IEEE IAVVC 2025

IEEE International Automated Vehicle Validation Conference
September 30 - October 2, 2025 | Baden-Baden, Germany

TABLE OF CONTENTS
WELCOME MESSAGE
TECHNICAL PAPERS
AUTHOR INDEX

IEEE INTERNATIONAL AUTOMATED VEHICLE VALIDATION CONFERENCE

CONFERENCE PROCEEDINGS

SPONSOR AND ORGANIZER

IEEE

IEEE 75 YEARS INSTRUMENTATION & MEASUREMENT SOCIETY

Part Number: CFP25DY9-ART
ISBN: 979-8-3315-2526-2

Technical Support
Conference Catalysts
Phone: +1 352 872 5544
cdyer@conferencecatalysts.com

2025 IEEE International Automated Vehicle Validation Conference (IAVVC) Proceedings

Additional copies may be ordered from:

IEEE Service Center
445 Hoes Lane
Piscataway, NJ 08855-1331 USA

+1 800 678 IEEE (+1 800 678 4333)
+1 732 981 1393
+1 732 981 9667 (FAX)
email: customer-service@ieee.org

IEEE Catalog Number: CFP25DY9-ART
ISBN: 979-8-3315-2526-2

This full-text paper was peer-reviewed at the direction of IEEE Instrumentation and Measurement Society prior to the acceptance and publication.

From Data-Compliance to Model-Introspection: Challenges in AV Rule Compliance Monitoring

Astrid Rakow*,Gustavo Gil Gasiola†,Dominik Grundt*, Joe Collenette‡, Eike Möhlmann*, Maike Schwammberger†
*Institute of Systems Engineering for Future Mobility
German Aerospace Center (DLR)
Oldenburg, Germany
<firstname.lastname>@dlr.de
†Karlsruhe Institute of Technology
Karlsruhe, Germany
<firstname.lastname>@kit.edu
‡University of Chester
School of Computer and Engineering Sciences

***Abstract*—Autonomous vehicles (AVs) are expected to comply with traffic laws, ensure safety, and provide transparent explanations of their decisions. Achieving these goals requires monitoring architectures that process large volumes of sensor, control, and contextual data. While real-time perception and decision-making are functionally indispensable, storing and using this data for auditing or improvement raises unresolved legal and technical challenges.**

Data protection regulations—such as the GDPR—mandate that personal data processing be limited to what is strictly necessary for specified purposes (Art. 5(1)(b), (c), and (e)). Yet, in practice, what counts as *"necessary"* remains ambiguous. This tension gives rise to the data-justification gap: the lack of systematic methods to determine which logged data is both sufficient to support compliance assessments and minimal under data protection constraints.

At the same time, aligning formalized rules with their legal intent poses a separate but interrelated challenge—the alignment problem. Legal norms are often ambiguous or context-dependent, and existing monitoring frameworks rarely guarantee that formal specifications faithfully reflect legal meaning.

This paper outlines a research agenda for bridging these gaps. We propose an integrated approach combining formal methods, legal reasoning, and runtime monitoring to develop data-justification frameworks. Such frameworks would enable developers to generate interpretable rule formalizations, synthesize minimally sufficient monitors, and justify data collection in a transparent and legally defensible manner.

***Index Terms*—autonomous vehicles, monitoring, GDPR, data minimization, formal methods, privacy**

I. Introduction

Autonomous vehicles (AVs) must meet stringent legal and safety requirements. To demonstrate compliance—whether for certification, liability assessment, or safety assurance—AVs log sensor data, control decisions, and environmental interactions. However, data protection frameworks such as the GDPR (Art. 5(1)(b), (c) and (e)) mandate that the processing of personal data be adequate, relevant and limited to what is necessary for pre-defined purposes. This creates a tension: insufficient logging can hinder the reproducibility of compliance assessments, whereas excessive logging may violate privacy regulations. We refer to this conflict as the *data-justification gap*—the lack of a means to determine which data is minimally sufficient to support compliance claims.

While significant progress has been made in formalising traffic rules and synthesising runtime monitors, existing approaches face two critical limitations. First, they rarely ensure that formal rule specifications faithfully capture the original legal intent, coined as the *alignment problem* [14]. Legal norms are often formulated in ambiguous or context-dependent natural language, which requires interpretative decisions when formalizing them. Without a comparison between the interpretation of the law, the formalized rules and real world behaviour, compliance with the regulations remains contestable. Without alignment between legal interpretation, formal models, and observed behaviour, rule compliance cannot be achieved. Second, current monitoring techniques offer no systematic way to justify the amount or type of data logged in view of data minimisation obligations.

This paper examines these two challenges in tandem. We argue that alignment and data justification are interdependent problems that must be addressed jointly to enable both auditable and legally defensible AV monitoring. To this end, we outline a research agenda that integrates formal methods, legal theory, and data efficiency techniques to develop data-justification frameworks that are compliant by design.

To lay the groundwork, Sect. II introduces the concept of self-monitoring in autonomous vehicles and outlines how

This work has received funding by KASTEL Security Research Labs (SRL), by the V&V4NGC project of the German Aerospace Center (DLR), as part of the project "FuturePorts" (DLR)

sensor data is collected and processed to support behavior evaluation, including rule compliance. Section III examines the legal foundations of data protection, particularly the GDPR, and highlights how its principles constrain monitoring practices. Section IV discusses the formalisation of traffic rules and outlines the alignment challenges that arise when translating legal norms into machine-readable specifications. Section V explores the limits of observability and categorises different types of traffic rules according to the data they require. Section VI presents a research agenda for bridging the data-justification gap, focusing on formal tools and frameworks for minimal, transparent, and legally defensible logging. Finally, Section VII concludes with a reflection on the implications for system design, accountability, and future regulatory alignment.

II. Monitoring in Autonomous Vehicles

Monitoring helps AVs assess whether they operate safely and lawfully. It involves collecting, processing, interpreting, and sometimes storing data about the vehicle's environment, internal state, and decisions. The basic steps of monitoring are illustrated in Fig. 1.

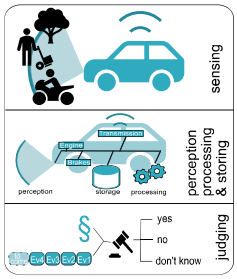

Figure 1: Basic Steps of Monitoring:

(1) collecting data via sensors,

(2) processing of percepts and logging (∼storing),

(3) evaluating compliance of requirements (∼rules)

To limit the scope of the discussion, we focus on self-monitoring, where an AV uses its onboard sensors such as lidar, cameras, or GPS to observe traffic situations. For example, to check compliance with a rule like *"Stop at red lights"*, the AV must detect the traffic light state and match it to its own behaviour. Monitoring extracts relevant sequences of events, such as *"light turned red"* followed by *"vehicle stopped"* and compares them with predefined requirements.

Some rules require monitoring of temporal aspects. For instance, to assess whether the AV has given way to a pedestrian, it must log monitoring data before and after the pedestrian crosses/attempts to cross.

This process can generate a high demand on system resources. Complex scenarios require complex perceptions, interpretation, and storage. AV developers must balance these demands with performance constraints, ensuring that systems remain responsive while collecting enough information to verify legal compliance.

III. Data Protection Issues in Monitoring

Data protection frameworks impose legal constraints on monitoring to the extent that it involves the processing of personal data. In the EU context, personal data is broadly defined and refers to any data relating to an identified or identifiable natural person (Art. 4(1) GDPR). When processing data linked to identifiable individuals (such as the vehicle identification number (VIN), or even technical and usage data of the vehicle and details of journeys made [2]), the monitoring will only be considered lawful if the conditions set out in the GDPR are met.

The regulatory structure of modern data protection frameworks, similar to or inspired by the GDPR, requires a proper justification for each processing of personal data carried out by monitoring activities. Under the principle of purpose limitation (Art. 5(1)(b) GDPR), the purposes of the processing—e.g. to ensure compliance with traffic rules or to provide ad hoc explanations—must be explicitly defined and, in general, communicated to the data subjects before the processing takes place. This is linked to the principles of data minimisation (Art. 5(1)(c) GDPR) and storage limitation (Art. 5(1)(e) GDPR), which require that the data collected and their storage be adequate, relevant and limited to what is necessary to fulfil this predefined purpose. In addition, the controller must be able to demonstrate that at least one of the legal grounds justifying the processing is present (Art. 6(1) GDPR), such as the consent of the data subject or a legal obligation. This regulatory structure therefore presupposes the feasibility of defining the purposes and the data necessary to carry out the monitoring.

Consider a monitoring module to ensure compliance with Art. 21(1)(b) of the Vienna Convention and its implementation in § 26(1) of the German Road Traffic Act (StVo), which regulates the driver's behavior towards pedestrians using or about to use crosswalks (as in the case of Figure 2). According to this rule, the driver must approach the crosswalk at a speed low enough not to endanger the pedestrian and, if necessary, stop to allow the pedestrian to cross. In the context of AV, adherence to this rule requires the collection of sufficient data, through multiple sensors, to automatically identify pedestrians on the sidewalk and their intentions ("about to use the crosswalk"), and to trigger the appropriate anticipatory behavior for that situation.

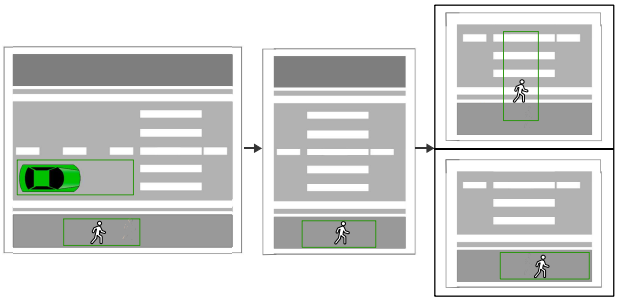

Figure 2: A pedestrian approaches a crossing. They might either pass the crossing and actually cross the road. An AV has to consider both possibilities.

Purpose and legal basis are provided by the traffic rule: the driver must identify pedestrian intentions and respond

accordingly. However, determining the data strictly necessary to monitor this rule remains a major challenge. It is essential to justify what data is collected and stored. Intransparencies—especially when using learning systems like neural networks—often reduce justification to observations such as *"Using this data collection works sufficiently"*. No established techniques exist for justifying the selection of training data. Only recently, a first approach to assess the relevance of knowledge was published by Grundt et al. [7]. In other words, the quantity and quality of collected data affect AV performance, so road safety and data protection—especially the principle of data minimisation [5]—can appear to conflict technically [10]. Alternatives have been proposed to anonymise collected data and thereby avoid data protection requirements. However, anonymisation measures such as automatic blurring of faces and number plates [5] may reduce data quality, making them unsuitable for some applications, such as detecting pedestrian intentions [5].

Let us consider the case that there may be a need to extend the data processed by the monitoring module to perform other functions, such as post-hoc analysis for determining liability, or even to further improve the performance of the automated driving system. *For example, in case a collision with a pedestrian is inevitable, the AV may need to mitigate the impact instead of avoiding the collision.* As the manufacturer might be required to proof that the mitigation, and not the avoidance, increased safety in this exceptional situation (see Annex III, Point 1.4.3.1.2, of the Commission Implementing Regulation (EU) 2022/1426), additional monitoring data must be collected to allow such explanation. The processing of data that is not strictly necessary to comply with the aforementioned transit rules must be justified by another legal ground, such as the public interest in road safety (Art. 6(1)(e) GDPR) or for the legitimate interest of the manufacturer (Art. 6(1)(f) GDPR).

The extent of data that must be processed by monitoring is directly impacted by the level of automation. In the context of a level 4 AV, which can handle all driving tasks but requires human takeover in exceptional situations, the monitoring of the driver's state (awareness, health, etc.) might be considered necessary in order to safely hand over control [10]. Naturally, this task would require collecting more data, including health data (special category data under Art. 9 GDPR), i.e., monitoring eye and head movements, heart or breathing rate, etc. However, it remains unclear whether such a monitoring system complies with data protection requirements and ensures an adequate level of protection.

While data protection laws impose significant legal constraints to protect fundamental rights, the monitoring of AVs presupposes the processing of personal data to realise other important rights and values, such as road safety or even the economic freedom of manufacturers to develop their products. The solution to this apparent conflict is unclear. Data protection rules assume that the manufacturer is able to clearly define which data are *"necessary"* to carry out the legitimate purposes of the monitoring. Nevertheless, as discussed in the next section, the challenges inherent in monitoring rules prevent an *a priori* definition of the necessary data without significantly compromising the efficiency of compliance monitoring.

IV. Formalising Legal Rules for Monitoring

Monitoring legal rules with AVs introduces unique difficulties beyond technical challenges like sensor noise or occlusions. A central challenge is interpreting and formalising traffic rules, which are often written in natural language and contain ambiguous, vague, or context-dependent terms. Without a formal basis, it is difficult to correctly compare observed behaviour with the intended specification. To monitor compliance, traffic rules must first be translated into formal representations. We summarize the main challenges of formalising traffic rules following Westhofen et al. [14].

Formalising traffic rules requires interpreting legal language to resolve ambiguity and make normative choices. Westhofen et al. identify the resulting congruence problem—the challenge of aligning interpretations of the legal rule—as fundamental. It affects both rule formalisation and monitoring (cf. Fig.3). The semiotic triangle [11], shown in Fig. 3, illustrates how legal terms like *"vehicle"* (a symbol) refer to abstract concepts (e.g. cars, trucks) which in turn relate to real-world entities (e.g. the car observed on Monday). Formalising and implementing a monitor each instantiate this triangle, possibly introducing differing interpretations at each stage.

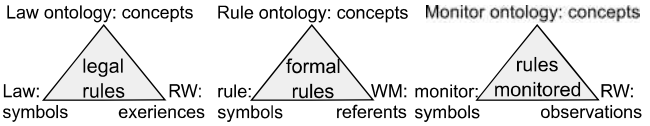

Figure 3: The semiotic triangle [11] relates symbols, concepts and referents. In practice, the triangles of legal rules, formal specifications and monitor interpretations may misalign—risking inconsistent compliance evaluation.

Following [14] thus the following challenges have to be solved: *Alignment*: Different stakeholders–legal experts, engineers, regulators–may interpret key concepts differently, hence shared interpretation must be ensured. *Observability*: Abstract legal concepts must be mapped to perceptible quantities/conditions. *Vagueness*: Many legal terms intentionally require interpretation. *Uncertainty*: Uncertainties (e.g. due to sensor limitations) must be dealt with. *Interrelations and Exceptions*: Legal rules are interrelated by exceptions, dependencies, or priority clauses. *Traceability and Justifiability*: Transparency, certification, and legal accountability require linking system behavior and formalized rules back to their legal origins.

V. Observability and the Data-Minimisation Dilemma

Even if rules are faithfully formalised, enforcing them in practice depends on what an AV can observe. This section discusses the limits of observability and introduces key categories of rules: *externally observable*, *introspective*, *predictive*, and *vague*. These categories pose distinct challenges for compliance monitoring and data minimisation, especially in the presence of non-explainable systems or GDPR constraints. We illustrate how partial observability can compromise both legal accountability and privacy compliance.

(i) *Externally observable rules* can be monitored via sensor data and external behaviour (e.g. *"Stop at red lights"*). (ii) *Introspective rules* require insight into the AV. For example, *"Drive in such a way that you can stop at any time"*, requires a risk assessment based on its own dynamics. (iii) *Predictive rules* require the AV to anticipate others' actions. For instance, *"Give way to pedestrians clearly intending to cross"* demands inferring intent from subtle behaviours like posture or gaze. (iv) *Vague rules* depend on interpretation and local context. *"Let pedestrians pass safely"* where the AV has to tell what *"safely"* means–an assessment that can vary by region. Many rules combine these aspects. The German StVO, for example, uses expressions like *"nach bestem Wissen und Gewissen"* (to the best of one's knowledge and belief) or "vorausschauendes Fahren" (anticipatory driving).

From an engineering perspective, introspective rules may require access to proprietary or non-explainable systems (e.g. neural networks), making external auditing difficult. Predictive rules are prone to be caught between performance and data protection requirements. For example, in order to recognise the intentions of pedestrians at a human level or better, highly sensitive biometric and behavioural data (glances, micro gestures, facial expressions) would have to be captured, which is generally restricted or completely prohibited under data protection law. Vague regulations add to the complexity, as their interpretation evolves with social norms and case law.

Figure 4 illustrates a situation where minimal data logging might fail. Suppose an AV brakes hard because it expects a pedestrian to step into the street. It may have noticed that a shop door is about to open while, at the same time, an e-scooter approaches. Without recording the scooter, the AV 's prediction may appear irrational. To avoid such gaps, AVs may need to store data beyond the immediate event. However, identifying which observations are relevant is particularly difficult—especially for systems whose internal reasoning is not explainable. This example highlights the core tension: storing enough data enables accurate compliance assessments, yet GDPR demands data minimisation. AVs must therefore collect sufficient context while justifying why each recorded element is necessary.

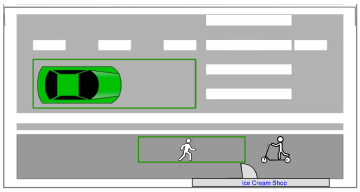

Figure 4: Why did the AV predict that the pedestrian enters the road? Capturing the relevant is a challenge

VI. From Rules to Logs: Toward a Data-Justification Framework

In this section we discuss next steps regarding two great challenges for monitoring traffic rules: the alignment problem (*"Does the rule formalization capture what was the legally intended?"*) and what we call the data justification gap (*"How to justify logging and processing of monitored data?"*)

The alignment problem is specific to rule-compliance monitoring: only properly formalized rules can be reliably monitored. We propose Traffic Sequence Charts (TSCs) as a formalism to address this gap. TSCs enable collaboration among domain experts during rule formalisation and audit, aligning real-world scenarios with their formal specifications. Our research agenda for rule formalisation and monitoring appears in Sect. VI-A. Another important issue is what we call the data justification gap. Manufacturers must justify what data is monitored and logged. We present a research agenda in Sect. VI-B.

A. TSCs: A Visual Rule Specification Formalism

1) The Role of Intuitive Rule Specification in AV Design: TSCs [4] provide a visual yet formal method to specify traffic scenarios. By composing symbolic scenes, they describe temporal evolutions in traffic, as illustrated in Fig. 2. For example, a pedestrian might approach a crossing, pause, and then cross or pass by.

Symbols appearing in each scene are declared in a symbol dictionary. Each symbol is linked to a formal concept, which in turn refers to world model objects (*referents*), thereby capturing the structure of the semiotic triangle [11]. For instance, the symbol *pedestrian* is associated with an object of the world model that moves.

This structure enables translation of TSCs into temporal logic formulas, from which runtime monitors can be synthesized [6], [13]. These monitors are aligned with the formal rules—they use the same symbols and ontology—and can be deployed during system development or operation to assess rule compliance. Figure 5 summarizes the TSC formalism.

Figure 3 illustrated how the interpretations of legal, formal, and monitored rules can diverge conceptually. Building on this, Figure 6 shows how TSCs help align these interpretations through shared symbols and ontologies, reducing ambiguity and improving traceability. While legal

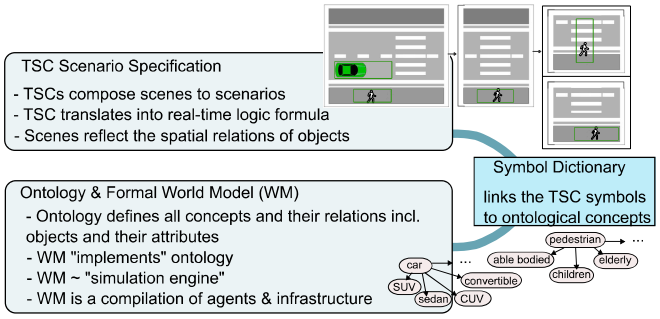

Figure 5: The TSC formalism. Scenes represent symbolic configurations of traffic. Rules are formalized visually and mapped to temporal logic monitors.

rules are derived from human-centered concepts, formal rules result from interdisciplinary interpretation, and monitors instantiate these concepts from sensor data. This alignment relies on a function obs_{WM} that maps real-world observations to instances of the formal world model [13, p. 13].

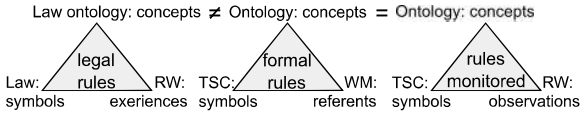

Figure 6: Monitoring TSCs: The semiotic triangles of legal, formalized, and monitored rules must be aligned. TSCs enable this alignment through shared symbols and ontology, supporting consistency across legal interpretation, specification, and observation [11], [14].

Because rule formalization and monitoring operate on the same symbolic foundations, feedback loops can be established: when monitoring exposes implausible system behavior (e.g. a pedestrian moving faster than expected), it may indicate gaps in the formal model. Such insights can improve rule formalization.

TSCs are particularly well suited for interdisciplinary auditing. Their visual representation aids communication between legal, engineering, and regulatory experts. TSCs allow users to: (p:a) visualize formal rule specifications intuitively, fostering cross-domain discussion; (p:b) generate monitors grounded in real-world observations, anchoring formalization in sensor data; (p:c) visualize deviations between formalization and observed behavior for post-hoc audits.

This makes TSCs a promising bridge between stakeholder comprehensibility and formal verifiability. Auditing simulated or real-world scenarios using TSCs supports alignment between the semiotic triangles of formalized and monitored rules and those held by auditing stakeholders.

2) Extending TSC Expressiveness and Monitor Synthesis: TSCs have already been used successfully to formalize various traffic rules [1]. Nevertheless, some aspects remain difficult to express, and extensions to the TSC formalism are being developed [3]. We briefly summarize that discussion here.

Current TSCs can already capture compliant and non-compliant behaviors *(e.g., maintaining a safety distance or obeying speed limits)*, but they lack explicit support for *obligations (O), permissions (P)* and *prohibitions (F)*. Extending TSCs with deontic modalities—by tagging scene subgraphs—will allow rules to be captured more accurately [15]. *For example, Fig. 7 illustrates how deontic TSCs could be used to capture the StVO's rules of main beam usage. Main beam may be used when visibility it requires and no one is dazzled (P); Drivers should switch on main beam in poor visibility to enhance safety (O); Main beam must not be used if it dazzles other road users (F).*

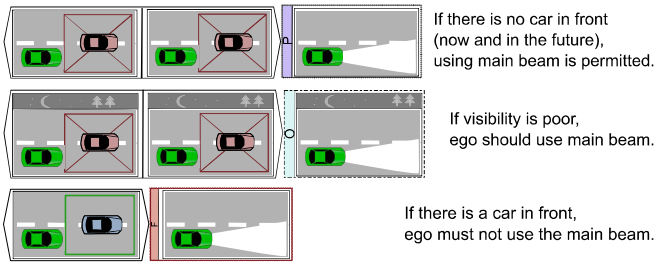

Figure 7: Sketch of deontic TSCs. Each is an "if–then" scenario chart: the (history|future)-premise is inside an elongated hexagon, and consequences carry deontic tags.

Currently, sequential TSC specifications support runtime monitoring via derived monitors [6], [13]. With appropriate sensor-to-attribute mappings, these monitors verify that observed behaviors conforms to the TSC. Planned next steps include: synthesizing runtime monitors for all TSC language features (e.g., all chart structures, time pins), developing a schema for orchestrating multiple runtime monitors, application of monitoring using automotive-grade hardware.

B. The Data–Justification Challenge

As discussed in Sect. III, the GDPR requires that personal data collection be limited to what is necessary for the stated purpose. In AVs, primary purposes include *real-time compliance monitoring* and triggering corrective actions, *post-hoc analysis* for liability and (safety) improvements and *continuous learning* of ML components. Yet, *"necessary"* remains legally qualitative, lacking an algorithmic definition. Collecting too much violates privacy while collecting to little risks functional deficiencies.

As argued in Sect. IV, for the above purposes we rely on a runtime monitor m derived from formal rules R that logs data. Let l denote the recorded log. The GDPR requires that the log contains the fewest recorded entries but allows logging of the data that is necessary for a monitor to produce its verdict, $m(l_{nec}) = m(l) \quad \wedge \quad records(l_{nec})$ *minimal*. In order to alleviate the

manufacturer's burden of justifying that they actually log the necessary data for judging compliance with rules R, approaches are required that log only the relevant records.

We believe that closing the data-justification gap hinges on answering a few core questions that span legal interpretation, formal methods, and systems design. In particular, any viable justification function $J(m, l) = l_{nec}$ must, for each runtime monitor m, pinpoint exactly which entries $l_{nec} \subseteq l$ are both *(a)* sufficient to reproduce m's verdict, and *(b)* minimal under GDPR-style data-minimization. Concretely, we identify the following interlocking research thrusts:

(1) *Standardized Scenario Recognition, Ontology & Performance Baseline:* Establish an industry- and regulator-endorsed standard that defines: *(a)* A core ontology of "must-recognize" entities and events *(e.g., pedestrians at crosswalks, cyclists overtaking, emergency vehicles). (b)* Quantitative sensing performance targets (e.g., *"$\geq$ 95% recall for pedestrian-in-crosswalk at 30 m under daylight"*). *(c)* A governance process for evolving the standard as sensor and AI capabilities advance. Anchoring the *"state of practice"* in a shared standard enables manufacturers to defend minimal logging schemes by conformance rather than proprietary claim.

(2) *Legal Semantics of Relevance and Minimality:* Precisely formalize the notion of "relevance" under Art. 5(1)(b), (c), and (e) GDPR: *(a)* Distinguish between raw and pre-processed data, retention periods, and levels of anonymization (e.g., pseudonymization, aggregation, noise addition). *(b)* Define purpose-specific criteria—for example, differentiating safety monitoring, post-hoc analysis, and continuous learning.

(3) *Log-Schema Synthesis and Certification (a)* Given a monitor m_R, synthesise a minimal *log schema* specifying which sensor streams and internal events reproduce all verdicts. *(b)* Provide a formal proof of sufficiency and GDPR-purpose minimality—i.e. that only personal-data elements strictly necessary for compliance are retained.

We envision employing minimal-witness and unsatisfiable-core methods [9] from temporal logic, enriched with relevance analysis [12] to pinpoint privacy-critical attributes per scenario. Since true minimality is NP-hard [8] and undecidable in unbounded time, we envision performing this calculation offline via a catalog of standard scenarios (which can also be used for certification). At runtime, the AV identifies the active scenario and executes a lightweight monitoring on its pre-defined, scenario specific data fields.

Answering these thrusts will yield a holistic *data-justification framework*, that prescribes not only *what* an AV must monitor but also *why* each recorded datum is both necessary and no more than necessary.

VII. Conclusion

Monitoring is vital for ensuring legal and safe behavior in autonomous vehicles. Yet current methods fail to ensure that formalized rules are aligned with legal intent, and it remains unclear what data must be stored to justify compliance without violating privacy regulations.

We identify two key gaps: the *alignment problem*, concerning the fidelity of formal rule representations, and the *data-justification gap*, concerning which data is necessary and sufficient for defensible monitoring. Bridging these gaps requires interdisciplinary tools that link legal norms, formal semantics, and real-world observability.

We proposed a research agenda centered on rule formalizations, minimal data logging strategies, and auditable justification frameworks. These are essential to balance safety, transparency, and data minimization—and to enable accountable, trustworthy AV systems.

Acknowledgement All sections of this work benefited from language and structuring feedback using the ChatGPT system (OpenAI, 2025).

References

[1] Jan S. Becker. A consistency analysis method for traffic sequence charts. In *VEHITS24 Doctoral Consortium*, Mai 2024.

[2] European Data Protection Board. Guidelines 1/2020 on processing personal data in the context of connected vehicles and mobility related applications, 2020.

[3] P. Borchers, W. Hagemann, D. Grundt, T. Werner, and J. Müller. Using Traffic Sequence Charts for Knowledge Formalization and AI-Application. In Kohei Arai, editor, *Intelligent Systems and Applications*, pages 198–220. Springer, 2024.

[4] W. Damm, S. Kemper, E. Möhlmann, T. Peikenkamp, and A. Rakow. Using traffic sequence charts for the development of havs. In *ERTS 2018*, Toulouse, France, 2018.

[5] David Fernández Llorca and Emilia Gómez. Trustworthy autonomous vehicles. *Publications Office of the European Union, Luxembourg,, EUR*, 30942, 2021.

[6] D. Grundt, A. Köhne, I. Saxena, R. Stemmer, B. Westphal, et al. Towards runtime monitoring of complex system requirements for autonomous driving functions. *EPTCS*, 371:53–61, 2022.

[7] D. Grundt, A. Rakow, P. Borchers, and E. Möhlmann. What does ai need to know to drive: Testing relevance of knowledge. *Science of Computer Programming*, 244:103297, 2025.

[8] C.-M. Li, Z. Zhu, F. Manyà, and L. Simon. Minimum satisfiability and its applications. In *Proc. of Conf. on Artificial Intelligence*, page 605–610. AAAI Press, 2011.

[9] Mark H. Liffiton and Karem A. Sakallah. Algorithms for computing minimal unsatisfiable subsets of constraints. *Journal of Automated Reasoning*, 40(1):1–33, Jan 2008.

[10] Trix Mulder and Nynke E Vellinga. Exploring data protection challenges of automated driving. *Computer Law & Security Review*, 40:105530, 2021.

[11] C.K. Ogden and I.A. Richards. *The meaning of meaning: A study of the influence of thought and of the science of symbolism.* Harcourt, Brace & World, 1923.

[12] Astrid Rakow. A Notion of Relevance for Safety Critical Autonomous Systems. In *Engineering Safe and TrustworthyCyber Physical Systems, ESTCPS*, volume 15471 of *LNCS*. Springer.

[13] R. Stemmer, I. Saxena, L. Panneke, D. Grundt, A. Austel, E. Möhlmann, and B. Westphal. Runtime monitoring of complex scenario-based requirements for autonomous driving functions. *Science of Computer Programming*, 244:103301, 2025.

[14] L. Westhofen, I. Stierand, J. S. Becker, E. Möhlmann, and W. Hagemann. Towards a congruent interpretation of traffic rules for automated driving - experiences and challenges. In G. Borges, K. Satoh, and E. Schweighofer, editors, *Proc. of LN2FR 2022*, pages 8–21, 2022.

[15] R. J. Wieringa and J.-J. Ch. Meyer. *Applications of deontic logic in computer science: a concise overview*, page 17–40. John Wiley & Sons, Inc., USA, 1994.