

Research Article

Open Access



Deep learning for autonomous driving systems: technological innovations, strategic implementations, and business implications - a comprehensive review

Laxmi Kant Sahoo, Vijayakumar Varadarajan

Department Swiss School of Business and Management, GBC - Geneva Business Center, Genève 1213, Switzerland.

Correspondence to: Dr. Laxmi Kant Sahoo, Department Swiss School of Business and Management, GBC - Geneva Business Center, Avenue des Morgines 12, Genève 1213, Switzerland. E-mail: laxmikanta.lk@gmail

How to cite this article: Sahoo LK; Varadarajan V Deep learning for autonomous driving systems: technological innovations, strategic implementations, and business implications - a comprehensive review. *Complex Eng Syst.* 2025, 5, 2. <https://dx.doi.org/10.20517/ces.2024.83>

Received: 24 Oct 2024 **First Decision:** 6 Dec 2024 **Revised:** 17 Jan 2025 **Accepted:** 5 Feb 2025 **Published:** 18 Feb 2025

Academic Editor: Zhaojing Wu **Copy Editor:** Ping Zhang **Production Editor:** Ping Zhang

Abstract

The rapid advancements in deep learning have significantly transformed the landscape of autonomous driving, with profound technological, strategic, and business implications. Autonomous driving systems, which rely on deep learning to enhance real-time perception, decision-making, and control, are poised to revolutionize transportation by improving safety, efficiency, and mobility. Despite this progress, numerous challenges remain, such as real-time data processing, decision-making under uncertainty, and navigating complex environments. This comprehensive review explores the state-of-the-art deep learning methodologies, including Convolutional Neural Networks (CNNs), Recurrent Neural Networks, Long Short-Term Memory networks, and transformers that are central to autonomous driving tasks such as object detection, scene understanding, and path planning. Additionally, the review examines strategic implementations, focusing on the integration of deep learning into the automotive sector, the scalability of artificial intelligence-driven systems, and their alignment with regulatory and safety standards. Furthermore, the study highlights the business implications of deep learning adoption, including its influence on operational efficiency, competitive dynamics, and workforce requirements. The literature also identifies gaps, particularly in achieving full autonomy (Level 5), improving sensor fusion, and addressing the long-term costs and regulatory challenges. By addressing these issues, deep learning has the potential to redefine the future of mobility, enabling safer, more efficient, and fully autonomous driving systems. This review aims to



© The Author(s) 2025. **Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, sharing, adaptation, distribution and reproduction in any medium or format, for any purpose, even commercially, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.



provide insights for stakeholders, including automotive manufacturers, artificial intelligence developers, and policymakers, to navigate the complexities of integrating deep learning into autonomous driving.

Keywords: Deep learning, autonomous driving, object detection, Advanced Driver Assistance Systems (ADAS), artificial intelligence in transportation, convolutional neural networks (CNNs)

1. INTRODUCTION

Road accidents remain a pressing global issue, with human error responsible for a significant 94% of incidents, as reported by the National Highway Traffic Safety Administration (NHTSA)^[1]. Leading causes of these accidents include impaired driving due to alcohol (40%), speeding (30%), and reckless driving (33%)^[2]. Distracted driving also plays a critical role in road fatalities. Autonomous vehicle (AV) technology emerges as a promising solution to mitigate these risks, either by augmenting the capabilities of human drivers or through full automation. Advanced Driver Assistance Systems (ADAS) and Automated Driving Systems (ADS) have been designed not only to enhance safety by preventing accidents but also to improve fuel efficiency, reduce emissions, and alleviate the mental strain of driving^[3]. Additionally, these technologies hold the potential for transforming mobility for people with disabilities, offering increased independence through driverless solutions.

AVs function as sophisticated decision-making systems, leveraging data streams from multiple onboard sensors such as cameras, radars, light detection and ranging (LiDAR), ultrasonic sensors, and GPS units to assess and respond to their environment. This sensor data is processed in real time by embedded computing systems, allowing AVs to make driving decisions autonomously. Effective autonomous operation requires not only environmental perception but also advanced path-planning algorithms and control over acceleration, braking, and steering. Decision-making in AVs is typically executed through either a modular pipeline-perception, planning, and action-or an End-to-End (End2End) learning model, where sensory inputs are directly converted into control commands.

This paper explores the role of deep learning in autonomous driving, focusing on its technological innovations, strategic implementations, and business implications. It will examine the impact of deep learning on key stakeholders, including automotive manufacturers, suppliers, service providers, and electric vehicle innovators, while addressing how this technology is shaping the future of transportation. By investigating the integration of deep learning into autonomous driving systems, this research seeks to highlight its transformative potential in enhancing road safety, revolutionizing mobility, and reshaping the automotive industry.

2. METHODS

Society of Automotive Engineers (SAE) defined six levels (L0 to L5) of automation for AVs^[4]. Level 0 vehicles are those under the full control of drivers. Level 1 allows automation of either the braking or steering system of the car and the rest of the control is with the human driver, e.g., adaptive cruise control. Level 2 cars can take some safety actions by automation of more than one system at a time, such as the smart pilot feature in XUV700, where the vehicle will do adaptive cruise control and automatic emergency braking at the same time. At level 3, the car can automatically drive in certain conditions by monitoring the surrounding environment. However, the human driver must still be on command to take control if the autonomous system fails. Daimler^[5] claimed that its S-class models featuring Automatic Lane change and Autobahn chauffeur have Level 3. In the case of Level 4, the car can safely take control and proceed accordingly if its request for human intervention is not responded to.

Level 4 cars are not recommended to be driven in uncertain weather conditions or unmapped areas. Lastly, level 5 vehicles cover full automation in all conditions and modes. Among deep learning techniques, Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long short-term memory (LSTM), and Deep Reinforcement Learning (DRL) are the most common methodologies applied to autonomous driving^[6].

CNNs are mainly used for processing spatial information, such as images, and can be viewed as image feature extractors and universal non-linear function approximators^[7]. Before the rise of deep learning, computer vision systems used to be implemented based on handcrafted features, such as HAAR-like features and Histograms of Oriented Gradients (HoG). Compared to these traditional handcrafted features, convolutional neural networks can automatically learn a representation of the feature space encoded in the training set. CNNs can be loosely understood as very approximate analogies to different parts of the human visual cortex^[8]. They are efficiently used for object and distance estimation^[9], vulnerable road user detection, lane detection and path prediction^[10], traffic sign recognition^[11], and visual localization^[12].

RNNs are especially good at processing temporal sequence data, such as text or video streams. Unlike conventional neural networks, an RNN contains a time-dependent feedback loop in its memory cell. The main challenge in using basic RNNs is the vanishing gradient encountered during training. LSTM networks are non-linear function approximators for estimating temporal dependencies in sequence data. As opposed to traditional recurrent neural networks, LSTM solves the vanishing gradient problem by incorporating three gates, which control the input, output, and memory state. RNN and LSTM networks are used for pose estimation^[13] and path planning^[14] in autonomous driving.

A significant advancement in this domain is the adoption of DRL models. Yang *et al.*^[14] introduced a decision-making framework for highway driving based on the Deep Deterministic Policy Gradient (DDPG) algorithm. This reinforcement learning (RL) approach enables the direct mapping of environmental observations to actionable driving decisions. By leveraging DDPG, AVs can learn optimal driving strategies in continuous action spaces, allowing them to effectively handle complex traffic scenarios, including lane changes and overtaking maneuvers. Additionally, by assessing the uncertainty of the learned policy at runtime, the system can detect unfamiliar situations and adjust its decisions, enhancing both safety and robustness.

In this section, the researchers shall briefly discuss different areas of autonomous driving development where deep learning is used or has the potential to be used.

2.1. Technological innovations

2.1.1. Driving scene understanding

In autonomous driving, scene understanding is a crucial element, particularly in urban environments where vehicles must navigate through diverse traffic participants, complex road layouts, and dynamic interactions. Urban areas present significant challenges due to the wide variety of object appearances, frequent occlusions, and unpredictable behaviors of pedestrians, cyclists, and other vehicles. For autonomous systems to function effectively, they must accurately detect, classify, and track traffic participants while identifying safe drivable areas in real time.

Deep learning-based perception systems, particularly CNNs, have emerged as the dominant approach for addressing these challenges. CNNs have demonstrated their superiority in object detection and scene recognition tasks, achieving outstanding performance in large-scale competitions such as the ImageNet

Large Scale Visual Recognition Challenge (ILSVRC)^[9]. This success has led to the widespread adoption of CNNs in autonomous driving, where their ability to process high-dimensional sensor data from cameras, LiDAR, and radar makes them ideal for identifying road features, obstacles, and traffic participants.

CNNs are especially well-suited for the complexities of urban driving, where occlusions and variations in object appearance are common. Through multi-layer feature extraction, they can generalize across diverse environmental conditions, allowing for robust object recognition and classification even in highly dynamic settings^[15]. Continuous advancements in CNN architectures, such as Mask Region-based CNN (R-CNN) and Faster R-CNN, have further improved their ability to accurately segment drivable areas and detect objects at varying scales and distances^[16].

The ability of CNNs to manage real-time perception tasks has positioned them as the backbone of modern autonomous driving systems. Their use in detecting vehicles, pedestrians, cyclists, road signs, and other infrastructure elements is critical to ensuring safe navigation in densely populated urban areas^[17]. Moreover, the development of specialized deep-learning models tailored to urban environments has led to significant enhancements in the performance and reliability of autonomous systems, enabling them to adapt to the unpredictable nature of urban driving scenarios^[18].

CNNs have revolutionized urban scene understanding in autonomous driving by enabling precise detection, classification, and tracking of traffic participants and road features in real time. Advancements in architectures such as Mask R-CNN and Faster R-CNN have enhanced the ability to handle occlusions, varying object appearances, and dynamic interactions, ensuring safer navigation in complex environments. CNNs excel in generalizing across diverse conditions, offering robust object recognition in dynamic scenarios. Tailored deep-learning models further boost performance, enabling autonomous systems to adapt to unpredictable urban settings, positioning CNNs as the backbone of perception systems and paving the way for safer, smarter mobility solutions.

2.1.2. Object detection

Object detection is essential in autonomous driving systems as it enables vehicles to identify and track various objects in their environment, such as vehicles, pedestrians, and road signs. Accurate detection and classification of these objects are critical for safe navigation and decision-making in AVs. Two primary architectures have emerged in object detection: single-stage and double-stage detectors, each with specific advantages regarding speed and accuracy.

Single-stage detectors, including You Only Look Once (YOLO)^[19] and Single Shot MultiBox Detector (SSD)^[20], perform object detection in one pass, combining object localization and classification into a single network. These detectors are renowned for their speed and computational efficiency, making them ideal for real-time applications in autonomous driving, where quick decision-making is crucial^[6]. For example, the ability of YOLO to detect multiple objects in real time with low latency makes it suitable for dynamic environments. Similarly, SSD uses a set of default bounding boxes of various aspect ratios and scales for fast and efficient object detection.

More recent single-stage detectors, such as CornerNet^[21] and RefineNet^[22], have further enhanced detection accuracy while maintaining fast processing capabilities. These models improve accuracy through techniques such as keypoint-based detection (CornerNet) and multi-path refinement (RefineNet). However, despite these improvements, single-stage detectors often lag behind double-stage detectors in terms of accuracy.

Double-stage detectors, such as Faster R-CNN^[23] and Region-based Fully Convolutional Networks^[24] (R-FCN), separate the object detection process into two stages: region proposal generation and object classification. In the first stage, region proposals are generated to identify areas likely to contain objects, and in the second stage, the model classifies these objects and refines their bounding boxes. This two-step approach allows for greater accuracy, as the model spends more time refining its predictions. For instance, Faster R-CNN uses a Region Proposal Network (RPN) followed by object detection in the second stage, achieving higher accuracy, albeit at the cost of speed^[6]. Similarly, R-FCN uses fully convolutional layers, reducing computational complexity while maintaining high accuracy.

Stereo images are often used for distance prediction in autonomous driving systems^[9]. Stereo vision provides depth information by calculating the disparity between two images captured from slightly different angles, allowing the system to estimate the distance to detected objects. Integrating stereo vision with object detection enhances the vehicle's perception and enables more precise navigation and obstacle avoidance. By using stereo images, autonomous systems can detect objects and estimate distances simultaneously, improving overall safety and decision-making^[17].

While double-stage detectors are generally more accurate, their computational complexity often makes them slower than single-stage detectors. Therefore, ongoing research focuses on hybrid models that combine the advantages of both approaches. For example, models such as YOLOv4^[25] and EfficientDet aim to balance the speed of single-stage detectors with the accuracy of double-stage models. Additionally, combining object detection with stereo image-based distance prediction enhances the comprehensive perception system, improving both safety and vehicle effectiveness^[16].

While significant advancements have been made in object detection for autonomous driving, a gap exists in balancing the speed of single-stage detectors with the accuracy of double-stage models. Current hybrid approaches, such as YOLOv4 and EfficientDet, still face challenges in optimizing this trade-off. Additionally, integrating stereo vision with object detection for simultaneous depth estimation and object classification remains underexplored, particularly in dynamic and low-visibility conditions. Future research should focus on developing more efficient hybrid models and robust perception systems that seamlessly combine speed, accuracy, and depth estimation, addressing computational complexity while improving real-time performance and reliability in diverse environments.

2.1.3. Semantic and instance segmentation

Semantic and instance segmentation are essential tasks in computer vision, playing a crucial role in achieving complete scene understanding for applications such as autonomous driving, indoor navigation, and virtual and augmented reality. Both tasks involve identifying and classifying objects within an image, but they serve different purposes. Semantic segmentation assigns a class label to each pixel in an image, grouping pixels that belong to the same object or region, while instance segmentation not only classifies objects but also distinguishes between multiple instances of the same class^[15].

In autonomous driving, understanding the scene in a detailed and granular manner is critical for making real-time decisions. Semantic segmentation helps the vehicle identify road elements, such as lanes, road boundaries, and traffic signs, while instance segmentation allows the system to differentiate between individual vehicles, pedestrians, and cyclists. This ability to distinguish and track multiple objects and road elements simultaneously is essential for safe navigation.

Several semantic segmentation networks, such as SegNet, IC-Net, ENet, AdapNet, and Mask R-CNN, have emerged as powerful tools for pixel-wise classification. These architectures are typically encoder-decoder networks, where the encoder extracts features from the input image and the decoder maps these features back to the pixel level to produce the segmentation mask. For example, SegNet and ENet are known for their efficiency in real-time applications, making them suitable for resource-constrained systems such as AVs^[26]. IC-Net focuses on achieving high-resolution segmentation results with minimal computation, addressing the challenge of processing large input images in real-time applications such as autonomous driving. Similarly, AdapNet is designed to adaptively handle different environments, making it a versatile choice for autonomous systems operating in diverse conditions^[27].

Mask R-CNN, one of the most popular frameworks for instance segmentation, extends the Faster R-CNN object detection framework by adding a branch for predicting segmentation masks. This allows the model to not only detect objects but also generate pixel-level masks for each instance, making it highly effective in tasks where instance-level precision is required, such as autonomous driving. However, deploying segmentation models across different environments poses significant challenges, especially when the model trained in one domain is applied to another, often referred to as the domain adaptation problem. This issue is particularly important in autonomous driving, where models may need to generalize across different cities, weather conditions, or lighting variations. Guan and Yuan (2023)^[28] propose an instance segmentation method that addresses the rapid deployment problem in autonomous driving applications. Their approach evaluates how models trained in a source domain can be adapted and deployed to multiple target domains with minimal performance degradation. This is crucial for ensuring that AVs can perform reliably in diverse driving conditions without the need for extensive retraining on new data^[28].

Semantic and instance segmentation are crucial for scene understanding in autonomous driving. While semantic segmentation labels pixels for road elements, instance segmentation distinguishes between individual objects such as vehicles and pedestrians. Despite advancements in models such as Mask R-CNN, a significant research gap remains in addressing domain adaptation challenges. Current models struggle to generalize across diverse environments, such as varying cities, weather, and lighting conditions, without extensive retraining. Future research should focus on developing robust segmentation models capable of adapting to new domains with minimal performance loss, ensuring reliable and efficient autonomous driving in dynamic, real-world scenarios.

2.1.4. Sensor fusion

Sensor fusion plays a pivotal role in autonomous driving by combining data from various sensors, such as cameras, LiDAR, and radar, to provide a comprehensive understanding of the vehicle's environment. Each sensor modality captures different data types: cameras capture perspective 2D views of the surroundings, while LiDAR collects 3D spatial data. This difference in data modalities introduces significant challenges, particularly in fusing them into a unified representation for multi-task perception. A well-integrated sensor fusion system is essential for enabling AVs to accurately perceive their environment, make decisions, and navigate safely.

One of the early approaches to sensor fusion involves projecting LiDAR point clouds onto camera images, resulting in RGB-D data that 2D CNNs can process. This method leverages the successes of 2D perception, especially in tasks such as object detection and segmentation^[29]. However, this LiDAR-to-camera projection suffers from severe geometric distortions, particularly when applied to tasks that require a high degree of geometric precision, such as 3D object recognition. The distortion arises because LiDAR data inherently captures depth and spatial information that cannot be accurately represented when projected onto 2D

images. This limits the effectiveness of this approach for tasks that rely heavily on accurate 3D information.

Another method to enhance sensor fusion involves augmenting the LiDAR point clouds with additional information, such as semantic labels^[29], CNN features^[30], or virtual points derived from 2D images^[31]. This approach improves the accuracy of 3D object detection by providing additional context to the LiDAR data, enabling more accurate predictions of 3D bounding boxes. However, these methods often fall short in semantic-oriented tasks, where understanding the meaning and context of objects is crucial. The camera-to-LiDAR projection used in these methods tends to be semantically lossy, as 2D camera images are not rich in spatial context, which is necessary for tasks such as semantic segmentation and scene understanding.

To address the limitations of previous fusion techniques, Liu *et al.*^[32,33] proposed BEVFusion—a multi-task, multi-sensor fusion framework that uses Bird's Eye View (BEV) representation to unify multi-modal features. BEVFusion effectively combines the geometric structure of LiDAR data with the semantic richness of camera data, allowing it to support a wide range of 3D perception tasks. By projecting sensor data into a common BEV representation, the system overcomes the distortions and semantic losses associated with previous methods, making it more effective for both geometric-oriented tasks, such as 3D object detection, and semantic-oriented tasks, including scene segmentation. This unified representation enables AVs to perceive their environment in greater detail and with higher accuracy, enhancing both object recognition and semantic understanding.

BEVFusion represents a significant advancement in the field of sensor fusion for autonomous driving, as it resolves the challenges posed by differing sensor modalities. By aligning the data from various sensors into a common BEV framework, this approach provides a richer, more detailed understanding of the environment, which is crucial for the development of robust perception systems. The ability to handle both geometric and semantic information effectively makes BEVFusion a versatile solution for addressing the multifaceted challenges of perception in autonomous driving.

Current sensor fusion methods, such as LiDAR-to-camera projection and LiDAR augmentation, face limitations due to geometric distortions and semantic losses, which hinder their effectiveness in tasks requiring precise 3D information and semantic context. Although these approaches improve object detection, they fall short in complex tasks such as semantic segmentation and scene understanding. The introduction of BEVFusion offers a promising solution by unifying multi-modal sensor data into a common framework, overcoming previous limitations. Future research should focus on refining and expanding this approach, ensuring more accurate and reliable AV perception in dynamic, real-world environments.

2.1.5. Localization

Visual Localization or Visual Odometry (VO) plays a critical role in autonomous driving, where it is responsible for determining the position of a vehicle by analyzing sequential images captured by onboard cameras. VO typically works by identifying key point landmarks in consecutive video frames and using these points as input for a perspective-n-point (PnP) mapping algorithm. This mapping algorithm computes the pose (i.e., the orientation and position) of the vehicle relative to the previous frame. Traditional approaches to VO, while effective, can suffer from inaccuracies due to the complexity of real-world driving environments, such as changing lighting conditions, occlusions, and dynamic obstacles.

Recent advances in deep learning have significantly improved the accuracy and robustness of VO by enhancing the key point detection process. Specifically, deep learning-based methods are able to identify more precise and reliable key points, which, in turn, lead to more accurate pose estimations. This has

proven particularly useful in Simultaneous Localization and Mapping (SLAM), a field that involves building a map of the environment while simultaneously keeping track of the vehicle's location within that map. By incrementally mapping the environment and calculating the camera's pose, SLAM techniques enable AVs to navigate even in unfamiliar or dynamic settings.

Neural networks have been increasingly adopted in this domain to estimate the 3D pose of a camera in an End2End fashion, where raw image data is directly fed into the model to output the vehicle's pose without the need for manual feature extraction. For instance, PoseNet^[34] was an early neural network designed for visual localization, utilizing deep learning to estimate the 6-DoF (degrees of freedom) camera pose. Further advancements, such as VLocNet++, integrate scene semantics with pose estimation, enhancing the vehicle's ability to understand not just its position but also the surrounding environment. Similarly, Sarlin *et al.* (2018)^[35] introduced an approach that leverages deep visual descriptors for hierarchical localization, allowing for more robust and accurate pose predictions in complex scenes.

More recent work has expanded beyond traditional image-based methods to incorporate other sensor modalities, such as LiDAR. For example, Charroud *et al.*^[36] proposed an explained deep learning LiDAR-based (XDLL) model that estimates the vehicle's position using only a minimal number of LiDAR points. This innovation not only reduces the computational load but also makes localization more efficient in environments where camera data might be unreliable or unavailable, such as during adverse weather conditions or in poorly lit areas. By leveraging LiDAR data, which provides highly accurate depth information, this approach enhances the robustness and precision of localization, particularly in 3D space.

Furthermore, these deep learning-based localization methods do not only focus on computing the vehicle's pose but also integrate scene semantics-information about the surrounding objects and environment. This combination of pose estimation and semantic understanding enables AVs to make more informed decisions, as they can recognize objects, pedestrians, and road signs while simultaneously determining their own position^[16].

Although deep learning has significantly improved visual localization and simultaneous localization and mapping (SLAM) in autonomous driving, challenges remain in ensuring robustness across dynamic environments with varying lighting and weather conditions. While recent approaches, including integrating LiDAR with deep learning, have enhanced localization accuracy, further research is needed to address computational efficiency and the fusion of diverse sensor modalities in real time. In conclusion, deep learning-based methods have revolutionized localization by improving accuracy and robustness, with continued advancements necessary to optimize performance across diverse and complex environments for fully AVs.

2.1.6. Perception using occupancy grid maps

Occupancy Grid Maps (OGMs) are a fundamental aspect of autonomous driving systems, providing a grid-based representation of the environment by dividing the driving space into cells that estimate the probability of occupancy. This method is crucial for real-time decision-making, particularly when navigating through environments that contain both static and dynamic objects^[37]. OGMs support tasks such as object detection, mapping, and contextual scene understanding, which are essential in complex urban driving environments.

Deep learning has significantly advanced OGM-based perception by enhancing dynamic object detection and the probabilistic estimation of the occupancy of each grid cell. By integrating sensor data from LiDAR,

cameras, and radar, deep learning models, such as RNNs and LSTM networks, enable the system to predict and track object movements, even when occlusions or incomplete sensor data are present^[17]. These models improve the robustness and real-time capabilities of OGMs by accumulating data over time, allowing better predictions of the vehicle's surrounding environment.

In addition to object detection, deep learning models assist in classifying driving environments. By continuously accumulating data, OGMs can categorize different driving contexts, such as highways, urban environments, or parking lots, based on the system's perception^[38]. This classification allows AVs to adjust their driving strategies to fit the environment, thereby enhancing safety and decision-making.

A key advancement in OGM-based systems is OGM completion, which addresses the problem of incomplete sensor data. Traditional OGMs are limited to real-time sensor inputs, leading to gaps when objects or structures block the view. Deep learning techniques, specifically OGM completion, extrapolate beyond sensor limitations to infer potential obstacles or structures in occluded areas, creating a more comprehensive and accurate map^[38].

Sensor fusion also plays a pivotal role in improving the functionality of OGMs. By combining multi-sensor data from LiDAR, cameras, and radar, Liu *et al.*^[32] have proposed multi-task, multi-sensor fusion using BEV representations. This approach enhances both geometric structure detection and semantic density estimation, boosting overall perception performance in 3D object detection and scene understanding.

While deep learning has substantially enhanced OGMs for autonomous driving by improving dynamic object detection and real-time environmental mapping, challenges remain in handling incomplete or occluded sensor data. Developing more advanced OGM completion techniques to fill in gaps caused by obstructions and integrating multi-sensor fusion more efficiently are crucial for improving perception robustness. In conclusion, deep learning advancements in OGMs significantly improve autonomous driving systems, offering enhanced decision-making, dynamic object tracking, and environmental understanding. Continued research into sensor fusion and OGM completion will be key to further optimizing safety and scalability.

2.1.7. Deep learning for path planning and behavior arbitration

Path planning and behavior arbitration are essential components in the development of autonomous driving systems, enabling vehicles to navigate complex environments while avoiding obstacles and interacting safely with other road users. Path planning involves finding an optimal route between a starting point and a desired destination, considering the vehicle's environment and dynamic obstacles. The goal is to ensure a collision-free trajectory that adapts to both static and dynamic elements, such as other vehicles, pedestrians, and road infrastructure. Deep learning, particularly through RL models, has become a promising approach for enhancing these capabilities.

Path planning requires the AV to continuously assess the environment and adjust its trajectory accordingly. Traditional rule-based methods, which rely on pre-defined algorithms to follow a set path, struggle to account for the dynamic and often unpredictable nature of real-world driving scenarios. Deep learning-based approaches, such as those discussed by Shalev-Shwartz *et al.*^[39], address these challenges by employing multi-agent systems that allow the host vehicle to negotiate interactions with other road users. For example, tasks such as overtaking, merging, or yielding require the vehicle to predict and respond to the behaviors of others, necessitating real-time adjustments to the planned route.

In addition to decision-making on highways, AVs must also handle more unstructured environments, such as urban areas where traffic rules may be ambiguous and pedestrian behavior more unpredictable. Hu *et al.*^[18] emphasize the need for behavior arbitration models that can predict and manage the behavior of other road users in such environments. Their work highlights how deep learning enables real-time adjustments to both path planning and behavior arbitration, as the vehicle must constantly adjust its trajectory based on evolving situations, such as pedestrians crossing unexpectedly or vehicles making unanticipated maneuvers.

Deep learning-based behavioral models not only enhance path planning but also optimize the decision-making process through end-to-end learning architectures. Liao *et al.*^[40,41] developed an integrated system for AVs that combines perception, prediction, and planning into a single neural network. This end-to-end model learns to identify safe trajectories directly from sensor data, bypassing the need for separate perception and planning modules. Such integrated architectures reduce the latency in decision-making, making the vehicle's responses faster and more adaptive in real-world driving conditions.

Moreover, model-based approaches such as BEVFusion, introduced by Liu *et al.*^[32], leverage BEV representations to unify multi-modal sensor data from LiDAR, radar, and cameras. This improves the system's ability to perform path planning and behavior arbitration by providing a comprehensive understanding of both the environment and potential obstacles. By fusing these sensor inputs into a coherent spatial representation, the vehicle can make more accurate predictions about the behavior of nearby objects and plan its path accordingly.

While deep learning, particularly RL, has significantly enhanced path planning and behavior arbitration for AVs, challenges remain in improving the system's adaptability in dynamic and unstructured environments, such as urban areas. Current models may struggle with unpredictable behaviors from pedestrians or other road users. Additionally, multi-modal sensor fusion, such as BEVFusion, offers potential but requires further refinement for seamless integration. In conclusion, deep learning has greatly advanced AVs' ability to plan paths and make real-time decisions. Continued research into improving adaptability, behavior arbitration, and multi-modal sensor fusion will be vital for enhancing safety and efficiency.

2.1.8. Safety of deep learning in autonomous driving

Safety in autonomous driving, particularly when utilizing deep learning techniques, is a critical concern as it directly influences the reliability and trustworthiness of self-driving systems. Safety, in this context, refers to the absence of conditions that may lead to dangerous outcomes or accidents. Ensuring that AVs operate safely is challenging because deep learning models are often opaque, making it difficult to predict how they will behave in novel situations. Varshney^[42] emphasizes that safety can be conceptualized in terms of risk, epistemic uncertainty, and the potential harm caused by unintended consequences, such as collisions or system failures. The nature of the cost function selected during model training plays a pivotal role in minimizing these risks, and care must be taken to ensure that the model generalizes well to real-world driving scenarios beyond the data it was trained on.

One of the significant challenges in ensuring the safety of deep learning systems in autonomous driving is the occurrence of accidents caused by unexpected behaviors of artificial intelligence (AI) models. Amodei *et al.*^[43] define accidents in machine learning systems as unintended and harmful behaviors that arise due to poorly designed AI systems. In autonomous driving, these accidents can stem from various factors, including incorrect object detection, faulty decision-making in complex environments, or the system's inability to handle edge cases. These harmful behaviors often occur because deep learning models, while

highly effective in many contexts, can fail in unpredictable ways when exposed to novel or rare driving situations. The black-box nature of deep learning models makes it particularly difficult to trace the root cause of such failures, further complicating efforts to ensure the safety of AV systems.

Baheri^[44] discusses the integration of RL in autonomous driving and highlights the difficulty of balancing performance with safety in real-world applications. The analysis focuses on the concept of reward hacking, where a system optimizes for short-term goals that may conflict with the broader goal of safety. For instance, an AV might optimize for speed or efficiency in a way that compromises safety, such as running through a yellow light to avoid delays. To mitigate these risks, the design of deep learning systems in autonomous driving must incorporate explicit safety constraints, ensuring that safety is always prioritized over performance metrics such as travel time or fuel efficiency.

Shalev-Shwartz *et al.*^[39] take a broader perspective, identifying autonomous driving as a multi-agent system where the vehicle must interact with other road users. This interaction introduces additional safety challenges, as the system must not only make safe decisions for itself but also anticipate the actions of pedestrians, cyclists, and other vehicles. Deep learning models must be trained to navigate these complex social dynamics safely, which requires robust datasets that account for a wide range of driving conditions and human behaviors. However, many current datasets are limited in scope, potentially leading to models that are ill-equipped to handle unusual or unexpected scenarios.

The concept of explainability is also crucial in enhancing the safety of deep learning systems in autonomous driving. As highlighted by Charroud *et al.*^[36], explainable AI (XAI) techniques are being developed to provide greater transparency into the decision-making processes of deep learning models. By making these models more interpretable, engineers can better understand why a system behaves in a certain way and identify potential safety issues before they result in accidents. Explainability not only improves model debugging and refinement but also increases stakeholder trust in the safety of autonomous driving systems, which is essential for widespread adoption.

While deep learning techniques have shown promise in autonomous driving, ensuring safety remains a significant challenge due to the opacity of these models and their unpredictable behavior in novel situations. Current models are limited in handling edge cases, unexpected behaviors, and the complex social dynamics of road interactions. The lack of explainability in AI models further complicates identifying and addressing safety issues. Future research needs to focus on integrating explicit safety constraints, developing more robust and diverse datasets, and improving model transparency through XAI techniques. These efforts are essential to making autonomous driving systems both reliable and safe.

2.1.9. Online vectorized high-definition map construction

The scalability of autonomous driving technology is heavily reliant on the availability, accuracy, and real-time update capability of high-definition (HD) maps. These maps offer comprehensive semantic information about road topology, traffic rules, and critical infrastructure, which is essential for the precise navigation and decision-making processes of AVs. The traditional approach to HD map creation involves manual processes that are not only time-consuming but also costly, limiting scalability. However, the emergence of deep learning-based solutions has revolutionized this space, enabling the real-time generation of vectorized HD maps.

Liao *et al.*^[41] introduced a significant advancement in this area with MapTRv2, a highly efficient end-to-end method for online vectorized HD map construction. Their deep learning model processes raw sensory data

from cameras and LiDAR systems to generate real-time HD map components, including road boundaries, pedestrian crossings, and lane dividers. Unlike traditional methods, MapTRv2 leverages the onboard high-processing GPUs of AVs, allowing HD map features to be generated dynamically while the vehicle is in motion. This innovation not only improves efficiency but also addresses the need for scalability, making it possible for AVs to operate across vast and dynamically changing environments.

Complementing this approach, Luo *et al.*^[45] developed a framework that integrates standard-definition (SD) maps into the HD map prediction process. This work introduces the SD Map Encoder, a Transformer-based model that enhances lane topology prediction by incorporating prior knowledge from SD maps. The model demonstrated a substantial improvement in the accuracy of lane detection and map precision, particularly in complex urban environments where road layouts can be intricate. By merging SD map data with real-time sensor input, this method enhances the predictive capability of deep learning models, resulting in more robust map construction for AVs.

Yuan *et al.*^[46] further refined online vectorized HD map creation by focusing on improving the temporal consistency and quality of map predictions. Their model utilizes a temporal fusion module with a streaming strategy that integrates information from multiple frames. This approach ensures smoother and more accurate HD map updates, addressing one of the critical challenges in autonomous driving: the need for map data to remain consistent as the vehicle moves through different environments. Temporal consistency is particularly important in urban settings, where dynamic changes, such as moving vehicles and pedestrians, require continuous updates to the map in real time.

Other researchers have contributed to this growing body of work. For example, Liu *et al.*^[47] developed VectorMapNet, an end-to-end system for vectorized HD map learning that builds on the concept of real-time map generation. Their model integrates camera and LiDAR data into a unified BEV representation, enhancing both the geometric accuracy and the semantic richness of the generated maps. This approach supports multiple perception tasks, such as object detection and lane segmentation, further extending the capabilities of autonomous driving systems.

The need for temporal and spatial integration in map construction has also been addressed by Wang *et al.*^[48] who proposed a method that augments LiDAR point clouds with CNN features derived from 2D images. This cross-modal fusion enhances the accuracy of 3D object detection, which is critical for precise map creation. Similarly, Yin *et al.*^[31] introduced a model that uses virtual points generated from 2D images to augment LiDAR-based HD maps, improving both detection and prediction tasks.

Despite significant advancements in online vectorized HD map creation for AVs, challenges related to scalability, real-time updates, and map consistency persist. Traditional manual methods are time-consuming and costly, limiting the widespread adoption of autonomous driving. Deep learning approaches, such as MapTRv2 and SD map integration, have improved efficiency; however, issues such as temporal consistency and complex urban environments remain under-addressed. Future research should focus on enhancing the real-time integration of dynamic map updates and improving the robustness of map features, ensuring the global scalability and reliability of AVs as they navigate diverse environments.

2.1.10. End-to-end autonomous driving

Traditionally, autonomous driving systems have relied on a modular architecture that divides the driving task into separate sub-modules, such as perception, planning, and control^[17]. Each of these modules processes specific aspects of the driving environment, sending outputs from one module to the next. While

this approach has been foundational in developing autonomous driving systems, it comes with several significant limitations. One major drawback is error propagation, where mistakes made in one module can adversely affect the performance of subsequent modules. For example, a misclassification in the perception module-such as incorrectly identifying a pedestrian as an inanimate object-can lead to incorrect planning decisions and, consequently, unsafe driving behavior. Additionally, managing these interconnected modules adds substantial computational complexity, as each module requires individual processing and data handling, making the system less efficient and more difficult to optimize as a whole.

To overcome these limitations, a newer approach called End2End Autonomous Driving has gained popularity^[49]. Unlike the modular approach, End2End driving simplifies the pipeline by directly mapping sensory input-such as data from cameras, LiDAR, and radar-into control outputs, bypassing the need for intermediate sub-tasks. This method leverages deep learning to handle the full spectrum of driving tasks in a single, unified model, which significantly reduces the risk of error propagation and improves overall system robustness. As a result, End2End systems can offer more streamlined and efficient performance, especially in dynamic and complex driving environments.

One of the key advancements in End2End driving has been the development of neural network-based models that can process large volumes of sensory data and make real-time decisions. For example, Shao *et al.* (2023)^[49] introduced a deep learning framework that improves decision-making for AVs by combining multiple sensor inputs in a more integrated fashion. Their model significantly enhances the vehicle's ability to make real-time adjustments in dynamic environments, such as urban areas with heavy traffic or unpredictable pedestrian movements. Similarly, Hu *et al.*^[18] demonstrated that an End2End approach could outperform traditional modular systems in terms of both safety and computational efficiency, particularly in complex scenarios including intersections and highway merging.

A key development in this space has been NVIDIA's Hydra-MDP model, introduced by Li *et al.*^[50]. This model uses a teacher-student knowledge distillation (KD) architecture, where the student model learns from a combination of human instructors and rule-based systems. The model can simulate various trajectory options optimized for different driving tasks, making it highly versatile in real-world driving conditions. This architecture enables the model to learn more efficiently and handle a wider range of scenarios, further solidifying the benefits of the End2End approach in autonomous driving. KD helps in maintaining high-performance levels, even as the model scales up to more complex driving situations, making the system more reliable and safer over time.

Another significant advantage of the End2End approach is its ability to simplify the training process. While modular systems require separate training for each module, End2End models can be trained holistically, reducing training time and computational resources^[14]. Shao *et al.*^[49] found that their End2End model required fewer computational resources to achieve the same level of accuracy as a comparable modular system, highlighting the efficiency gains of this approach. Furthermore, NVIDIA developed an End2End model that adapts to real-time changes in the environment more effectively than modular systems, demonstrating the approach's adaptability in dynamic driving conditions.

However, despite its advantages, End2End autonomous driving is not without its challenges. One of the most significant hurdles is ensuring that these systems can generalize well across different environments and driving conditions. For example, while an End2End model may perform well in one region, it might struggle when deployed in a different geographic location with varying traffic laws, weather conditions, or road structures. Wang *et al.*^[48] identified this as a key challenge for scaling End2End models, suggesting that

further research is needed to improve the generalization capabilities of these systems.

While End2End autonomous driving has made significant strides in improving system efficiency and robustness by simplifying the architecture, challenges related to generalization across different environments and driving conditions remain. Despite reducing error propagation and computational complexity, these models may still struggle when deployed in diverse geographic locations with varying traffic laws, weather conditions, and road structures. Further research is needed to enhance the generalization capabilities of End2End models to ensure their reliable and scalable deployment across a wider range of driving scenarios. Innovations such as NVIDIA's Hydra-MDP model demonstrate the potential for End2End systems to adapt and improve in real-world conditions; however, more work is required to address the diverse complexities of global autonomous driving applications.

2.1.11. Computational hardware and deployment

Deploying deep learning algorithms on target edge devices is not a trivial task. The main limitations when it comes to vehicles are the price, performance issues and power consumption. Therefore, embedded platforms are becoming essential for integration of AI algorithms inside vehicles due to their portability, versatility, and energy efficiency. The market leader in providing hardware solutions for deploying deep learning algorithms inside autonomous cars is NVIDIA®. NVIDIA DRIVE Hyperion™ is a production-ready platform for AVs. This AV reference architecture accelerates development, testing, and validation by integrating DRIVE Orin™-based AI compute with a complete sensor suite that includes 12 exterior cameras, three interior cameras, nine radars, 12 ultrasonics, and one front-facing lidar, plus one lidar for ground truth data collection. DRIVE Hyperion features the full software stack for autonomous driving (DRIVE AV) and driver monitoring and visualization (DRIVE IX), which can be updated over the air, adding new features and capabilities throughout the life of the vehicle, and is an energy-efficient computing platform, with 254 trillion operations per second, while meeting automotive standards such as the ISO 26262 functional safety specification. The scalable DRIVE Orin product family lets developers build, scale, and leverage one development investment across an entire fleet, from Level 2+ systems all the way to Level 5 fully AVs. NVIDIA is also building The DRIVE Thor super chip that leverages the latest CPU and GPU advances to deliver an unprecedented 2,000 TFLOPS of performance, while reducing overall system cost, targeting 2025 vehicles. Renesas also provides a similar SoC, called R-Car H3^[51] which delivers improved computing capabilities and compliance with functional safety standards. Equipped with new CPU cores (Arm Cortex-A57), it can be used as an embedded platform for deploying various deep learning algorithms, compared with R-Car V3H, which is only optimized for CNNs.

2.1.12. Dataset for autonomous driving

The most relevant datasets researchers use for developing autonomous driving systems are publicly available and summarized in Table 1. These datasets cover a wide range of problem spaces, addressing key challenges in autonomous driving such as object detection, scene understanding, and path planning. Detailed information on the sensor setups, geographic locations, and traffic conditions for each dataset can be found in Table 1, offering comprehensive insights into the environments and conditions under which the data was collected. These datasets serve as critical resources for advancing the accuracy and robustness of deep learning models in real-world autonomous driving scenarios.

2.2. Adopting deep learning in autonomous driving: strategic implementations

The integration of deep learning into autonomous driving has emerged as a transformative strategy in the automotive industry, enabling vehicles to process vast amounts of data in real time to make informed decisions. Strategic implementations of deep learning in autonomous driving are not only technical but also

Table 1. Publicly available datasets for autonomous driving research

Dataset	Problem space	Sensor set up	Location	Traffic condition
NuScenes	3D object detection, tracking, online vectorized map creation	Camera, radar, lidar, GPS, IMU	Boston, Singapore	Urban
KITTI	3D object detection, tracking, SLAM	Camera, lidar, GPS, IMU	Karlsruhe, Germany	Urban, Rural
Udacity	3D object detection, tracking	Camera, lidar, GPS, IMU	Mountain View, USA	Rural, Urban
Cityscapes	Semantic segmentation	Camera, lidar, GPS, IMU	Switzerland, France	Urban
Ford	3D object detection, tracking	Camera, lidar, GPS, IMU	Michigan	Urban
Daimler pedestrian	Pedestrian detection, classification, segmentation, path prediction	Mono and stereo camera	Europe, China	Urban
BDD	2D/3D object detection, tracking, semantic segmentation	Camera	USA	Urban, Rural
Oxford	3D tracking, 3D object detection	Camera, lidar, GPS, IMU	Oxford	Urban, Highway

SLAM: Simultaneous localization and mapping; IMU: inertial measurement unit; BDD: berkeley deep drive dataset; KITTI dataset: Karlsruhe Institute of Technology (KIT) and Toyota Technological Institute (TTI).

organizational, requiring companies to adapt their business models, resources, and long-term goals to harness the full potential of this technology.

One of the critical strategic considerations is the role of deep learning in enhancing perception and decision-making. Autonomous driving systems rely on deep learning models to interpret sensory inputs from various sources, such as cameras, LiDAR, radar, and ultrasonic sensors. These models can identify and categorize objects, predict movements, and determine safe routes. Organizations must invest in building robust sensor fusion frameworks to integrate data from multiple modalities and create a cohesive understanding of the driving environment. For instance, Tesla uses a camera-based deep learning approach that allows its vehicles to detect and react to traffic conditions more effectively than traditional rule-based systems

However, the adoption of deep learning for autonomous driving also brings forth challenges that require strategic planning, particularly in areas such as data infrastructure and computational resources. Deep learning algorithms are data-hungry, requiring continuous access to high-quality, labeled datasets for training and refinement^[52]. This places significant demands on organizations to invest in large-scale data collection, storage, and processing systems. Autonomous driving companies including Waymo and NVIDIA have recognized this and have built extensive data pipelines to support the development of their deep learning models.

Additionally, companies adopting deep learning face the challenge of scalability. Traditional automotive manufacturers, such as BMW and General Motors, have had to reconfigure their production processes to accommodate the integration of AI-driven components in their vehicles. This involves a rethinking of manufacturing strategies, workforce training, and collaboration with external AI research firms^[49]. The successful implementation of deep learning technologies also requires a significant shift in organizational culture, as companies must cultivate expertise in AI and machine learning to stay competitive. Upskilling existing teams and hiring AI specialists are common strategies that automotive companies use to meet the demands^[51].

Partnerships and collaborations have also emerged as essential strategies for integrating deep learning into autonomous driving. Companies often collaborate with research institutions, technology providers, and even competitors to share resources and knowledge. For example, Ford has collaborated with Argo AI to enhance its self-driving technology, leveraging Argo's expertise in deep learning^[53]. These partnerships allow companies to overcome resource constraints and accelerate innovation by tapping into specialized knowledge and cutting-edge technologies.

Moreover, regulatory compliance plays a critical role in the strategic implementation of deep learning for autonomous driving. Governments around the world are developing regulations to ensure the safety and reliability of AI-driven vehicles, which requires companies to build deep learning models that not only meet performance standards but also adhere to safety protocols^[39]. Companies such as Waymo and Cruise have been at the forefront of working with regulatory bodies to ensure their deep learning systems comply with evolving safety standards, especially concerning object detection, collision avoidance, and ethical decision-making in edge cases^[43].

The implementation of deep learning in autonomous driving also involves strategic decision-making around data privacy and security. AVs collect vast amounts of data about their environment, much of which includes personal and sensitive information. Companies must develop policies and technologies to ensure this data is securely stored and processed, while also maintaining transparency with consumers and regulators about how data is used.

In summary, the strategic adoption of deep learning in autonomous driving is multifaceted, requiring careful planning and execution in areas such as data management, scalability, workforce development, partnerships, regulatory compliance, and data security. These strategic elements are critical to ensuring the successful deployment of deep learning technologies, which are essential for achieving the long-term vision of fully AVs.

2.3. Adopting deep learning in autonomous driving - business implications

The adoption of deep learning technologies in autonomous driving has profound business implications, influencing various aspects of the automotive industry, from operational efficiency to competitive advantage. Deep learning, a subset of AI, enables AVs to process large amounts of sensory data, improving their ability to make real-time decisions. As automotive companies race toward achieving fully autonomous driving, the strategic integration of deep learning is reshaping the industry's economic and business landscape.

One of the primary business implications of adopting deep learning is its potential to drastically improve operational efficiency. Deep learning models, particularly those used for perception tasks such as object detection, scene segmentation, and lane tracking, enable vehicles to navigate complex environments with minimal human intervention. This automation reduces the need for manual input, which in turn lowers labor costs and enhances productivity. Additionally, autonomous fleets powered by deep learning can operate around the clock, providing opportunities for cost savings in industries such as logistics and ride-hailing services^[47].

Moreover, the adoption of deep learning technologies offers a competitive advantage to companies that effectively integrate AI into their autonomous driving platforms. Firms such as Mercedes-Benz and NVIDIA have collaborated to accelerate AI innovation for self-driving technologies. Mercedes-Benz plans to introduce Level 3 autonomy, allowing drivers to relinquish full control under certain conditions, with

deep learning models playing a crucial role in ensuring real-time decision-making and safety. Such innovations enable manufacturers such as Mercedes-Benz to differentiate their offerings in the premium automotive market by combining advanced technology with luxury.

While the potential benefits are significant, adopting deep learning also presents substantial challenges from a business perspective. One of the key challenges is the high cost of implementation. Developing, testing, and deploying deep learning algorithms require substantial investment in data infrastructure, computational resources, and skilled personnel^[54]. For many traditional automotive companies, these costs pose a barrier to entry, particularly when compared to tech companies that have historically had more experience and resources in AI development. Moreover, the continuous need for data collection and model retraining increases operational expenses, which can influence profit margins.

Another business implication of adopting deep learning in autonomous driving is the shift toward partnerships and collaborations. Given the complexity of deep learning technologies, many automotive companies are forming strategic alliances with technology firms, AI startups, and research institutions. Mercedes-Benz, for instance, has partnered with Bosch and NVIDIA to develop AVs with Level 4 capabilities, enabling fully driverless cars in controlled environments (Mercedes-Benz, 2023). Such collaborations allow traditional automakers to benefit from advanced AI capabilities, accelerating the path to autonomous driving.

In addition to partnerships, regulatory compliance is a significant business consideration when adopting deep learning in autonomous driving. As governments introduce stricter regulations to ensure the safety and security of AI-driven vehicles, automotive companies must invest in compliance measures. This involves ensuring that deep learning models are robust enough to handle edge cases-uncommon but potentially dangerous driving scenarios^[39]. Companies that fail to meet regulatory standards risk delays in product deployment, legal liabilities, and reputational damage. Therefore, adhering to evolving regulations is not only a legal requirement but also a strategic imperative for business sustainability.

Finally, the adoption of deep learning in autonomous driving has implications for workforce management. As AI systems become more integrated into vehicle production and operations, there is a growing need for workers with expertise in machine learning, data science, and robotics^[51]. Automotive companies must invest in upskilling their existing workforce or hiring specialized talent to manage and maintain deep learning systems. This shift in the required skill set represents both a challenge and an opportunity for businesses. While the demand for AI talent may increase labor costs in the short term, it also offers the potential for long-term efficiency gains through automation and AI-driven decision-making.

In conclusion, the adoption of deep learning in autonomous driving has far-reaching business implications, affecting operational efficiency, cost structures, competitive dynamics, and workforce development. While companies such as Mercedes-Benz, Tesla, and Waymo that successfully implement deep learning can achieve significant strategic advantages, they must also navigate challenges related to cost, regulatory compliance, and talent acquisition. As the autonomous driving industry continues to evolve, businesses will need to balance innovation with strategic planning to fully realize the potential of deep learning technologies.

3. DISCUSSION

This survey provides a comprehensive analysis of the application of deep learning to autonomous driving systems, emphasizing its transformative potential across multiple aspects, from improving road safety to

driving innovation in the automotive industry. The key findings of this study highlight the immense impact of deep learning on enhancing the perception, decision-making, and control mechanisms of AVs. With advancements in technologies such as CNNs, RNNs, and RL, the development of autonomous systems has significantly progressed, especially in areas such as object detection, path planning, and scene understanding. This survey also outlines how deep learning plays a pivotal role in overcoming the limitations of traditional systems, particularly through End2End learning models, which enhance efficiency by simplifying the complex modular pipeline of perception, planning, and control.

The strengths of this study lie in its exploration of the full spectrum of deep learning techniques applied to various challenges in autonomous driving, such as sensor fusion, object detection, and real-time decision-making. Additionally, the study discusses the importance of deep learning in enhancing business models and strategic implementations, enabling traditional automotive manufacturers to adopt cutting-edge AI-driven solutions. The study also underscores the potential societal benefits, such as increased road safety, reduced emissions, and improved mobility for disabled individuals.

However, the study also has limitations. It primarily focuses on existing deep learning applications and may not fully capture the rapid pace of innovation and novel techniques emerging in the field. Moreover, while the study provides valuable insights into the technical and business implications of deep learning, it is limited in addressing the ethical and regulatory challenges associated with fully AVs. The complexity of ensuring safety in deep learning-based systems, particularly in rare and unpredictable driving conditions, requires more detailed exploration. Additionally, the discussion on workforce transformation and the costs associated with integrating deep learning into existing business structures is relatively underdeveloped and needs further investigation.

This study also raises certain controversies, particularly in the trade-offs between performance and safety in deep learning models for AVs. The black-box nature of many deep learning algorithms raises concerns about explainability, accountability, and trust in the technology. Additionally, the potential for algorithmic bias in decision-making processes remains an area that requires attention, especially when addressing ethical considerations around life-critical decisions made by autonomous systems.

Future research directions are necessary to address these limitations and challenges. There is a need for further exploration of XAI techniques to enhance the transparency and accountability of deep learning models in autonomous driving. Additionally, future studies should focus on integrating various sensor modalities beyond the current reliance on cameras and LiDAR, enabling more robust systems that can operate safely in diverse weather conditions and environments. The development of generalized models capable of handling unseen driving scenarios will be essential to achieving Level 5 autonomy. Research into the long-term business implications of deploying deep learning in autonomous driving, including cost analysis, regulatory frameworks, and workforce development, will also be crucial for scaling these technologies.

In conclusion, while deep learning has already demonstrated its significant potential in revolutionizing autonomous driving systems, several technical, ethical, and business-related challenges need to be addressed. Future research should focus on advancing the robustness and transparency of these systems, ensuring their safe and widespread adoption, while also considering the long-term societal and economic impacts.

4. CONCLUSIONS

The integration of deep learning into autonomous driving systems marks a paradigm shift in transportation, unlocking safer, more efficient, and intelligent mobility solutions. This research underscores the transformative impact of deep learning technologies-such as CNNs, RNNs, and RL-in enhancing the core functionalities of AVs. From object detection and path planning to sensor fusion and scene understanding, deep learning has redefined the way autonomous systems perceive, interpret, and interact with their environment. The introduction of end-to-end learning models has further streamlined traditional workflows, improving efficiency while pushing the boundaries of innovation in the automotive industry.

Beyond technical advancements, this research provides key insights into the broader implications of deep learning. It highlights the potential to revolutionize not only how vehicles operate but also how they influence society-promising improved road safety, lower emissions, and expanded mobility for underserved populations such as individuals with disabilities. Moreover, the adoption of deep learning opens new avenues for automotive businesses, enabling them to transition toward AI-driven strategies and remain competitive in an evolving industry landscape.

However, the study also sheds light on the challenges that remain, particularly in achieving full autonomy at SAE Level 5. Real-time decision-making in unpredictable scenarios, integrating diverse sensor modalities, and building systems capable of operating safely in uncharted environments represent critical hurdles. Furthermore, ethical and regulatory complexities-such as ensuring explainability, minimizing algorithmic bias, and addressing safety concerns-emphasize the need for greater transparency and accountability in deep learning applications.

The research also reveals gaps in understanding the long-term business implications, including the costs of infrastructure, workforce transformation, and navigating regulatory frameworks across different regions. These challenges highlight the necessity for collaborative efforts between researchers, policymakers, and industry stakeholders to realize the full potential of autonomous driving.

By providing a comprehensive analysis of the current capabilities and limitations of deep learning, this study offers valuable insights into its transformative role in autonomous driving. It emphasizes the need for future research to address unresolved technical, ethical, and economic challenges, ensuring that the benefits of these systems are maximized while fostering trust, scalability, and societal impact.

DECLARATIONS

Authors' contributions

Made substantial contributions to the conception and design of the study, performed the survey and analysis, and identified research Gap: Sahoo LK

Provide technical guidance and support: Varadarajan V

Availability of data and materials

Not applicable.

Financial support and sponsorship

Not applicable.

Conflicts of interest

Varadarajan V is an Editorial Board member of the journal *Complex Engineering Systems*, and he is not involved in any steps of editorial processing, notably including reviewer selection, manuscript handling, or decision-making. The other author declares that there are no conflicts of interest.

Ethical approval and consent to participate

Not applicable.

Consent for publication

Not applicable.

Copyright

© The Author(s) 2025.

REFERENCES

1. Singh S. Critical reasons for crash investigation in the national motor vehicle crash causation survey. (Traffic safety facts crash•stats. report No. DOT HS 812 506). Washington, DC: National Highway Traffic Safety Administration. 2018. Available from: <https://crashstats.nhtsa.dot.gov/Api/Public/Publication/812506>. [Last accessed on 12 Feb 2025].
2. Lana I, Del Ser J, Velez M, Vlahogianni EI. Road traffic forecasting: recent advances and new challenges. *IEEE Intell Transport Syst Mag.* 2018;10:93-109. DOI
3. Crayton TJ, Meier BM. Autonomous vehicles: developing a public health research agenda to frame the future of transportation policy. *J Transp Health.* 2017;6:245-52. DOI
4. Goldfain B, Drews P, You C, et al. AutoRally: an open platform for aggressive autonomous driving. *IEEE Control Syst.* 2019;39:26-55. DOI
5. First internationally valid system approval for conditionally automated driving. Mercedes-Benz Group. Available from: <https://group.mercedes-benz.com/innovation/product-innovation/autonomous-driving/system-approval-for-conditionally-automated-driving.html>. [Last accessed on 12 Feb 2025].
6. Khanum A, Lee C, Yang C. Involvement of deep learning for vision sensor-based autonomous driving control: a review. *IEEE Sensors J.* 2023;23:15321-41. DOI
7. Lecun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. *Proc IEEE.* ;86:2278-324. DOI
8. HUBEL DH, WIESEL TN. Shape and arrangement of columns in cat's striate cortex. *J Physiol.* 1963;165:559-68. DOI PubMed PMC
9. Song JG, Lee JW. CNN-based object detection and distance prediction for autonomous driving using stereo images. *Int J Automot Technol.* 2023;24:773-86. DOI
10. Lee D, Liu J. End-to-end deep learning of lane detection and path prediction for real-time autonomous driving. *SIViP.* 2023;17:199-205. DOI
11. Li Q, Wang Y, Wang Y, Zhao H. Hdmapnet: an online HD map construction and evaluation framework. In 2022 International Conference on Robotics and Automation (ICRA); 2022 May 23-27. Philadelphia, PA, USA. IEEE; 2022. pp. 4628-34. Available from: <https://ieeexplore.ieee.org/abstract/document/9812383>. [Last accessed on 12 Feb 2025]
12. Ghintab S, Hassan M. CNN-based visual localization for autonomous vehicles under different weather conditions. *ETJ.* 2022;41:1-12. DOI
13. Hoque S, Xu S, Maiti A, Wei Y, Arafat MY. Deep learning for 6D pose estimation of objects - a case study for autonomous driving. *Expert Syst Appl.* 2023;223:119838. DOI
14. Yang K, Tang X, Qiu S, Jin S, Wei Z, Wang H. Towards robust decision-making for autonomous driving on highway. *IEEE Trans Veh Technol.* 2023;72:11251-63. DOI
15. He K, Gkioxari G, Dollár P, Girshick R. Mask R-CNN. In: Proceedings of the IEEE international conference on computer vision; 2017. pp. 2961-9. DOI
16. Radwan N, Valada V, Burgard W. Vlocnet++: deep multitask learning for semantic visual localization and odometry. *IEEE Robot Autom Lett.* 2018;3:4407-14. DOI
17. Chen C, Seff A, Kornhauser A, Xiao J. Deepdriving: learning affordance for direct perception in autonomous driving. In Proceedings of the IEEE international conference on computer vision; 2015. pp. 2722-30. DOI
18. Hu Y, Yang J, Chen L, et al. Planning-oriented autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition; 2023. pp. 17853-62. Available from: https://openaccess.thecvf.com/content/CVPR2023/html/Hu_Planning-Oriented_Autonomous_Driving_CVPR_2023_paper.html. [Last accessed on 12 Feb 2025].
19. Redmon J, Divvala S, Girshick R, Farhadi A. You only look once: unified, real-time object detection. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR); 2016. pp. 779-88. DOI

20. Liu W, Anguelov D, Erhan D, et al. SSD: single shot MultiBox detector. In: Leibe B, Matas J, Sebe N, Welling M, editors. *Computer Vision - ECCV 2016*. Cham: Springer International Publishing; 2016. pp. 21-37. DOI
21. Law H, J. Deng J. CornerNet: detecting objects as paired keypoints. In: *Proceedings of the European conference on computer vision (ECCV)*; 2018, pp. 734-50. DOI
22. Lin G, Milan A, Shen C, Reid I. RefineNet: multi-path refinement networks for high-resolution semantic segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*; 2017. pp. 1925-34. DOI
23. Ren S, He K, Girshick R, Sun J. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE Trans Pattern Anal Mach Intell.* 2017;39:1137-49. DOI PubMed
24. Dai J, Li Y, He K, Sun J. R-FCN: object detection via region-based fully convolutional networks. *Adv Neural Inf Process Syst.* 2016;29. Available from: <https://proceedings.neurips.cc/paper/2016/hash/577ef1154f3240ad5b9b413aa7346a1e>. [Last accessed on 12 Feb 2025]
25. Wang R, Wang Z, Xu Z, et al. A real-time object detector for autonomous vehicles based on YOLOv4. *Comput Intell Neurosci.* 2021;2021:9218137. DOI
26. Badrinarayanan V, Kendall A, Cipolla R. SegNet: a deep convolutional encoder-decoder architecture for image segmentation. *IEEE Trans Pattern Anal Mach Intell.* 2017;39:2481-95. DOI PubMed
27. Valada A, Vertens J, Dhall A, Burgard W. Adapnet: adaptive semantic segmentation in adverse environmental conditions. In: *2017 IEEE International Conference on Robotics and Automation (ICRA)*; 2017 May 29-Jun 3; pp. 4644-51. DOI
28. Guan L, Yuan X. Instance segmentation model evaluation and rapid deployment for autonomous driving using domain differences. *IEEE Trans Intell Transport Syst.* 2023;24:4050-9. DOI
29. Vora S, Lang AH, Helou B, Beijbom O. PointPainting: sequential fusion for 3d object detection. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*; 2020. pp. 4604-12. DOI
30. Wang C, Ma C, Zhu M, Yang X. PointAugmenting: cross-modal augmentation for 3d object detection. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*; 2021. pp. 11794-803. Available from: https://openaccess.thecvf.com/content/CVPR2021/html/Wang_PointAugmenting_Cross-Modal_Augmentation_for_3D_Object_Detection_CVPR_2021_paper.html?utm_campaign=Akira%27s+Machine+Learning+News+++&utm_medium=email&utm_source=Revue+newsletter. [Last accessed on 12 Feb 2025].
31. Yin T, Zhou X, Krähenbühl P. Multimodal virtual point 3d detection. *Adv Neural Inf Process Syst.* 2021;34:pp.16494-507. Available from: <https://proceedings.neurips.cc/paper/2021/hash/895daa408f494ad58006c47a30f51c1f>. [Last accessed on 12 Feb 2025]
32. Liu Z, Tang H, Amini A, et al. BEVFusion: multi-task multi-sensor fusion with unified bird's-eye view representation. In: *2023 IEEE International Conference on Robotics and Automation (ICRA)*; 2023 29 May-2 Jun. London, United Kingdom. IEEE; 2023. pp. 2774-81. DOI
33. Liu Z, Lin Y, Cao Y, et al. Swin transformer: hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF international conference on computer vision*; 2021. pp. 10012-22. DOI
34. Kendall A, Grimes M, Cipolla R. PoseNet: a convolutional network for real-time 6-dof camera relocalization. In: *Proceedings of the IEEE international conference on computer vision*; 2015. pp. 2938-46. DOI
35. Sarlin PE, Debraine F, Dymczyk M, Siegwart R, Cadena C. Leveraging deep visual descriptors for hierarchical efficient localization. In: *Conference on Robot Learning*; 2018. pp. 456-65. Available from: <https://proceedings.mlr.press/v87/sarlin18a.html>. [Last accessed on 12 Feb 2025].
36. Charroud A, El Moutaouakil K, Palade V, Yahyaouy A. XDLL: explained deep learning LiDAR-based localization and mapping method for self-driving vehicles. *Electronics.* 2023;12:567. DOI
37. Thrun S. Robotic mapping: a survey. 2002. Available from: <https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=cbe046f24f31aace8b61b36e41392a93225029e0>. [Last accessed on 12 Feb 2025].
38. Stojcheski J, Nürnberg T, Ulrich M, Michalke T, Gläser C, Geiger A. Self-supervised occupancy grid map completion for automated driving. In: *2023 IEEE Intelligent Vehicles Symposium (IV)*; 2023, pp. 1-7. DOI
39. Shalev-Shwartz S, Shammah S, Shashua A. Safe, multi-agent, reinforcement learning for autonomous driving. 2016. Available from: <https://arxiv.org/abs/1610.03295>. [Last accessed on 12 Feb 2025].
40. Liao B, et al. Maptrv2: an end-to-end framework for online vectorized hd map construction. 2024. Available from: <https://arxiv.org/abs/2308.05736>. [Last accessed on 18 Feb 2025].
41. Liao B, Chen S, Wang X, et al. Maptr: structured modeling and learning for online vectorized hd map construction. 2022. Available from: <https://arxiv.org/abs/2208.14437>. [Last accessed on 12 Feb 2025].
42. Varshney KR. Engineering safety in machine learning. In: *2016 Information Theory and Applications Workshop (ITA)*; 2016 Jan 31-Feb 5. La Jolla, CA, USA. IEEE; 2016. pp. 1-5. DOI
43. Amodei D, Olah C, Steinhardt J, Christiano P, Schulman J, Mané D. Concrete problems in AI safety. 2016. Available from: <https://arxiv.org/abs/1606.06565>. [Last accessed on 12 Feb 2025].
44. Baheri A. Safe reinforcement learning with mixture density network, with application to autonomous driving. *Results Control Optim.* 2022;6:100095. DOI
45. Luo Z, Gao L, Xiang H, Li J. Road object detection for HD map: full-element survey, analysis and perspectives. *ISPRS J Photogramm Remote Sens.* 2023;197:122-44. DOI
46. Yuan T, Liu Y, Wang Y, Wang Y, Zhao H. StreamMapNet: streaming mapping network for vectorized online HD map construction.

- In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision; 2024, pp. 7356-65. Available from: https://openaccess.thecvf.com/content/WACV2024/html/Yuan_StreamMapNet_Streaming_Mapping_Network_for_Vectorized_Online_HD_Map_Construction_WACV_2024_paper.html. [Last accessed on 12 Feb 2025].
47. Liu Y, Yuan T, Wang Y, Wang Y, Zhao H. Vectormapnet: end-to-end vectorized hd map learning. In: International Conference on Machine Learning; 2023. pp. 22352-69. Available from: <https://proceedings.mlr.press/v202/liu23ax.html>. [Last accessed on 12 Feb 2025].
 48. Wang S, Yang D, Wang B, et al. UrbanPose: a new benchmark for VRU pose estimation in urban traffic scenes. In: 2021 IEEE Intelligent Vehicles Symposium (IV); 2021 July 11-17. Nagoya, Japan. IEEE 2021. pp. 1537-44. DOI
 49. Shao H, Wang L, Chen R, Li H, Liu Y. Safety-enhanced autonomous driving using interpretable sensor fusion transformer. In: Proceedings of The 6th Conference on Robot Learning; 2023. Available from: <https://proceedings.mlr.press/v205/shao23a.html>. [Last accessed on 18 Feb 2025].
 50. Li Z, Li K, Wang S, et al. Hydra-MDP: end-to-end multimodal planning with multi-target hydra-distillation. Available from: <https://arxiv.org/abs/2406.06978>. [Last accessed on 18 Feb 2025].
 51. R-Car-V3H: SoC optimized for automotive application in stereo front cameras | Renesas. 2022. Available from: <https://www.renesas.com/us/en/products/automotive-products/automotive-system-chips-socs/r-car-v3h-system-chip-soc-designed-intelligent-camera-deep-learning-capabilities>. [Last accessed on 12 Feb 2025].
 52. Caesar H, Bankiti V, Lang AH, et al. nuScenes: a multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR); 2020. pp. 11621-31. DOI
 53. Wilson B, Qi W, Agarwal T, et al. Argoverse 2: next generation datasets for self-driving perception and forecasting. 2023, Available from: <https://api.semanticscholar.org/CorpusID:244906596>. [Last accessed on 12 Feb 2025].
 54. Veiga A, Astakhova LV, Botha A, Herselman M. Defining organisational information security culture-perspectives from academia and industry. *Comput Secur.* 2020;92:101713. DOI