1

The AI-Defined Vehicle: Navigating the Convergence of AI and Autonomous Systems

Ovidiu Vermesan[1], Valerio Frascolla[2], Patrick Pype[3], and Vahid Hashemi[4]

[1]SINTEF AS, Norway
[2]Intel Deutschland GmbH, Germany
[3]NXP Semiconductors, Belgium
[4]AUDI AG, Germany

Abstract

The automotive landscape is experiencing a paradigm shift, driven by the pervasive integration of artificial intelligence (AI) across all functional layers, the availability of previously unattainable data-processing capabilities, and an increasingly tight convergence of sensors, actuators, and advanced communication technologies. This perspective article explores the evolution from connected vehicles to Software (SW)-defined vehicles (SDVs), focusing on the emerging frontier of AI-defined vehicles (AIDVs) as a self-evolving architecture that integrates generative AI (GenAI) for scenario synthesis and agentic AI for autonomous decision-making. This article discusses the technology evolution and the role of AI and vehicle-to-everything (V2X) communication in enabling this transformation. The architecture of SDVs and the conceptual framework of AIDVs are examined, highlighting the transition from Hardware (HW)-centric to SW- and AI-centric designs. Recent advances and future directions are also discussed, presenting a multifaceted view of the opportunities and challenges by synthesising insights from research and industry trends and offering a forward-looking perspective

on the technological, societal, and ethical factors shaping the evolution of automated intelligent transportation systems (ITSs).

Keywords: AI-defined vehicle, software-defined vehicle, autonomous vehicles, vehicle communication AI, V2X, edge computing, edge AI.

1.1 Introduction

The concept of the automobile is being fundamentally redefined, with recent vehicles evolving into cyber-physical, interconnected, and intelligent platforms. This transformation is underpinned by three key technological pillars: widespread embedding of sensing and actuation, advanced communication largely through V2X technologies, and pervasive integration of AI. In addition, as modern vehicles are increasingly SW-defined, the control and compute infrastructure using AI becomes a strategic asset, with central domain and zonal controllers replacing fragmented Electronic Control Units (ECU) architectures [1]. This shift increases reliance on high-performance microprocessors, AI, machine learning (ML), deep learning (DL), GenAI, agentic AI (AI systems with the capacity to perceive, reason, make decisions, act autonomously, and take multi-step actions to achieve specific goals with minimal human oversight), accelerators, and real-time controllers. The journey began with the connected vehicle, a concept that has steadily evolved into today's advanced automated vehicle systems [2, 3].

Advanced autonomous vehicles leverage novel ML/AI-based computational methods combined with new architectures and platforms [4]. The evolution of these vehicles enables ITS, which is seen as the future of transportation systems integrating sensors, information, and computing and communication technologies. AI applications, which are steadily enhancing the human-like intelligence into ITS, and advanced autonomous vehicles share common challenges, such as demanding real-time requirements, cybersecurity threats, and network bandwidth constraints [5]. As vehicle architectures advance, key issues including trustworthiness, intellectual property (IP), and insurance liability underscore the need for robust regulation to support the integration of autonomous vehicles in transportation systems and the smart cities of the future [6].

Technology developments in ITS are firmly evolving in the era of the SDV, where vehicle functions are increasingly decoupled from HW and updatable over-the-air (OTA), and real-time sensor notification and warning systems enhance overall safety [7], a shift that analysts see as the future of

Table 1.1 The evolution of the of vehicle architecture.

Traditional Architecture	**Software-Defined Vehicle (SDV) Architecture**	**AI-Defined Vehicle (AIDV) Architecture**
HW-Defined: Functionality tied to specific ECUs.	**SW-Defined:** Decoupling of SW from HW via HW Abstraction Layer (HAL). Centralized computation and processing.	**AI-Defined:** AI embedded in vehicle operation and user experience.
Inflexible: Difficult and costly to update or add new features.	**Flexible and Scalable:** OTA updates for new features, performance improvements and upgrades.	**Adaptive and Learning:** Continuous learning and personalization based on real-time data.
Siloed Systems: Limited communication and data sharing between ECUs.	**Service-Oriented Architecture (SOA):** Enables integration and communication between SW components.	**Data-Driven and Predictive:** Proactive maintenance, personalized experiences, and anticipatory actions.
Focus on Mechanics: Core value in the vehicle physical engineering.	**Focus on SW and Services:** Value shifts to the user experience and connected services including new business models.	**Focus on Intelligence and Autonomy:** Core value in the vehicle ability to perceive, infer, and act intelligently.

the automotive industry [8]. Looking ahead, the horizon is dominated by the concept of the AIDV, a system in which AI, at the core of operations and user experience, is the main driver of the evolution of the vehicle, as sketched in Table 1.1.

V2X communication serves as the nervous system of this intelligent transportation ecosystem and can be considered as a general concept that encompasses vehicle-to-vehicle (V2V), vehicle-to-infrastructure (V2I), vehicle-to-pedestrian (V2P), vehicle-to-network (V2N), vehicle-to-grid (V2G), vehicle-to-home (V2H), vehicle-to-maintenance (V2M), vehicle-to-owner (V2O), vehicle-to-users (V2U), and other interactions [9, 10]. It provides the real-time data streams necessary for advanced driver-assistance systems (ADAS), cooperative autonomous driving, and efficient traffic management. Concurrently, AI provides vehicles with the equivalent of a brain, enabling vehicles to perceive the environment, make complex decisions, and learn from experience, building on a long history of development in autonomous systems.

The development of SDVs and AIDVs requires architecture and technology solutions that integrate cooperation and real-time communication into their design specifications. The collaboration and exchange of data and information depend on trust among humans, vehicles, and infrastructure. In advanced autonomous systems, trust is earned through the dependability of vehicle systems and subsystems, as reflected in the trustworthiness concept [11]. Complex AI systems integrated into SDVs and AIDVs deepen the existing challenge of safety assurance of autonomous driving. A solution to mitigate this challenge is to utilize a set of techniques like spatio-temporal prediction [1, 12], explainable AI (XAI) [13] and interpretable AI (IAI) [14] methods for safe and trustworthy autonomous driving, focusing on key aspects such as data, model, and agency in the defined operational design domain (ODD) of the vehicle.

The autonomous functions required by SDVs and AIDVs are executed by the vehicle modules integrated into the vehicle architecture to perform perception, planning, and control and communicate with other vehicles, road users or the road infrastructure.

Communication technologies play an essential role in SDV Network (SDVN) architectures, added to the SDV architectures by integrating novel electrical and electronic (E/E) architectures, operating systems (OS), open-source HW, SW tools, processing at the edge, OTA updates/upgrades techniques, features-on-demand, and the deployment of Digital Twins (DT) and immersive triplets technologies for modelling, simulation, and operations. SDV refers to transforming vehicles from primarily HW-driven to embedded and SW-centric platforms. This shift implies that SW drives much of the functionality of a vehicle functionality, customization, and performance enhancements rather than mechanical or HW changes. This approach allows for continuous updates, upgrades and improvements of the features and functions of vehicles and more in general of mobile autonomous systems.

Vehicle architectures have evolved to achieve full autonomous functions, as illustrated in Figure 1.1.

The evolution includes domain-based, body-zonal, cross-domain zonal, and consolidated computing architectures, where zonalization and consolidation of the computing tasks drive the developments of the vehicle architecture. The domain-based functions in the vehicles are grouped progressively into zonal controllers. While several applications still require isolation for performance, safety, or security requirements, in the future, complete consolidation of the electronic functions into a central vehicle computer is expected. In the consolidated compute vehicle architecture, sensors and actuators data are collected in I/O aggregator units placed in their proximity. While reducing the system complexity, zonalization and consolidation introduce new safety and

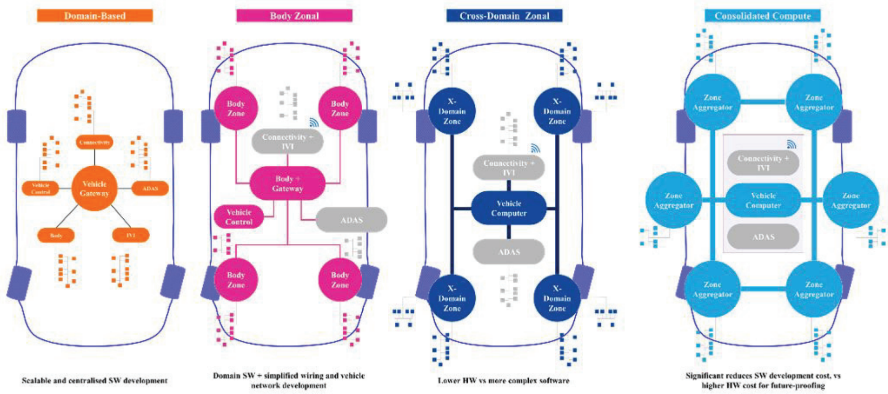

Figure 1.1 Evolving vehicle architectures.

security concerns. Mixed-critical applications share HW resources, such as communication links or vehicle processors, potentially interfering with each other. The system must separate the execution of safety-critical, real-time applications from the execution of non-critical ones. Isolation, safety, and security mechanisms must be taken into consideration in the very early phases of the architectural design of new vehicles, as they will be key elements for commercializing future mobile autonomous and robotic systems.

The transition from conventional vehicles to SDVs illustrates a fundamental shift in automotive perspective, architecture, and capability. Traditional vehicles are primarily HW-defined, with their functions and performance inherently linked to the physical and electromechanical components installed in the vehicle during manufacturing. Upgrades or new features generally require physical modifications or component replacements, which are often performed at a dealership. A distributed network of multiple ECUs characterizes the vehicle architecture, each with a specific, usually isolated function, connected by intricate wiring harnesses that add weight and limit integration possibilities. Innovation cycles are tied to model years and HW changes, resulting in slower evolution.

The integration of sensors, actuators, connectivity, and AI into the vehicle-edge-cloud continuum requires managing data effectively, reliably, securely, and privately via various data-, model-, and code-centric architectures, combining proprietary and open-source building blocks. The evolution of the E/E architecture for SDV to AIDVs is envisioned in Figure 1.2.

One of the core elements of autonomous SDVs and AIDVs lies in the decision-making and planning systems that interpret sensor data,

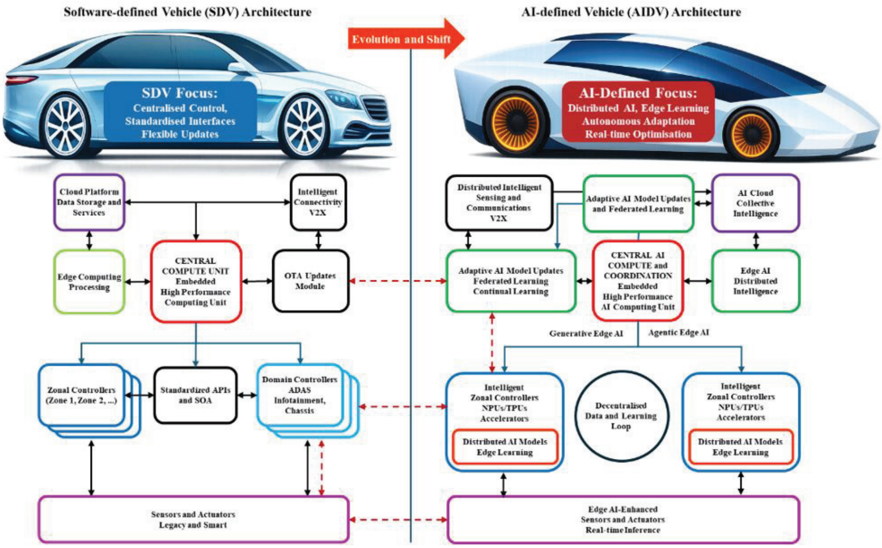

Figure 1.2 E/E Architecture types: SDVs vs. AIDVs.

predict traffic behaviour, and determine optimal driving actions. Traditional rule-based systems, while offering transparency, often lack the flexibility needed for unpredictable real-world scenarios. End-to-end AI models that integrate perception, prediction, and planning have shown promise in this domain. Research emphasizes the use of reinforcement learning (RL), imitation learning, and advanced neural network architectures that combine these elements for real-time decision-making [15].

Emerging techniques that leverage deep RL have the potential to scale decision-making by dynamically adapting to new driving environments. Incorporating multi-modal data, such as real-time weather updates, V2X information, and historical traffic patterns, can significantly enhance prediction accuracy and planning reliability. Improving these models requires rigorous simulation and field testing, along with advanced synthetic data generation methods, to replicate rare yet critical driving scenarios and provide explainable solutions for autonomous vehicles [16].

In this context, challenges remain in effectively combining multimodal, multisource, internal, and external vehicular data to support real-time decision-making, as well as in addressing synchronisation, formatting, and management issues. Integrating vehicular system intelligence with human knowledge requires robust frameworks that ensure secure, private, yet efficient data management pipelines, focusing on adaptive techniques that integrate newly emerging data modalities and sources.

The deployment of ML, DL, and GenAI solutions requires optimised energy and resource use for processing and communication. As the models can run in vehicles using specialised HW accelerators, task prioritisation, management, and optimisation are necessary for effective decision-making. Decision-making in real-time using large-scale sensor data analytics and GenAI relies on novel algorithmic innovations, including federated learning (FL), neural networks, and agentic AI. As a result of specific tasks performed by SDVs and AIDVs, the models must be personalised in the vehicular environment based on human responses.

Personalisation can be critical for the driving experience, where environments and road events vary constantly. Creating models for every vehicle, again and again, is inefficient, as (re)training is resource-intensive and vehicles are constrained environments that require lightweight frameworks capable of handling large-scale datasets as part of the data processing pipeline [17].

This article presents a perspective on the synergistic evolution of AI by the pervasive integration across all functional layers in shaping the future of autonomous systems and ITSs, by exploring architectural shifts from traditional vehicles to SDVs and AIDVs, discussing the current state of the art and future research directions, and offering a view of the extensive potential and significant challenges that lie ahead.

1.2 Architectural Evolution: From HW to AI-Defined

The transition from HW-centric to AI-defined architectures is a multi-stage evolution, with each phase building upon the capabilities of the previous one.

The autonomous vehicle architecture began with a modular pipeline approach [18, 19] used in the design of the driving systems, separating each part of the system into distinct SW and HW components, split which created significant challenges in optimising and synchronising the system. The modular pipeline approach breaks down the task of the autonomous driving system into perception, prediction, planning, and control. In this approach, each module is developed separately and is responsible for a specific functionality in the whole system. To address the development of autonomous vehicles, an end-to-end pipeline approach [19, 20, 21] was introduced to integrate separate components into a unified system and optimise the entire system in a differentiable manner. The end-to-end pipeline manages autonomous driving as a single learning task using imitation and RL, taking raw sensor data as input and directly outputting the control signal, optimising final planning performance as its primary objective, thereby providing greater safety and reliability than a modular pipeline approach. These approaches are part of the evolution towards new architectural concepts for SDVs and AIDVs.

1.2.1 The Rise of the SW-Defined Vehicle

Traditionally, vehicle functionalities have been tightly coupled with dedicated ECUs. The ECU-based architecture, while reliable, is limited in its ability to update and introduce new features, leading to a complex, costly upgrade process. An SDV architecture marks a significant departure from this model by centralising computing resources and abstracting SW from the underlying HW, a concept that has been thoroughly surveyed in the context of SDVNs [22] and AI-defined wireless networking [23, 24].

SDVs represent a shift from HW-centric design to a model where SW governs functionality, performance, and user experience. As vehicles integrate increasing numbers of sensors, ECUs, and connectivity features, they generate large volumes of data related to performance, driver behaviour, and system health. SDVs leverage this data to continuously improve vehicle capabilities, enabling manufacturers to refine features, diagnose issues, and deploy enhancements throughout the vehicle lifecycle.

A key technical enabler of SDVs is the use of DTs and immersive triplets, which act as virtual representations of real-world vehicles [25]. By transmitting operational data to the cloud, vehicles can provide detailed insights into battery health, performance of ADAS, and feature usage under real driving conditions. This continuous feedback loop allows original equipment manufacturers (OEMs) to accelerate development cycles, identify recurring faults, and implement corrective actions before issues become widespread, particularly in complex domains such as autonomous driving.

V2X communication further extends the value of SDVs by enabling secure data exchange between vehicles, infrastructure, and other road users. Information such as vehicle speed, position, and lane departure warnings can be shared in real time to enhance situational awareness and improve traffic safety. This interconnected environment relies on robust data handling and secure communication protocols, both of which are central to SDV platforms.

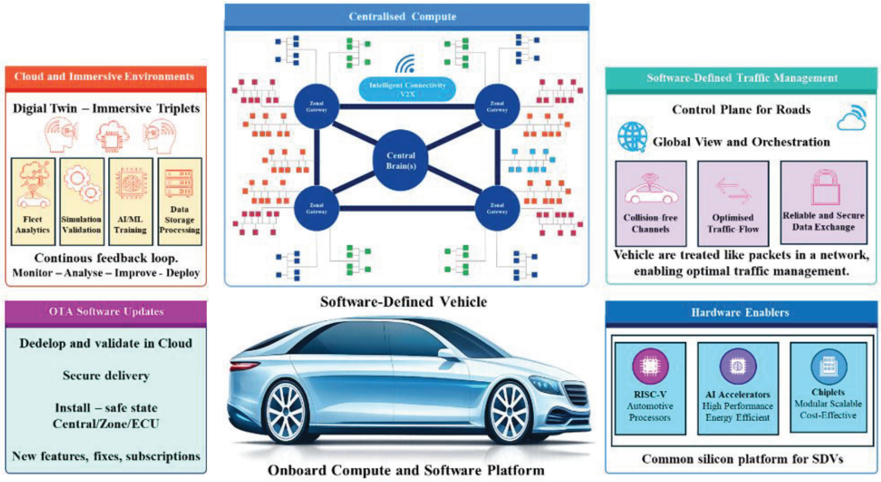

Figure 1.3 SDV architecture.

A defining feature of SDVs is the ability to update and expand functionality through OTA SW updates, which are developed and validated in cloud environments before being securely delivered to vehicles. Once downloaded, updates can be installed across central computing systems, zone controllers, or edge ECUs, typically requiring a system restart in a safe vehicle state.

OTA capabilities allow OEMs to introduce new features, fix SW defects, and offer subscription-based services without requiring physical access to the vehicle.

The transition to SDVs is closely tied to the evolution of electrical and electronic E/E architectures. Traditional distributed and domain-based architectures are being replaced by zone-based networks that reduce wiring complexity and support centralised SW control. In a zone architecture, the vehicle is divided into physical regions, each managed by a zone control module that aggregates inputs from sensors and actuators within its area.

Within these architectures, OEMs adopt different SW deployment strategies. In a fully centralised approach, a single high-performance computer controls most vehicle functions, simplifying SW management and enabling consistent updates. However, this model introduces challenges related to real-time control latency and functional safety, particularly if communication links between the central unit and peripheral components are disrupted.

A hybrid approach distributes SW responsibilities between a central computer and zone controllers. High-performance applications such as ADAS and infotainment are typically centralised, while time-sensitive control functions may remain closer to the HW in zone modules or edge ECUs. This balances computational efficiency with responsiveness and safety requirements.

A more distributed model retains several domain controllers alongside zone modules, allowing a gradual transition from legacy systems. In this configuration, different zones may handle varying combinations of functions, such as body control, lighting, or chassis systems, depending on design choices. This flexibility enables OEMs to tailor architectures to specific vehicle platforms while incrementally adopting SDV principles.

The design of SDV architectures requires careful trade-offs between latency, network performance, functional safety, cybersecurity, and SW scalability. OEMs must align HW and SW strategies to ensure reliable real-time control while enabling the flexibility and continuous innovation that define SDVs.

From a road perspective, a collection of individual vehicle decisions is inevitably sub-optimal, and it is essential that roads directly intervene (control/assist) in vehicle operation to enable optimal traffic management. Solutions such as SW-defined traffic management (SDTM), which incorporate SW-Defined Network (SDN) concepts, are being used in the transportation environment. In SDTM, each vehicle is handled like a packet in an SDN, and its movement can be assisted/controlled using a global view of the road. In this concept, the SDTM enables reliable interaction between vehicles and infrastructure by allocating collision-free channels to vehicles, reducing their messaging interval by fully utilising available channel resources, and redistributing sensitive information from the infrastructure [26].

The maturity of open-source SW (e.g., Linux-based solutions) has demonstrated how shared platforms can reduce duplication, lower entry barriers, foster innovation, and enhance interoperability. This success needs to be replicated in the domain of silicon by enabling the co-development of a common European platform of RISC-V-based automotive processors and AI-accelerators. Chiplets are small, specialised integrated circuits (ICs) designed to perform specific functions within a larger system. This is considered a strategic approach to mitigate the substantial design costs associated with newer, smaller process geometries and to address the challenges posed by large, monolithic die sizes of 300 mm and above, which often result in poor yields. Additionally, providing diverse processing elements for both low- and high-power loads is a game changer in the continuously evolving market for commercial processing elements (CPUs, GPUs, AI accelerators, etc.).

1.2.2 The Dawn of the AI-Defined Vehicle

AIDVs represent a decisive shift in automotive system design, moving beyond the SDV paradigm toward architectures where AI is the primary driver of functionality, behaviour, and evolution. While SDVs introduce abstraction, modularity, and updatable SW layers, AIDVs embed intelligence as a foundational property of the vehicle. In this model, AI does not merely support isolated features but defines how the vehicle perceives, decides, adapts, and interacts across all domains of operation.

At the architectural level, AIDVs build on centralised and zonal computing frameworks that consolidate previously distributed ECUs into high-performance computing platforms as illustrated in Figure 1.4, where AI is the primary driver for functionality, behaviour and evolution.

These platforms integrate CPUs, GPUs, NPUs, and dedicated AI accelerators, enabling the execution of complex ML models with strict real-time constraints. The shift toward centralised compute reduces latency, improves data coherence, and supports advanced sensor fusion. Robust sensor fusion involves integrating data from cameras, LiDAR, radar, and other

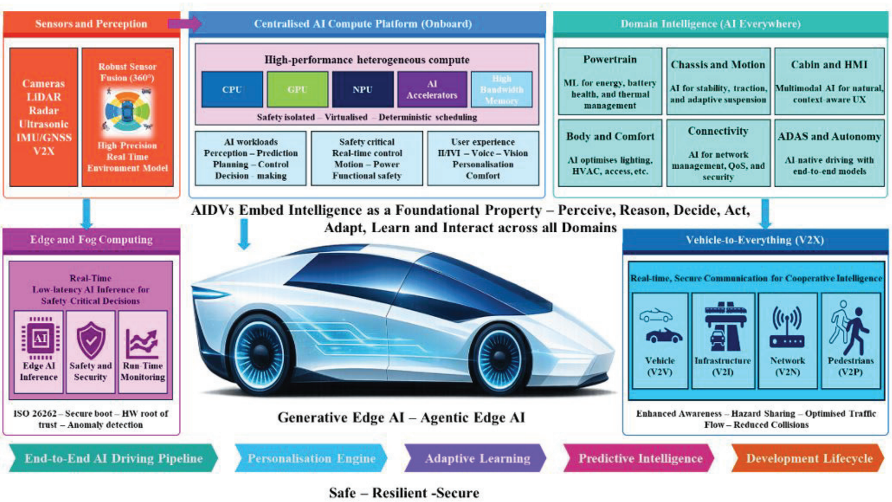

Figure 1.4 AIDV architecture.

sensing devices to create a detailed and reliable picture of the environment and enhancing DL methods for object detection and localisation to handle complex, dynamic environments. Enhancing these technologies is critical to optimising traffic flow and reducing casualties during unforeseen events [27, 63].

AI workloads are deployed across a heterogeneous computing fabric, where safety-critical tasks coexist with high-level perception, planning, and user-experience functions, under strict isolation and scheduling policies.

A defining characteristic of AIDVs is pervasive AI integration. Instead of limiting intelligence to ADAS, AI is embedded across the entire vehicle stack. In the powertrain domain, ML models optimise energy efficiency, battery health, and thermal management in real time. In chassis and motion control, AI enhances stability, traction, and adaptive suspension behaviour by continuously learning from driving conditions and driver intent. Within the cabin, multimodal AI systems fuse voice, vision, and contextual data to deliver adaptive human–machine interfaces that respond naturally to occupants. This pervasive intelligence is enabled by tightly coupled sensor networks and high-bandwidth in-vehicle communication systems, often based on automotive Ethernet.

Fog and edge computing are critical enablers of AIDVs, providing low-latency processing for safety-critical decisions [28]. AI models deployed at the edge process raw sensor data from cameras, radar, LiDAR, ultrasonic sensors, and inertial units in real time. These systems must meet stringent functional safety requirements, often aligned with standards such as ISO 26262 [29], while also addressing cybersecurity concerns through secure boot, HW roots of trust, and runtime anomaly detection. The vehicle effectively becomes a secure, high-performance computing node capable of operating autonomously even in degraded connectivity scenarios.

Continuous learning is another core pillar of the AIDV concept. Vehicles continuously collect data from onboard sensors, V2X communications, and user interactions. This data is selectively transmitted to cloud infrastructure, where large-scale training and validation pipelines refine AI models. Updated models are then deployed back to the fleet via OTA updates. This closed-loop lifecycle allows the vehicle to improve over time, adapting to new environments, edge cases, and usage patterns. Importantly, mechanisms for data governance, privacy preservation, and dataset curation are integral to ensuring both regulatory compliance and model robustness.

Personalisation in AIDVs extends beyond static user profiles to dynamic, context-aware adaptation. AI systems learn driver and passenger preferences across multiple dimensions, including driving style, seating position, climate control, infotainment choices, and navigation habits. These preferences are continuously refined using behavioural data and contextual cues such as time of day, location, and occupancy. The result is an in-cabin experience that evolves with the user, delivering a level of customisation that approaches individualised mobility services. This personalisation is often supported by multimodal user identification and cloud-synchronised profiles that persist across vehicles.

Predictive and prescriptive capabilities distinguish AIDVs from reactive systems. By integrating real-time sensor data with historical and fleet-level insights, AI models can anticipate hazards, optimise routes, and predict component degradation. Predictive maintenance algorithms analyse vibration patterns, thermal signatures, and usage data to forecast failures before they occur, reducing downtime and maintenance costs. Similarly, predictive energy management systems optimise battery usage, charging strategies, and regenerative braking to maximise efficiency and range. These capabilities rely on both onboard inference and cloud-based analytics, forming a hybrid intelligence model.

A major technological trend underpinning AIDVs is the shift toward continuous end-to-end AI pipelines, particularly in autonomous driving. Traditional autonomy stacks rely on modular pipelines that separate perception, localisation, prediction, planning, and control. In contrast, end-to-end models use deep neural networks to map raw sensor inputs directly to control outputs such as steering, acceleration, and braking. These models are trained on large-scale datasets collected from vehicle fleets and augmented with synthetic data generated through high-fidelity simulation. While end-to-end approaches can reduce system complexity and improve adaptability, they also introduce challenges in interpretability, validation, and safety assurance.

The development lifecycle of AIDVs is tightly coupled with continuous integration and continuous deployment practices. AI models are iteratively trained, validated, and deployed across a computing continuum that spans cloud and edge environments. Cloud platforms provide the computational resources for large-scale training and scenario simulation, while edge platforms execute optimised models in real time. Emerging techniques in GenAI and agentic systems are being integrated into these pipelines, enabling automated scenario generation, data augmentation, and even autonomous system tuning. This accelerates development cycles while improving coverage of rare and safety-critical scenarios.

AIDVs represent a convergence of advanced AI, high-performance computing, and connected vehicle ecosystems. By embedding intelligence into every aspect of the vehicle and enabling continuous evolution through data-driven pipelines, AIDVs redefine the vehicle as an adaptive, learning system. This transformation has implications not only for vehicle performance and safety but also for the broader mobility landscape, where vehicles become intelligent agents operating within a dynamic, interconnected environment.

1.2.3 The Transition to Self-Evolving Vehicular Architectures

While SDVs decouple SW from HW, AIDVs introduce a paradigm shift toward "Self-Evolving Architectures" integrating GenAI and agentic AI technologies. The core challenge lies in moving from static OTA updates to dynamic, self-supervised learning loops. The defining characteristic of the future AIDV architecture is its "self-evolving" capability, a departure from static, rule-based systems toward a dynamic, continuous evolving ecosystem. This evolution is driven by the cyclical interaction between the AI paradigms, such as GenAI and agentic AI [30].

The synergy allows the system to move beyond supervised learning on pre-labelled datasets into a continuous, self-supervised learning cycle, where GenAI functions serve as a knowledge generator, and foundation models synthesise high-fidelity virtual environments by ingesting vast streams of real-world driving data, including telemetry, sensor fusion logs, and V2X communication packets.

The trajectory toward AIDVs is accelerating as advances in connectivity, compute, foundation models, GenAI, and agentic AI begin to converge into an AI-centric vehicle stack. Beyond incremental improvements, the field is shifting toward architectures in which perception, planning, communication, and user interaction are co-optimised through shared AI models and continuous learning loops spanning fleet, edge, and cloud.

AI-powered V2X communication is evolving from static protocol optimisation to fully adaptive, learning-driven network orchestration. Recent trends integrate RL and graph neural networks to model highly dynamic vehicular topologies, enabling predictive scheduling, interference mitigation, and semantic-aware communication while transmitting only task-relevant information. This reduces bandwidth pressure while improving latency bounds for safety-critical services. The transition toward 5G-Advanced and early 6G concepts further introduces integrated sensing and communication (ISAC), allowing the communication infrastructure itself to act as a distributed sensor [31]. In this context, AI-native air interfaces and cross-layer optimisation are expected to support sub-10 ms end-to-end latency and ultra-reliable low-latency communication (URLLC) at scale, which is essential for cooperative autonomy [32].

Cooperative perception and manoeuvring are moving beyond simple data sharing toward collective intelligence. Instead of exchanging raw sensor feeds, vehicles increasingly share compressed feature maps or object-level abstractions generated by deep neural networks, reducing communication overhead while preserving semantic richness. Emerging approaches use multi-agent RL and distributed consensus algorithms to coordinate manoeuvres such as platooning, merging, and intersection negotiation. This enables intent-aware driving, in which vehicles not only perceive the environment but also anticipate others' actions [33]. Progress in uncertainty quantification and trust-aware fusion is critical here, ensuring robustness against noisy, delayed, or malicious inputs in open environments.

Edge and cloud computing are being redefined by the rise of heterogeneous AI accelerators and AI-defined compute fabrics. Modern AIDV platforms integrate GPUs, TPUs, NPUs, neuromorphic, and domain-specific accelerators optimised for transformer-based workloads, enabling real-time execution of increasingly large foundation models and agentic AI workflows. A key trend is the emergence of split computing paradigms, in which inference pipelines are partitioned across the vehicle, roadside edge, and cloud based on latency, privacy, and energy constraints. FL and continual learning frameworks enable fleets to collaboratively improve models without centralising raw data, addressing both scalability and regulatory concerns. At the same time, AI-defined infrastructure enables dynamic deployment of AI services OTA, effectively turning vehicles into nodes of a distributed learning system.

GenAI in the cabin is advancing from conversational assistants to context-aware, multimodal copilots. These systems integrate speech, vision, and vehicle telemetry to provide proactive assistance, such as anticipating driver needs, explaining vehicle decisions, or adapting interfaces in real time. The underlying models are increasingly domain-specialised small language models (SLMs) and vision-language models (VLMs) fine-tuned for automotive safety and reliability [34]. On-device inference is becoming feasible through model compression and distillation, reducing reliance on cloud connectivity. A key research direction is aligning generative models with safety constraints

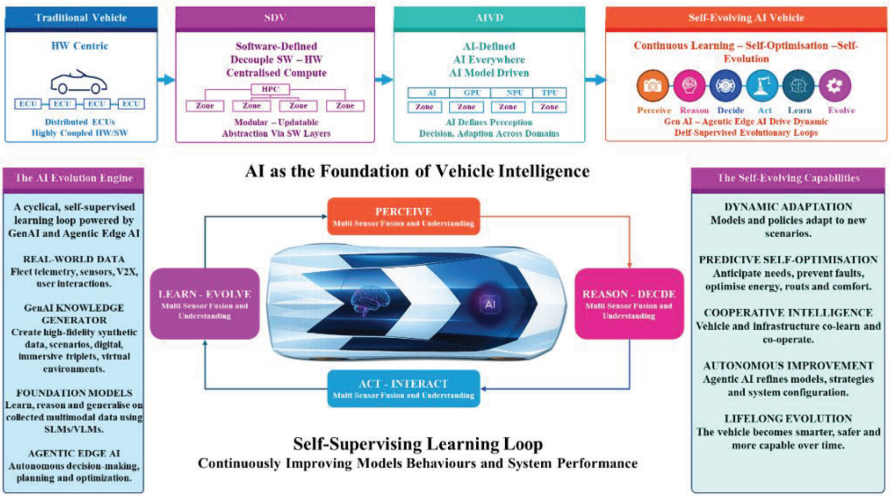

Figure 1.5 The vision of self-evolving vehicular architecture.

and human factors, ensuring that natural interaction does not compromise driver attention or trust.

DTs are evolving into high-fidelity, continuously synchronised replicas that allow for not only validation but also real-time decision support and lifecycle optimisation. Advances in simulation realism, driven by neural rendering and physics-informed AI, allow DTs to capture complex environmental and behavioural dynamics. The concept of immersive triplets extends this by integrating spatial computing and mixed reality interfaces, enabling engineers and even vehicles themselves to interact with virtualised environments. Emerging frameworks incorporate secure attestation mechanisms that validate AI models and OTA updates against their DT counterparts before deployment, enhancing safety and cybersecurity [35, 36]. Looking ahead, closed-loop integration between fleet data, simulation, and model retraining is enabling a "simulation-driven development at scale" paradigm, significantly shortening innovation cycles.

These trends indicate a shift from isolated intelligent functions to fully integrated, learning-driven vehicular ecosystems. The defining characteristic of next-generation AIDVs will not be any single capability, but the ability to continuously adapt, collaborate, and improve through tightly coupled AI, communication, and compute infrastructures.

Looking ahead, the integration of 6G communication technologies is expected to be a significant catalyst for AIDVs. The ultra-low latency, high bandwidth, massive connectivity, integrated sensing, precise localisation, and embedded AI of 6G will enable even more evolved V2X applications.

Recent developments have significantly expanded the scope of V2X research by integrating generative models and advanced game theory into trajectory planning, particularly at the intersection of aerial and ground vehicular networks. In path planning, Generative Adversarial Networks (GANs) have emerged as a tool for navigating complex, high-dimensional spaces. Visual data is utilised to autonomously generate energy-efficient flight paths, highlighting the potential for generative models to solve optimisation problems in dynamic, three-dimensional traffic environments where traditional heuristic methods struggle [37].

Complementing these generative approaches are advancements in modelling the interactive decision-making processes between vehicles during critical manoeuvres, such as on-ramp merging. Moving beyond static rule-based logic, an integrated motion planning framework based on Stackelberg Game modelling can be used [38].

Advancements in neuromorphic and cognitive computing, which mimic the structure and function of the human brain, could lead to more efficient and powerful AI HW for the various electronic components integrated into diverse vehicle domains.

Security solutions in ITS standards, based solely on Key Performance Indicators (KPI), leave several areas for reconsideration, and new approaches and solutions are required for SDVs and AIDVs to ensure vehicular communication is indeed secure so that the overall objective to make the roads safer and reduce road accidents can be achieved, rather than providing a new target for cyberattacks. The development of SDVs and AIDVs towards autonomous cyber-physical systems, utilising vehicular communication as a key domain for exchange between vehicles, infrastructure, and traffic participants, is crucial to ensure that the connectivity system is not suboptimal, thereby avoiding physical damage. To prevent such losses, new relevant security and privacy aspects of vehicular communication for SDVs and AIDVs must be addressed [39].

1.3 Opportunities and Challenges

The shift toward AIDV represents a fundamental overhaul of automotive architecture, moving from distributed ECUs to centralised, high-performance computing platforms. This AI-centric approach allows for continuous OTA updates and predictive maintenance, turning vehicles into evolving intelligent agents rather than static HW. However, this reliance on complex AI models requires immense processing power and data management, challenging manufacturers to build robust, scalable systems that support real-time decision-making without latency.

As vehicles become hyper-connected, trustworthiness and security emerge as a paramount challenge; a single vulnerability in the SW and AI stack could compromise fleet safety or user privacy. To mitigate these risks, the industry must navigate a complex landscape of evolving regulation and standardisation. Establishing universal protocols is essential not only for preventing malicious cyberattacks but also for ensuring interoperability between different manufacturers and smart infrastructure, preventing a fragmented ecosystem of incompatible safety standards.

The widespread adoption of autonomous capabilities will trigger a profound transformation of jobs, disrupting traditional roles in logistics, transportation, and ride-hailing services. While the demand for human drivers may decline, a new spectrum of opportunities can emerge in SW engineering, teleoperations, and fleet oversight. The challenge lies in managing this workforce transition equitably, requiring significant investment in upskilling to ensure that the creation of high-value technical roles balances the displacement of manual labour.

The deployment of these vehicles introduces complex ethical dilemmas regarding algorithmic accountability and decision-making during unavoidable accidents. Programming a machine to make life-or-death value judgments raises difficult questions about moral priority and legal liability. Furthermore, ensuring that these AI systems are free from data bias, such as detecting pedestrians of all demographics with equal accuracy, is critical to gaining public trust and preventing the convenience of automation from coming at the cost of equitable safety.

The key opportunities and challenges of AIDV vehicles are presented in Table 1.2, further described in the subparagraphs below and illustrated in Figure 1.6.

1.3.1 AI

AI is likely to disrupt the way vehicles are developed and built, the solutions used, the intrinsic characteristics of automotive components and systems, and how consumers interact with them. The first stages of AI development, from ML and DL to generative and agentic [40], are already evident in new vehicle

Table 1.2 Key opportunities and challenges of AI-defined vehicles.

Opportunities	Challenges
Enhanced Safety: Drastic reduction in accidents caused by human error [39].	**Cybersecurity:** Increased attack surface for malicious actors [3, 45].
Improved Traffic Efficiency: Reduced congestion and travel times through optimized routing and traffic flow.	**Data Privacy:** Collection and use of vast amounts of personal and location data.
Increased Accessibility: Greater mobility for the elderly, disabled, and those unable to drive.	**Ethical Dilemmas:** Programming "trolley problem" scenarios and ensuring fairness in AI decision-making.
New Business Models: Emergence of new services in areas like in-vehicle commerce, entertainment, and logistics [8].	**Regulatory Hurdles:** Lack of standardized regulations and liability frameworks for autonomous vehicles.
Environmental Benefits: Optimized driving patterns and platooning can reduce fuel consumption and emissions.	**Infrastructure Investment:** Significant investment required in V2X infrastructure and 6G networks.
Enhanced User Experience: Personalized and intuitive in-vehicle experiences [1].	**Public Trust and Acceptance:** Overcoming scepticism and building confidence in the safety and reliability of AI-driven systems.

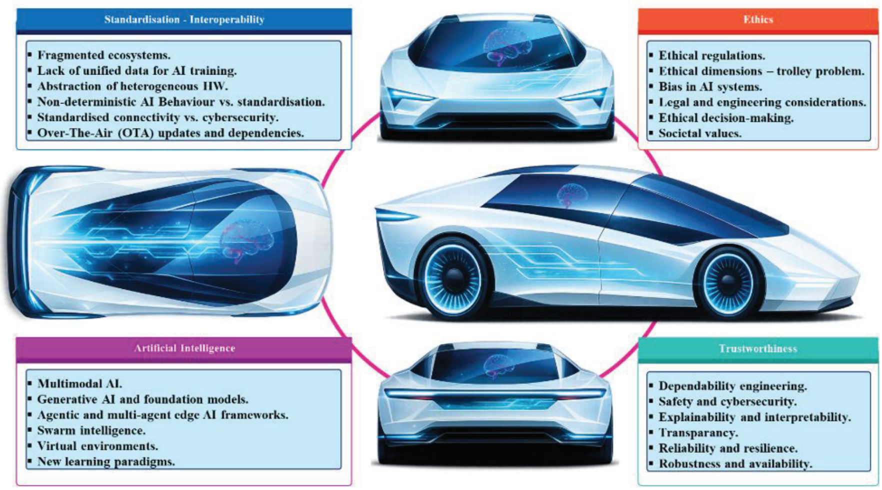

Figure 1.6 Self-evolving vehicles opportunities and challenges.

concepts and architectures, and their evolution will transform every aspect of the automotive value chain.

Future developments in AIDVs are likely to focus on improved perception through multimodal sensing, combining cameras, radar, LiDAR, and contextual data for better situational awareness. Advances in real-time decision-making and edge computing allow vehicles to process information faster and operate safely without relying heavily on cloud connectivity. AI models become more efficient, reducing power consumption while increasing reliability.

V2X-connected intelligence can reduce congestion, improve safety, and enable coordinated driving behaviours. Continuous learning systems may allow fleets to share knowledge, so improvements made by one vehicle can benefit others almost instantly.

The advantages of autonomous vehicle development include increased road safety by reducing human error, which is a leading cause of accidents. AI-driven vehicles can react faster, maintain consistent attention, and make data-driven decisions under pressure. They also offer improved mobility for people who cannot drive, such as the elderly or disabled.

Autonomous vehicles can optimise traffic flow, reduce fuel consumption, and lower emissions by enabling smoother driving and more efficient route planning. Over time, they may reshape urban design by reducing the need for parking and enabling more efficient transportation systems. Economically, they can reduce logistics and transportation costs while enabling new services such as autonomous delivery and ridesharing.

As both modular and end-to-end driving approaches face challenges such as causal confusion, explainability, interpretability, generalization, and robustness, a solution is to embed LLMs [41, 42], VLMs [19], and agentic AI [40, 43] into the vehicle architecture pipeline, alongside scene understanding, reasoning, zero-shot recognition, in-context learning, and interpretability. The use of a unified multimodal input and vocabulary for vision, language, and action can unify vision- and action-related tasks, including, among others, scene understanding, planning, and control, by using agentic AI approaches to implement interactions with the vehicle decision-making support.

The concept of agentic AI shifts the focus from isolated vehicle intelligence to cooperative behaviour among multiple autonomous systems. In the future, vehicles will not only make decisions individually but will also coordinate with each other through multi-agent learning systems and swarm intelligence frameworks. This cooperation is particularly essential for managing dense traffic situations, achieving efficient route planning, and ensuring collective safety during complex manoeuvres.

Developing robust multi-agent frameworks involves research into FL paradigms, where each vehicle processes local data while contributing to a shared global model without compromising data privacy. Moreover, integrating these intelligent agents with V2X communication networks enables real-time information exchange, enhancing overall traffic management and coordination. The challenges lie in ensuring system scalability, reducing latency, and safeguarding cybersecurity across distributed platforms [27].

These research directions open the door to a future in which urban mobility is revolutionised by interconnected, cooperative vehicle systems.

Swarm intelligence involves the coordination of multiple autonomous vehicles working together as a collective to optimise traffic flow and reduce congestion. Future research in this domain should focus on developing algorithms that facilitate cooperative behaviour among vehicles, enabling them to make joint decisions in real time. Such systems must be robust enough to handle dynamic changes in traffic and varying environmental conditions. The development of cooperative multi-agent systems will rely on novel RL techniques and advanced communication protocols that enable vehicles to reliably share sensor data and decision-making parameters. Leveraging swarm intelligence could revolutionise urban mobility by enabling distributed coordination among vehicles, reducing travel times, lowering emissions, and enhancing overall traffic stability. Research should also evaluate the potential cybersecurity risks inherent in large-scale inter-vehicle communication and develop protocols to mitigate these risks [27].

The prospect of a future dominated by AIDVs is both stimulating and challenging, given the profound societal implications of this technological shift. The potential of AI to dramatically improve road safety is perhaps the most significant benefit, as the World Health Organization estimates that over 1.3 million deaths occur each year from road traffic crashes [44]. Preliminary data for the first half of 2024, which covers 26 countries, show an improvement compared to the same period in 2023. Road fatalities decreased in 16 countries, while increasing in ten. On average, road deaths across these countries declined by 2% [45]. It is expected that by removing the human element, which is a factor in most accidents, AIDVs have the potential to save many lives.

1.3.2 Trust

Trust in AI-driven vehicles is grounded in the belief that the autonomous systems behave safely, predictably, and in alignment with human expectations. For users and regulators, trust develops when the vehicle consistently demonstrates correct decisions across a wide range of driving conditions, including rare and uncertain situations.

Trustworthiness is closely tied to dependability properties such as reliability, safety, availability, and robustness. Reliability refers to the system performing its intended functions without failure over time. Safety ensures that even when failures occur, the system minimises harm. Availability means the vehicle can operate when needed without unexpected downtime, while robustness reflects its ability to handle noise, uncertainty, and changing environments.

Research into robust, resilient autonomous systems focuses on reducing cumulative error propagation within multi-module architectures and ensuring redundancy across critical components. By integrating emergency response protocols and actively learning from near-miss incidents, these systems can improve overall reliability and public trust.

The emerging of new safety concepts emphasises deep world understanding and introspective AI models that simulate human-like intuition. Traditional scenario modelling, which relies on rigid safety frameworks, falls short in addressing the "long-tail" of rare and complex emergencies. Future research must develop adaptive, introspectable models that align with international safety standards such as ISO 26262 [29] and SOTIF [46]. The challenge is not only to predict and prevent injuries but also to implement fail-safe measures that allow vehicles to respond gracefully to unexpected conditions. Recent studies indicate that leveraging GenAI for predictive world modelling can enhance vehicle safety, reduce accidents, and improve situational awareness in densely populated or challenging environments [47].

Other key aspects of trustworthiness are transparency, explainability, and interpretability. Users are more likely to trust AI-driven vehicles if they can understand, at least at a high level, why certain decisions are made. This is particularly important in situations involving sudden manoeuvres or accident avoidance, where opaque behaviour can reduce confidence even if the outcome is safe.

XAI has emerged as a crucial requirement for the next generation of autonomous vehicles. The inherent "black-box" nature of many DL models challenges the ability of engineers, regulators, and end-users to understand AI-driven decisions. Techniques that combine vision, language, and action (such as LLMs and VLMs) offer potential solutions by providing interpretable narratives that can be monitored both during operation and through post-trip analysis. Developing robust XAI frameworks is also critical for regulatory compliance and user accountability. With nearly 94% of road accidents attributed to human error, a significant motivation for the adoption of autonomous vehicles, the need for clear, interpretable AI decisions is essential to bridging the trust gap between AI systems and human stakeholders [48].

Validation and verification are central to building trust. Unlike traditional SW, AI systems learn from data, which makes their behaviour harder to predict and formally verify. Ensuring that models perform safely across diverse real-world scenarios, including rare edge cases, remains a major challenge.

One of the main challenges in developing trust is handling uncertainty in complex environments. Weather conditions, unpredictable human behaviour, and incomplete sensor data can all affect decision-making. AI systems must not only perform well under ideal conditions but also degrade gracefully when conditions worsen.

Another challenge is bias and data limitations. If the training data does not adequately represent all driving environments or populations, the system may perform unevenly, posing safety risks and eroding trust. Continuous data collection and updating are necessary, but they introduce further complexity in maintaining consistency and safety.

Human-machine interaction plays a significant role in AIDVs. Misunderstandings about system capabilities, overreliance, or lack of proper user awareness can lead to misuse. Building intuitive interfaces and clearly communicating system limits is necessary to ensure that trust is appropriate and not misplaced.

Cybersecurity is a paramount concern for AIDVs. AI-driven vehicles are connected systems that can be vulnerable to attacks, potentially compromising safety and eroding public trust. Ensuring secure communication, resilient architectures, and rapid response to threats is essential.

As vehicles become more connected and reliant on SW, they also become more vulnerable to cyberattacks. A successful attack could have devastating consequences, ranging from data theft to remote control of vehicles. Comprehensive surveys highlight the multi-layered security and privacy challenges in V2X systems that must be addressed [39, 49].

When addressing security vulnerabilities in an AIDV, several approaches can be employed to strengthen data exchange against common threats. Trust Execution Environments (TEEs), both SW- and HW-based, and secure enclaves, more HW-oriented, can be used to isolate sensitive code or data from the main OS and applications [50]. These measures can enhance system robustness against potentially compromised applications, thereby addressing one of the major V2X threats: OTA update manipulation. Additional security risks arise from man-in-the-middle and side-channel attacks, which can be mitigated using advanced cryptographic approaches such as secure multi-party computation or homomorphic encryption [51].

In scenarios utilising heterogeneous sensor analytics, e.g., different kinds of cameras, privacy-preserving protection is essential [52]. Anonymisation algorithms can be deployed to de-identify faces of individuals, household interiors or license plates of cars detected around moving AIDVs. FL offers another solution by sharing only model updates (gradients) rather than raw data, dramatically reducing privacy exposure in communications [53].

1.3.3 Ethics

Regulatory and ethical challenges further complicate the development of trust. There is still no universal agreement on safety standards, liability in case of accidents, or acceptable risk levels. These uncertainties make it harder for the public to fully trust the introduction of new technologies.

AIDVs' concept and the implementation of autonomous functions that allow the vehicles to operate independently raise questions about how much control should remain with humans and when the system should override human decisions. These issues are interlinked with the AIDV data and AI models' dependence, since these vehicles rely on massive amounts of data from cameras, Global Navigation Satellite System (GNSS), and other sensors, as well as AI models that process the data to enable the vehicles to function safely.

Safety is a central ethical issue. While AI can reduce human error, it can also fail in unpredictable ways. Determining who is responsible in the event of an accident (the manufacturer, the SW developer, or the user) is a complex legal and moral challenge.

Bias in AI systems is another concern. If the data used to train these systems is incomplete or skewed, the vehicle may make unfair or unsafe decisions, such as misidentifying pedestrians or behaving differently in certain environments.

Privacy is also at stake. AIDVs collect and process large amounts of personal and location data, raising concerns about surveillance, data misuse, and consent. Users may not fully understand how their data is being stored or shared.

AIDVs' deployment can have broader societal impacts as widespread adoption could affect jobs in driving professions and change urban infrastructure. Ensuring that these technologies are accessible and beneficial to all, rather than only to a privileged group, is an ongoing ethical challenge.

The ethical dimensions of AI in autonomous vehicles are a subject of intense debate. The "trolley problem," where the vehicle must make a split-second decision in an unavoidable accident; the thorny issue of liability attribution; the primary concern of data privacy; the chilling prospect of hacking vulnerabilities; and the potential for overall job displacement are well-known examples [54]. SDVs and AIDVs must make decisions that carry ethical dimensions that are increasingly significant and safety-critical, as choosing a specific trajectory determines how risks are distributed among traffic participants [55].

Programming these ethical choices into machines is a complex and contentious issue that requires broad societal input and careful consideration by developers. Policymakers, standardisation organizations and vehicle producers must conceptualize what (shall) constitute(s) ethical decision-making for SDVs and AIDVs and integrate ethical, legal and engineering considerations into the development process by defining approaches on computational ethics (particularly in autonomous driving) while offering practitioners in the automotive sector a decision-making process for SDVs and AIDVs that is technically viable, legally permissible, ethically grounded and adaptable to societal values [55].

1.3.4 Architecture

The vehicle architecture has evolved from the domain architecture, a distributed system defined by function, to the zonal architecture, which addresses the limitations of the domain approach and provides the physical backbone for SW-defined mobility, to the SDV architecture, which is a paradigm shift that leverages zonal architecture to decouple SW from underlying HW opening the way for AI-defined architecture that builds upon the SDV and zonal foundation by deeply integrating AI, ML, DL, GenAI (SLMs, LLMs, VLMs, etc.) and agentic AI into core operations.

The evolution of computing platform design in the ITS domain is ideally based on the chiplets concept and RISC-V architectures, as these are two key ingredients that can foster faster productisation, lower costs, and quicker market adoption of newly developed solutions, especially those devised within the European market.

Chiplet architectures facilitate heterogeneous integration and high-bandwidth interconnects, such as the Universal chiplet interconnect express (UCIe), enabling higher processing density and improved power efficiency for demanding applications like ADAS and AI-driven electronic units. The flexibility inherent in Chiplet designs helps facilitate the transition towards centralised, zonal E/E architectures emerging in the AIDV, which can be upgraded and improved over the lifespan of a vehicle [56, 57, 58].

The feasibility of the AIDV architecture relies heavily on an open, scalable HW foundation, a role increasingly filled by RISC-V, which has reached a level of industrial maturity sufficient for safety-critical automotive applications, including those requiring the highest ISO 26262 ASIL-D integrity levels. The AIDV functions are defined by SW and AI, and the underlying HW must ensure the necessary foundation. RISC-V brings openness and flexibility to HW and offers standardisation and an ecosystem to support development toward certified-ready RISC-V Microcontroller Units (MCUs) for multiple Automotive Safety Integrity Levels (e.g., ASIL-B, ASIL-D) and toward compliance with automotive security standards (e.g., ISO21434) [59, 60].

The introduction of heterogeneous platforms incorporating conventional processing units (CPUs and GPUs) and more energy efficient and effective architectures (TPUs, NPUs, HW accelerators, neuromorphic computing-based devices) has the potential to drastically reduce energy consumption, e.g., an order of magnitude in the case of neuromorphic-based applications, and improve the capability to deliver complex real-time value-added services running on dedicated HW (HW accelerators), capable of processing the needed AI workloads at the edge.

1.3.5 Job transformation

The transition to AIDVs will have a significant impact on the workforce as well. Jobs in areas such as trucking, taxi services, and delivery will be impacted. Proactive measures, such as retraining programs and social safety nets, will be needed to mitigate the social and economic consequences of this shift.

Engineering careers will also be transformed by AIDVs, requiring a re-focus of skills: mechanical engineers will evolve into mechatronics specialists, electrical engineers into SW-HW integration experts, and service technicians into SW diagnostics and remote troubleshooting specialists, among other role transformations. In summary, most of the required skills will have to be based on AI-focused methods, combining elements of SW and HW design, AI frameworks and Data [61].

1.3.6 Regulation and Standardisation

The fragmented global regulatory landscape is a hindrance to the broader development of AIDVs. Alignment among regulatory agencies and agreements on world-recognised legislation, such as the EU AI Act, will enable a much smoother and more effective deployment of AIDVs [62]. Similar fragmentation exists within standardisation bodies, where multiple international entities develop technical specifications that often overlap and occasionally conflict on key aspects [63]. Greater alignment and clearer responsibility distribution are essential to tackle the complex challenge of AIDV standardisation.

Finally, it is worth noting that a balance between regulation and free-market rules, and between safety and security, is to be found; otherwise, the market entry of potentially breakthrough innovations and life-saving new features may be hindered by overly strict rules required to fulfil the regulations.

The convergence of AI and advanced communication is steering the automotive industry into an era of innovation. The AIDV, with its promise of enhanced safety, efficiency, and personalisation, represents a transformative vision for the future of both personal mobility and ITSs. Realising this vision will require a concerted effort from researchers, technologists, policymakers, and society to address the significant technical, ethical, and societal challenges that lie ahead. The road to the AI-defined future is still under construction, but its direction is clear [64].

1.4 Future Research Directions

The evolution of DTs into immersive triplets offers a transformative approach for predictive modelling and real-time operations in AIDVs. Future research should prioritise creating highly accurate, scalable virtual representations that seamlessly integrate vehicles with edge devices and cloud platforms. Establishing interoperability standards for these simulations is essential to ensure they function across various manufacturers. This enables rigorous AI testing in virtual environments, significantly reducing costs and the need for physical prototypes while accelerating development timelines.

As AIDV's architecture becomes more complex, new research is needed to develop virtual validation, verification, and benchmarking pipelines. This involves developing AI-based development tools that combine SW workflows with GenAI and agentic AI to safely implement automotive functions. High-fidelity simulations must be advanced to recreate rare, hazardous, or complex traffic scenarios, enabling automakers to train and validate integrated systems spanning HW, SW, and AI algorithms without the risks of real-world testing.

GenAI and synthetic data approaches are expected to accelerate the training and evaluation of AIDVs' autonomy across diverse scenarios, while multi-agent systems and collaborative frameworks foster cooperative behaviour among autonomous fleets, optimising urban mobility. Improvements in operational workflows that facilitate human–machine collaboration, along with robust V2X communications and distributed predictive maintenance systems, enhance overall vehicle performance. Further research on swarm intelligence is needed to develop scalable, coordinated transportation systems that ultimately redefine how future cities and vehicles interact.

End-to-end AI models that integrate RL and imitation learning will strengthen decision-making and planning. These approaches allow systems to learn both from experience and from observed behaviour, improving adaptability in complex situations. GenAI and synthetic data are increasingly important for training robust systems. They enable the simulation of rare edge cases that are difficult to capture in real-world datasets. Agentic AI and multi-agent systems will support cooperative driving and swarm intelligence. These capabilities can help optimise traffic flow and improve overall urban mobility.

To address the growing energy demands of intensive AI models, energy efficiency and sustainability must become central research areas. This includes developing energy-efficient algorithms using model compaction techniques and optimising edge computing to process data locally rather than in the cloud. Future investigations should conduct full AI life-cycle assessments and explore renewable energy integration to minimise the ecological footprint of AIDVs.

Interdisciplinary collaboration is crucial for addressing challenges in interoperability, cybersecurity, explainability, and ethics. Research must extend beyond technical specifications to address the ethical decision-making of AI, focusing on bias and accountability, while also safeguarding data privacy across the V2X ecosystem. Partnerships between automakers, semiconductor companies, edge, and cloud providers are necessary to standardise edge computing communications and create lightweight, scalable architectures that can operate efficiently on resource-constrained devices.

The development of end-to-end unified AI compute platforms and safety-certified operating systems is vital for the autonomy continuum. Future work envisions advancing vehicle architectures through RISC-V System-on-Chips (SoCs), heterogeneous chiplets interfaces, safety-certified MCUs, and the integration of neuromorphic, quantum sensing, and communication technologies.

Future research is needed to define and develop AIDV's full technology stack, a comprehensive safety system that unifies vehicle architecture, AI models, chips, SW, tools, and services to ensure the safe development of autonomous vehicles across the computing and communication continuum from vehicle to edge and cloud.

Finally, research must focus on interdisciplinary approaches to develop components, systems, tools, workflows, and platforms to be integrated into a robust AIDV ecosystem that supports the entire design, training, simulation, and in-vehicle processing loop, ensuring that real-time AI operates with end-to-end reliability from the vehicle to the edge and the cloud.

Acknowledgements

This publication has received funding through the Chips JU MOSAIC project. The Chips JU MOSAIC "A MOSAIC of Essential Electronic Components and Systems (ECS) for our Automated Digital Future in Industry and Mobility" project is supported by the Chips Joint Undertaking and its members, including top-up funding by Austria, Belgium, Denmark, France, Germany, Greece, Israel, Italy, Latvia, Netherlands, Norway, Poland and Turkey under Grant Agreement No. 101194414.

References

[1] B. Zheng *et al.*, "Online Location Planning for AI-Defined Vehicles: Optimizing Joint Tasks of Order Serving and Spatio-Temporal Heterogeneous Model Fine-Tuning," *IEEE Transactions on Mobile Computing*, vol. 25, no. 2, pp. 1777-1794, Feb. 2026, https://www.doi.org/10.1109/TMC.2025.3608179.

[2] L. Zhao *et al.*, "Vehicular Communications: Standardization and Open Issues," *IEEE Communications Standards Magazine*, vol. 2, no. 4, pp. 74-80, Dec. 2018, https://doi.org/10.1109/MCOMSTD.2018.1800027.

[3] S. E. Shladover, "Connected and automated vehicle systems: Introduction and overview," *Journal of Intelligent Transportation Systems*, vol. 22, no. 3, pp. 190–200, Jun. 2018, https://doi.org/10.1080/15472450.2017.1336053.

[4] A. Mohanty *et al.*, "Age of Computational AI for Autonomous Vehicles," *Artificial Intelligence for Autonomous Vehicles: The Future of*

Driverless Technology, Wiley, 2024, pp. 25-54, https://doi.org/10.1002/9781119847656.ch2.

[5] T. Gong *et al.*, "Edge Intelligence in Intelligent Transportation Systems: A Survey," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 9, pp. 8919–8944, Sep. 2023, https://doi.org/10.1109/tits.2023.3275741.

[6] D. Greenbaum, "Autonomous vehicles: innovations, challenges, and implications for society and law," *Oxford University Press eBooks*, May 2025, https://doi.org/10.1093/9780198972877.003.0034.

[7] K. Yadav *et al.*, "Design Considerations and Framework Analysis for Software-Defined Autonomous Vehicles," in *2024 IEEE 99*$^{\mathrm{th}}$ *Vehicular Technology Conference (VTC2024-Spring)*, Singapore, 2024, pp. 1-5, https://doi.org/10.1109/VTC2024-Spring62846.2024.10683354.

[8] Morgan Stanley. "Software-Defined Vehicles: The Future of Cars," *Morgan Stanley Research*. https://www.morganstanley.com/ideas/software-defined-vehicles-outlook, 2024.

[9] W. Tong *et al.*, "Artificial Intelligence for Vehicle-to-Everything: A Survey," *IEEE Access*, vol. 7, pp. 10823-10843, 2019, https://doi.org/10.1109/ACCESS.2019.2891073.

[10] O. Vermesan *et al.*, "Automotive Intelligence Embedded in Electric Connected Autonomous and Shared Vehicles Technology for Sustainable Green Mobility," *Frontiers in Future Transportation*, vol. 2, Aug. 2021, https://doi.org/10.3389/ffutr.2021.688482.

[11] P. Koopman *et al.*, "Ethics, Safety, and Autonomous Vehicles", Virtual roundtable. Computer. Dec. 2021. pp 28-37. [Online] Available: https://users.ece.cmu.edu/~koopman/pubs/koopman21_Ethics_Safety_AVs_IEEE_Roundtable.pdf.

[12] Y. Mao *et al.*, "A Survey on Spatio-Temporal Prediction: From Transformers to Foundation Models," *ACM Computing Surveys*, Sep. 2025, https://doi.org/10.1145/3766546.

[13] A. Kuznietsov *et al.*, "Explainable AI for Safe and Trustworthy Autonomous Driving: A Systematic Review," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 12, pp. 19342-19364, Dec. 2024, https://doi.org/10.1109/TITS.2024.3474469.

[14] E. Haque *et al.*, "Towards an Interpretable AI Framework for Advanced Classification of Unmanned Aerial Vehicles (UAVs)," in *2024 IEEE 21*st *Consumer Communications & Networking Conference (CCNC)*, Las Vegas, NV, USA, 2024, pp. 644-645, https://doi.org/10.1109/CCNC51664.2024.10454862.

[15] J. Huang *et al.,* “Reinforcement Learning Based Resource Management for 6G-Enabled mIoT With Hypergraph Interference Model,” *IEEE Transactions on Communications*, vol. 72, no. 7, pp. 4179-4192, July 2024, https://doi.org/10.1109/TCOMM.2024.3372892.

[16] S. Atakishiyev *et al.,* “Explainable Artificial Intelligence for Autonomous Driving: A Comprehensive Overview and Field Guide for Future Research Directions,” *IEEE Access*, vol. 12, pp. 101603-101625, 2024, https://doi.org/10.1109/ACCESS.2024.3431437.

[17] E. Peltonen *et al.,* “Towards Data-Centric and Context-Aware Decision Making in Software-Defined Vehicles,” in *2025 IEEE 22nd International Conference on Software Architecture Companion (ICSA-C)*, Odense, Denmark, 2025, pp. 574-577, https://doi.org/10.1109/ICSA-C65153.2025.00085.

[18] E. Yurtsever *et al.,* “A Survey of Autonomous Driving: Common Practices and Emerging Technologies,” *IEEE Access*, vol. 8, pp. 58443-58469, 2020, https://doi.org/10.1109/ACCESS.2020.2983149.

[19] X. Zhou *et al*., “Vision Language Models in Autonomous Driving: A Survey and Outlook,” *IEEE Transactions on Intelligent Vehicles*, pp. 1–20, 2024, https://www.doi.org/10.1109/TIV.2024.3402136.

[20] L. Chen *et al*., “End-to-End Autonomous Driving: Challenges and Frontiers,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 46, no. 12, pp. 10164-10183, Dec. 2024, https://www.doi.org/10.1109/TPAMI.2024.3435937.

[21] A. Tampuu *et al*., “A Survey of End-to-End Driving: Architectures and Training Methods,” *IEEE Transactions on Neural Networks and Learning Systems*, vol. 33, no. 4, pp. 1364-1384, April 2022, https://doi.org/10.1109/TNNLS.2020.3043505.

[22] N. Hameed Hussein *et al*., “SDN-Based VANET Routing: A Comprehensive Survey on Architectures, Protocols, Analysis, and Future Challenges,” *IEEE Access*, vol. 13, pp. 126801-126861, 2025, https://doi.org/10.1109/ACCESS.2024.3355313.

[23] S. D. A. Shah, M. A. Gregory, F. Bouhafs and F. D. Hartog, “Artificial Intelligence-Defined Wireless Networking for Computational Offloading and Resource Allocation in Edge Computing Networks,” *IEEE Open Journal of the Communications Society*, vol. 5, pp. 2039-2057, 2024, https://doi.org/10.1109/OJCOMS.2024.3382265.

[24] E. Kartsakli *et al*., “AI-Powered Edge Computing Evolution for Beyond 5G Communication Networks,” in *2023 Joint European Conference on*

Networks and Communications & 6G Summit (EuCNC/6G Summit), Gothenburg, Sweden, 2023, pp. 478-483, https://doi.org/10.1109/EuCNC/6GSummit58263.2023.10188371.

[25] O. Vermesan and V. Frascolla, Eds., "Edge IoT Industrial Immersive Technologies and Spatial Computing Continuum," AIOTI White Paper, Alliance for Internet of Things Innovation, Brussels, Belgium, 2024. [Online]. Available: https://aioti.eu/wp-content/uploads/AIOTI-Paper-Edge-AI-IoT-Immersive-Technologies-Published.pdf.

[26] J. Hwang, K. -W. Lim and D. Kim, "What If Vehicles are Controlled by the Road?," *IEEE Transactions on Intelligent Vehicles*, vol. 10, no. 4, pp. 2219-2233, Apr. 2025, https://doi.org/10.1109/TIV.2024.3372007.

[27] K. Kandali and S. Nouh, "AI-Native V2X and Internet of Vehicles Systems: Architectures, Learning Paradigms, and Open Challenges," Feb. 2026, https://doi.org/10.36227/techrxiv.177006069.92717813/v1.

[28] Xu et al., "A Low-Latency and Massive-Connectivity Vehicular Fog Computing Frame-work for 5G," in *2018 IEEE Globecom Workshops (GC Wkshps)*, Abu Dhabi, United Arab Emirates, 2018, pp. 1-6, https://doi.org/10.1109/GLOCOMW.2018.8644428.

[29] ISO 26262:2018, "Road vehicles — Functional safety," International Organization for Standardization, Geneva, Switzerland, 2018. [Online]. Available: https://www.iso.org/publication/PUB200262.html.

[30] D. Hepp and M. Harrysson. "Beyond CES 2025: How gen AI is revolutionizing vehicle tech." *McKinsey & Company*. [Online]. Available: https://www.mckinsey.com/featured-insights/themes/beyond-ces-2025how-gen-ai-is-revolutionizing-vehicle-tech.

[31] Jesus Gutierrez *et al.*, "Seamless Integration of Efficient 6G Wireless Technologies for Communication and Sensing Enabling Ecosystems," *IFIP advances in information and communication technology*, pp. 190–201, Jan. 2024, doi: https://doi.org/10.1007/978-3-031-63227-3_13.

[32] X. Li *et al.,* "MultiX: Advancing 6G-RAN Through Multi-Technology, Multi-Sensor Fusion, Multi-Band and Multi-Static Perception," *IEEE Wireless Communications*, 2025, https://doi.org/10.1109/mwc.2025.3629589.

[33] D. Firmani *et al.,* "INTEND: Intent-Based Data Operation in the Computing Continuum," in *International Conference on Advanced Information Systems Engineering (CaiSE) 2024*, Limassol, Cyprus, 03-07.07.2024, pp. 43-50.

[34] Z. Xi *et al.,* “The rise and potential of large language model based agents: a survey,” *Science China Information Sciences*, vol. 68, no. 2, Jan. 2025, https://doi.org/10.1007/s11432-024-4222-0.
[35] W. Raslan *et al.*, “Smart Vehicle Safety: AI-Driven Driver Assistance and V2X Communications,” in *2024 International Telecommunications Conference (ITC-Egypt)*, Cairo, Egypt, 2024, pp. 787-792, https://doi.org/10.1109/ITC-Egypt61547.2024.10620463.
[36] K. Agrawal *et al.*, “Advancing Software-Defined Vehicles: An End-to-End Framework with Digital Twin Based Attestation for OTA Updates,” in *2025 17$^{\text{th}}$ International Conference on COMmunication Systems and NETworks (COMSNETS)*, Bengaluru, India, 2025, pp. 514-522, https://doi.org/10.1109/COMSNETS63942.2025.10885560.
[37] M. Eskandari *et al.*, “Visual GANs for End-to-End UAV Trajectory Generation in RIS-Assisted Energy-Efficient Wireless Vehicular Networks,” *Green Energy and Intelligent Transportation*, pp. 100365–100365, Sep. 2025, https://doi.org/10.1016/j.geits.2025.100365.
[38] L. Zhang *et al.*, “Integrated Motion Planning for On-Ramp Merging Based on Stackelberg Game Modeling Considering Interactive Characteristics,” *IEEE Transactions on Vehicular Technology*, vol. 74, no. 8, pp. 11762-11776, Aug. 2025, https://doi.org/10.1109/TVT.2025.3554978.
[39] T. Yoshizawa *et al.*, “A Survey of Security and Privacy Issues in V2X Communication Systems,” *ACM Computing Surveys*, Aug. 2022, https://doi.org/10.1145/3558052.
[40] Satyabrata Pradhan, “Agentic AI for Software-Defined Vehicles: A Generative Testing Framework for Autonomous Feature Assurance,” *Journal of Electrical Systems*, vol. 21, no. 01, pp. 993–1001, Apr. 2025, https://doi.org/10.52783/jes.9181.
[41] Mao *et al.*, “GPT-Driver: Learning to Drive with GPT,” in *37th Annual Conference on Neural Information Processing Systems (NeurIPS)*, New Orleans, Dec. 2023, https://doi.org/10.48550/arXiv.2310.01415.
[42] J. Wei *et al.*, “OccLLaMA: An Occupancy-Language-Action Generative World Model for Autonomous Driving,” *arXiv (Cornell University)*, Sep. 2024, https://doi.org/10.48550/arxiv.2409.03272.
[43] J. Mao, J. Ye, Y. Qian, M. Pavone, and Y. Wang, “A Language Agent for Autonomous Driving,” *arXiv (Cornell University)*, Nov. 2023, https://doi.org/10.48550/arxiv.2311.10813.
[44] World Health Organization. *Global status report on road safety 2023*. https://assets.bbhub.io/dotorg/sites/64/2023/12/WHO-Global-status-report-on-road-safety-2023.pdf.

[45] International Transport Forum. Road Safety Annual Report 2024. https://www.itf-oecd.org/sites/default/files/docs/irtad-road-safety-annual-report-2024.pdf.

[46] ISO/PAS 21448:2019, "Road vehicles-Safety of the intended functionality," International Organization for Standardization, Geneva, CH, 2019. [Online]. Available: https://www.iso.org/standard/70939.html.

[47] Wayve. "Pioneering a New Paradigm in AV Safety: Solving Long-Tail Challenges," Jul., 2024. https://wayve.ai/technology/safety-framework/

[48] S. Atakishiyev *et al.*, "Explainable Artificial Intelligence for Autonomous Driving: A Comprehensive Overview and Field Guide for Future Research Directions," *IEEE Access*, vol. 12, pp. 101603-101625, 2024, https://doi.org/10.1109/ACCESS.2024.3431437.

[49] S. E. Shladover, "Connected and automated vehicle systems: Introduction and overview," *Journal of Intelligent Transportation Systems*, vol. 22, no. 3, pp. 190–200, Jun. 2018, https://doi.org/10.1080/15472450.2017.1336053.

[50] Y. Bai *et al.*, "SecPaging: Secure Enclave Paging with Hardware-Enforced Protection against Controlled-Channel Attacks," pp. 1–6, Jun. 2024, https://doi.org/10.1145/3649329.3658241.

[51] Y. Liang *et al.*, "Secure Multi-Party Computation on the Encrypted Network: a Lattice-based Multi-Party Homomorphic Encryption Protocol," in *2024 9$^{\text{th}}$ International Conference on Communication, Image and Signal Processing (CCISP)*, Gold Coast, Australia, 2024, pp. 112-117, https://doi.org/10.1109/CCISP63826.2024.10765563.

[52] W. Abbasi *et al.*, "Privacy vs Accuracy Trade-Off in Privacy Aware Face Recognition in Smart Systems," *2022 IEEE Symposium on Computers and Communications (ISCC)*, Rhodes, Greece, 2022, pp. 1–8, doi:10.1109/ISCC55528.2022.9912465.

[53] L. Skovajsova *et al.*, "A Review of Multi-Objective and Multi-Task Federated Learning Approaches," in *2025 IEEE 23rd World Symposium on Applied Machine Intelligence and Informatics (SAMI)*, Stara Lesna, Slovakia, 2025, pp. 000035-000040, https://doi.org/10.1109/SAMI63904.2025.10883172.

[54] K. Shah and E. Guven, "Exploring Ethical Issues and Challenges of Autonomous Vehicles in Different Countries and Cultures," in *2025 59$^{\text{th}}$ Annual Conference on Information Sciences and Systems (CISS)*, pp. 1–6, Mar. 2025, https://doi.org/10.1109/ciss64860.2025.10944682.

[55] F. Poszler *et al.*, "Ethical Decision-Making for Self-Driving Vehicles: A Proposed Model & List of Value-Laden Terms that Warrant (Technical)

Specification," *Science and Engineering Ethics*, vol. 30, no. 5, Oct. 2024, https://doi.org/10.1007/s11948-024-00513-0.

[56] P. Onufryk and S. Choudhary, "UCIe: Standard for an Open Chiplet Ecosystem," *IEEE Micro*, vol. 45, no. 1, pp. 16-25, Jan.-Feb. 2025, https://doi.org/10.1109/MM.2024.3451532.

[57] M. Odema *et al.*, "Performance Implications of Multi-Chiplet Neural Processing Units on Autonomous Driving Perception," in *2025 Design, Automation & Test in Europe Conference (DATE)*, Lyon, France, 2025, pp. 1-7, https://doi.org/10.23919/DATE64628.2025.10993228.

[58] T. Uehling *et al.*, "Chiplet Package for Automotive and Edge Processors," in *2025 IEEE 75*$^{\text{th}}$ *Electronic Components and Technology Conference (ECTC)*, Dallas, TX, USA, 2025, pp. 1259-1263, https://doi.org/10.1109/ECTC5168.

[59] S. Roy et al., "Exception Coverage on Automotive Processors," *IEEE Embedded Systems Letters*, vol. 18, no. 1, pp. 27–30, Feb. 2026, https://doi.org/10.1109/LES.2025.3571040.

[60] L. Cuomo *et al.,* "Towards a RISC-V Open Platform for Next-generation Automotive ECUs," in *2023 12*$^{\text{th}}$ *Mediterranean Conference on Embedded Computing (MECO)*, 2023, pp. 1-8, https://doi.org/10.1109/MECO58584.2023.10154913.

[61] K.-H. Wang and W.-C. Lu, "AI-induced Job impact: Complementary or substitution? Empirical Insights and Sustainable Technology Considerations," *Sustainable Technology and Entrepreneurship*, vol. 4, no. 1, 2024, https://doi.org/10.1016/j.stae.2024.100085.

[62] European Parliament and Council of the European Union, "Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) and amending Regulations (EC) No 300/2008, (EU) No 167/2013, (EU) No 168/2013, (EU) 2018/858, (EU) 2018/1139 and (EU) 2019/2144 and Directives 2014/90/EU, (EU) 2016/797 and (EU) 2017/2394," *Official Journal of the European Union*, vol. L 1689, July 12, 2024. [Online]. Available: https://artificialintelligenceact.eu/the-act/.

[63] B. Raaf *et al.,* "Key technology advancements driving mobile communications from generation to generation," *Intel Technology Journal*, vol. 18, no. 1, 2014.

[64] U. Rajasekaran *et al.,* "Artificial Intelligence in Autonomous Vehicles - A Survey of Trends and Challenges," pp. 1–24, Feb. 2024, https://doi.org/10.1002/9781119847656.ch1.