. and M.Sc. degrees in Computational Engineering Science from RWTH Aachen University, Germany, in 2015 and 2017, respectively. He currently pursues his Ph.D. degree at the same university. From 2018 to 2021, he worked as a research associate at the Institute for Automotive Engineering (ika), RWTH Aachen University. His research on safety assurance of automated vehicles focuses on test methods for the perception subsystem.

Maike Scholtes received her B.Sc. and M.Sc. (Dean's List Award) degrees in Computational Engineering Science from RWTH Aachen University, Germany in 2015 and 2017, respectively. During semesters abroad she, inter alia, studied at the University of California, San Diego. After working in the area of driving simulation at fka SV Inc. in the Silicon Valley, she started her Ph.D. at the Institute for Automotive Engineering (ika) at RWTH Aachen University in 2018. She works on the assessment of automated

driving and ADAS with a focus on machine perception.

Lutz Eckstein studied mechanical engineering at Stuttgart University. He received the Ph.D. degree from Stuttgart University in 2000. During his time in the industry, he worked on active safety systems and human-machine interactions for Daimler AG and BMW AG. He was appointed to his professorship and current position at RWTH Aachen University in 2010. He is currently the Chair of the Institute for Automotive Engineering (ika), RWTH Aachen University.