9

Towards Automated Liability Determination for Autonomous Vehicles in Road Accidents

Rohit Bohara[1], Omkar Joglekar[1], Mirko Ross[1], and Rob van Kranenburg[2]

[1]asvin GmbH, Germany
[2]Martel Innovate, Switzerland

Abstract

The rapid deployment of autonomous and semi-autonomous vehicles introduces fundamental challenges in liability assessment, as existing legal and regulatory frameworks are primarily designed for human drivers and do not directly apply to autonomous decision-making systems. Traditional investigation methods are manual, slow, subjective, and insufficient to capture the complex interactions among vehicles, infrastructure, and evolving mobility ecosystems. This work proposes an automated liability assessment framework that integrates edge intelligence, machine learning technology and graph theory to address the gaps and objectively and transparently assign responsibility in self-driving vehicle accident scenarios. Sensor data from vehicles, weather and roadside units is processed locally on embedded platforms using lightweight AI models optimized for real-time inference and low power consumption. The framework constructs event identity graphs to represent causal relationships among accident entities, factors, cross-referenced with insurance codes, traffic regulations, and ethical policies. The proposed framework integrates mechanisms that inherently support transparency, fairness, and explainability throughout the decision-making process. By rethinking liability assessment for autonomous vehicles, this approach reduces investigation delays, enhances transparency, and supports scalable, secure, and ethically aligned decision-making in the era of self-driving vehicles.

Keywords: Autonomous vehicle, Liability assessment, Incident Recording, Event Identity Graph.

9.1 Introduction

Road traffic injuries continue to represent one of the most critical global public health and socio-economic challenge of the modern era. Despite technological advancements in vehicle safety and the implementation of road safety policies, the absolute number of fatalities resulting from road crashes has continued to rise worldwide. According to [1], approximately 1.35 million people died in road traffic accidents in 2016, making road crashes a more significant cause of mortality than HIV/AIDS, tuberculosis, or diarrhoeal diseases. In Europe, 25,150 fatalities and approximately 135,000 serious injuries were recorded on roads in 2018 [2]. Beyond the humanitarian dimension, the economic burden of road crashes within the EU is estimated at €280 billion annually, equivalent to roughly 2% of total GDP, representing a significant drain on societal and economic resources. In 2022, approximately 891,831 reported road crashes resulted in 20,634 fatalities and 1,14,189 injured people in the EU member state as illustrated in Figure 9.1 [3].

The persistent burden of road traffic injuries underscores the need for transformative approaches to mobility safety. One of the most promising developments in this regard is the emergence of autonomous and semi-autonomous vehicle (AV) technologies, which leverage advanced sensors, machine learning algorithms, and real-time decision-making systems to

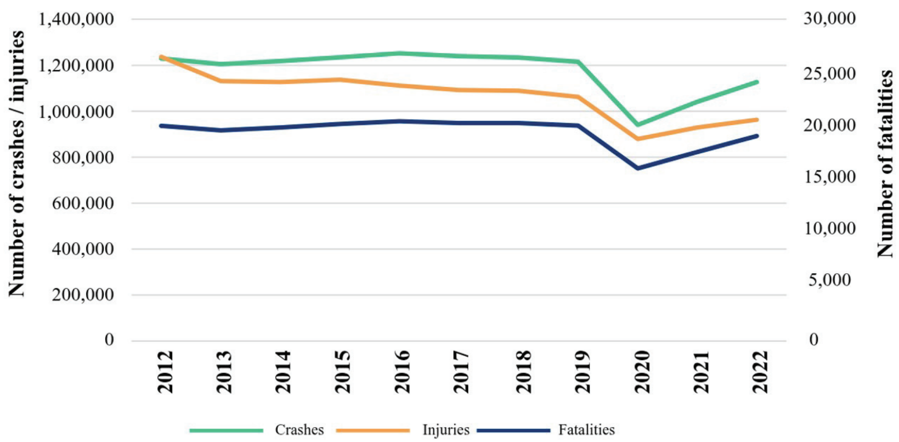

Figure 9.1 Number of road crashes, fatalities and injured people in the EU [3].

enhance driving performance and reduce human error. Human factors, such as distraction, fatigue, impaired judgment, reaction delay, alcohol, and drugs contribute to more than 90% traffic accidents [4–6], making them a primary target for intervention. AVs, by continuously monitoring the driving environment, predicting potential hazards, and executing precise control manoeuvres, have the potential to significantly mitigate collisions attributable to driver mistakes. Early studies and pilot deployments suggest that even partial automation can improve situational awareness, optimize vehicle spacing, and reduce collision severity, thereby lowering both fatalities and serious injuries [7–9].

Beyond improving road safety, AV technologies offer significant societal and economic benefits by transforming the way people and goods move. By optimizing traffic flow and enabling coordinated vehicle interactions, AVs can reduce congestion, shorten travel times, and increase overall productivity in urban and intercity transport networks. They also enhance mobility access for populations traditionally limited by age, disability, or the inability to obtain a driver's license, including the elderly, teenagers and individuals with certain medical conditions. By combining enhanced safety with improved accessibility and efficiency, AVs have the potential to create more inclusive, resilient, and productive transportation systems.

Despite their potential to improve safety and mobility, AVs introduce complex and unresolved challenges in liability and accountability assessment. Traditional legal and regulatory frameworks are designed for human drivers and do not directly apply to software-based decision-making. In accidents involving AVs, fault may arise from multiple sources, including vehicle manufacturers, software developers, fleet operators, infrastructure managers, or even end users, making causal attribution highly complex. Current approaches are manual, slow, subjective, and unable to capture the interactions among vehicles, infrastructure, and adaptive mobility systems. These challenges are compounded by privacy concerns and the proprietary nature of AV data, which can limit access to crucial evidence for liability determination.

The contribution of this work is a conceptual framework for automated liability assessment in accidents involving AVs. The framework integrates three strands of research that are rarely combined in this domain. First, it employs lightweight artificial intelligence models on embedded platforms to enable real-time, low-power processing of vehicle and roadside sensor data at the edge. Second, it employs graph theory to capture causal relationships of entities involved and connect them to insurance codes, traffic regulations, and ethical guidelines. Finally, the framework embeds fairness auditing to systematically evaluate and mitigate bias in liability outcomes, keeping transparency and explainability at the core. Together these elements form a structured, scalable, and ethically grounded approach to liability assessment that addresses both the technical limitations of traditional investigation methods for software-driven vehicles.

9.2 Background and Related Work

Autonomous vehicles represent a transformative advancement in modern transportation, combining sophisticated sensing, perception, and decision-making systems to operate with minimal or no human intervention. The Society of Automotive Engineers (SAE) classifies vehicle automation into six levels (0-5), starting from no automation (Level 0), driver assistance (Level 1), partial automation (Level 2), conditional automation (Level 3), high automation (Level 4) to full automation without human oversight (Level 5) [10].

Levels 0-2, associated with Advanced Driver Assistance Systems (ADAS), are categorized as driver support features in which the human driver retains full responsibility for monitoring the driving environment and maintaining vehicle control. In contrast, Levels 3–5, representing Automated Driving Systems (ADS), encompass automated driving features capable of performing the entire dynamic driving task under specific (Level 3–4) or all (Level 5) conditions. This classification framework underscores the technological evolution from assistive automation (using ADAS) to high and full automation (using ADS), reflecting a fundamental shift toward autonomous mobility [10, 11].

AV technologies rely on a combination of heterogeneous sensors, including light detection and ranging (LiDAR), cameras, ultrasonic sensors, Global Positioning System (GPS) etc, to perceive the vehicle's surrounding environment [12]. Artificial intelligence (AI) lies at the core of processing these data streams: sensor fusion techniques, powered by machine learning models, integrate multiple inputs to generate robust and accurate representations of the driving environment [13]. Perception and prediction systems, typically based on deep learning and computer vision algorithms, not only detect and classify objects such as vehicles, pedestrians, and traffic signals but also predict the behaviour and intent of surrounding agents. AI enables vehicles to anticipate the actions of other drivers, pedestrians, or cyclists, and supports motion planning algorithms that determine safe and efficient trajectories under diverse weather, lighting, and traffic conditions. AI empowers AVs to perceive their environment, predict events, plan actions, and control the vehicle. Vehicle-to-Everything (V2X) communication further enhances situational awareness by allowing vehicles to exchange information with other vehicles, roadside infrastructure, and cloud services. This connectivity supports coordinated manoeuvres, real-time traffic management, and early hazard detection, which are particularly critical in urban and high-density traffic environments.

Traditional approaches to road accident analysis rely primarily on post-incident investigation, including police reports, witness statements, and expert reconstruction of crash events [14, 15]. These standards say that all possibilities for clarifying the course of a traffic accident must be exploited, including use of digital evidence from vehicles, event data recorders, smartphones, etc. Also, when the police investigate an accident, they must reserve traces at the accident scene and secure evidence at the scene: witness data, spontaneous statements from involved parties, photographs of the scene, the positions of vehicles, etc [15]. Each report so submitted should include, at a minimum, the following information relating to the crash: location, date, time, identification of drivers, owners, pedestrians, passengers, motor vehicles, direction of travel each motor vehicle, a narrative description of the events and circumstances leading up to the time of the crash and immediately after the crash [14].

While these methods provide valuable insights, they are often time-consuming, subjective, and limited in capturing the complex interactions among vehicles, road infrastructure, environmental conditions, and human behaviour. Accident reconstruction frequently depends on manual interpretation of physical evidence, which may be incomplete or degraded, leading to delayed or inconsistent assessments. Moreover, conventional analyses are typically designed for scenarios involving human drivers, and are ill-suited for multi-agent, automated, and connected traffic environments.

Road traffic crash analysis literature can be broadly divided into pre-crash predictive models and post-crash severity models, both offering valuable insights for safety and liability assessment. Pre-crash models estimate the likelihood and potential severity of crashes using traffic, environmental, and driver-related factors. [16] showed that traffic congestion significantly influences crash severity, highlighting how human driver behaviour interacts with traffic conditions. It is demonstrated in [17] that interactions between vehicles affect severity outcomes through a copula-based model, emphasizing the interdependence of driver actions. [18] found that real-time weather conditions significantly impact freeway crash risk, and [19] successfully predicted motorcycle crash outcomes using multivariate Bayesian models. While these studies are effective for forecasting crash likelihood and risk exposure, they focus on human-driven scenarios and cannot fully account for autonomous system interventions, algorithmic decision-making, or control handovers, which are critical in AV environments. Thus, these pre-crash models cannot alone determine actual liability in post-crash scenarios involving autonomous or mixed-traffic vehicles.

Post-crash models provide evidence of actual crash dynamics, occupant injuries, and environmental impacts, essential for liability determination. [20] linked kinetic energy dissipation to injury severity, illustrating how crash physics predict harm. Vehicle and crash characteristics strongly influence injury outcomes in side-impact collisions [21], while [22] demonstrated that fixed roadside objects exacerbate injury severity. Correlation was established between the extent of vehicle damage and occupant injury in head-on crashes in [23]. While these models accurately capture outcomes in human-driven crashes, they do not account for AV-specific factors such as automated braking, emergency manoeuvres, or interaction with mixed human-autonomous traffic. Traditional post-crash severity models, such as [20] and [21], rely on detailed vehicle dynamics, occupant injury data, and environmental factors to reconstruct crashes and assess liability. In contrast, the model in [24] analyses statistical trends in crashes and AV disengagements, making it a macro-level post-crash analysis rather than a micro-level reconstruction of individual crash mechanics.

The transition to autonomous and semi-autonomous vehicles further complicates liability determination. In AV-related accidents, fault may arise from a combination of factors, including vehicle software, sensor performance, manufacturer design choices, operator intervention, or infrastructure deficiencies. Current regulatory and legal frameworks are largely unprepared to address these distributed sources of responsibility, creating ambiguity in fault attribution and insurance claims. In addition, the increasing volume of sensor and telematics data generated by AVs introduces both opportunities and challenges: while these datasets can provide high-resolution evidence of accident dynamics, they also raise privacy, security, and data integration issues. These limitations underscore the need for automated, data-driven frameworks that can systematically analyse accident events, incorporate causal reasoning, and provide transparent and equitable liability assessments in the context of autonomous mobility.

9.3 Event Identity Graph Framework

The paper proposes the Event Identity Graph Framework (EIGF) which provides an automated, transparent, and fairness-aware approach to liability assessment in autonomous vehicle accident scenarios. It integrates edge intelligence, machine learning, and causal graph modelling to process distributed sensor data securely and interpret complex interactions among vehicles, infrastructure, and environmental factors. The framework is designed to capture the multimodal and dynamic nature of AV ecosystems, where responsibility can no longer be attributed solely to human drivers but must be inferred from a network of interconnected agents and decision-making systems.

The EIGF operates through four primary steps: (i) event context data ingestion and pre-processing at edge; (ii) Event Identity Graph (EIG) construction using graph theory, which represents causal relationships among entities (vehicles, humans, objects, regulators, insurer, police etc.) and contextual factors (vehicle condition, telemetry, environment conditions, insurance codes, traffic rule) and (iii) liability assessment and report generation using graph theory and large language model (LLM) and (iv) fairness, transparency and ethicality assessment. The proposed EIGF models accident analysis and liability assessment as a structured flow of information from data acquisition to explainable decision outcomes. Figure 9.2 illustrates the overall architecture of the proposed framework, organized according to the classical **input–processing–output** paradigm. The **input layer** represents multimodal data sources, including sensor streams, contextual metadata, and regulatory texts. The **processing layer** captures causal reasoning and entity-relationship modelling within the event graph, while the **output layer** delivers structured liability assessments, interpretability reports, and fairness evaluations.

An Event Identity Graph is a directed graph, illustrated in Figure 9.3, representing events and the entities involved, functioning as an advanced, machine-readable equivalent of a traditional crash report. Unlike conventional reports that primarily describe incidents in narrative form, the EIG captures the full context of a crash, including the vehicles, drivers, sensors, environmental conditions, and any preceding or resulting events. The EIG not only facilitates the explanation of the crash by making explicit how different factors contributed to the incident but also provides a foundation for privacy-preserving, and ethically aligned systematic liability analysis, bridging the gap between technical performance and legal accountability in autonomous mobility. Moreover, its machine-interpretable format ensures interoperability

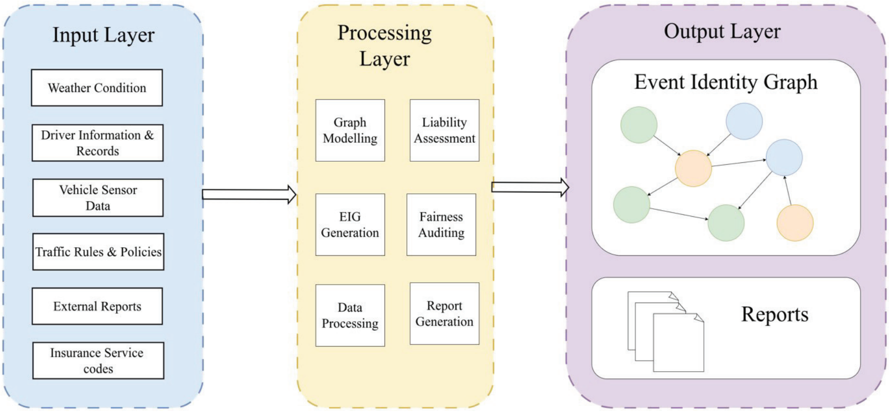

Figure 9.2 Event Identity Graph Framework Architecture.

with AI-based assessment tools, predictive analytics, and fairness auditing frameworks, making it an essential component in modern crash investigation and liability determination. The subsequent sub-sections elaborate on the fundamental processes and architectural elements that constitute the proposed framework.

9.3.1 Data Acquisition and Preprocessing

Accurate and reliable data acquisition forms the foundation of the EIGF, as liability assessment in AV accidents depends on the integrity and diversity of multimodal information. The input layer of the framework acquires and pre-process heterogeneous and distributed data streams originating from vehicles, infrastructure, and Roadside Units (RSUs) to create a unified representation of an event. The data includes onboard vehicle sensors data (e.g., LiDAR, camera, GPS), vehicle telemetry data (speed, braking force, steering angle, system state), human-related factors (driver profile, reaction time), traffic flow from RSUs. Supplementary data such as traffic signal states, environment, weather conditions, traffic laws, insurance codes, manufacturer guidelines and map-based infrastructure information are also integrated to provide a comprehensive situational view.

Each input is temporally stamped and semantically categorized to ensure accurate event reconstruction. The preprocessing pipeline includes data validation, temporal synchronization, across heterogeneous sensors, noise reduction and filtering for signal integrity, and feature extraction to derive semantically meaningful representations such as vehicle trajectories, relative

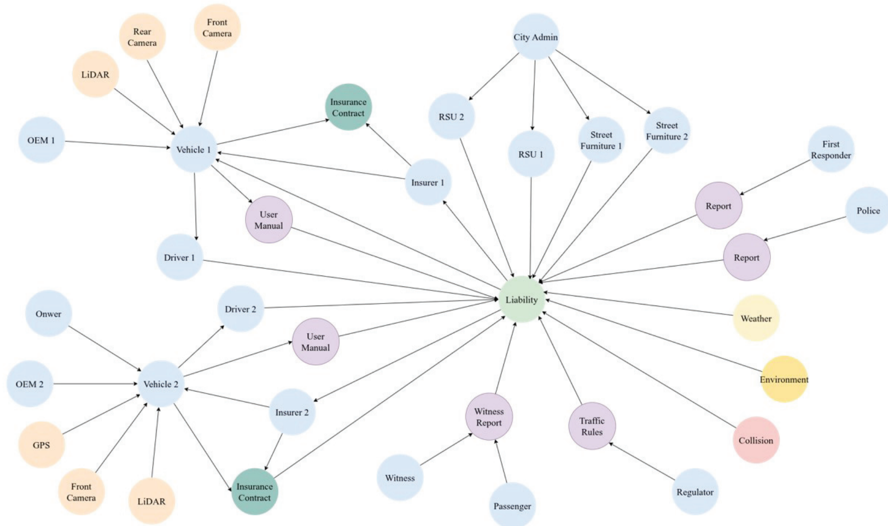

Figure 9.3 Entity Identity Graph.

distances and context tagging before being structured into a graph-compatible schema for downstream analysis. This multi-source approach ensures that both micro-level driving behaviours and macro-level environmental interactions are captured, enabling holistic and high-fidelity reconstruction of event with causal reasoning and supporting granular liability attribution.

9.3.2 Event Identity Graph Construction

An EIG serves as the core analytical representation within the proposed framework, translating complex accident scenarios into structured causal relationships. Each graph as shown in Figure 9.3, is composed of nodes representing entities such as vehicles, drivers, passenger, infrastructure components, environmental conditions, and regulators, and edges denoting the interactions and dependencies among them. These relationships capture both temporal (e.g., sequence of braking or acceleration events) and causal (e.g., vehicle A's manoeuvre causing vehicle B's reaction) connections, allowing for an interpretable reconstruction of the event chain. Machine readable JSON and XML formats of an EIG can be found at [25]. Data streams collected at the edge layer are abstracted into event descriptors and contextual attributes, which are then mapped to graph components through a combination of rule-based logic, temporal correlation, and causal discovery algorithms.

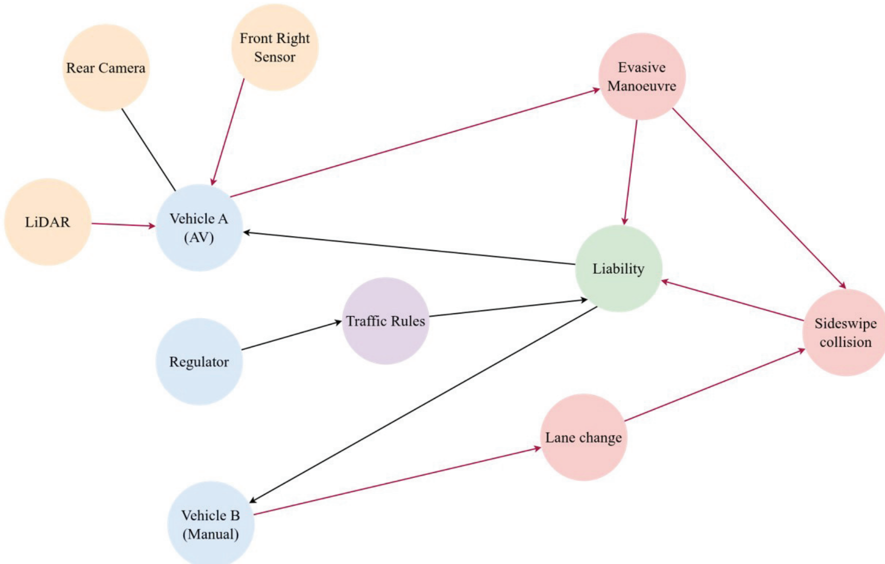

Figure 9.4 Entity Identity Graph Example Walkthrough.

Once constructed, the EIG becomes a dynamic knowledge graph capable of representing both micro-level interactions and macro-level contextual factors that influence accident outcomes. By encoding cause-effect relationships, the EIG provides a foundation for explainable reasoning, enabling the system to trace back the origin of each contributing factor and assess its relative significance. The graph structure also facilitates multi-source integration, linking sensor data with regulatory databases and insurance codes for comprehensive liability assessment. This approach ensures that the resulting interpretations are not only data-driven but also context-aware and legally grounded, bridging the technical and normative dimensions of AV accountability.

Let us explore this further with help of a simple example scenarios where a manually driven vehicle (Vehicle B) performs an unsignalled lane change into the path of an AV (Vehicle A). The AV detects the intrusion, initiates an evasive manoeuvre, and a minor sideswipe occurs. This would result in an event identity graph shown in Figure 9.4. The resulting EIG captures a minimal causal chain: Lane change (Vehicle B) $\rightarrow$ detection by Front Right Sensor and LiDAR (Vehicle A) $\rightarrow$ evasive manoeuvre (Vehicle A) $\rightarrow$ collision, with Vehicle B's rule-violating manoeuvre as the primary causal factor, and Vehicle A's evasive action as a mitigating response.

9.3.3 Liability Inference

The Liability Inference constitutes the reasoning core of the EIGF, transforming the structured causal information encoded in the EIG into transparent, legally grounded, and ethically defensible responsibility assessments. Each node and edge in the EIG is evaluated to determine the contribution of vehicles, pedestrians, infrastructure, and other entities to the sequence of events leading to an accident. Established principles of causal inference and graph-based influence propagation algorithms are employed to quantify influence, identify proximate, estimate causal impact and distal causes, and ensure that liability determinations are consistent with traffic regulations, and insurance policies. The EIGF also draws on principles from neuro-symbolic AI [28] and knowledge-graph reasoning [29], which demonstrate how structured causal representations can support explainable, rule-consistent inference in safety-critical domains.

To enhance interpretability and contextual understanding, the module incorporates an LLM-augmented reasoning layer. The LLM interprets graph-based relationships alongside unstructured textual data-such as regulation documents, motor vehicle laws, police reports, and witness statements, to generate human-readable explanations of liability assignments. It also cross-validates outcomes against traffic codes and regulatory guidelines, providing a semantic audit of decisions and highlighting potential inconsistencies or biases. Outputs include structured liability assessments, confidence scores, and traceable decision pathways, accompanied by natural-language reports for investigators, insurers, or regulatory authorities as depicted in Figure 9.2. By combining causal graph reasoning with LLM-augmented interpretive capabilities, EIGF delivers a scalable, transparent, and traceable methodology for automated liability assessment in autonomous vehicle ecosystems.

9.3.4 Fairness and Transparency, and Compliance Assessment

Ensuring fairness, transparency, and legal compliance is essential for any automated liability assessment system, particularly in the context of autonomous vehicles where decisions directly impact human safety, legal outcomes, and insurance claims. The proposed framework incorporates a dedicated auditing layer that evaluates both algorithmic bias and procedural fairness. It ensures that every decision or liability attribution is transparent, auditable, and ethically defensible.

Fairness is achieved by accompanying each inference or liability assignment by contextual metadata, such as the underlying data source, model confidence score, and reasoning chain derived from causal inference models. An EIG visually encodes these relationships, making it possible to trace how specific actions, sensor readings, or environmental factors contributed to the final liability outcome. Transparency is achieved through graph-based explainability and LLM-augmented natural-language reports, which provide human-interpretable rationales for each liability decision. Decision pathways, causal weights, and contributing factors are logged in traceable and auditable formats, enabling investigators, insurers, or regulatory authorities to verify the reasoning behind each outcome. Regulatory compliance is further reinforced by integrating traffic laws, insurance codes, and normative guidelines directly into the liability inference process.

9.4 Discussion

By representing incidents as structured causal graphs, the EIGF allows investigators and insurers to quickly identify contributing factors, assign responsibility, and generate interpretable reports. This structured approach reduces reliance on manual reconstruction and subjective judgment, thereby shortening investigation times and improving consistency in liability assessments.

Beyond operational efficiency, the EIGF enhances transparency and accountability in decision-making. Stakeholders, including vehicle manufacturers, regulatory agencies, and insurance companies, can trace the chain of events leading to a decision, fostering trust in automated systems.

From a technical perspective, the framework introduces several innovations. Edge intelligence allows real-time event detection on embedded platforms, while graph-based causal modelling captures complex interactions among vehicles, infrastructure, and environmental factors. The integration of LLM-augmented reasoning adds semantic understanding of unstructured data, generating interpretable explanations and validating outcomes against regulations. Collectively, these features provide a scalable, explainable, and robust mechanism for automated liability assessment.

9.4.1 Limitations and Challenges

Despite its conceptual strengths, the EIGF faces several limitations. Its effectiveness depends on high-quality, synchronized, and comprehensive data from vehicles, infrastructure, and environmental sensors. Gaps in sensor coverage, data loss, or inaccurate inputs could compromise graph construction and liability inference.

Another challenge lies in algorithmic misattribution. While causal inference and LLM reasoning provide structured analysis, errors in model predictions or incomplete knowledge could lead to incorrect liability assignments. Human oversight may still be required to validate and correct outcomes in complex or ambiguous cases.

One significant challenge in developing and evaluating liability assessment frameworks is the scarcity of publicly available, high-fidelity datasets on autonomous vehicle collisions. Because such accidents are relatively rare and often subject to proprietary, regulatory, or privacy constraints, researchers lack large-scale, multimodal datasets that capture the full sensor, telemetry, contextual, and legal data needed to validate causal and liability models. For instance, the California Department of Motor Vehicles maintains an Autonomous Vehicle Collision Reports database in which AV testing entities must report any collision involving property damage, bodily injury, or death within 10 days [26]. As of October 10, 2025, the DMV has received 875 Autonomous Vehicle Collision Reports, with older reports archived and available upon request.

Similarly, the National Highway Traffic Safety Administration (NHTSA) has issued a Standing General Order [27] requiring manufacturers and operators of vehicles equipped with ADS or Level 2 ADAS to report certain crashes involving these vehicles. This initiative aims to provide timely and transparent notification of real-world crashes associated with ADS and Level 2 ADAS technologies, enabling NHTSA to respond to crashes that raise safety concerns through further investigation and enforcement. While both DMV and NHTSA are valuable initiatives, the dataset remains limited in scope, coverage, and granularity, especially concerning raw sensor logs or internal system states. Without extensive annotated datasets, concept frameworks such as the Event Identity Graph Framework will face difficulty in rigorous empirical validation and benchmarking, potentially slowing progress toward real-world deployment.

The adoption of such systems also raises ethical, legal, and societal challenges. Trust in automated liability decisions may vary across stakeholders, and regulatory acceptance will depend on demonstrable accuracy, fairness, and transparency. Public perception and the willingness of insurers, regulators, and the legal system to rely on AI-driven assessments are critical factors.

9.4.2 Future Directions

Future research will focus on developing a functional prototype of the EIGF framework to demonstrate its practical feasibility. For edge computing platforms NVIDIA Jetson Nano and Raspberry Pi will be employed to host lightweight AI models for real-time event detection and local inference. Graph databases such as Neo4j will be used to model accident events and relationships among entities, supporting efficient traversal, causal reasoning, and graph analytics.

Additionally, LLM-based models such as LegalBERT or GBERT will be integrated to extract and interpret regulatory, traffic, and insurance-related textual information, enabling the framework to link unstructured legal and policy data to structured graph representations. Since such LLMs exceed the computational capabilities of edge devices, the LLM-augmented reasoning layer described earlier will be executed on a cloud or backend service. This allows the edge to remain lightweight and real-time while the cloud handles the heavier semantic reasoning tasks.

Ensuring fairness and transparency will continue to be a central objective. Toolkits such as AI Fairness 360 will be utilized to audit liability assessments, detect potential biases, and enforce fairness across sensitive attributes such as driver demographics, vehicle types, or geographic regions. Hybrid human–AI workflows will be developed to combine automated causal reasoning with expert oversight, enhancing reliability and trust.

9.5 Conclusion

This paper has proposed Event Identity Graph Framework for automated liability assessment in autonomous vehicle accident scenarios. The framework integrates causal graph modelling, edge intelligence concepts, and LLM-augmented reasoning to provide a structured, interpretable, and scalable approach for analysing complex interactions among vehicles, infrastructure, and environmental factors. It emphasizes the integration of traffic regulations, insurance policies, and safety guidelines, while incorporating fairness auditing and explainability mechanisms to ensure that liability assessments are socially and legally defensible. Future work will explore implementing the framework for multi-agent traffic networks, integrating regulatory, insurance, and urban planning considerations, and incorporating semantic analysis of legal and policy texts using models such as LegalBERT or GBERT.

Acknowledgements

This work has been supported by the EdgeAI-Trust "Decentralized Edge Intelligence: Advancing Trust, Safety, and Sustainability in Europe" project which has received funding from Chips Joint Undertaking (Chips JU) under grant agreement No 101139892. The Project website, https://www.edgeai-trust.eu/ provides more information.

References

[1] World Health Organization, "Global Status Report on Road Safety 2018," 2018 Geneva, Switzerland.
[2] European Commission, "EU Road Safety Policy Framework 2021–2030 – Next Steps Towards Vision Zero," 2019, Brussels, Belgium.
[3] H. Atasayar, M. Fleischer, M. Donabauer "Annual statistical report on road safety in the EU 2024," Feb 2024.
[4] A. Smiley and K. Brookhuis, "Alcohol, Drugs and Traffic Safety. Road Users and Traffic Safety," 1987.
[5] National Motor Vehicle Crash Causation Survey. National Highway Traffic Safety Administration, U.S. Department of Transportation.
[6] Bureau of Infrastructure, Transport and Regional Economics "Fatal Road crashes in Australia in the 1990s and 2000s: crash types and major factors," BITRE Canberra ACT.
[7] Mark Mario Morando, Long T. Truong, Hai L. Vu, "Investigating safety impacts of autonomous vehicles using traffic micro-simulation," Australasian Transport Research Forum 2017, Auckland, New Zealand
[8] P. S. Oruganti, E. Deosthale and C. L. Jimenez, "Assessing safe autonomous vehicle behavior via large scale traffic simulation," 2023 IEEE Intelligent Vehicles Symposium (IV), Anchorage, AK, USA, 2023.
[9] N. Novat, E. Kidando, B. Kutela, and A. E. Kitali, "A comparative study of collision types between automated and conventional vehicles using Bayesian probabilistic inferences," Journal of Safety Research, Feb. 2023.
[10] SAE International Recommended Practice, "Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles," SAE Standard J3016_202104, Revised April 2021

[11] UNECE, GRVA-WS06-03 "Global Technical Regulation on Automated Driving System (ADS)," 6th GRVA Workshop on Automated Driving System (ADS), Geneva, Switzerland 2025.
[12] M. Martínez-Díaz and F. Soriguera, "Autonomous vehicles: theoretical and practical challenges," Transportation Research Procedia, 2018.
[13] A. E. Sallab, M. Abdou, E. Perot, and S. Yogamani, "Deep Reinforcement Learning framework for Autonomous Driving," ei, vol. 29, no. 19, pp. 70–76, Jan. 2017.
[14] National Highway Traffic Safety Administration, "Highway Safety Program Guideline No. 18: Motor Vehicle Crash Investigation and Incident Reporting" Washington, DC, 2017.
[15] Deutscher Verkehrssicherheitsrat (DVR), Standards zur Verkehrsunfallaufnahme, Vorstandsbeschluss vom 28.04.2025, May 2025, Berlin
[16] M. A. Quddus, C. Wang, and S. G. Ison, "Road Traffic Congestion and Crash Severity: Econometric Analysis Using Ordered Response Models," J. Transp. Eng., May 2010.
[17] T. A. Rana, S. Sikder, and A. R. Pinjari, "Copula-Based Method for Addressing Endogeneity in Models of Severity of Traffic Crash Injuries: Application to Two-Vehicle Crashes," Transportation Research Record, Jan. 2010.
[18] Q. Zeng, W. Hao, J. Lee, and F. Chen, "Investigating the Impacts of Real-Time Weather Conditions on Freeway Crash Severity: A Bayesian Spatial Analysis," IJERPH, Apr. 2020.
[19] W. Cheng, G. S. Gill, T. Sakrani, M. Dasu, and J. Zhou, "Predicting motorcycle crash injury severity using weather data and alternative Bayesian multivariate crash frequency models," Nov. 2017
[20] A. Sobhani, W. Young, D. Logan, and S. Bahrololoom, "A kinetic energy model of two-vehicle crash injury severity," Accident Analysis & Prevention, May 2011.
[21] C. M. Farmer, E. R. Braver, and E. L. Mitter, "Two-vehicle side impact crashes: The relationship of vehicle and crash characteristics to injury severity," Accident Analysis & Prevention, May 1997.
[22] J. M. Holdridge, V. N. Shankar, and G. F. Ulfarsson, "The crash severity impacts of fixed roadside objects," Journal of Safety Research, Jan. 2005.
[23] C. Conroy et al., "The influence of vehicle damage on injury severity of drivers in head-on motor vehicle crashes," Accident Analysis & Prevention, July 2008.

[24] A. Sinha, V. Vu, S. Chand, K. Wijayaratna, and V. Dixit, "A Crash Injury Model Involving Autonomous Vehicle: Investigating of Crash and Disengagement Reports," Sustainability, July 2021.
[25] R. Bohara, M. Ross, and O. Joglekar, "Event Identity Graph Framework - Supplementary information," https://github.com/asvin-io/eigf2025.
[26] California Department of Motor Vehicles, "Autonomous Vehicle Collision Reports," DMV.
[27] National Highway Traffic Safety Administration, "Standing General Order on Crash Reporting," Washington, DC.
[28] A. d'Avila Garcez, M. Gori, L. C. Lamb, L. Serafini, M. Spranger, and S. N. Tran, "Neural-Symbolic Computing: An Effective Methodology for Principled Reasoning and Learning," May 2019.
[29] M. Nickel, K. Murphy, V. Tresp and E. Gabrilovich, "A Review of Relational Machine Learning for Knowledge Graphs," Jan 2016.