Towards an Operational Design Domain That Supports the Safety Argumentation of an Automated Driving System

Magnus Gyllenhammar^{*,‡}, Rolf Johansson^{*,∥}, Fredrik Warg[†], DeJiu Chen[‡], Hans-Martin Heyn[§] Martin Sanfridson[§], Jan Söderberg[¶], Anders Thorsén[†], Stig Ursing^{**}

> *Zenuity AB, magnus.gyllenhammar@zenuity.com †RISE Research Institutes of Sweden, {fredrik.warg, anders.thorsen}@ri.se [‡]KTH Royal Institute of Technology, chendj@kth.se [§]Volvo Technology AB, {hans-martin.heyn, martin.sanfridson}@volvo.com ¶Systemite AB, jan.soderberg@systemite.se **Semcon Sweden AB, stig.ursing@semcon.com ¶Autonomous Intelligent Driving, rolf.johansson@aid-driving.eu

Abstract-One of the biggest challenges for self-driving road vehicles is how to argue that their safety cases are complete. The operational design domain (ODD) of the automated driving system (ADS) can be used to restrict where the ADS is valid and thus confine the scope of the safety case as well as the verification. To complete the safety case there is a need to ensure that the ADS will not exit its ODD. We present four generic strategies to ensure this. Use cases (UCs) provide a convenient way providing such a strategy for a collection of operating conditions (OCs) and further ensures that the ODD allows for operation within the real world. A framework to categorise the OCs of a UC is presented and it is suggested that the ODD is written with this structure in mind to facilitate mapping towards potential UCs. The ODD defines the functional boundary of the system and modelling it with this structure makes it modular and generalisable across different potential UCs. Further, using the ODD to connect the ADS to the UC enables the continuous delivery of the ADS feature. Two examples of dimensions of the ODD are given and a strategy to avoid an ODD exit is proposed in the respective case.

Index Terms—ADS, automated driving system, functional safety, ODD, operational design domain

I. INTRODUCTION

Automated driving systems (ADS) show great potential to disrupt the transport industry and change the way we travel drastically. One of the remaining challenges in realising this potential is to show that the ADSs are sufficiently safe to be released on public roads. If not done adequately it would incur an unreasonable risk of serious loss events, including fatalities. It is not feasible to prove that the ADS is sufficiently safe predeployment by driving on public roads. Even a fleet of 100 cars would have to be driven continuously for five centuries to accumulate enough exposure hours to prove the ADS is safer than human drivers when it comes to fatalities [1]. Instead, the focus should be on minimising the residual risk that is

This research has been supported by Vinnova via the project ESPLANADE (ref 2016-04268).

left when launching the ADS. The entire process of ensuring sufficient safety of the ADS is discussed in Sec. II.

Within the automotive industry, it has been agreed that following the processes outlined in ISO 26262 [2] yields a system of sufficient functional safety. However, this standard was conceived with traditional automotive systems in mind, where the driver with a certain probability is able to compensate for a failure that emanates from within the computer system. The key process is still valid for developing an ADS, but two of the main tasks of ISO 26262, namely showing completeness of the hazard analysis and risk assessment (HARA) as well as the completeness of the verification, turns out to be significantly more difficult for an ADS compared to a traditional automotive system. One reason for this difficulty being that an ADS, in order to achieve an SAE automation level 3-5 [3], tends to require a complex set of interwoven functions and sub-systems to perform the dynamic driving task (DDT) on its own. A second reason being that the ADS needs to cope with the complete uncertainty of its environment. Traditional automotive systems have an implicit operational design domain (ODD) (everything that it can be exposed to on the road), but for an ADS the ODD can be used to confine the content of the HARA and subsequently the scope of the verification. An efficient format and content of the ODD can thus support showing completeness of the HARA as well as limiting the required verification.

A. An Example ADS

Assume we are tasked with developing an ADS that is able to drive autonomously between the cities of Gothenburg and Malmö in Sweden with a maximum speed of 110 km/h. It is restricted to only operate on the E6 motorway and can only be active during daytime between April and October and in dry weather (< 1mm rain per hour). This information gives a quite good indication of what needs to be handled by the ADS, but it does not directly help in the development process. Almost all the relevant information is not explicitly stated, and it will require an elaborate process to go from this list of information to one that can be connected to the HARA and the requirements of the ADS. It is crucial to conduct this process in order to get a set of requirements on the ADS that is possible to verify.

One explicitly stated dimension is the speed limit. This is easy to design for and it is easy to connect to a technical requirement. The restriction to spring-autumn gives us some additional information. It tells us that we most likely do not need to handle snow and probably only on really rare cases temperatures will reach below zero degrees Celsius. But we are not sure if this needs to be within the capabilities of the ADS or not. There are probably a few good reasons for imposing this restriction in the first place, all of which take into account the customer value and the ease of implementation and design. However, these reasons are wrapped up in a "use casestatement", like the one in the first paragraph of this section, that is easy for the customer and business side to understand, but which obscures the impact to the design. The limitation of rain is likely one of those examples. It is not the rain itself that is the problem to the development of the ADS, but the loss of sight and reduction of friction on the road.

Similarly, the daytime restriction is probably there to ensure that there is enough illumination, but this is not what we have to account for. There is still the rare event of eclipses. Do we have to handle those? Or are they outside of the scope even if they occur during daytime? Additionally, what does it mean to be restricted to the E6 motorway between these two cities? Are there barriers present? What are the speed limits? What curvatures and banking can we expect? These roadway characteristics are fairly straight forward to determine and model, as they are static (except for repairs and road works). We just need to go for a drive on this stretch of road and note down the interesting details. However, what details should we note down? We do not know what to design for without the notes and we also do not know what will be important for us during the design process. We might realise that other aspects entirely would impact the ADS or facilitate the design process.

The largest implicit category of dimensions in this "use case-statement" is the traffic behaviour. We know that the driving should adhere to Swedish law, but there are many more nuances to traffic behaviour. Furthermore, traffic behaviour can be modelled with probabilistic methods. For example, it is unlikely that we encounter pedestrians on the E6 motorway, although it is not impossible. Without considering these nuances we will end up with an ADS which will either have unnecessarily large or (worse) insufficient safety margins. The unreasonably large safety margins would occur if we place too much weight on accounting for the "worst cases" when in reality they are extremely unlikely. Places along the road that lie far from areas trafficked by pedestrians and which rarely see vehicles at standstill will have a very low occurrence of pedestrians. On the other hand, the insufficient safety margins might be because, on other parts of the road, cars break down often and thus pedestrians appear frequently there. Both cases are difficult to estimate without data.

The examples above show the need for a process to classify and quantify the operating conditions (OCs) of the ADS such that what is imposed by the UCs match with what is possible to implement and verify in the ADS. In this paper, we propose to use the ODD to model and collect these conditions to make them explicit and to ensure that they confine what is required from the UC. That said, the ODD should still be refined throughout the development process, as visualised in Fig. 3.1 in [4]. However, clarifying the OCs early in the development phase will greatly support achieving a performant and safe ADS.

B. Contributions and outline

In this paper we discuss how the ODD of the ADS can be used to define what external and internal OCs that the ADS needs to handle and subsequently what it should be designed and verified for. Each such OC is associated with at least one of four presented strategies to ensure that the ADS does not experience an ODD exit during operations. It is acknowledged that UCs constitute prominent and useful examples of the spatial and temporal strategy to ensure the ADS remains in its ODD for collections of OCs. Further, we present a framework to enable mapping between the ODD and a UC and suggest that writing the ODD on this format facilitates the comparison to different UCs.

One result of viewing the ODD in this way is that instead of explicitly defining geographic regions within the ODD we propose that geographic limitations show up in the values of other OCs or as a potential strategy to ensure that the OCs are fulfilled. We believe this to be in line with the SAE definition [3] and it further emphasises the modularity of the ODD. This modularity supports continuous deployment and continuous delivery of the ADS feature, as well as handling product portfolios containing a large number of product variants.

Each strategy to avoid ODD exits needs to be constructed with sufficient integrity. One way of achieving this for UCs is by using driving data. Further, such quantification of the requirements from the UC not only brings meaning to the dimensions of the ODD for the design and implementation of a performant ADS but also reduces unnecessary safety margins.

The process of achieving a safe ADS before deployment is discussed in Sec. II and the process of mapping a UC to the ODD is visualised. The four strategies to ensure that the ADS remains in its ODD are presented in Sec. III. Further, a framework for categorising the OCs of the ODD is presented in Sec. IV-A and how the categories can be quantified for UCs is discussed in Sec. IV-B. In Sec. V the role of the ODD when designing and implementing the ADS is outlined. Related work is discussed in Sec. VI, and in Sec. VII conclusions and potential future work are presented. To highlight the impact of the ODD to the design process two examples of ODD dimensions are given in Sec. IV-B1 and Sec. III-A respectively. Further suitable strategies to ensure the OC for each example are discussed. Problem Domain



UC1 UC4 UC3 E

Fig. 2. Use cases (UCs) and the ODD. The ADS can only be released for UCs contained within the ODD, i.e. not UC3.

Fig. 1. A use case (UC) being quantified into internal and external operating conditions (OCs) through a world model. The OCs of the UC are contained within the OCs of the ODD.

II. THE PROCESS OF ACHIEVING A SAFE ADS

The safety of the ADS can be argued in several ways, but it is important to consider that the main purpose is to ensure that the released product does not expose its users or other traffic participants to unreasonable risks. Koopman and Wagner [5] present a validation framework that aims to do just this by considering the residual risk at each level of testing. Junietz et al. [6], however, note that the safety validation of ADSs cannot be "solved" mathematically.

Safety validation is just one aspect. Before the validation comes the problem formulation, the HARA, the requirement breakdown, the design, the implementation, and the verification of the posed requirements. In this section, a potential process to achieve a safe deployable ADS is outlined. Note that this process would benefit from being used in conjunction with a validation framework such as the one proposed by [5] and potentially coupled with an analysis of the validation results using extreme value theory [7].

The ODD helps to confine the HARA as well as the design, development, and verification processes. However, if the ADS is designed and verified towards an ODD it is paramount that the ADS does not leave that ODD. For each dimension in the ODD, there should be a strategy to ensure that this dimension is not exceeded during operations. One prominent and convenient way of viewing these strategies is bundling them into a UC. This has the further advantage to clarify the negotiation between user benefits and ADS design, as highlighted in Sec. I-A.

In Fig. 1 the process of modelling a UC through a world model is depicted. These pieces need to exhaustively model the UC and to quantify what the external conditions are around the ADS. If this model is based on a representative set of data, it can be argued that the model itself is adequately complete. Further, it can then be argued that the residual risk of events not captured by the model is small enough. For example, if the model is populated with representative data from T hours of driving, it can be ensured that any events that would require a change of the model have a mean time between the events of

$$\bar{T_E} \ge \frac{2}{\chi^2_{CL,2}} \cdot T,\tag{1}$$

using the connection between the Poisson process arrivals and the chi-squared distribution as discussed in [8]. (This assumes that the events are independent and occur with exponentially distributed times between the events.) The denominator is the Chi-squared distribution with confidence level CL. For CL = 75%, the factor in front of T becomes approximately 0.72. This can be used to estimate the risk that the model is incomplete and also to quantify the incurred residual risk. Note that in most of the modelled dimensions it will be possible to extrapolate, which makes the model useful beyond the collected data.

To allow the ADS to operate within a UC the ODD needs to encapsulate all the OCs that the quantified model of the UC requires. How this might look like for a few different UCs simultaneously is depicted in Fig. 2. Given an ODD the HARA can be conducted and the requirements, to fulfil the safety goals as well as the remaining OCs of the ODD, can be broken down and defined. These requirements together with the requirements from strategies to ensure that ADS remains within the ODD make up the system specification. The ADS can subsequently be implemented and verified against this set of requirements, as depicted in Fig. 3. Following the ISO 26262 [2], development processes starting from the ODD will ensure that there is a sufficiently low risk contribution following the HARA, the requirement breakdown and the implementation and verification.



Fig. 3. Depicts how the safety goals (SGs) of the HARA and the nonsafety/quality (QM) requirements are valid within the ODD and how the requirements from the strategies to remain within the ODD ensure that the ADS stays within the ODD. The ADS is verified towards these three classes of requirements.

Finally, once released, the environment should be monitored to ensure that the UC does indeed follow the estimated OCs. Further, the ADS should also be monitored in the field to capture any potential violations of the safety requirements before an event of significant loss occurs. One additional benefit with having the field monitoring is that the models of the UC, that used collected data, can be further verified and refined with the added data from the field vehicles.

III. STRATEGIES FOR GUARANTEEING NOT TO EXIT ODD

As we discuss in later sections of the paper, there are different kinds of OCs that can constitute an ODD. We need to guarantee that the ADS feature will never experience an ODD exit while enabled. It is thus necessary to have strategies that are able to ensure this property for all OCs. Four such strategies are presented in Table I. Some conditions, like the internal having a maximum speed of 60 km/h, might be possible to directly guarantee for the ADS as part of the basic feature definition (strategy I), requiring no direct triggering condition at run-time. Some external conditions can be fulfilled by only accepting trips (strategic tasks) that beforehand can be guaranteed to be possible to complete without exiting the ODD (strategy II). As an example, only trips which are possible to complete by driving on road segments that are ensured to adhere to the ODD (i.e. UCs which are contained by the ODD) should be accepted. If point A and B are separated by a stretch of road that has not been characterised and ensured to be within the ODD, this trip should not be accepted.

Some other external conditions can be guaranteed by expressing geographical or temporal triggering conditions (strategy III). Many of the external OCs are connected to the UCs which they have been estimated from. Thus, if we assume that the UCs intended for the ADS is connected to a region. the geographic boundaries used for triggering ODD exits of these OCs might well be given already in the UC. In addition to geographic boundaries there are also restrictions in time: time of day, time of year, etc. Strategy III is closely linked to II, which in many cases will use the same regions, but to accept/reject a strategic task rather than triggering on potential exits during run-time, which is the scope of strategy III. Finally, some external conditions can be guaranteed by dynamically measuring properties related to the OC (strategy IV), e.g. road friction or rain. The example of rain intensity being the OC is developed further in Sec. III-A below.

These strategies help in bridging the UC definition and the OCs of the ODD. Recall the original example, from Sec. I-A, where we had daytime as a restriction. Here this restriction is reintroduced, but only as a means for ensuring that the OCs, illumination or road characteristics, will remain within the defined limits.

A. Example - Rain intensity

Fig. 4 shows a hypothetical frequency of rain with different intensity as estimated for a certain UC (the distribution is made up for the sake of this example and has no real bearing). If we want to include the entire UC the line corresponding to the ODD needs to be set using appropriate statistical modelling including confidence bounds based on the data from the UC. By acknowledging that different intensities occur with different frequencies it is possible to provide solutions, with different integrity, to the different cases. Alternatively, one could do a more coarse division of the rain intensities if this differentiation is not needed.

There are multiple reasons why one might want to restrict the amount of rain that the ADS should be able to handle, for example, to avoid handling aquaplaning and to ensure sufficient perception capabilities. In addition to selecting UCs that have a low probability of rain (strategy III), it might be possible to devise methods to predict the future risk of high rain intensity to make use of strategy IV. If such a method exists, e.g. by coupling a weather forecast with a sensor, measuring the current rain intensity, and a rain radar, it is possible to remove these high rain intensities from the ODD. The magenta line in Fig. 4 corresponds to the ODD where the rain intensity is limited by a combination of strategy III, taking care of the left part of the dashed vertical line, and IV, ensuring that the truncation is valid.



Fig. 4. A hypothetical distribution of rain. The solid black line being the "real" estimated distribution and the magenta one corresponding to the decided ODD. Note the truncation at the dashed vertical line.

The ODD for the rain intensities in Fig. 4 defines what the ADS can expect during operations and it is also against this that the ADS needs to be designed and verified. Considering rain on its own might result in conservative assumptions being made with respect to other dimensions. However, if there is a need to find a less conservative solution for the combination of, let's say overtaking speed and rain, then the analysis can be expanded to take that correlation into account, as is discussed in Sec. IV-B.

IV. CONNECTING A USE CASE TO THE ODD

When defining the strategies to avoid ODD exits in Sec. III it was acknowledged that UCs provide a natural example of strategy III and that it will be useful for many different OCs. In the following section, we present a framework to categorise the OCs of the UC to ensure that it is sufficiently modelled. Further, the possibility of quantifying the OCs of a UC is discussed in Sec. IV-B.

A. Categorising the Operating Conditions

Ulbrich et al. [9] present a terminology in which a UC is defined by scenario(s), functional range, desired behaviour and functional system boundaries. This structure is also selected

 TABLE I

 Strategies for guaranteeing not to exit ODD.

Strategies		Need to estimate inside ODD in design-time.	Need to define triggering cond. for DDT-fallback	Need for reliable map info.	Need for sensors capable of measuring condition	
Ι	Internal	Inherent in ADS feature definition	N	N	N	N
II		Checking mission when accepting strategic task	Y	N	Y	N
III	External	Statically defined, spatial and temporal triggering conditions	Y	Y	Y	N
IV		Run-time measurable triggering cond. related to OC	N	Y	N	Y

as basis for the taxonomy of use case, scene, and scenario in the ongoing standardisation work for ISO/PAS 21448 (Safety of the intended functionality, SOTIF). As the framework is geared towards testing and validation, and not modelling and requirement definition, there are some changes that can be made to make it better support the task of categorising the OCs. In Fig. 5 the framework with our amendments is shown. There are two main categories of OCs, the internal and the external ones. The internal OCs are the conditions pertaining to the ADS itself and its user. They are defined by the UC directly or they follow from the requirements on the interaction with the user of the ADS. For example, the speed restriction of the ADS (which would be located in Functional Range) is likely given by the UC definition, whereas the time for user to take back control after a request by the ADS (located in Fallback Ready User) can be estimated through user profiling and usage statistics. The external conditions generally need to be modelled and estimated. This estimation process is discussed further in Sec. IV-B.

The aim is to make sure that the UC is sufficiently modelled. Since we make a separation of internal and external OCs, it is beneficial to add one level of granularity by separating other actors and ego vehicle underneath the Actions and Events category. Only the actions and events of other actors are relevant for the analysis of the OCs, which is why the ego vehicle box is greyed out in Fig. 5. It should be noted that there will potentially be an interaction between the ego vehicle and the actions from the other actors. However, the allowed actions and events of the ego vehicle are captured within the outer layer within the Functional Range rather than as part of the scenario. Further, the Self-Representation of the different actors present in the scene is important when simulating and recreating test cases, but for modelling the UC it brings little value to know the other vehicles' own estimates of their abilities. Thus, this box is also greyed out. The Functional System Boundaries and the Desired Behaviour are also of no use when modelling the OCs of the world surrounding the ADS. The functional boundaries are proposed to be defined by the underlying OCs of each of the other categories and the desired behaviour will be limited by the same OCs.

In Fig. 5 the *Goals and Values* category has been split into the sub-categories *Permanent* and *Transient* as discussed by [9]. We decide to make this separation explicit in the figure to be able to split the categories into internal and external OCs.

In addition to the original categories from [9] there is a



Fig. 5. Categories of operating conditions (OCs). Reworked based on [9]. The greyed out boxes are not applicable as categories of OCs. Further, the OCs are split into external and internal conditions pertaining to the environment or the ADS itself and its user.

need to add one for the *Fallback Ready User*, within which the conditions assumed/required of the user are located, e.g. should the user be capable of resuming control of the ADS? Does (s)he have a driver's licence? Is the ADS a dedicated vehicle without a user? Etc. The category of *Connectivity* is also proposed to be added in order to capture the constraints related to the availability of different connectivity-based support functions. Examples of such include global position system (GPS), vehicle to vehicle (V2V) communication, mobile network capabilities (such as 5G) or access to remote fleet management systems.

The relevant categories OCs are described in Table II together with some examples and potential data sources to contribute to the estimation of their values within a UC.

B. Quantifying the Operating Conditions

The categories presented in Sec. IV-A all need to be exhaustively modelled. The scenario databases proposed by the PEGASUS project [10], as well as by TNO [11], aim at capturing the challenges imposed by the world around the ADS as test cases. The majority of the works cited in Sec. VI are, along the same lines, trying to categorise the world in order to outline the tests that the ADS should be able to handle. These frameworks would be a good place to start in order to populate the categories of Fig. 5. The potential sources for estimation of the categories can differ greatly, as can be seen from the third column of Table II, and there will consequently be a need for different models to capture the different data.

Czarnecki [12], [13] develops an operational world model (OWM) ontology which he argues can, at least in part, be used for defining the ODD of an ADS. This OWM for the roads,

TABLE II

EXAMPLES OF OPERATING CONDITIONS WITHIN EACH CATEGORY (FROM FIGURE 5) AND POTENTIAL STRATEGIES TO QUANTIFY THEM.

Category	Description and examples	Potential source(s) for estimation
Dynamic Elements	Elements that move or are able to move. The types of such objects, e.g. • vehicles • pedestrians • emergency vehicles and their properties, e.g. • colour • shape Also including road works, including trucks, workers, etc. but excluding stationary items like cones.	 institutional reports accident databases measurements from driving
Scenery	Geo-spatially stationary elements and quasi static phenomena • vegetation • buildings • road properties, e.g. - barriers - curvature • weather	 map data institutional reports, guidelines, etc. measurements from driving historical weather data (however, global warming is expected affect the validity of the historical data)
Connectivity	Availability of • GPS • vehicle to infrastructure communications • remote fleet management systems • weather forecast data	 need from system implementation related to strategies to remain within ODD need for run-time information from these sources satellite coverage maps
Actions & Events, Other Ac- tors	 The actions associated with the dynamic elements e.g. cars changing lanes car turning on the hazard lights animal crossing the road motorcycle driving in between lanes in a queue 	 institutional reports accident databases measurements from driving
Goals & Values - Permanent	• traffic rules, laws and policies, e.g. speed limits	• law texts etc.
Goals & Values - Transient	 operator input change of mission objectives	• user/usage mapping/profiling
Functional Range	 speed restrictions maximum brake allowed	• internal feature specification
Fallback Ready User	time needed for user takebackcapability of user to conduct takeback	limited by usage requirementsuser mapping/profiling

as presented in [12], gives a holistic view of how to classify as well as quantify road segments. Starting from there would greatly help in defining the *Scenery* category of Fig. 5. The OWM ontology for road users [13] lists the dynamic elements that the ADS might be exposed to. However, their actions are not exhaustively discussed and here we see an additional need for developing more in-depth models, such as the one presented by e.g. [14].

Our vision is to use models that collectively exhaust the UC. The pieces of such a model should, on the lowest level, capture the entirety of the UC, albeit with a granularity that might be next to useless when constructing the ADS feature. However, the pieces can be analysed and combined in more elaborate formations to increase the level of granularity and

thus achieve an ADS with sufficiently low unnecessary margins. For example, assume we want to model the longitudinal movement of other vehicles. In a simple model the parts within this category could be braking, accelerating, and standstill. There is no overlap between the pieces, and they can be parameterised and quantified individually. In the first-level analysis, these parameters will be considered independently of any other parameters. However, in a more refined analysis, one might find that convertible cars accelerate more aggressively, whereas there is a higher probability of trucks to be at standstill at the road shoulder. Thus, if there are issues with fulfilling the OCs resulting from the initial analysis a separation could be done to reduce the unnecessary margins for non-trucks with respect to standstills at the road shoulder. Having a world model built in such a way would create a foundation for a requirement analysis that can be just as granular and elaborate as needed. In the first instance, one could assume complete independence between the pieces, which would yield a conservative set of requirements since the worst cases from each piece need to be combined with each other, even though this might not at all be present in the real world. However, the granularity of the analysis could be increased by incorporating the correlation between the pieces that are most limiting for the ADS. Thus, using a world model in this way not only provides an efficient way of exhausting the OCs of the UC in the first place, but it also provides an efficient way of refining them when necessary.

There is significant work left to be done to compile a complete set of models for the purpose of exhaustively quantifying the UC. This is left as a suggestion for future work, but we give two examples of what we believe could be the format of the OCs resulting from such modelling. One has been presented in Sec. III-A already and the second one is presented below.

1) Velocity difference during overtake: A vehicle overtaking the ego-vehicle with ADS is an example of a piece of the model for Actions and Events, Other Actors. This piece of the model needs to be parameterised to quantify how overtakes might occur. The velocity difference during the overtake could be one such parameter. Note that in this granularity we do not account for the different vehicle types that might conduct the overtake. Thus, in the first instance of the analysis, it might be assumed that all vehicle types (modelled in the Dynamic Elements category) conduct overtakes in the same manner. This analysis can subsequently be improved to model the conditional probability or distribution for the different vehicle types (a convertible car being, for example, more likely to conduct an overtake than a truck).

The frequency of velocity differences between the own car and the overtaking vehicle as measured from around 500 hours of driving across Europe, the US and Asia is presented in Fig. 6. From here one could make a statistical model of the frequency distribution (including statistical confidence) and pose the frequencies for each bin as OCs that the UC requires. Alternatively, one could continue the analysis and, for example, classify these events according to severity, as defined in ISO26262 [2]. Assume that the assessed severity is for when the own vehicle swerves into the adjacent lane. In that case, the delta velocity can quite directly be linked to different severity levels, as indicated in Fig. 6. The choice is now whether to include the outcome of the severity classification as the OC or to stick with the fine-grained histogram. The larger the bins used, the more potentially unnecessary margins are needed. For example, half the events of the S1 category in Fig. 6 occur infrequently whereas the other half is quite common. However, if it is not possible to make separate solutions within the ADS to these different cases there is not use to make a more finegrained separation, because the solution needs to solve the most limiting case anyways.

The velocity difference makes up one parameter of the *overtake* within the *Actions and Events* category. To exhaust



Fig. 6. The frequency of measured delta velocity of overtaking vehicles while they pass the ego car together with potential classification of severity, as defined in ISO26262 [2], with S1 being the least severe and S3 corresponding to a fatality. The data is from roughly 500 hours of motorway driving across the US, Europe, and Asia.

this category there is a need for both more parameters for overtakes, as well as more pieces to complete the model and capture all relevant aspects. Similar to this example the additional pieces and parameters could be estimated for a certain UC in order to judge what is needed from the ADS.

The example of velocity differences during an overtake is an example of an external OC that can be assured through the third strategy *Statistically defined, spatial and temporal triggering conditions* from Table I. This is possible since the frequencies in Fig. 6 are estimated from data collected in a geographically and temporally restricted space. As long as the ADS remains within this space there is a high confidence that the velocity differences that the ADS experiences during overtakes will follow these estimated frequencies (including the statistical modelling with confidence).

V. THE ROLE OF THE ODD

As highlighted by the example from Sec. I-A it is beneficial if there is a clear distinction between the problem domain, outlining what should be solved, and the solution domain, which contains the design, requirements breakdown and system's solution. In making this distinction, it is necessary to ensure that what ends up in the solution domain is confined in such a way that it is (1) implementable and (2) verifiable. We suggest that this is the primary purpose of the ODD, to confine the necessary activities of the HARA as well as the verification of the requirements and make them independent of the specific UC considered. Fig. 3 depicts the ODD on the boundary between the problem and solution domain and how that relates to what needs to be done on in the solution domain.

By positioning the ODD in the border between the solution and problem domain, the solution domain will contain two parts to solve the requirements posed by the ODD.

- (i) The part of the ADS solving the operating conditions outlined in the ODD;
- (i) the part of the ADS ensuring that the OCs of the ODD are not surpassed, i.e. ensuring the system remains in an environment where the defined ODD limits are valid.

These two also correspond to the two categories of risk that Wittman et al. [15] discuss. Strategies to fulfil the second part are presented in Sec. III.

One prominent and useful strategy to ensure many OCs is UCs. It is important that the definition of the problem domain (i.e. the ODD) completely encapsulates the UCs that the ADS sets out to solve. Meaning that it is possible to compare a defined ODD with a UC and argue that the UC is completely contained within the ODD, as depicted in Fig. 2. The ADS can consequently be designed, implemented, and verified against the ODD rather than against the UC. Note further, that there is no direct link between the UC and the solution domain except that which is provided by the ODD. If the models of the UC are correct, this means that the need for testing the system on the roads for verification purposes is limited, as also suggested in Sec. II. Decoupling the design and verification of the ADS from the specific UC gives the further advantage that the same ADS can be used for several UCs as long as the UCs can be mapped towards the ODD of the ADS. As an example, releasing the ADS on (a) highways in Sweden require us to estimate the OCs of this UC. Once this is done an ODD that encapsulates these OCs can be defined and an ADS can be implemented according to this ODD. Further, when implementing the ADS, it might be turn out to be easy to design solutions for some of the operating conditions. It is thus decided to increase the ODD and design a system that is more capable than what is required for (a). In consecutive development we might want to also release the ADS on (b) highways in Germany. Again, we need to estimate what the OCs are within this region. If it turns out that our ODD of the already implemented ADS encapsulates the OCs of (b) the ADS is already implemented and verified, and we can go ahead releasing it in (b). However, if the OCs of (b) are not encapsulated by the ODD already, there is a need to update the ODD as well as the design, implementation, and verification of the ADS. The independence between the ADS and its UCs is a key feature that allows for continuous deployment and expansion of the ADS. How to construct an ODD in relation to the potential UCs such that they allow for modularisation of the ADS feature itself has not been investigated further, but is identified as future work.

Viewing the ODD as a medium to bridge the UC and the ADS means that we suggest that geographic region is not a direct dimension of the ODD itself, but rather

- captured as OCs within, for example the *permanent goals and values* category of Fig. 5, and
- as a factor that impacts where the defined ODD is applicable.

In this way, the "zones" proposed as one of the main categories in the ODD checklist presented by Thorn et al. [16] should also be kept out of the ODD definition. Instead, each of the implications should be contained within the proposed categories of Fig. 5. With a holistic modelling, this should be no issue since, for example, the interference zones (e.g. tunnels, parking garage, limited GPS) can be captured within the *Scenery* and *Connectivity* categories. The behaviour related to the traffic participants would then be modelled within the *Actions and Events* category. We suggest that the ODD is quantitatively defined for all applicable OCs. For more complex ADSs this facilitates the comparison between the possible regions for release of the ADS as well as during the implementation and verification phases. Even a level 5 function, where the ADS can operate the vehicle under all driver-manageable road conditions within its region of the world [3], will benefit from having an explicit ODD to express what is meant with "all driver-manageable conditions". When considering "smaller" features, it is obviously useful to define an explicit ODD to show the constraints of the feature. Two examples of such projects with limited ODD in the Gothenburg region are *VERA* by Volvo trucks [17] and *Born to Drive* [18].

- The ODD can be constructed from either of two directions:
- from defining and quantifying a UC that the ADS should solve and subsequently set an ODD that completely encapsulate this UC, or
- (2) by removing known challenges of the design and solution from the ODD and then match a UC to these criteria.

In practice, the ODD will be decided using a combination of both (1) and (2). There will be a negotiation between the business side aiming for a large UC and the engineers of the system focusing on finding a problem that is possible to solve and verify. The ODD can act as the medium through which this negotiation is made. In this light we can also view UC3 from Fig. 2 as a UC which was difficult to handle and thus we decided to draw the ODD to exclude it. To give an example, we might have trouble handling sun blinding the sensors during a certain incidence angle if there is rain on the road (an OC within the Scene category from Sec. IV-A). Knowing this we might opt to remove this challenge from the problem domain. The negotiation becomes a question of balancing the cost for solving this issue versus the loss of customer value. However, it is not possible to omit something from the ODD unless there are strategies to ensure the ADS will never experience it, as discussed in Sec. III. For example, restricting the feature to only be active during the time of day (strategy III from Table I) when the angle of the sun will not fall within the critical interval of incidence (for a specific set or roads since their slope will also impact the angle) could be one strategy. Alternatively, we can devise a strategy to ensure the ADS is only active when the road surface is dry (a combination of strategies II to IV of Table I might suffice).

VI. RELATED WORK

The view of the ODD presented in this paper takes the current specification from literature one step further and makes its uses more explicit. J3016 defines the ODD as consisting of the "Operating conditions under which a given driving automation system or feature thereof is specifically designed to function[...]" [3]. This is exactly what the proposed ODD is as well, we just define the OCs such that they are more detailed and explicit in order to facilitate the subsequent development process steps.

It should be noted that so far, literature does not provide a uniform definition of the ODD. The closest one gets to a common standard is SAE's recommended practice J3016 [3]. To give an example, SAE states "environmental, geographical, and time-of-day restriction, and/or [...] traffic or roadway characteristics" as possible condition criteria. This paper suggests a framework to categorise the operating conditions (OCs) making up the ODD but does not make any claims of providing a complete set of conditions. There are many other works that could be used in conjunction with ours for this purpose. Koopman and Fratrik [19] provide a comprehensive, and in eight categories organised, list of criteria that should at least be included in an ODD. The list includes (1) operational terrain, (2) environment/weather, (3) operational infrastructure, (4) interaction with environment, (5) geographical region (in regards to traffic rules). (6) communication modes. (7) availability of (map) data, and (8) expected distributions of rare elements (e.g. toll booths). The U.S. Department of Transportation recommends a minimum set of five items that include roadway types, geographic areas, speed range, environmental conditions and "other constraints" [20]. For reasons explained in Sec. V we, however, believe that geographic region should not be a direct ODD dimension but rather that it will impact the OCs required of the ODD. We further believe that the structure of Fig. 5 is able to capture all of the categories albeit in slightly different configurations. The benefits from having one structure for both the validation framework as well as for the definition of the ODD we would argue turns the scales to favour the categories based on the framework from [9] outlined in Sec. IV-A.

A report by Thorn et al. [16], focusing on testing of ADS, includes a chapter describing the identification of attributes that can be used to define the ODD for an ADS. Based on a literature survey they define and categorise the ODD into a taxonomy with six top-levels categories; (1) Physical infrastructure, (2) operational constraints, (3) (dynamic) objects, (4) connectivity, (5) environmental conditions and (6) zones (i.e. geographical restrictions or special road zones like construction zones). Each of the categories contains two to five subcategories. Additionally, they suggest an ODD checklist and they complete the checklist for three simplified ADS examples. We strongly suggest that the ODD should not be viewed as only a checklist, as this destroys the main purposes of the ODD in facilitating the development of the ADS. The OCs of the ODD can contribute with a detailed description of what the ADS needs to handle, which we believe we should use to the largest extent possible.

As part of the PEGASUS project, to provide simulation scenarios, Bagschik et al. [21] propose a knowledge-based ontology for defining the functional system boundaries of an automated vehicle. They include five layers: Road, traffic infrastructure, temporary changes to road/traffic, objects, and environment. Wittmann et al. [15] also set out to describe the system boundaries of an ADS. They present a specification space with four categories: Static environment, Traffic dynamics, Environmental conditions, and State of ego vehicle and passengers. They further differentiate between static and dynamic items which define the system boundaries. It is

similar to the definition of a scene with three main categories in [9], upon which the framework presented in Sec. IV-A is based. Further, the state of the ADS's passengers fits within the *Fallback Ready User* of 5. Reschka et al. [22] give a list of items that fit into the category self-representation of the vehicle from [9]: Position accuracy, grip value, viewing area, system operation status, and system reaction time.

All of the above frameworks capture the same information, but in different flavours. It is clear, however, that the detailed modelling is still left unexplored, maybe with the OWM ontology for road structures presented by Czarnecki [12] being the exception, as discussed before. This road ontology might further be quantified using the hierarchical tree-structured model with belief propagation recommended by Töpfer et al. [23].

Wittmann et al. [15] argue that the safety of an automated vehicle can be assured defining a functional system boundary, similar to how we propose to use the ODD, and a twopart surveillance system to ensure that the boundaries are not violated. Similar to the process proposed in this paper they argue that scenarios can be used to set the functional system boundary. The surveillance concept they propose does not give room for a fallback to the user if the boundary is approached. Additionally, neither when defining the functional system boundary with scenarios nor when discussing the surveillance concept, do they account for the exposure of the scenarios. The sheer number of scenarios are of no interest to the safety argumentation, but the risk exposure they contribute with is. Without accounting for the exposure their surveillance concept does not help to assess the residual risk of the ADS. Hörwick et al. [24] also suggest a function boundary monitor to ensure that the ADS (they call it fully automated driving assistance system) operates inside the defined functional boundaries, e.g. by an ODD. Rather than a monitor we suggest that there exist four generic strategies, as presented in Sec. III to ensure that the ADS will remain in its ODD. Colwell et al. [25] note that the ODD not only defines the restriction on the ADS, but also defines the functional requirements for systems responsible for ODD monitoring. The authors introduce a restricted operational domain (ROD) which, in contrast to the ODD, also captures a degraded functionality of the ADS. They propose a runtime monitoring system that continuously determines if the ADS is still acting inside the restriction of the ROD, similar to the safety supervisor architecture proposed by [26]. For many of the OCs, which will be required to exhaust the categories of Fig. 5, it will not be possible to make such a restriction, for others it will be highly tractable. Thus, the ROD should be considered when continuing the work with detailing the dimensions of the framework presented in this paper.

It is clear from the plethora of ODD definitions that there is plenty of work left to align on a complete list of dimensions to exhaust the framework presented in this paper. Further, it is also evident that there exist several visions of what the purpose(s) of the ODD should be. We hope that this paper sheds some light on this and shows the potential the ODD has when designing and developing an ADS.

VII. CONCLUSIONS AND FUTURE WORK

The operational design domain (ODD) of an automated driving system (ADS) aims to confine where the ADS is valid and restrict the scope of the safety case as well as the verification efforts. To ensure that the ADS never experience an ODD exit while activated we present a set of four strategies for the ADS to remain in its ODD and examples highlighting the use of the strategies are given. It is acknowledged that use cases (UCs) provide a convenient strategy for a collection of OCs. A framework to categorise the OCs of a UC is presented and how the OCs can be quantified is discussed. Writing the ODD using these same OCs further facilitate the mapping between different UCs towards the ODD. Since using the ODD to model the UC makes the ADS independent of the specific UC the ODD supports both continuous deployment of the ADS features as well as to accommodate a wide variety of ADS variants. Additionally, having the ODD on the proposed format makes it modular and generalisable across multiple UCs.

Detailed models of each of the OCs, within the categories presented in this paper, is suggested for future work. Further, analysing the impact of the proposed strategies to the ADS architecture is an important next step. Additionally, investigating the modularisation of the ADS itself and how that relates to potentially different UCs or ODDs is another interesting avenue for future work.

VIII. LIST OF ACRONYMS

ADS	automated driving system (fulfilling SAE		
	automation levels 3-5)		
DDT	dynamic driving task		
HARA	hazard analysis and risk assessment [2]		
OC	operating condition		
ODD	operational design domain		
OWM	operational world model [12], [13]		
ROD	restricted operational domain [25]		
UC	use case		

ACKNOWLEDGEMENT

The authors would like to thank the participants in the ESPLANADE project ¹ for valuable discussions and input.

REFERENCES

- N. Kalra and S. M. Paddock, "Driving to safety: How many miles of driving would it take to demonstrate autonomous vehicle reliability?" Transportation Research Part A: Policy and Practice 94, pp. 182–193, 2016.
- [2] ISO, "ISO 26262:2018 Road vehicles Functional safety," 2018.
- [3] SAE, "SAE J3016:201806 SURFACE VEHICLE RECOMMENDED PRACTICE - Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," 2018.
- [4] Colwell, Ian, "Runtime restriction of the operational design domain: A safety concept for automated vehicles," 2018. [Online]. Available: http://hdl.handle.net/10012/13398
- [5] P. Koopman and M. Wagner, "Toward a framework for highly automated vehicle safety validation," SAE Technical Paper, Tech. Rep., 2018.

¹See: http://esplanade-project.se/

- [6] P. Junietz, W. Wachenfeld, K. Klonecki, and H. Winner, "Evaluation of different approaches to address safety validation of automated driving," in 21st International Conference on Intelligent Transportation Systems (ITSC). IEEE, 2018, pp. 491–496.
- [7] D. Åsljung, J. Nilsson, and J. Fredriksson, "Using extreme value theory for vehicle level safety validation and implications for autonomous vehicles," *IEEE Transactions on Intelligent Vehicles*, vol. 2, no. 4, pp. 288–297, 2017.
- [8] A. Gorski, "Chi-square probabilities are poisson probabilities in disguise," *IEEE transactions on reliability*, vol. 34, no. 3, pp. 209–211, 1985.
- [9] S. Ulbrich, T. Menzel, A. Reschka, F. Schuldt, and M. Maurer, "Defining and substantiating the terms scene, situation, and scenario for automated driving," in *Proceedings of IEEE 18th International Conference on Intelligent Transportation Systems*, Las Palmas, Spain, Sep. 2015.
- [10] A. Pütz, A. Zlocki, J. Bock, and L. Eckstein, "System validation of highly automated vehicles with a database of relevant traffic scenarios," in *12th ITS European Congress*, Strasbourg, France, 2017.
- [11] H. Elrofai, J.-P. Paardekooper, E. de Gelder, S. Kalisvaart, and O. Op den Camp, "Streetwise scenario-based safety validation of connected and automated driving," TNO, Tech. Rep., July 2018.
- [12] K. Czarnecki, "Operational world model ontology for automated driving systems-part 1: Road structure," Waterloo Intelligent Systems Engineering Lab (WISE) Report, University of Waterloo, 2018.
- [13] —, "Operational world model ontology for automated driving systems-part 2: Road users, animals, other obstacles, and environmental conditions," *Waterloo Intelligent Systems Engineering Lab (WISE) Report, University of Waterloo*, 2018.
- [14] H. Weber, J. Bock, J. Klimke, C. Roesener, J. Hiller, R. Krajewski, A. Zlocki, and L. Eckstein, "A framework for definition of logical scenarios for safety assurance of automated driving," *Traffic Injury Prevention*, vol. 20, no. sup1, pp. S65–S70, 2019, pMID: 31381437. [Online]. Available: https://doi.org/10.1080/15389588.2019.1630827
- [15] D. Wittmann, C. Wang, and M. Lienkamp, "Definition and identification of system boundaries of highly automated driving," 7. Tagung Fahrerassistenz, 2015.
- [16] E. Thorn, S. Kimmel, and M. Chaka, "A framework for automated driving system testable cases and scenarios," Report No. DOT HS 812 623. National Highway Traffic Safety Administration, Sep. 2018.
- [17] Volvo Trucks Global. (2019) "Vera's first assignment: Volvo Trucks presents an autonomous transport between a logistics centre and port". [Online]. Available: https://www.volvotrucks.com/en-en/news/ press-releases/2019/jun/pressrelease-190613.html
- [18] RISE Research Institute of Sweden. (2019) "Born to drive". [Online]. Available: https://www.viktoria.se/projects/born-to-drive
- [19] P. Koopman and F. Fratrik, "How many operational design domains, objects, and events?" in *Proceedings of AAAI Workshop on Artificial Intelligence Safety*, Honolulu, USA, Jan. 2019.
- [20] NHTSA, "Automated driving systems 2.0 a vision for safety," Transportation Research Part A: Policy and Practice 94, 2018.
- [21] G. Bagschik, T. Menzel, and M. Maurer, "Ontology based scene creation for the development of automated vehicles," in *Proceedings of 2018 IEEE Intelligent Vehicles Symposium (IV)*, Changshu, China, Jun. 2018.
- [22] A. Reschka, J. Böhmer, T. Nothdurft, P. Hecker, B. Lichte, and M. Maurer, "A surveillance and safety system based on performance criteria and functional degradation for an autonomous vehicle," in *Proceedings* of *IEEE Conference on Intelligent Transportation Systems*, Anchorage, USA, Sep. 2012.
- [23] D. Töpfer, J. Spehr, J. Effertz, and C. Stiller, "Efficient road scence understanding for intelligent vehicles using compositional hierarchical models," *IEEE Transactions on Intelligent Transport Systems*, vol. 15, no. 1, pp. 1–10, 2015.
- [24] M. Hörwick and K.-H. Siedersberger, "Strategy and architecture of a safety concept for fully automatic and autonomous driving assistance systems," in *Proceedings of 2010 IEEE Intelligent Vehicles Symposium*, San Diego, USA, Jun. 2010.
- [25] I. Colwell, B. Phan, S. Saleem, R. Salay, and K. Czarnecki, "An automated vehicle safety concept based on runtime restriction of the operational design domain," in *proceedings of 2018 IEEE Intelligent Vehicles Symposium (IV)*, Changshu, China, Jun. 2018.
- [26] M. Törngren, X. Zhang, N. Mohan, M. Becker, L. Svensson, X. Tao, D.-J. Chen, and J. Westman, "Architecting safety supervisors for high levels of automated driving," in 2018 21st International Conference on Intelligent Transportation Systems (ITSC). IEEE, 2018, pp. 1721–1728.