



HAL
open science

Monitoring Operational Design Domain Compliance in Intelligent Vehicles

Thibault Charmet, Véronique Cherfaoui, Javier Ibanez-Guzman, Alexandre Armand

► **To cite this version:**

Thibault Charmet, Véronique Cherfaoui, Javier Ibanez-Guzman, Alexandre Armand. Monitoring Operational Design Domain Compliance in Intelligent Vehicles. 36th IEEE Intelligent Vehicles Symposium (IV 2025), IEEE, Jun 2025, Cluj-Napoca, Romania. pp.754-760, <10.1109/IV64158.2025.11097647>. <hal-05085630>

HAL Id: hal-05085630

<https://hal.science/hal-05085630v1>

Submitted on 26 May 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



HAL Authorization

Monitoring Operational Design Domain Compliance in Intelligent Vehicles

Thibault Charmet^{1,2}, Véronique Cherfaoui¹, Javier Ibanez-Guzman², Alexandre Armand²

Abstract—Advanced driver assistance systems (ADAS) and automated driving functions are becoming integral to modern vehicles. Ensuring their safety and reliability requires validating their operation within well-defined Operational Design Domains (ODD). Monitoring ODD compliance is crucial to determine when these functions can operate safely. This paper presents a systematic approach to ODD monitoring using a formalized, machine-readable ODD description and fuzzy logic. The method evaluates compliance and provides explanations for non-compliance. The approach introduces a two-level hierarchical ODD representation, a membership score quantifying compliance, and an explanation mechanism for identifying the primary factors contributing to non-compliance. The monitoring results are integrated into a Conditional Activation Control System (CACS), which governs function activation based on ODD compliance. The proposed system was implemented within a production vehicle and validated using real-world data, demonstrating its feasibility for deployment. By enabling clear, real-time assessments of ODD adherence, this approach supports safer and more reliable automated driving, promoting user confidence and regulatory compliance.

I. INTRODUCTION

This paper builds on prior research by leveraging a formalized, machine-understandable description of the ODD and utilizing fuzzy logic to assess ODD compliance. It also provides an explanation for detected non-compliances by identifying the primary contributing factors in the ODD description. This enhances the transparency of intelligent vehicle systems usage, a critical aspect for ensuring user confidence in autonomous and highly automated driving technologies.

Several works have addressed the importance of ODD monitoring in the context of intelligent vehicles. For example, Wittmann et al. [1] discuss a methodology to define specific functional system boundaries (i.e., an ODD), Colwell et al. [2] discuss the restriction of the ODD based on the degraded capabilities of the system, while Gyllenhammar et al. [3] emphasize the need for a structured approach to ODD monitoring. Alsayed et al. [4] use machine learning algorithms to monitor the ODD of a 2D laser-based localization algorithm, while Woerter et al. [5] propose an Operational Domain Monitor (ODM) to assess the service quality of a capacity. Schwalb et al. [6] employ a statistical approach to identify risky situations by computing their

level of compliance with fuzzy requirements. Input validation techniques such as anomaly detection and outlier detection are also relevant for ODD monitoring, as they can help identify unexpected or erroneous data that may indicate a departure from the expected operational domain [7].

Some works have proposed formalized descriptions of the ODD. For example, Whiteside et al. [8] and Rohne et al. [9] use domain-specific languages to describe the ODD, while Schwalb et al. [10] propose a two-level language for ODD description. The ASAM OpenODD project [11] introduces a standardized ODD description format.

Existing approaches either focus on monitoring specific systems or subsystems (e.g., perception, localization) using their specificities, or on formalizing the ODD without integrating it into a general monitoring framework. However, to our knowledge, no prior work effectively bridges the gap between ODD formalization and comprehensive ODD monitoring based on such descriptions. This paper aims to help fill this gap by establishing a direct link between the formalized ODD description and the monitoring process, enabling systematic and practical ODD compliance evaluation.

The paper details the systematic approach, starting with the formulation of the ODD using a two-level hierarchical description, enriched with examples related to vehicle automation. A key contribution of our work lies in the methodology for monitoring ODD compliance via a membership score calculation. Additionally, we introduce a novel mechanism for inferring explanations from the membership scores, providing insights into the reasons behind non-compliance. Furthermore, we integrate the results of the monitoring process into the vehicle to conditionally control the system activation through a purpose-built mechanism called the Conditional Activation Control System (CACS). We detail the design and implementation of CACS and demonstrate its effectiveness by integrating it into the data structure of a production-level vehicle. The proposed system is designed to function using real-time data generated by existing vehicle sensors and onboard systems, ensuring feasibility for deployment in production-level prototypes. The remainder of this paper is structured as follows: Section II describes the ODD formulation and the membership score calculation. Section III presents the ODD monitoring methodology and explains how compliance explanations are inferred. Section IV details the experimental setup and results, using data from an experimental vehicle.

*This work is performed within the SIVALab, a common laboratory between Ampere and Heudiasyc (Université de technologie de Compiègne, CNRS), and funded by Ampere and the Association Nationale de la Recherche et de la Technologie.

¹Université de Technologie de Compiègne, CNRS, Heudiasyc laboratory, Compiègne, France name.surname@hds.utc.fr

²Ampère, Guyancourt, France name.surname@ampere.cars

II. ODD DESCRIPTION

This section presents the formal specification of an ODD using the proposed description method. Since no publicly available implementation of ODD descriptions is currently available, we developed a custom description language tailored to our specific use case. The complete specification and implementation of our ODD description language will be made publicly accessible in the future to support the standardization of ODD descriptions and their implementations. The description of the ODD is based on a taxonomy of the elements of the environment, similar to ISO 34503 [12] and other works compared in Charmet et al. [13], ensuring that all relevant attributes are systematically defined and used consistently. The rationale behind splitting the description into two levels is explained, detailing their structures and differences, and explaining how the high-level description is translated into a low-level description.

The purpose of the ODD is to facilitate communication among various stakeholders involved in the design and implementation of autonomous driving systems, including engineers, lawyers, insurers, and authorities. To achieve this, a two-level description of the ODD is necessary. The high-level description is intended for human readers and is designed to be concise and easily understandable. In contrast, the low-level description will be read by programs. Therefore, it is designed to be as uniform and standardized as possible.

The ODD can be represented as a large logical expression made up of sub-expressions, organized in a tree structure. Each leaf of the tree is an atomic expression that imposes constraints on an attribute of the ODD. For example, the formal expression describing a heavy rain condition could be `rain.intensity >= heavy`. Here, `rain.intensity` is an attribute of the ODD defined in a taxonomy or ontology of significant elements of the driving scene, along with other metadata (name, description, measurement unit, etc.). The term `heavy` refers to a categorical preset for rain intensity, corresponding to a range of 7.6 to 50 mm/h [14]. The branches of the tree connect their children using fuzzy logical operators [15]. Consider the example of tree-structure shown in Fig. 2 and described below. The low-level description is in a JSON-based format. However, to limit the size of the example, it is presented in a YAML-like format. Accordingly, only the end of the full taxonomic name of the attributes is provided.

This ODD contains two expressions of depth 1. The first expression is a conjunction of two atomic expressions of depth 2, while the second expression is a single atomic expression. The ODD is a conjunction of these two expressions. Each expression can be composed of multiple sub-expressions. The list of logical operators “ops” is used to combine the sub-expressions. The “not” field is used to negate the expression. The “att” field specifies the attribute of the ODD, the “interval” field specifies the range of values for the attribute, and the “set” field specifies the set of values for the attribute.

This description would be equivalent to the following

```

odd:
  exprs:
    - name: Cannot exceed 25 km/h when
      there are 10 or more entities
      in the scene.
      exprs:
        - att: target_speed
          interval: "[ (25,inf) ]"
          not: false
        - att: entity_number
          interval: "[ (10,inf) ]"
          not: false
      ops: [AND]
      not: true
    - name: Cannot drive when there are
      cyclist or pedestrians in sight.
      exprs:
        - att: vulnerable.in_sight
          set: [cyclists, pedestrians]
          not: false
      ops: []
      not: true
  ops: [AND]
  not: false

```

Fig. 1. Low-level ODD Description in YAML-like format.

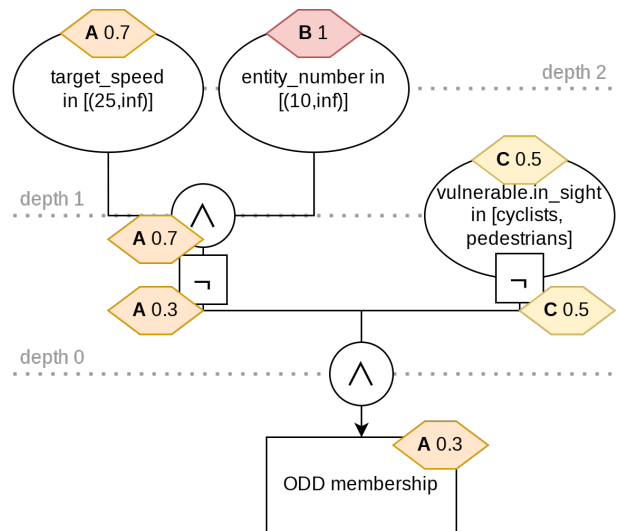


Fig. 2. Logical tree representation of the ODD description. Atomic expressions are represented by the leaves at the top, and the root at the bottom represents the ODD membership. The logical operators are depicted by the nodes using the symbols of Table I. The colored hexagons contain ODD membership scores.

high-level description.

```

forbid target_speed in [25, inf]
and entity_number in [10, inf]
forbid vulnerable.in_sight in [cyclists,
pedestrians]

```

This description is permissive, meaning that everything not explicitly forbidden is required. Thus, a target speed of 20 km/h or the presence of other vulnerable road users in sight would be implicitly accepted.

The description would also be equivalent to the following restrictive description.

```

require target_speed in [0, 25]
  or entity_number in [0, 9]
require vulnerable.in_sight not in
  [cyclists, pedestrians]

```

In restrictive descriptions, any condition not explicitly required is forbidden. For instance, a target speed of 30 km/h when there are 2 entities in the scene would be implicitly forbidden.

In the high-level description, each line or ‘statement’ is equivalent to an expression of depth 1. When an expression is required, it is not only accepted, but also expected to be true. When an expression is forbidden, it is expected to be false¹.

In both description modes, any ODD attribute that is not defined in the ODD description will be ignored. Implicitly forbidding undefined attributes in restrictive descriptions would result in the ODD always evaluating to false (0). This would necessitate defining every ODD attribute in the taxonomy within the ODD description, which would be cumbersome and could lead to unexpected behavior when new attributes are added to the taxonomy. However, during the editing phase, it would be beneficial to warn the user that the ODD is incomplete and that some attributes have not yet been defined.

The ODD can be represented as a logical tree of expressions, making permissive and restrictive descriptions equivalent. The choice between them depends on the use case and the number of constraints to be defined. A permissive description is more concise when there are few restrictions, while a restrictive description is more concise when there are few requirements. At low-level, the difference is that permissive descriptions will have depth 1 expressions negated (the negation at depth 1 in Fig. 2), while restrictive descriptions will have depth 1 expressions as they are. Since all statements should be true, at low-level, they are combined with an *AND* operator at depth 0. This format for describing the ODD will allow for effective monitoring.

III. ODD MONITORING

A. ODD membership score evaluation

ODD monitoring involves observing the attributes of the ODD in the vehicle environment (i.e. the operational domain) and calculating an ODD membership score. Presented in Charret et al. [16], this score is a measure of the compliance of the environment with the ODD. It ranges from 0 to 1, and can be between 0 and 1 when measurements are uncertain. The closer the score is to 1, the more confident we are that the environment is compliant with the ODD. The closer the score is to 0, the more confident we are that the environment is not compliant with the ODD.

Two types of attribute values are considered: continuous attribute values with imprecise measurements and discrete attribute values with confidence scores. Continuous attribute

values with imprecise measurements have measurement uncertainty represented as a normal distribution. For those attributes, the membership function is defined as the sum of the areas under the curve for each sub-interval within the interval specified by the ODD, divided by the total area under the curve in the attribute’s definition interval.

Let $x \sim \mathcal{N}(\mu, \sigma^2)$ represent the value of an attribute, with f_x as its probability density function. The membership function of the expression $\mu_{expr} : \mathbb{R}^2 \rightarrow [0, 1]$ is defined by:

$$\mu_{expr}(x) = \frac{\sum_{i=0}^n \int_{a_i}^{b_i} f_x(t) dt}{\int f_x} \quad (1)$$

With $[a_i, b_i]$ the sub-intervals of the interval given by the ODD.

Let’s illustrate this with an example. Consider the expression “target_speed < 30” which evaluates whether the target speed of the system is less than 30 km/h. If the target speed is exactly 25 km/h, the membership score will be 1. However, if the target speed has a measurement uncertainty of 5 km/h, the membership score will be approximately $\mu_{target_speed}(x) = \int_0^{30} f_{att}(x) dx = 0.84$.

Discrete attribute values with confidence scores denote categorical attributes that can take on a finite number of values, each with a corresponding confidence score. For those attributes, the membership function μ_{expr} is defined as follows:

$$\mu_{expr}(X) = \begin{cases} 0 & \text{if } X = \emptyset \text{ or } C = \emptyset \\ \max\{X[k] \mid k \in C\} & \text{otherwise} \end{cases} \quad (2)$$

where:

- X is a dictionary of confidence values, with keys representing classes and values representing confidence scores.
- C is the set of classes described in the expression.
- $X[k]$ is the confidence value for class k .
- $\{X[k] \mid k \in C\}$ is the set of confidence values for the classes in X that are present in C .

This function represents the maximum classification confidence for each class specified in the expression set.

The expression “vulnerable.in_sight in [cyclists, pedestrians]” evaluates whether cyclists or pedestrians are in sight. Detected road users have both an existence score (indicative of the presence or absence of an object) and a classification score (class probability), and the confidence is considered as the product of these scores [17]. For example, if the system detects a road user with an existence probability of 0.9 and a classification probability of 0.8 as a cyclist and 0.2 as a pedestrian, the confidence values will be 0.72 for cyclists and 0.18 for pedestrians. Therefore, the membership score will be 0.72.

Once each atomic expression in the ODD has been evaluated, the membership score is propagated from the leaves to the root to obtain the overall ODD membership score. See Fig. 2 for an example of the membership score evaluation (the values inside the colored hexagons). This evaluation uses the fuzzy logic operators listed in Table I

¹Compared to Charret et al. [16], the statement keywords have been changed to ‘forbid’ and ‘require’ to better reflect their logical meaning.

| Boolean | Fuzzy | Symbol |
|---------------|--------------|--------------|
| NOT(x) | $1 - x$ | $\neg x$ |
| AND(x, y) | $\min(x, y)$ | $x \wedge y$ |
| OR(x, y) | $\max(x, y)$ | $x \vee y$ |

TABLE I

FUZZY LOGIC OPERATORS SORTED BY PRECEDENCE. \neg HAS PRIORITY OVER \wedge WHICH HAS PRIORITY OVER \vee .

to combine the membership scores of the sub-expressions into the membership score of the parent expression.

B. Explaining ODD membership score

The benefit of using logical rules to describe the ODD is that it allows for explaining why an ODD is not being met. Similar to the evaluation of the membership score, it is possible to determine for each expression which of its sub-expressions contributed most to its score. This expression is referred to as “the explanation” in the following.

When the logical operators used are $\{1 - x, \min(x, y), \max(x, y)\}$, the contribution of each sub-expression to the parent expression can be determined as follows. Let AB be the parent expression of sub-expressions A and B , with their respective memberships denoted as μ_{AB} , μ_A , and μ_B , such that $\mu_A < \mu_B$.

- For the *and* operator, the explanation will be the sub-expression(s) with the smallest membership value(s) (A). This expression(s) can be considered to contribute 100% to the membership of the parent expression, as no infinitesimal change in the membership of other sub-expressions will affect the parent expression’s membership.
- For the *or* operator, the explanation will be the sub-expression(s) with the largest membership value(s) (B).
- For the *not* operator, the explanation will be the single sub-expression.

Given the structure of the ODD, each parent expression contains n sub-expressions and $n - 1$ operators. To evaluate the membership and determine the explanation for each expression accurately, the operations will be performed according to the precedence of the operators, grouping the sub-expressions by their respective operators.

The benefit of this explanation is that it helps identify which expression contributes most to lowering the ODD membership score. This understanding can enable the implementation of preventive or corrective measures to avoid leaving the ODD, such as reducing the target speed. For example, in Fig. 2, the target speed is the principal factor affecting the ODD membership score. By addressing this specific expression (e.g., by reducing the target speed), the overall ODD membership score can be increased from 0.3 to 0.5.

C. Conditional activation control system

The computed membership score can be used to decide when to activate or deactivate a driving system. This decision

algorithm is referred to as the Conditional Activation Control System (CACS). The CACS is based on a hysteresis control system and time-weighted moving average filters.

A hysteresis control system is a type of control mechanism that helps to prevent frequent switching or oscillation around a setpoint [18]. It achieves this by introducing a range or “band” around the setpoint within which no switching occurs. In the CACS, this is equivalent to having two thresholds: a low threshold bl and a high threshold bh . The system is activated when the membership score is greater than bh and deactivated when the membership score is less than bl . This approach helps to limit the state changes of the driving system, preventing it from oscillating around the setpoint.

Additionally, the sliding window introduces “inertia” to the membership score, although at the cost of latency. This is implemented as a time-weighted moving average over n seconds, as detailed in Charmet et al. [16]. In signal processing, a moving average is a type of low-pass finite impulse response filter [19]. It attenuates high-frequency components of the signal by averaging it over a window of length n . This filter is time-weighted to ensure independence from the sampling rate.

At the ODD level, it can act like an accumulator with a threshold on the membership score. For instance, with a sliding window of n seconds, if the membership score is 0 at time t_0 , it will take at least $0.95 \times n$ seconds for the membership score to reach 0.95 or higher, assuming all subsequent values are 1. Generally, with a low threshold bl , a high threshold bh , and a sliding window of n seconds, the minimum time between deactivation and reactivation is $(bh - bl) \times n$ seconds.

IV. EXPERIMENTS AND RESULTS

A. Experimental setup

We used an experimental vehicle to 1) record data and 2) embed the system that was developed. The vehicle is an original Renault Megane Etech that is used by several R&D projects for prototyping and testing next generation innovative ADAS (Fig. 3). We used the original perception system of the vehicle (fusion camera and radar) as well as an automotive GNSS based localization system. All information is available on the vehicle CAN bus, which is interfaced with the ROS2 Humble middle-ware.



Fig. 3. The vehicle used for the experimentation.

The ODD monitoring system is developed as a Python package running inside a ROS 2 node. This setup allows for real-time data processing and integration with other vehicle systems. Although the current implementation is a prototype and not optimized for performance, it operates at a few hertz, making it suitable for real-time applications. The ROS 2 node subscribes to relevant topics, processes the incoming data to evaluate ODD compliance, and publishes the results for use by other systems in the vehicle. The low-level ODD description and attributes taxonomy are loaded from JSON and YAML files, respectively. The system is designed to be easily configurable and adaptable to different scenarios and vehicles.

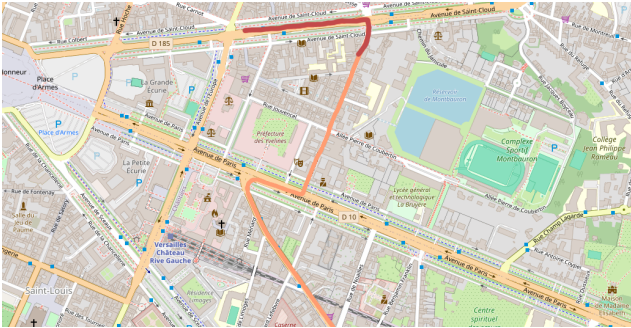


Fig. 4. Map showing full scenario trajectory (Orange) and the final segment (Red) in Versailles. © OpenStreetMap contributors.

The use cases are based on a 300-second recording made in Versailles (Orange trajectory in Fig. 4). The ODD has been monitored on the whole sequence and is displayed in Fig. 7. To better illustrate the method, specific sequences will be detailed in the following section corresponding to the end of the recording (Red trajectory).

B. Conditional activation control system

Let's focus on the final segment of the recording, between 190 and 300 seconds, with the ODD description defined in Fig. 5. This ODD description is detailed enough to demonstrate how the membership score changes in response to various expressions and logical operators. Since the vehicle does not drive autonomously, we will assume that the measured speed corresponds to the target speed of the hypothetical driving system.

Fig. 6 shows the activation and deactivation of the driving system based on the ODD membership score. The hysteresis control system is configured with a low threshold of 0.4 and a high threshold of 0.95 (red lines). The moving average for

```

forbid target_speed in [(25, inf)]
  and entity_number in [(10, inf)]
  or visibility in [(0, 100)]
  or GNSS.service_health in [(0, 0.5)]
forbid vulnerable.in_sight in ['cyclists']
  and vehicles.in_sight in ['large_trucks']
forbid restricted_zones in ['no_adas']
  
```

Fig. 5. High-level ODD description for the use case 1.

each expression is set to 1 seconds, while the sliding window for the overall ODD is set to 0 second. A low threshold of 0.4 allows the driving system to be deactivated when there is enough evidence that the ODD is not met. Conversely, a high threshold of 0.95 permits the driving system to be reactivated when there is strong confidence that the ODD is respected. As the membership score for the ODD evolves (blue line), the hysteresis control system determines whether the driving system should be activated (white background) or deactivated (colored background).

In this segment of the circuit, we approach a junction with a red light, where cyclists, trucks, and numerous other road users are frequently present. Given our ODD, we anticipate that the driving system will be unavailable most of the time. Between 212 and 218 seconds, the vehicle's visibility is artificially reduced to simulate fog, causing the driving system to become temporarily unavailable. Upon entering the intersection, cyclists (at 239 seconds) and trucks (at 246 seconds) are detected. The vehicle exceeds 25 km/h while more than ten objects are detected (at 258 seconds) then slows down (at 283 seconds), but other bicycles and trucks remain present in the scene.

This experiment also teaches us that not all ODD expressions are adapted for ODD monitoring. For instance, expressions related to very dynamic attributes, such as the number of entities in sight, should not be used directly as-is in the ODD description. In Fig. 7 for example, the confident but wavering detection of a truck around 120 seconds might cause an indecisive activation behavior of the driving system. Thus, for this kind of attributes, using smarter expressions, such as average traffic density over a few seconds, would be more appropriate.

This experiment highlights that not all ODD expressions are well-suited for monitoring. For example, expressions involving highly dynamic attributes, such as the number of entities in sight, should not be used directly in the ODD description. In Fig. 7, the confident yet fluctuating detection of a truck around the 120-second mark could lead to indecisive activation behavior of the driving system. For such attributes, it would be more effective to use refined expressions, such as the average traffic density calculated over a few seconds, to ensure smoother and more reliable system behavior.

These results demonstrate the effectiveness of the Conditional Activation Control System in managing the activation and deactivation of the driving system based on the ODD membership score. The system helps to prevent frequent switching and ensures that the driving system is only activated when there is high confidence in ODD compliance.

C. ODD membership explanation

Using the previous scenario, we can explain why the ODD is not respected. In Fig. 6, the background color indicates the expression that most significantly contributes to the ODD non-compliance. Each expression is represented by a different color. When the ODD is respected, the background color is white.

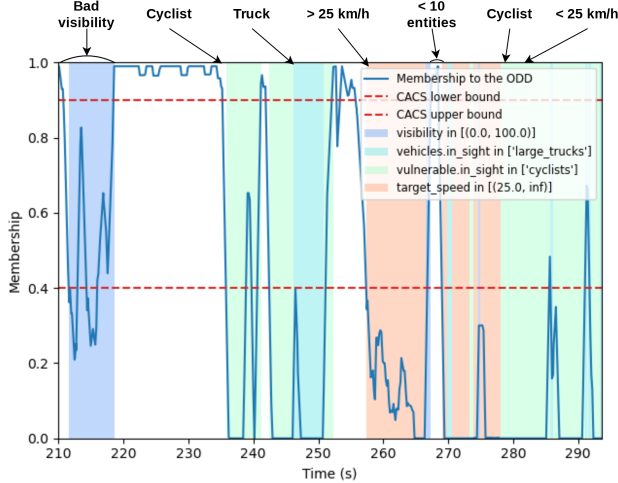


Fig. 6. ODD compliance and ODD membership explanation for the use case 1. Background color represents the expression that contributed most to the non-conformity of the ODD when outside the ODD. White background represent ODD compliance

Fig. 7 illustrates the evolution of the ODD membership score and the corresponding explanations over an extended period. This visualization helps in analyzing the vehicle’s operational domain in relation to its ODD. Notably, between seconds 15 and 33, the vehicle passes near the Versailles women’s prison, which is considered a restricted area for ADAS in this scenario. This example demonstrates the application of geofencing-based ODD constraints.

In this second usecase, the same scenario will be used, but only 2 elements will be considered: visibility and the presence of cyclists. The ODD description defined in Fig. 8 will be used. To illustrate the monitoring properties, we will use an artificially set visibility value, incorporating noise and uncertainty. The visibility is set to 105 ± 5 meters until 212 seconds, and linearly increase to 125 ± 5 meters until 218 seconds. The probability of existence of the detected cyclists is also made more uncertain when there is a detection to reflect the greater perception uncertainty in low visibility conditions. The moving average is set to 1 second at the expression level and 1.5 seconds at the ODD level.

The ODD description specifies that either poor visibility or the presence of cyclists will result in non-compliance with the ODD. Using the $\{min, max, 1 - x\}$ operators, the values of expressions linked by an “OR” generally affect the lowest ODD membership independently. For example, with an expression “A or B”, if $A = 0.2$, no value of $B \leq 0.2$ will affect the ODD membership. This is why other fuzzy logic operators could prove useful, such as $\{x \times y, x + y - xy, 1 - x\}$, where every value has a direct impact on the membership score, and which tends to push the membership score towards 0 or 1.

However, with the $\{min, max, 1 - x\}$ operators, the expressions can still have combined effects on the membership score. For instance, in Fig. 9, if cyclists are detected, the membership score will be low. If the visibility was also poor

before detecting the cyclists, a membership score using an ODD moving average will be even lower because the low visibility was already contributing to lowering the average membership score.

A second combined effect is possible due to the properties of the CACS. Initially, since the visibility is not poor enough, the driving system can remain active. However, as soon as a cyclist is detected, the membership score falls below the deactivation threshold. Even after the cyclist is no longer detected, the driving system will not be reactivated until the visibility has improved to an acceptable level with minimal uncertainty.

V. DISCUSSION

The proposed ODD monitoring system is designed to be flexible and adaptable to various scenarios and vehicles. The use of a two-level ODD description allows for clearer communication among stakeholders while ensuring that the system can be easily integrated into existing vehicle architectures. The use of fuzzy logic operators for evaluating ODD membership scores provides a robust framework for handling uncertainty in sensor measurements.

Thought, using a general ODD monitoring system has limitations. The system relies on the observability of ODD attributes and the correctness of the detections. In cases where sensor data is unreliable or unavailable, the system may not be able to accurately assess ODD compliance. Even though solutions might exist to mitigate this issue, such as monitoring some perception system using different perception systems (which might be explored in future works), ODD monitoring systems will always have some reliance on the perception and situation assessment systems. This means that other works focused on improving perception robustness and reliability remain essential to ensure the overall safety and performance of autonomous driving systems.

VI. CONCLUSION

In this paper, we extended our previous work on ODD monitoring with uncertain observations. We presented a comprehensive approach for monitoring Operational Design Domain (ODD) compliance in intelligent vehicles. We introduced a formalized, machine-readable ODD description method and demonstrated how to evaluate ODD membership scores using fuzzy logic. Additionally, we proposed a Conditional Activation Control System (CACS) to manage the activation and deactivation of driving systems based on ODD compliance.

Our experimental results, using real-world data from a production-level vehicle, validated the practical feasibility of the proposed system. The ODD monitoring system effectively evaluated compliance, provided explanations for non-compliance, and ensured that the driving system was activated only when there was high confidence in ODD compliance. The explanations helped identify the factors contributing to non-compliance, enabling the implementation of preventive or corrective measures.




Fig. 7. ODD membership and out-of-ODD explanation for a 5 minutes scenario, illustrating the behavior of the system over time. Background color represents the expression that contributed most to the non-conformity of the ODD when outside the ODD.

```

forbid visibility in [(0, 100)]
or vulnerable_in_sight in ['cyclists']

```

Fig. 8. High-level ODD description for the use case 2.

REFERENCES

- [1] D. Wittmann, C. Wang, and M. Lienkamp, "Definition and identification of system boundaries of highly automated driving," in *7. Tagung Fahrerassistenzsysteme*, 2015.
- [2] 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 *2018 IEEE Intelligent Vehicles Symposium (IV)*. IEEE, 2018, pp. 1910–1917.
- [3] M. Gyllenhammar, R. Johansson, F. Warg, D. Chen, H.-M. Heyn, M. Sanfridson, J. Söderberg, A. Thorsen, and S. Ursing, "Towards an Operational Design Domain That Supports the Safety Argumentation of an Automated Driving System," in *10th European Congress on Embedded Real Time Systems (ERTS 2020)*, 2020.
- [4] Z. Alsayed, P. Resende, and B. Bradai, "Operational Design Domain Monitoring at Runtime for 2D Laser-Based Localization algorithms," in *2021 20th International Conference on Advanced Robotics (ICAR)*, 2021, pp. 449–454.
- [5] F. Woerter, A. Kreutz, A. Salvi, M. Dreiser, and G. Weiss, "Enhanced System Awareness as Basis for Resilience of Autonomous Vehicles," in *CARS 2021 6th International Workshop on Critical Automotive Applications: Robustness & Safety*, 2021.
- [6] E. Schwalb, A. Richter, and D. Rohne, "Validating Autonomous Behaviors against Partially Specified Ambiguous Requirements," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2022, pp. 1342–1349.
- [7] D. Bogdoll, M. Nitsche, and J. M. Zollner, "Anomaly Detection in Autonomous Driving: A Survey," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. New Orleans, LA, USA: IEEE, Jun. 2022, pp. 4487–4498.
- [8] Iain Whiteside, "ASAM OpenXOntology Proposal Workshop 5 Ontologies and ODDs at Five," 2020.
- [9] D. Rohne, A. Richter, and E. Schwalb, "Implementing ODD as single point of knowledge to support the development of automated driving," in *2022 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2022, pp. 1364–1370.
- [10] E. Schwalb, P. Irvine, X. Zhang, S. Khastgir, and P. Jennings, "A Two-Level Abstraction ODD Definition Language: Part II," in *2021 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2021, pp. 1669–1676.
- [11] ASAM, "OpenODD: Concept Paper," 2021. [Online]. Available: <https://www.asam.net/standards/detail/openodd>
- [12] International Organization for Standardization (ISO), "ISO 34503:2023 : Road Vehicles — Test scenarios for automated driving systems — Specification for operational design domain," 2023.
- [13] T. Charmet, V. Cherfaoui, J. Ibanez-Guzman, and A. Armand, "Overview of the Operational Design Domain Monitoring for Safe Intelligent Vehicle Navigation," in *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, 2023, pp. 5363–5370.
- [14] B. S. Institution, "PAS 1883:2020 Operational design domain (ODD) taxonomy for an automated driving system (ADS) - Specification." 2020.

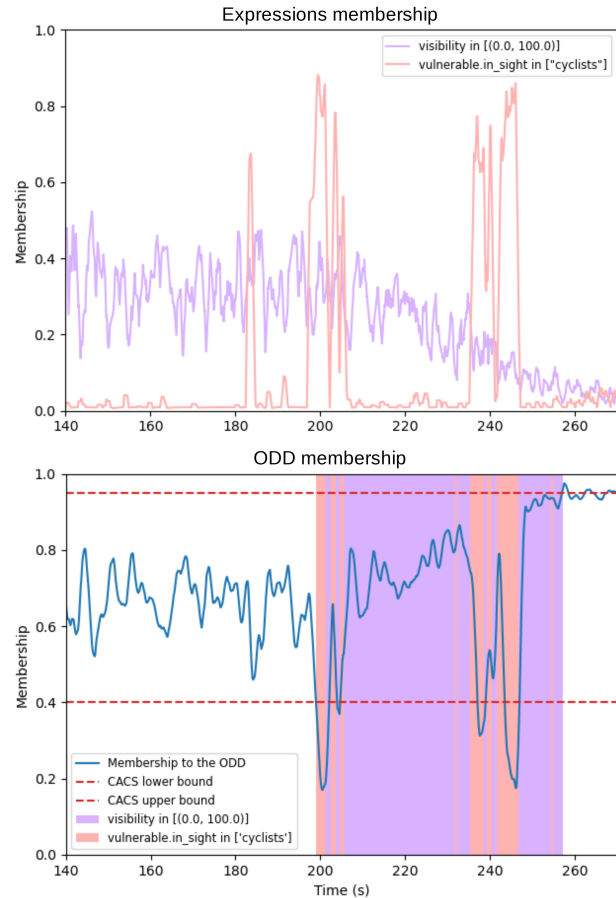


Fig. 9. Top: ODD membership explanation for the use case 2. Background color represents the expression that contributed most to the non-conformity of the ODD when outside the ODD. Bottom: expression membership for the visibility and cyclists detection.

- [15] D. Dubois and H. Prade, *Fuzzy Sets and Systems: Theory and Applications*. Academic Press, 1980.
- [16] T. Charmet, V. Cherfaoui, J. Ibanez-Guzman, and A. Armand, "Operational Design Domain Monitoring with Uncertain Measurements," in *2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC)*, 2024.
- [17] M. Hussain, "YOLO-v1 to YOLO-v8, the Rise of YOLO and Its Complementary Nature toward Digital Manufacturing and Industrial Defect Detection," *Machines*, vol. 11, no. 7, p. 677, Jul. 2023.
- [18] D. Liberzon, *Switching in systems and control*. Springer, 2003, vol. 190.
- [19] S. W. Smith *et al.*, "Moving average filters," *Digital Signal Processing*, pp. 277–284, 2003.