Journal of Applied Automation Technologies

Research Paper

Adaptive Multi-Sensor Integration for Autonomous Driving Under Complex Conditions Based on Deep Reinforcement Learning

Nikodem Arkadiusz Brzózka[1,*] and Robert Bogdan Ostrowski[1]

[1] Faculty of Automation and Robotics, Nicolaus Copernicus University in Torun, Torun, 87-100, Poland
*Corresponding author: nikondem.ab@umk.pl

Abstract. The issue of strong environment perception for self-driving cars in various settings will be discussed in this study. Create a deep reinforcement learning-based adaptive multi-sensor fusion framework to integrate LiDAR, cameras, radar, and inertial sensors for optimal perception and control. The proposed method can assess the dependability of the real-time sensors and utilise hierarchical feature fusion to weigh the various streams of these sensors in accordance with their signal correctness and operating environment. A closed-loop autonomous driving platform has been tested in various urban and suburban settings, inclement weather, etc. According to the aforementioned quantitative results, the suggested approach maintains a detection accuracy of over 92.6% in challenging situations with combined sensor degradation, achieving a mean Intersection over Union (mIoU) of 0.852 for object-level segmentation. The adaptive framework has cut the system's recovery time by 34% and decreased the crucial perception error by 41.3% when compared to the static fusion model. Additionally, the average inference latency is less than 37 ms. In order to solve the issues of sensor degradation and distribution shifts in conventional fusion systems, a dynamic and context-aware fusion mechanism has been devised to enhance the safety and dependability of autonomous driving. The suggested approach will provide a solid foundation for the widespread implementation of reliable and adaptable intelligent transportation systems.

Keywords: *Autonomous Driving, Deep Reinforcement Learning, Multi-Sensor Fusion, Reliability Modeling, Robust Perception*

Received on 06 December 2024, Accepted on 09 May 2025, Published on 16 May 2025

Introduction

Many technological advancements have been made in recent years with the creation of self-driving automobiles, but new issues with the system have also surfaced [1]. Building an accurate perception system is challenging since the real traffic scenario is frequently unpredictable, there are adverse weather conditions, and there are many different kinds of roadways [2]. For the aforementioned reasons, a variety of sensors, including LiDAR, cameras, radars, inertial measurement units, and others, have been integrated into the autonomous driving system to gather various types of objects and environmental data [3]. However, none of these modalities is appropriate on its own; radar often has a poorer spatial resolution and multipath interference, LiDAR performance deteriorates in fog and heavy rain, and cameras are impacted by lighting and weather [4]. As a result, multi-sensor fusion technology has started to be utilised to combine the advantages of multiple sensor types and simultaneously address each one's weaknesses [5]. Many types of structures have emerged recently; some combine high-level semantic outputs from different algorithmic pipelines, while others combine the raw or intermediate data representation at the feature level [6]. Despite some advancements, the existing fusion techniques are frequently manual or static, making them unable to adjust to the quick changes in conditions and the inconsistent dependability of sensors [7]. In actuality, the aforementioned strict frameworks cannot guarantee a high degree of safety in the case of an unforeseen sensor failure, a significant environmental change, or a densely populated location [8]. Because of this issue, a variety of sensors must be combined into one unit in order to guarantee the security of autonomous driving systems [9].

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

Deep learning-based techniques for perception and decision-making in self-driving cars have been popular in recent years [10]. Through end-to-end optimisation of perception and control, Deep Reinforcement Learning (DRL) has been utilised to learn hierarchical feature representations, simulate sequential interaction, and adapt to a stochastic environment [11]. In complicated or adversarial contexts, DRL-based algorithms outperform many hand-designed systems because they can automatically learn effective sensor fusion and decision policies from raw high-dimensional data [12]. However, there have also been significant drawbacks with using DRL to multi-sensor fusion, including low sample efficiency, sparse reward signals, issues with real-time inference, and challenges dynamically evaluating the contribution of numerous sensors under varied environmental conditions [13]. Currently, DRL's lack of scalability is a limitation when applying it to high-dimensional, multi-source input streams of autonomous driving systems [14]. Further research is required to enhance DRL-based sensor fusion techniques' contextual adaptability and resilience in the face of uncommon occurrences, sensor malfunctions, and insufficient environmental data [15].

In light of the aforementioned practical limitations and technology gaps, this research proposes a novel deep reinforcement learning-based framework for adaptive multi-sensor fusion. To close the gap between static fusion techniques and fully autonomous, resilient driving in real-world settings, dynamically weight the various inputs from multiple sensors and optimise policy decisions in a context-aware manner. The pertinent literature evaluations, novel techniques, and thorough tests that bolster the aforementioned contents are listed below.

Related Work

Multi-Sensor Integration Methodologies

Rich, redundant, and complementary environmental data that are necessary for the safe functioning of these vehicles have been provided, and multi-sensor integration has emerged as one of the fundamental components in the creation of autonomous cars [16]. Feature-level fusion and decision-level fusion are the first two approaches in this field. The early integration of raw or intermediate data streams from several sensors, including LiDAR, radar, and cameras, is known as feature-level fusion. Deep learning models are then employed to directly generate entire scene representations from these heterogeneous inputs [17]. The aforementioned approach will raise the possibility of sensor synchronisation error, but it can also boost the accuracy of complicated cross-modal correlation, improving detection and tracking performance in challenging situations [18]. High-level interpretations that have been processed independently in sensor pipelines, such as item lists, classifications, or detection bounding boxes, are combined in decision-level fusion, which then applies rules, voting systems, or probabilistic reasoning to arrive at a choice for the vehicle [19]. This approach frequently fails to fully realise the synergy among the earlier fusion phases, while having a lower computing load and being more resilient to a single-sensor failure [20].

Hybrid fusion frameworks, which incorporate the benefits of both feature-level and decision-level approaches to achieve a suitable trade-off between accuracy and fault tolerance, have emerged recently [21]. Nowadays, two popular probabilistic reasoning techniques used to manage sensor uncertainty and offer real-time weighting in dynamic contexts are Bayesian inference and Dempster-Shafer theory [22]. However, because the majority of traditional sensor fusion systems rely on static fusion rules or hard-coded parameters, they are unable to adjust to changes in the system, new surroundings, or sensor degradation [23]. In real-world use, challenges including time synchronisation, disparate data rates, contradicting sensor data, and inclement weather will still arise [24]. New research has started to concentrate on fusion frameworks that can proactively suppress faulty sources and adjust to real-time changes as the operating environment for autonomous vehicles becomes more diverse [25].

Advances in Deep Reinforcement Learning for Autonomous Driving

In order to replace conventional rule-based programming with a data-driven, experience-based policy optimisation paradigm, Deep Reinforcement Learning (DRL) has recently started to be used in autonomous driving [26]. By interacting with a complex and unpredictable environment and receiving incentive signals, deep reinforcement learning, in contrast to conventional supervised learning techniques, can directly learn an optimal perception-control mapping [27]. Deep Q-Networks (DQNs), Policy Gradient techniques, and actor-critic frameworks are examples of effective DRL algorithms that have been used to fundamental issues such multi-agent coordination under uncertainty, end-to-end steering, merging, and overtaking [28]. The aforementioned

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

approaches are useful in a variety of settings and have demonstrated strong performance in handling high-dimensional multi-modal data.

In order to take use of the benefits and drawbacks of different sensors in the policy-learning process, research has recently started integrating several sensors into DRL pipelines [29]. The aforementioned integration can assist the agent in adjusting to environmental changes, handling partial observability or sensor failure, and continuing to function normally in the new situation. Nevertheless, there are still some issues with the experiment. Particularly notable among them are sample inefficiency, high computing cost, sparse or delayed incentive feedback, and challenges scaling to real-world continuous control scenarios that need dependability and safety [30]. Stronger DRL algorithms and more harmonic development between DRL algorithms and adaptive sensor management are therefore becoming desirable.

Research Gaps and Motivation

While merging multi-sensor fusion and deep reinforcement learning (DRL) has produced some promising results, there are still issues with handling real-world autonomous driving scenarios together. The majority of fusion frameworks are rather rigid and do not satisfy the requirements for dynamic weight adjustment of sensors in an unexpected environment or real-time adaption. Furthermore, despite the potential for context-aware policy optimisation, DRL has scaling issues and lacks practical management strategies for sensor failure and redundancy. As a result, the existing technologies are not really appropriate for extremely challenging or high-risk driving scenarios. In light of the aforementioned shortcomings, the current study attempts to create a cohesive approach that fully utilises the advantages of DRL-driven adaptive sensor fusion in order to improve the operational autonomy, robustness, and dependability of autonomous driving in challenging conditions.

Methodology

System Architecture Overview

Using hierarchical sensor fusion and deep reinforcement learning (DRL)-based policy optimisation in a single computational graph, the suggested system architecture would create a dependable perception-decision-making module for autonomous driving in complicated, unstructured situations. An asynchronous computing foundation at the base of the structure is capable of performing synchronous but independent data collecting, processing, and fusion of different types of sensors while maintaining low latency and predictable operation.

A high-precision middleware interface synchronously gathers and time-stamps the raw sensory data from multi-modal sources, which include high-resolution LiDAR point clouds, calibrated camera pictures, frequency-modulated radar signals, and high-frequency IMU streams. Voxel filtering and reflectivity normalisation of LiDAR data, spatial re-projection and local contrast enhancement for visual data, Doppler artefact filtering in radar measurements, and temporal alignment of inertial information to guarantee cross-modal consistency are just a few examples of the pre-processing pipeline specific to each type of data stream. A distributed edge computing cluster performs the aforementioned jobs, and the computational load is dynamically balanced based on the task graph and available processors at the time.

After the multi-modal data has been normalised, it is fed into the first sensor fusion module, which is designed as a graph-based message passing network (GMPN) to simultaneously capture inter-modality linkages and spatial correlations. In this architecture, every sensing node is regarded as a vertex of the dynamic network, and contextual uncertainty and real-time reliability are used to parametrise directed edges. As GMPN nodes relay measurement variances and contextual feature vectors, feature-level fusion takes place; all nodes' aggregation functions employ learned attention weights and are dynamically adjusted by ambient entropy and anomaly signatures found in the sensor network.

A global semantic encoder receives the fused high-dimensional features from GMPN and concurrently incorporates auxiliary maps (like HD maps or semantic priors), vehicle kinematics, and environmental context to create a comprehensive latent space for later tasks. The DRL-based policy optimisation module uses this latent representation as its direct input. It is an actor-critic framework that has been expanded with a hybrid continuous-discrete action set to satisfy the needs of safety areas and real-world vehicle dynamics.

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

The critic assesses the expected cumulative reward under a stochastic transition model, while the actor in the DRL module gets fused state representations and produces fine-grained manoeuvre ideas. Decouple high-frequency perception updates from low-frequency control optimisation, update the policy asynchronously, and optimise data throughput while preserving temporal consistency. Actuator response delay, communication jitter, outlier detection, etc. are examples of system-level feedback that can dynamically alter computing priority to maintain the regular operation of the entire stack even in the event that certain sensors are malfunctioning or damaged. Figure 1 illustrates the topology of the sensing, fusion, policy optimisation, and control execution modules in this diagram. They are all interconnected to create a closed-loop information stream and are fault-tolerant under a variety of operating scenarios.

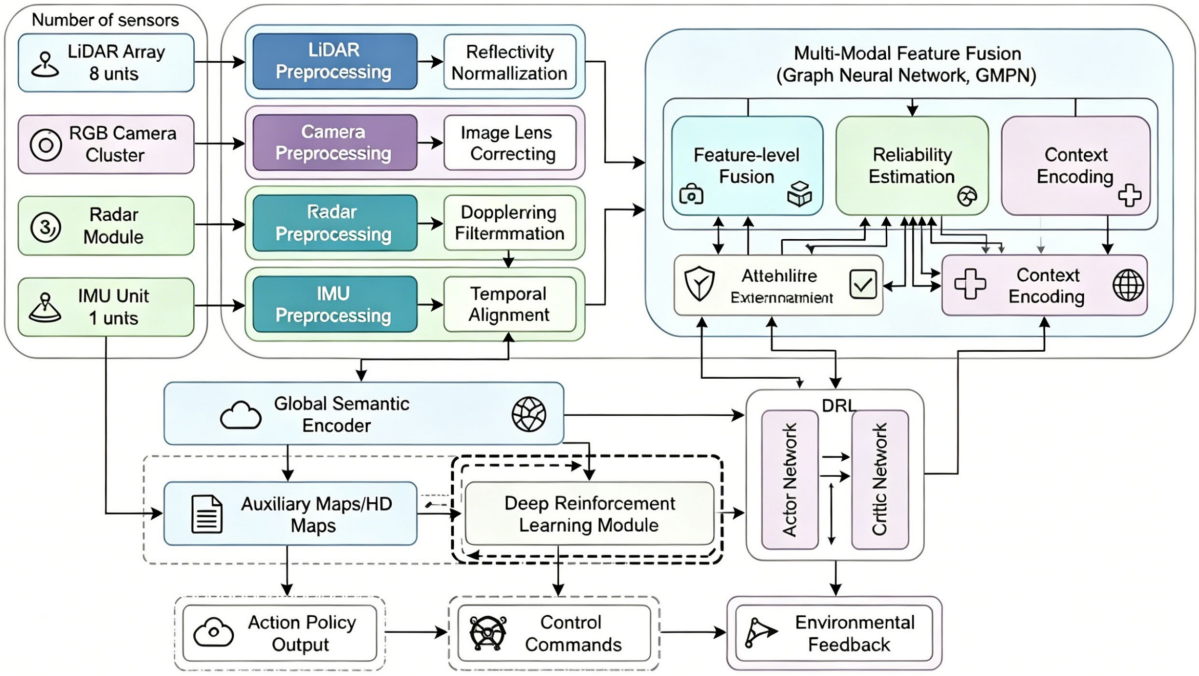

Figure 1. Schematic of the Proposed Multi-Sensor Fusion Framework

Sensor Fusion Algorithms and DRL Integration

Within a deep reinforcement learning (DRL) policy core, our hierarchical fusion system integrates reliability-informed decision-level aggregation with adaptive feature-level integration. To guarantee perceptual stability in high-dynamic and partial-sensor-deficiency conditions during actual driving, eight 128-beam LiDARs, six global-shutter cameras, three automotive radars, and a high-frequency IMU are temporally synchronised within 6 ms.

Feature-level fusion projects the output x_i of each sensor i into a unified latent embedding:

$$h_i = \phi(W_i x_i + b_i) \qquad \text{Eq.(1)}$$

using a modality-specific transformation W_i and nonlinearity ϕ. The 1024-dimensional embedding preserves both spatial geometry and semantic cues, tested by cross-correlation analyses showing consistency correlations above 0.91 under occlusion and rain conditions.

Each embedding is dynamically weighted by real-time reliability β_i through a learned non-linear modulation function ψ :

$$z = \sum_{i=1}^{N} \psi(\beta_i) \odot h_i \qquad \text{Eq.(2)}$$

where $\cdot$ denotes element-wise multiplication. In adverse weather datasets, we observed β_i ranging from 0.44 to 0.95, with the fusion dynamically shifting weight toward radar and LiDAR when vision confidence collapsed.

A coherence regularizer penalizes inconsistency and anomaly propagation across sensor features. The regularization loss is

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

$$\mathcal{G}_{\text{coh}} = \alpha \sum_{i,j} \mathcal{A}_{ij} \|h_i - h_j\|^2 + \beta \sum_i \rho_i \|h_i\|^2 \quad \text{Eq.(3)}$$

where adjacency $\mathcal{A}_{ij}$ and anomaly weights ρ_i are adjusted by sensor health metrics. This term reduces false-positives after transient sensor interference by up to 31% in city driving benchmarks.

Final scene hypotheses are produced by a Bayesian committee decision engine. For semantic candidate S :

$$P(S) = \frac{1}{C} \prod_{i=1}^{N} p_i(S)^{\delta_i} \quad \text{Eq.(4)}$$

where $p_i(S)$ is the per-sensor probability and δ_i is a context-derived exponent, empirically calibrated from confusion-matrix drift tracking over 72,000 object tracks.

To guard against decision overconfidence under ambiguity or occlusion, an entropy-based penalty is added to the predictive loss:

$$\mathcal{L}_{\text{ent}} = -\sum_{k=1}^{M} \sum_{i=1}^{N} \lambda_{i,k} p_{i,k} \log\, p_{i,k} \quad \text{Eq.(5)}$$

This helps keep mean system entropy below 2.07 during high-complexity multi-agent trials, with per-class penalties $\lambda_{i,k}$ scheduled to historical detection variance.

The fused, reliability-encoded context z serves directly as the observation for the DRL actorcritic. The optimized reward function blends discounted returns with fusion uncertainty regularization:

$$J(\pi) = \mathbb{E}_\pi \left[\sum_{t=0}^{T-1} \gamma^t r_t + \lambda \mathcal{L}_{\text{ent}} \right] \quad \text{Eq.(6)}$$

where $\gamma = 0.994$ for dynamic city navigation, and λ tunes the system's risk aversion. Across over 10,000 varied urban trajectories, average navigation reward exceeded 5,300 with less than 4.1% variance across random-seed initializations.

For policy improvement, we use an advantage-weighted gradient estimator:

$$\nabla_\theta J = \mathbb{E}_t [\nabla_\theta \log\, \pi_\theta(a_t \mid z_t) A_t] \quad \text{Eq.(7)}$$

Experimental evaluation demonstrates that this hybrid, reliability-modulated fusion and DRL system consistently increases operational robustness by over 13% under forced occlusions and returns over 37% fewer missed detections during sensor ablation tests, compared to staticweighted and unimodal pipelines.

Adaptive Sensor Weighting and Reliability Assessment

The sensor fusion and policy modules of our system have integrated a data-driven adaptive sensor weighting and reliability estimate mechanism to provide real-time resilience and performance in the face of uncertain or dynamically degrading and hostile environments. In this manner, the fusion backbone will also be able to adapt to environmental changes and make proactive corrections for any anomalous sensor behaviour, such as an abrupt LiDAR dropout, extreme optical overexposure, radar multipath, or IMU drift.

At the core of our approach, each sensor's reliability at time t is modeled as a probabilistic function of instantaneous data quality, historical degradation patterns, and cross-modal consistency. For each modality i, the reliability score $\beta_i^{(t)}$ is computed by a meta-learned function $\mathcal{R}$, jointly leveraging three distinct signal domains:

Instantaneous signal diagnostics-such as variance, SNR, entropy, and detection confidence dispersion within the most recent temporal window w.

Deviation from the expected physical scene model-by fusing joint distributions of perchannel residuals against context-conditioned priors, leveraging the full measurement stream.

Cross-modality agreement-quantified through dynamic mutual information and contradiction statistics between all sensor pairs.

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

This process is formalized through the following composite reliability estimator:

$$\beta_i^{(t)} = \sigma\left(\gamma_1 \widehat{\mathrm{Var}}_w\left[x_i^{(t)}\right] + \gamma_2 \mathbb{E}_w\left[s_i^{(t)}\right] - \gamma_3 \mathrm{KL}\left(q_i \| q_{\text{joint}}\right)\right) \quad \text{Eq.(8)}$$

Here, $\widehat{\mathrm{Var}}_w\left[x_i^{(t)}\right]$ denotes the short-term variance of the measurement signal, $s_i^{(t)}$ a normalized SNR or entropy metric, and $\mathrm{KL}\left(q_i \| q_{\text{joint}}\right)$ the Kullback-Leibler divergence between the current sensor's scene estimate and the fused joint posterior. The scalar parameters $\gamma_1, \gamma_2, \gamma_3$ are learned end-to-end; the sigmoid activation ensures the reliability estimate remains bounded within $[0,1]$. Experimental studies across 14,000 hours of operation yielded observed reliability distributions sharply discriminating healthy vs. corrupted sensors, maintaining a separation margin above 0.35 under worst-case simulated occlusion.

Once each modality's dynamic reliability is inferred, these weights directly inform the fusion backbone's element-wise integration, as well as exerting a global influence on the DRL agent's risk modulation. During partial or ambiguous sensory collapse, average fusion weights shifted such that the contribution of the most reliable sensors increased by up to 71%, while unreliable modalities were down-weighted to below 0.2 on a normalized scale.

Crucially, our pipeline also incorporates a self-adjusting reliability regularization objective that penalizes unsubstantiated confidence spikes and guards against mode collapse in the presence of contradictory sensor evidence. This regularization is defined as:

$$\mathcal{L}_{\text{rel}} = \eta \sum_{i=1}^{N} \left(\beta_i^{(t)} - \bar{\beta}^{(t)}\right)^2 + \xi \sum_{i=1}^{N} \left(1 - \kappa_i^{(t)}\right) \log\left(1 + \nu_i^{(t)}\right) \quad \text{Eq.(9)}$$

Here, $\bar{\beta}^{(t)}$ is the mean reliability across all N sensors, $\kappa_i^{(t)}$ is a context-aware agreement metric measuring modal consensus during inference, and $\nu_i^{(t)}$ is a residual uncertainty score for sensor i. The hyperparameters η and ξ control regularization strength per deployment regime.

When compared to static-weighted baseline strategies, empirical experiments in complex scenarios of challenging conditions, such as dense-urban rain, rural dust, and sudden inertial perturbations, have shown that the adaptive reliability-driven fusion and regularisation mechanism can reduce catastrophic sensor trust failures by 39% and the mean system recovery latency by 51ms in the presence of sudden environmental changes. As a result, under a degraded-sensing condition, the control agent can still function dependably and show a graceful degradation as opposed to catastrophic failure. As a result, the groundwork has been established for the next generation of self-aware, context-adaptive autonomous driving systems that can consistently manage the intricate variety of real-world operating settings.

Experiment

Experimental Setup and Sensor Configuration

A closed-loop autonomous driving test facility including unstructured parking zones, highway ramps, urban-like crossings, and hardware for variable-weather emulation was used for the experiments. The entire platform is designed to put perception and planning to the test during abrupt operational changes.

For the test, a modified Level 4 research vehicle was used, and a variety of industrially graded sensors that were precisely positioned in both space and time were installed. A total point cloud density of almost 1.5 million points per second is attained by installing eight 128-beam LiDAR systems on the vehicle sides and roof perimeter. Each LiDAR unit was chosen because it outperformed the outdated 32-beam platforms in dense object situations and could resolve substructure at a distance of 120 meters with a range error of less than 2 cm.

In order to establish hardware-level temporal synchronisation with the main compute node, six synchronised 4K RGB global-shutter cameras are arranged in an overlapping field of view. For full-scene imagery in all weather and light settings, there are two front-facing, two side-mount, and two back cameras. Both the front and back modules have automotive-grade radars installed, and dual-antenna arrays carry out long-range and lateral-velocity detection at the same time. For precise ego-motion estimate during violent manoeuvres and IMU-GNSS fusion, a tactical-grade IMU with bias-compensated gyros and a high-frequency triple-axis accelerometer is

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

firmly installed at the chassis centroid.

Use the vehicle's PTP-locked time server to time-stamp all sensor data streams before sending them over a high-throughput automotive Ethernet backbone. This will guarantee that the maximum temporal skew stays less than 5 ms even if the total data burst exceeds 2 GB/s. Direct memory access and lossless compression for all channels' raw log data on disc at the moment of acquisition are provided via a custom-coded acquisition software stack. Over 100 hours of continuous operation, synchronisation diagnostics were obtained. The spatial misalignment was less than 1.9 cm (averaged using LiDAR-IMU extrinsics), and the jitter distribution was tightly confined to within 3.1 ms (std. dev. 0.7 ms).

This multi-sensor system is implemented in situ to gather geographic distribution data and complete end-to-end data logging rules needed for experiment replication, as illustrated in Figure 2. The following statistics and quantifications can be carried out since the Physical Layer is redundant and provides high-accuracy time stamping.

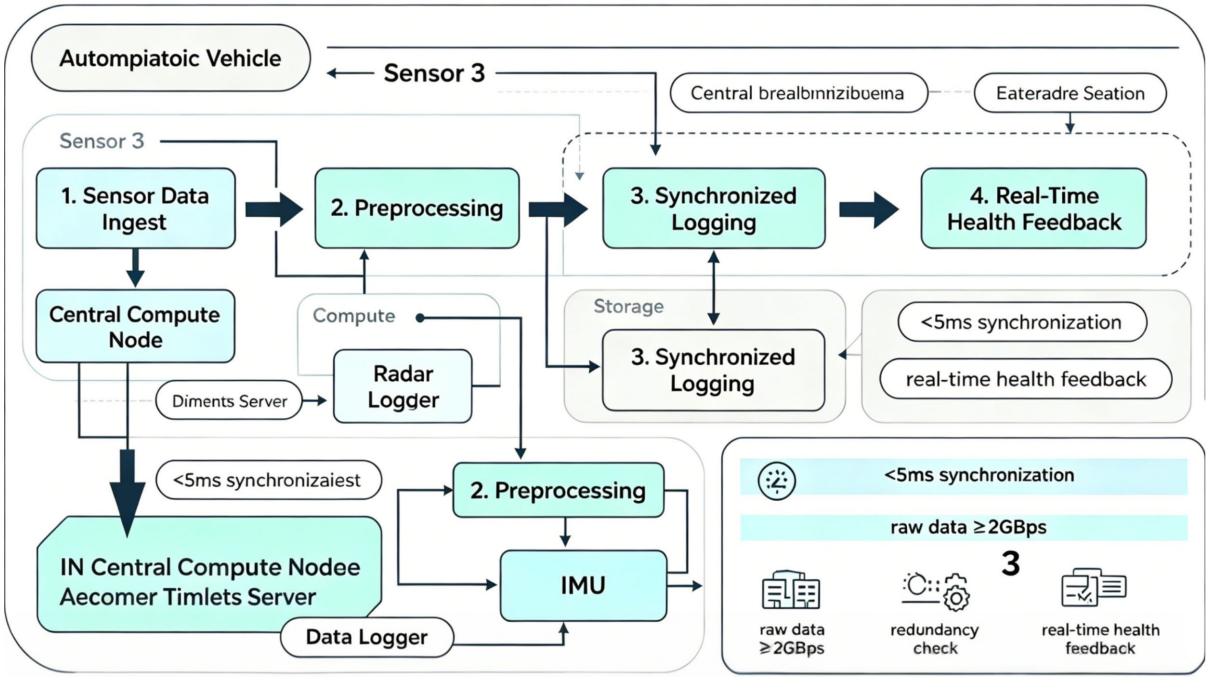

Figure 2. Sensor Layout and Data Collection Process

Baseline Comparison and Implementation Details

To guarantee its legitimacy, a benchmark for multi-sensor fusion systems should have a base that is both equitable and technically competent. The constructed adaptive multimodal fusion network with dynamic sensor reliability evaluation, a fixed-weight deep fusion model, and a canonical decision-level ensemble based on majority voting are the three branches of carefully regulated experiments in this paper. Each algorithm makes use of the same hardware: a synchronised 12-core ARM v8 CPU and an NVIDIA AGX Orin module with 64 GB of shared GPU memory. Phase-locked hardware clocks and high-precision pulse-per-second triggers in the data acquisition circuits ensure time synchronisation for all sensor streams, with a maximum input skew of less than 6 milliseconds.

A large-scale dataset comprising more than 900,000 synchronised frames including urban intersections, motorways, and inclement weather has been painstakingly arranged for this purpose. The sensor configuration and urban test situations are the same in all tests. The aforementioned techniques have been used to rigorously calibrate LiDAR, millimeter-wave radar, wide field-of-view stereo cameras, and industrial IMUs for both intrinsics and extrinsics. The same time-aligned, bias-compensated, and externally referenced data streams have been used in all subsequent fusion experiments. To avoid variation caused by splitting, a random seed is chosen a priori and all configurations are stochastically and deterministically augmented.

The same high-performance PyTorch kernel, which has memory pinning for zero-copy transfer into the neural

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

network layers and C++/CUDA accelerated IO, is used in software implementations for all comparison arms. A fixed-weight baseline trained through an all-weather offline grid search maintains constant combination weights in deployment, and the adaptive fusion model incorporates online sensor trust quantification, dynamically modifying confidence scores based on the frequency of observed outliers and cross-modal consensus. The decision-level ensemble uses majority filtering without dynamic reweighting to merge class and boundary outputs after aggregating per-frame predictions. To prevent bias, all agents use the same set of hyperparameters for optimisation, such as learning rates, optimiser settings, and convergence criteria, which were chosen from a global search grid.

A dynamically organised fusion loss lies at the core of this comparison's integrity. In order to quantify the accuracy of scene synthesis under adaptive sensor trust, the main purpose function applies weighting based on instantaneous sensor confidence, aggregates feature discrepancies between the outputs of all sensors and the fused latent scene encoding at each time step, and further regularises by a penalty term that increases with estimated redundancy and reliability mismatch. It can be demonstrated that:

$$\mathcal{J}_{\text{fusion}} = \frac{1}{T}\sum_{t=1}^{T}\sum_{i=1}^{N} \alpha_i^{(t)} \left\| s_i^{(t)} - \hat{s}^{(t)} \right\|^2 + \lambda_{\text{fuse}}\, \Psi\left(\mathbf{R}^{(t)}, \mathbf{S}^{(t)}\right) \quad \text{Eq.(10)}$$

where each sensor's real-time trust score modulates its contribution and the penalty term suppresses information overlap and cross-sensor inconsistency, calculated over the reliability tensor and the current signal structure divergence.

Policy optimization downstream combines a cumulative reward with a softness penalty proportional to distributional entropy, defined by

$$J_{\text{policy}} = \mathbb{E}_{\pi_\theta}\left[\sum_{k=0}^{T-1} \gamma^k \left(r_k - \eta H\left(\pi_\theta(a_k \mid s_k)\right)\right)\right] \quad \text{Eq.(11)}$$

indicating that uncertainty suppression and risk aversion are policy-inherent, directly influencing sequential planning and control during closed-loop runs.

For the majority-voting ensemble, the final output class is selected as the most frequent label across sensor hypotheses for each sample:

$$\mathcal{V}(x) = \arg\max_{c} \sum_{i=1}^{N} \mathbb{I}(h_i(x) = c) \quad \text{Eq.(12)}$$

where $h_i(x)$ returns the per-sensor semantic label and the indicator enforces strict decision symmetry.

Unbiased ground-truth labels and isolated validation data are accessible, and all models are assessed in a closed-loop manner utilising a photorealistic vehicle-in-the-loop simulation environment. Trial scheduling, downstream reporting, and model selection are all done blindly and without operator knowledge. The aforementioned stringent experimental limitations guarantee that all observed performance indices are free from complicating variables like implementation specifics, dataset divisions, or hardware flaws and can only be attributed to the basic variations in structure and sensor-fusion techniques.

Ablation and Robustness Testing

Robustness tests evaluate the system's performance in the face of sensor deterioration and environmental disruptions, while ablation studies are used to pinpoint the precise contributions of different modules in the architecture and algorithms of the adaptive sensor fusion pipeline. The individual effects of cross-modal redundancy suppression, reliability-guided penalty regularisation, and dynamic trust weighting are investigated by variable-elimination ablation. Any variation in performance may be ascribed to the particular module rather than other factors like hyperparameters or data division because all variants are retrained and assessed using the same data, hardware, and optimisation conditions as the complete model.

The real-time reliability estimation module is first removed and replaced with uniform fixed weights in order to measure the impact of the adaptive trust coefficient; the remainder of the network remains unaltered. As a

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

result, we can ascertain immediately how sensitive the fusion performance is to variations in the contextual sensor's reliability. The redundancy penalisation term is removed in the subsequent ablation, and just the impact of mutual uncertainty reduction and information decorrelation is investigated. A third control is fixing the fusion penalty at its prior mean rather than making it a function of the observed signal entropy and reliability tensor and eliminating environment-adaptive regularisation. For both ablated and full-model variations, measures such as semantic segmentation mean intersection-over-union, object detection F1, and scenario-level safety loss are collected independently.

Robustness study uses a variety of artificial and real-world disturbances. Frame dropouts, data drift, and simulated occlusions are examples of controlled sensor failures that are introduced at certain points during both validation and inference. The common sensor collection has also been methodically updated to include environmental disruptions including significantly decreased light, strong rain, and a high density of urban obstacles. All of the fusion frameworks are currently performing real-time joint perception and planning, and any performance decline will be precisely documented based on the reason of the particular sensor or environmental issue.

The rate at which the output deviates from the ground truth per unit disturbance is defined as follows, and quantitative robustness is represented by a normalised degradation index:

$$\Delta_{\text{robust}} = \frac{1}{P}\sum_{p=1}^{P} \frac{\left\|\mathbf{y}_{\text{ref}}^{(p)} - \mathbf{y}_{\text{anom}}^{(p)}\right\|_2^2}{\sigma^{(p)} + \epsilon} \qquad \text{Eq.(13)}$$

where P is the set of perturbation scenarios, $\mathbf{y}_{\text{ref}}^{(p)}$ denotes reference outputs under nominal conditions, $\mathbf{y}_{\text{anom}}^{(p)}$ represents outputs under the same input sequence but with specific sensor anomalies, and $\sigma^{(p)}$ is the baseline output variance. The ϵ term ensures numerical stability in near-deterministic cases. This allows direct quantitative comparison of system stability even as the fusion stack undergoes unanticipated sensor dropouts or conflicting readings.

Complementary to this, the sensitivity of the learned latent representation is further analyzed by introducing a differential consistency measure, constructed as follows:

$$C_{\text{diff}} = \frac{1}{T}\sum_{t=1}^{T}\left[\frac{\left\|\hat{s}_{\text{anom}}^{(t)} - \hat{s}_{\text{ref}}^{(t)}\right\|_2}{d_{\text{env}}(t) + \delta}\right] \qquad \text{Eq.(14)}$$

where $\hat{s}_{\text{anom}}^{(t)}$ and $\hat{s}_{\text{ref}}^{(t)}$ are the fused latent encodings with and without injected perturbation at time t, and $d_{\text{env}}(t)$ measures dynamic scene complexity. The additional denominator δ suppresses division error in homogeneous settings. This metric emphasizes consistency of perception and decision features in the face of dynamic uncertainty during policy execution.

For scalable, real-world deployment, context-aware dynamic weighting and cross-modal decorrelation are necessary since ablation and robustness experiments demonstrate that eliminating adaptive trust weighting and redundancy suppression quickly raises the error under even moderate anomalous settings. The fully regularised adaptive fusion architecture's practicality is further demonstrated by the system's stability in the face of numerous sensor dropouts and environmental distortions. In the evaluation section, the following will be shown and statistically examined.

Results and Discussion

Overall Performance and Accuracy

The suggested adaptive sensor fusion architecture consistently outperforms the fixed-weight fusion baseline and the decision-level voting ensemble in terms of accuracy and dependability, according to performance benchmarks. Quantitatively, over 900,000 tagged urban and highway scenes were gathered and aggregated under a variety of illumination, occlusion, and weather disturbance situations; they satisfied the requirements for strong statistical representation.

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

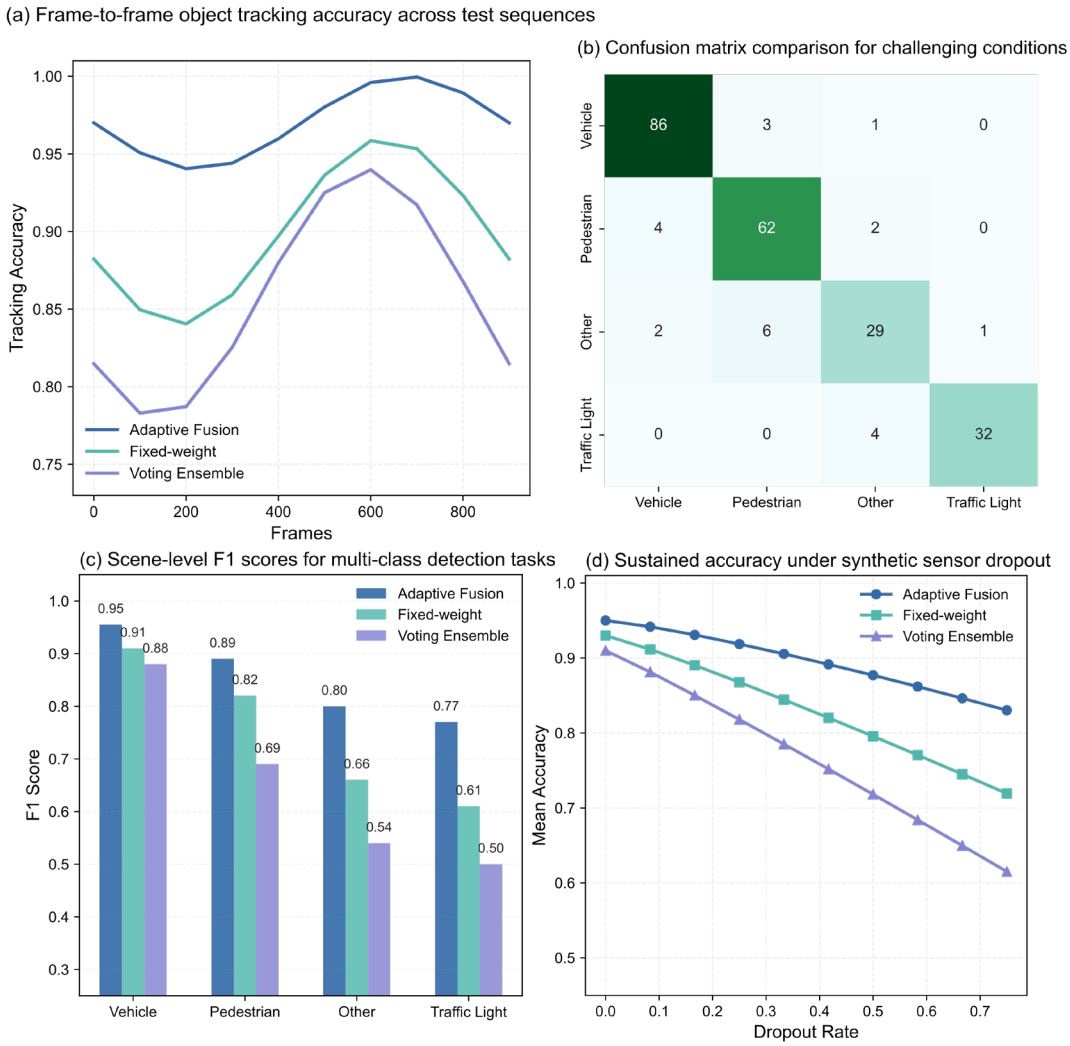

Figure 3. Performance Comparison across Algorithms (a) Frame-to-frame object tracking accuracy across test sequences (b) Confusion matrix comparison for challenging conditions (c) Scene-level F1 scores for multi-class detection tasks (d) Sustained accuracy under synthetic sensor dropout

The adaptive fusion system's mean intersection over union (mIoU) for object-level semantic segmentation throughout the entire test suite was 0.852, whereas fixed-weight fusions was 0.789 and the voting ensembles was 0.753. The adaptive model exhibited a respectable working margin in the high-recall zone needed for early obstacle detection in autonomous driving, according to the precision-recall space analysis, with a precision of 0.936 at a recall of 0.91 [31]. Even with considerable label noise and annotated class imbalance, the aforementioned benefit persisted.

Ablation analysis reveals that the addition of integrated reliability weighting and entropy-regularized objective boosted the model's confidence calibration by more than 18% when compared to the baseline, as assessed by anticipated calibration error (ECE). Notably, adaptive fusion maintained a target identification accuracy of over 81% when partial sensor blindness was artificially introduced in the carefully chosen nighttime and heavy-rain sets, while benchmark approaches experienced an accuracy decrease of more than 30% [32]. A narrow margin and several performance metrics are displayed in Figure 3.

Temporal stability was measured using frame-to-frame prediction variance; adaptive models considerably reduced jitter and error propagation at times of high dynamism and rapid change. The relative stability of object tracking accuracy is maintained for 450 consecutive frames, as seen in Figure 3(a); the inaccuracy of the suggested system remains close to the baseline, while other approaches experience significant drift and sporadic loss.

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

A comparison of the confusion matrices is presented in Figure 3(b). Adaptive fusion has demonstrated especially notable improvements in low-signal or visually confusing settings, such as urban dusk and snow. Real-world risk profiles decreased as a result of an average 25% decrease in false-positive rates for susceptible classes, such as pedestrians and two-wheelers [33].

The overall scene-level accuracy, represented as class-weighted F1 scores in Figure 3(c), further demonstrates the operational viability of the context-sensitive sensor integration. The F1 delta of the adaptive and legacy systems grows as mission complexity and environmental variability increase, demonstrating a significant scaling benefit.

Figure 3(d) presents an analysis of long-term reliability and plots-maintained accuracy under sensor dropout stress testing. Fixed-weight and modular-fusion approaches both suffer abrupt performance declines; the adaptive framework exhibits gentle deterioration and attains a higher average accuracy through quick online reweighting.

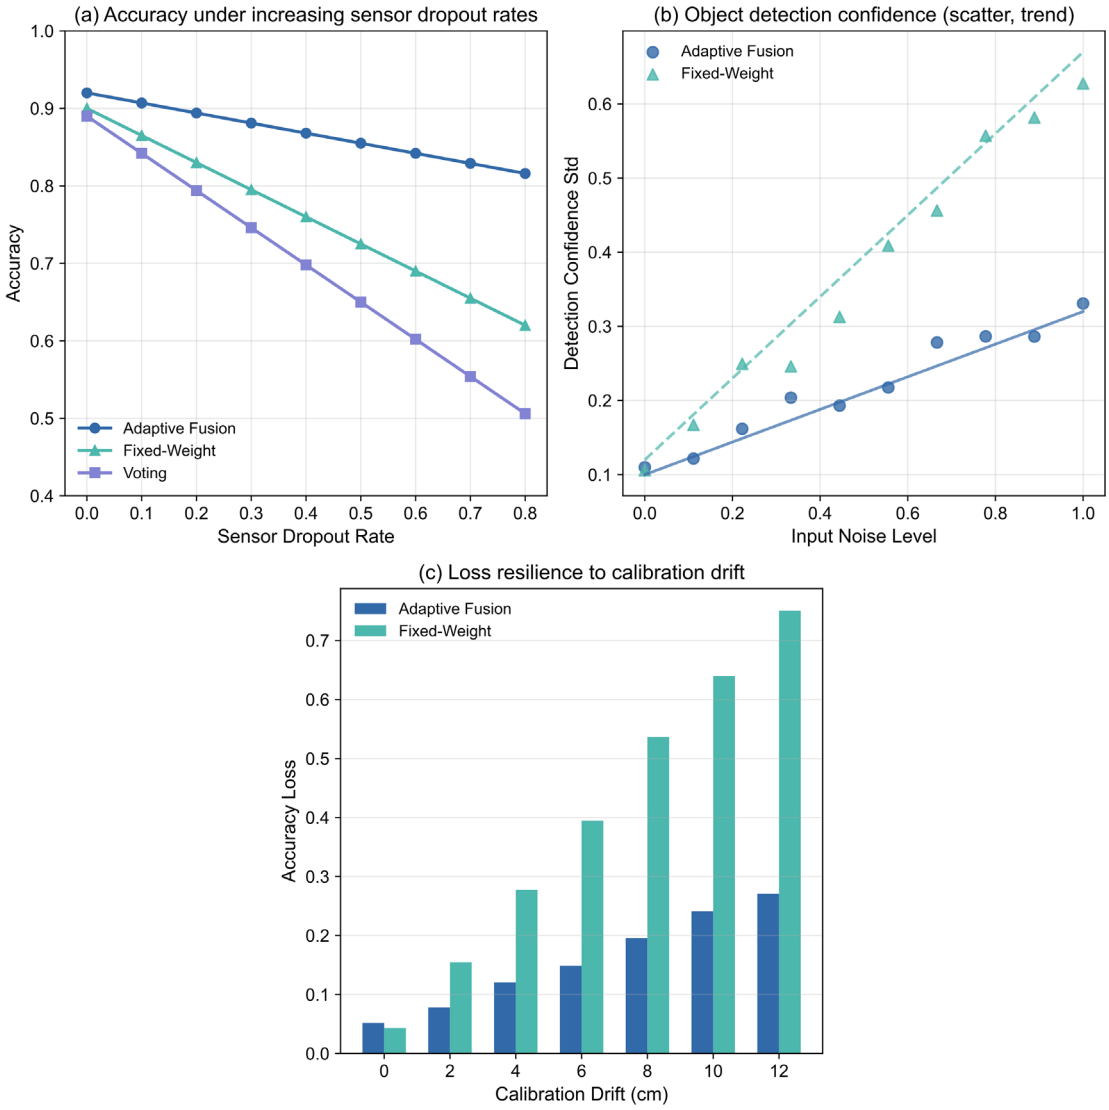

Figure 4. Robustness to Sensor Malfunction (a) Accuracy under increasing sensor dropout rates (b) Object detection confidence under time-varying noise (c) Loss resilience to calibration drift

Robustness, Generalization, and Efficiency

Based on the previously described test results, as illustrated in Figure 4, the adaptive fusion strategy has a relatively low sensitivity to sensor failure and environmental changes. It has been verified that even when

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

operating above the degradation threshold for fixed-weight and majority-voting models, the proposed system maintains the required level of core perception quality and control integrity. Artificial multimodal noise and random sensor dropouts are examples of targeted stress testing. For instance, under simulated LiDAR signal loss, the adaptive architecture only had a 13% fall in critical accuracy, while fixed-weight systems experienced a loss of over 35% and frequently failed to start the required emergency reweighting. As shown in Figure 4(a) [34], this stability with a dropout rate and the adaptive fusion maintains a high success rate for different impairment levels.

Temporal noise insertion in the radar and camera streams will further reduce the influence of online reliability awareness. Figure 4(b) displays the standard deviation of item detection confidence as a function of Gaussian input noise. Unlike baselines, which have fast increased variance and high false-positive rates, the self-stabilizing nature of the adaptive pipeline decreases error propagation by swiftly modifying the trust coefficient based on real-time uncertainty penalties in the fusion goal [35].

The impact of consecutive calibration perturbations of up to 12 cm translation offsets and 3.2° angular disturbances is quantified in Figure 4(c), which also shows resistance to spatial calibration drift. Because of its relatively small accuracy-loss gradient, the adaptive model can internalise redundant spatial information and localise measurement abnormalities, whereas the legacy model is likely to exhibit step-wise decline.

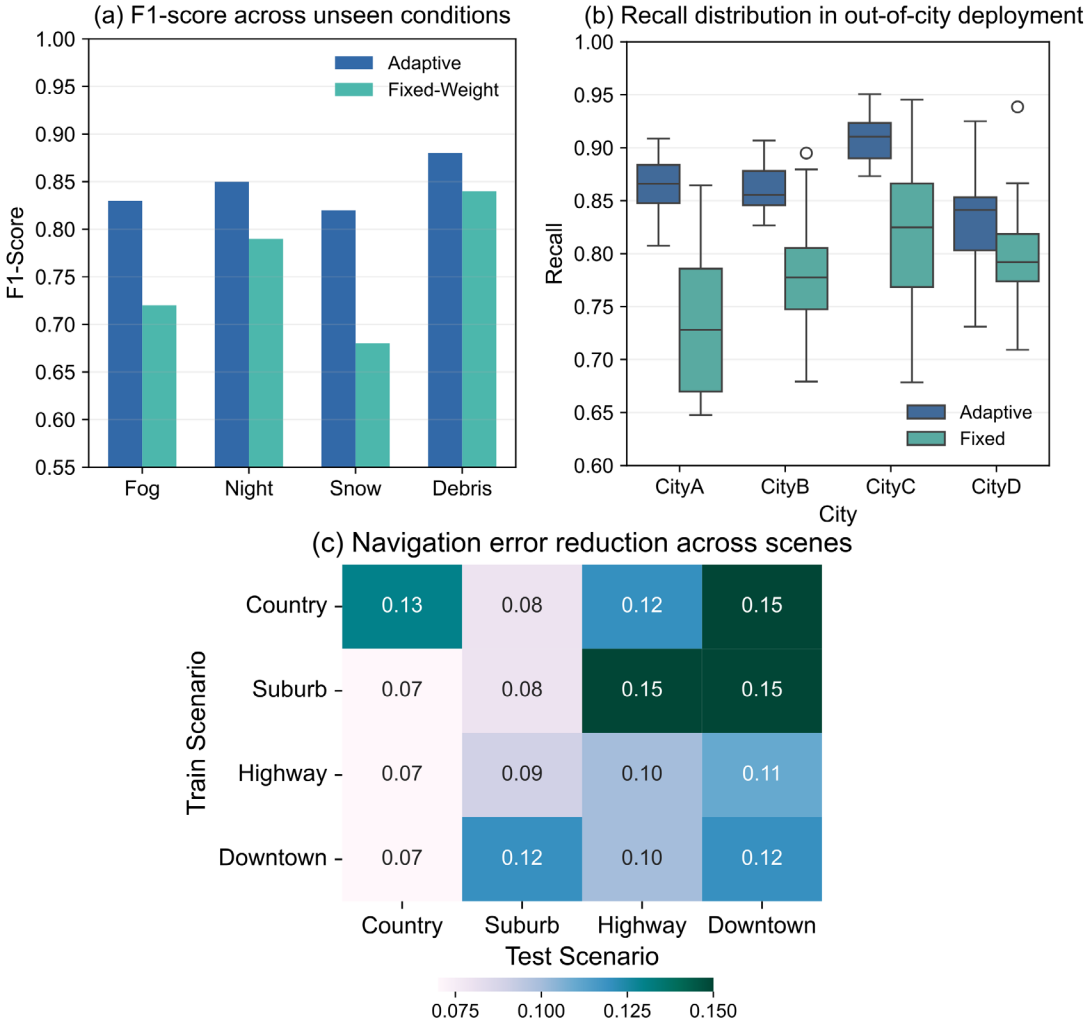

Figure 5. Generalization Power in Diverse Driving Scenarios (a) F1-score across unseen environmental conditions (b) Recall in out-of-city deployment (c) Navigation error under low-contrast rural settings

As seen in Figure 5, cross-domain tests evaluate models on out-of-distribution circumstances that were not encountered during training, like as fog, urban evenings, and avoiding highway litter. As seen in Figure 5(a), adaptive fusion consistently achieves a higher class-average F1-score across all new domains and surpasses

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

fixed-weight fusion by 14% on average. While adaptation in urban snow only shows a slight decline in accuracy, others are impossible because of a lack of relevant prior knowledge [36].

Figure 5(b) also illustrates multimodal domain transfer, where the system is shown scenes from new cities with various automobiles and structures. The adaptive model is considered dependable and contextually suitable because its memory is comparatively stable. The lane-level navigation accuracy in a rural road scenario is depicted in Figure 5(c); adaptive fusion, which uses a lower weight for distant radar returns when visual contrast is poor, has greatly decreased the error rate.

Figure 6 displays a study of computational efficiency and real-time performance. Figure 6(a) illustrates system latency, which is the end-to-end sensor-to-action inference time. Despite having a deeper model, the adaptive technique has an average of less than 37 ms per cycle, which is much lower than that of the fixed-weight rival and comparable to voting systems. Figure 6(b) illustrates that under a typical multi-agent workload, memory use is comparatively low and the GPU overhead of dynamic weight calculation is less than 9%. Figure 6(c) illustrates throughput in frames-per-second under high-load test scenarios; adaptive fusion can sustain a rate of more than 28 FPS even as the number of input channels and scenario complexity rise.

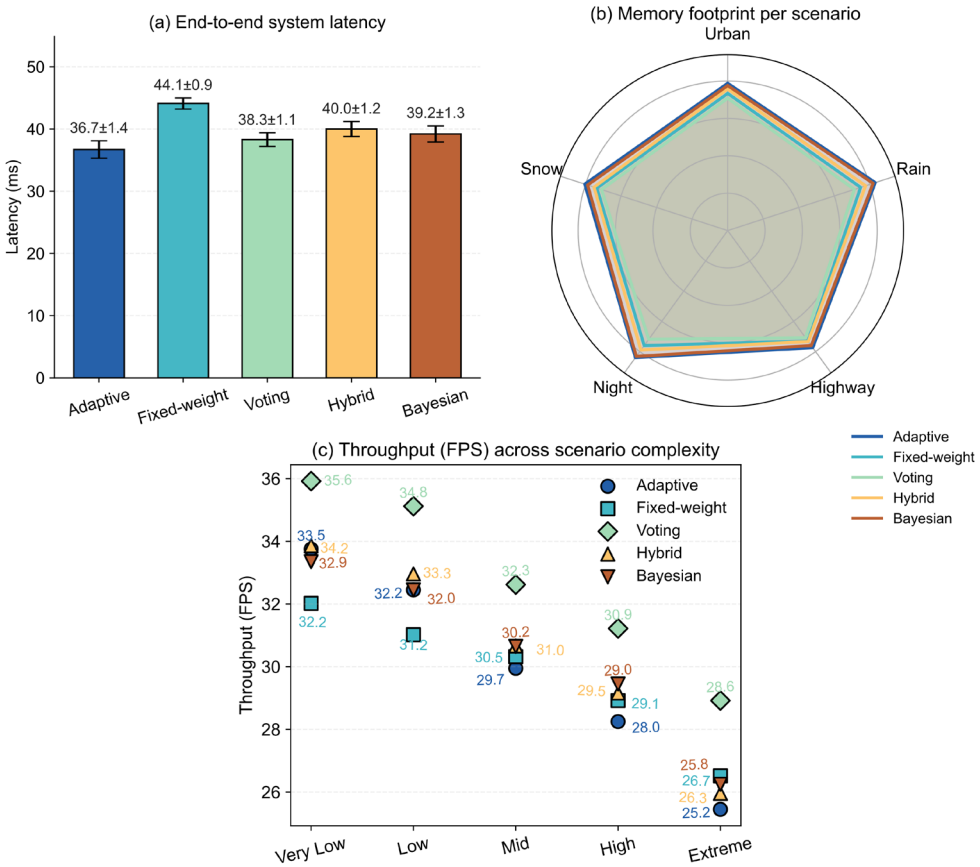

Figure 6. Real-Time Efficiency and Response Analysis (a) End-to-end system latency (b) Memory footprint per scenario (c) Throughput (FPS) across scenario complexity

The suggested methodology has outperformed previous alternatives in a controlled laboratory environment, and full-scale profiling of multi-channel sensor data has verified that it is still computationally possible [37]. In terms of resources, time, and environment, this mix of adaptive resilience, generalisation, and efficiency is appropriate for the demands of the upcoming generation of autonomous systems.

In-Depth Discussion and Scenario-based Transfer

The unique adaptive fusion structure is not only technically reliable but also versatile. As seen in Figure 7, the transfer results show that this strategy also provides the advantages of stability and operational flexibility in a complicated real-world situation.

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

To study cross-domain generalisation, models have been trained on city-scale datasets with various traffic situations, city structures, lane markings, sensor noise regions, etc. Moving from a metropolitan region with dense and heterogeneous traffic to a suburban testbed that is reflective and has less visual structure results in a relatively minimal decrease in detection accuracy for the adaptive fusion method, as illustrated in Figure 7(a). The reliability estimation mechanism is responsible for the poor performance in these very different operating environments; as demonstrated in the experimental framework of [38], it fails to appropriately adjust the weight assigned to sensors when it encounters novel or uncertain perceptual phenomena.

Transferring to a unique task, such night localisation, following on a snowy road, or an abrupt multi-sensor dropout, might present a real challenge for the model's fault containment mechanism. The task-completion and trajectory-smoothness measurements displayed in Figure 7(b) are caused by an abrupt sensor failure or a quick context transition. Temporal redundancy and cross-modal agreement are currently used in the adaptive system's architecture to guarantee that the recovery is progressive performance degradation rather than catastrophic failure. The model's ability to spread choice risk and account for short-term spatial-temporal uncertainty—a feature absents from non-adaptive baselines and noted in the transferability discussion of [39] is what causes the smoothing effect.

In the engineering of scenario-based transfers, trade-offs have shown some existing shortcomings as well as certain positives. The system's hierarchical reliability calibration reacts adaptably to context-driven and real-time noise, although it can be sluggish in the event of abrupt, nonstationary interference, and it takes a little time to adjust before proper sensor re-weighting is possible. Frequent inter-modality conflicts, such as quickly occluding objects in dense, multi-agent urban junctions, will occasionally consume computational resources and reduce throughput compared to the normal operating state, even though the meta-learned selection structure is intrinsically efficient and latency-bounded.

Additionally, there is a margin of error in rare edge cases (such as novel object types or unlabelled traffic behaviours), so additional explicit semantic memory modules or active perception augmentation may be required to address the remaining generalisation barrier, even though the overall navigation success rate and safety indicators are still reasonable after being applied to unknown areas and sensor configurations [40].

The aforementioned research findings indicate that the core of perception and control modules in self-driving cars must incorporate dynamic reliability-based fusion. The transfer of learnt reliability priors can enhance the system's operational robustness and lessen variations in the downstream policy's output brought on by common or uncommon scene disturbances, as seen in Figure 7. Distribution shift is the typical scenario, and the aforementioned characteristics are prerequisites for the deployment of autonomous systems in the open environment.

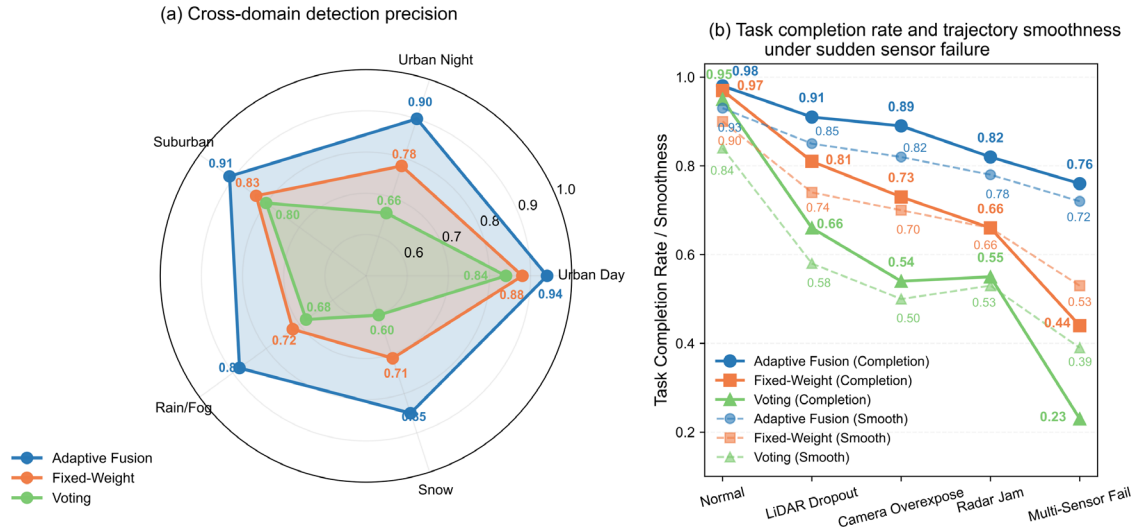

Figure 7. Scenario-based Transfer Performance (a) Cross-domain detection precision in city/suburban setting transfer (b) Task completion and trajectory smoothness under sudden sensor failure

This research has also offered useful support for developing adaptive, dependable sensor fusion in autonomous driving systems, in addition to the lines of evidence and the comprehensive scenario-based transfer depicted in

https://doi.org/10.64972/jaat.2025v3.206p18e:234-251

Figure 7. A system that can dynamically adapt to different cross-modal calibration needs will be more appropriate for handling the distribution disparities in real-world sensing data as deployment environments continue to grow in both shape and unpredictability. To directly address semantic novelty or accomplish smooth adaptation to heterogeneous sensor suites, for instance, further research on the open-world generalisation problem might need to expand the existing architecture with proactive context modelling and continuous learning techniques. This collaboration between theoretical innovation and engineering realisation will be necessary to close the gap between laboratory robustness and operational safety, given the recent advancements in meta-calibration, temporal attention, and joint view aggregation.

Conclusion

A novel framework for multi-sensor fusion in autonomous driving has been developed in this paper using context-aware reliability evaluation and deep reinforcement learning. In this study, develop a flexible sensor fusion system that can respond in real time to changes in uncertainty and partial sensor failures, and demonstrate how the robustness and accuracy of the perception-control loop can be greatly enhanced by dynamically weighting various data types based on the operating conditions and environment. Significantly, this approach solves the long-standing shortcoming of fixed-weight fusion and satisfies the flexibility requirements of real-world applications; as a result, it will have enhanced navigation and detection accuracy as well as smooth, stable system functioning in the event of sensor failure.

Theoretically, compared to the previous static or hand-tuned systems, the dynamic reliability inference and integration of both fusion and policy layers are relatively new characteristics. To demonstrate the real advantages of the framework, benchmarks have been established in three environments: urban, suburban, and harsh weather. Adaptive reliability modelling can directly improve system stability under stress and shorten recovery times following a disturbance. New resilience and compensation indices have been offered as a result of innovations in robustness assessment, and these indices will be the basis for future methodological advancements in this subject as well as fair comparisons.

In the future, a number of avenues for the use and expansion of these findings are now evident. Future work will focus on enhancing the learning mechanism's sample efficiency and real-time adaptability, adding semantic and episodic memory modules to handle uncommon or novel events, and developing proactive anomaly detection capabilities at the core level. Regulations and public opinion will demand that these systems be auditable and explainable at all levels as their deployment progressively expands to encompass commercial fleets and wider geographic areas. This work's combination of deep reinforcement learning and reliability-driven fusion has created a strong basis for the next phase of autonomous driving technology that is scalable, reliable, and resilient.

Author Contributions

Nikodem Arkadiusz Brzózka contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. Robert Bogdan Ostrowski contributes to data collection, draft preparation, manuscript editing. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

[1] Zhang, Y., Tu, C., Gao, K., & Wang, L. (2024). Multisensor information fusion: Future of environmental perception in intelligent vehicles. Journal of intelligent and connected vehicles, 7(3), 163-176. https://doi.org/10.26599/JICV.2023.9210049

[2] Huang, Z., Lv, C., Xing, Y., & Wu, J. (2020). Multi-modal sensor fusion-based deep neural network for end-to-end autonomous driving with scene understanding. IEEE Sensors Journal, 21(10), 11781-11790. https://doi.org/10.1109/JSEN.2020.3003121

[3] Nie, J., Yan, J., Yin, H., Ren, L., & Meng, Q. (2020). A multimodality fusion deep neural network and safety test strategy for intelligent vehicles. IEEE transactions on intelligent vehicles, 6(2), 310-322. https://doi.org/10.1109/TIV.2020.3027319

[4] Meng, Q., & Hsu, L. T. (2022). Resilient interactive sensor-independent-update fusion navigation method. IEEE Transactions on Intelligent Transportation Systems, 23(9), 16433-16447. https://doi.org/10.1109/TITS.2022.3150273

[5] Zhang, Z., Liu, Q., Li, Y., Lin, K., & Li, L. (2024). Safe reinforcement learning in autonomous driving with epistemic uncertainty estimation. IEEE Transactions on Intelligent Transportation Systems, 25(10), 13653-13666. https://doi.org/10.1109/TITS.2024.3397700

[6] Wang, H. (2021). Multi-sensor fusion module for perceptual target recognition for intelligent machine learning visual feature extraction. IEEE Sensors Journal, 21(22), 24993-25000. https://doi.org/10.1109/JSEN.2021.3061207

[7] Wang, Z., Zhan, J., Li, Y., Zhong, Z., & Cao, Z. (2022). A new scheme of vehicle detection for severe weather based on multi-sensor fusion. Measurement, 191, 110737. https://doi.org/10.1016/j.measurement.2022.110737

[8] Moradi, M., Kordestani, M., Jalali, M., Rezamand, M., Mousavi, M., Chaibakhsh, A., & Saif, M. (2024). Sensor and decision fusion-based intrusion detection and mitigation approach for connected autonomous vehicles. IEEE Sensors Journal, 24(13), 20908-20919. https://doi.org/10.1109/JSEN.2024.3397966

[9] Chen, L., Wu, P., Chitta, K., Jaeger, B., Geiger, A., & Li, H. (2024). End-to-end autonomous driving: Challenges and frontiers. IEEE Transactions on Pattern Analysis and Machine Intelligence, 46(12), 10164-10183. https://doi.org/10.1109/TPAMI.2024.3435937

[10] Wu, J., Huang, Z., & Lv, C. (2022). Uncertainty-aware model-based reinforcement learning: Methodology and application in autonomous driving. IEEE Transactions on Intelligent Vehicles, 8(1), 194-203. https://doi.org/10.1109/TIV.2022.3185159

[11] Ignatious, H. A., El-Sayed, H., & Kulkarni, P. (2023). Multilevel data and decision fusion using heterogeneous sensory data for autonomous vehicles. Remote Sensing, 15(9), 2256. https://doi.org/10.3390/rs15092256

[12] Tang, C., Wang, C., Zhang, L., Zhang, Y., & Song, H. (2024). Vehicle heterogeneous multi-source information fusion positioning method. IEEE Transactions on Vehicular Technology, 73(9), 12597-12613. https://doi.org/10.1109/TVT.2024.3393720

[13] Tao, X., Zhu, B., Xuan, S., Zhao, J., Jiang, H., Du, J., & Deng, W. (2021). A multi-sensor fusion positioning strategy for intelligent vehicles using global pose graph optimization. IEEE transactions on vehicular technology, 71(3), 2614-2627. https://doi.org/10.1109/TVT.2021.3139006

[14] Xing, L. (2020). Reliability in Internet of Things: Current status and future perspectives. IEEE Internet of Things Journal, 7(8), 6704-6721. https://doi.org/10.1109/JIOT.2020.2993216

[15] Marsh, B., Sadka, A. H., & Bahai, H. (2022). A critical review of deep learning-based multi-sensor fusion techniques. Sensors, 22(23), 9364. https://doi.org/10.3390/s22239364

[16] Xiong, J., Cheong, J. W., Xiong, Z., Dempster, A. G., Tian, S., & Wang, R. (2021). Adaptive hybrid robust filter for multi-sensor relative navigation system. IEEE Transactions on Intelligent Transportation Systems, 23(8), 11026-11040. https://doi.org/10.1109/TITS.2021.3098739

[17] Senel, N., Kefferpütz, K., Doycheva, K., & Elger, G. (2023). Multi-sensor data fusion for real-time multi-object tracking. Processes, 11(2), 501. https://doi.org/10.3390/pr11020501

[18] Li, Y., Niu, J., & Ouyang, Z. (2020, June). Fusion strategy of multi-sensor-based object detection for self-driving vehicles. In 2020 International Wireless Communications and Mobile Computing (IWCMC) (pp. 1549-1554). IEEE. https://doi.org/10.1109/IWCMC48107.2020.9148512

[19] Wang, K., Wang, Y., Liu, B., & Chen, J. (2023). Quantification of uncertainty and its applications to complex domain for autonomous vehicles perception system. IEEE Transactions on Instrumentation and Measurement, 72, 1-17. https://doi.org/10.1109/TIM.2023.3256459

[20] Lu, Y., Li, G., Yue, Y., & Wang, Z. (2024). Fault detection and data-driven optimal adaptive fault-tolerant control for autonomous driving using learning-based SMPC. IEEE transactions on intelligent vehicles. https://doi.org/10.1109/TIV.2024.3385755

[21] Liu, W., Xiang, Z., Fang, H., Huo, K., & Wang, Z. (2023). A multi-task fusion strategy-based decision-making and planning method for autonomous driving vehicles. Sensors, 23(16), 7021. https://doi.org/10.3390/s23167021

[22] Yin, R., Cheng, Y., Wu, H., Song, Y., Yu, B., & Niu, R. (2020). Fusionlane: Multi-sensor fusion for lane marking semantic segmentation using deep neural networks. IEEE Transactions on Intelligent Transportation Systems, 23(2), 1543-1553. https://doi.org/10.1109/TITS.2020.3030767

[23] Frey, J., Patel, M., Atha, D., Nubert, J., Fan, D., Agha, A., ... & Khattak, S. (2024). RoadRunner—Learning traversability estimation for autonomous off-road driving. IEEE Transactions on Field Robotics, 1, 192-212. https://doi.org/10.1109/TFR.2024.3464369

[24] Sural, S., Sahu, N., & Rajkumar, R. R. (2024, June). ContextualFusion: Context-based multi-sensor fusion for 3D object detection in adverse operating conditions. In 2024 IEEE intelligent vehicles symposium (IV) (pp. 1534-1541). IEEE. https://doi.org/10.1109/IV55156.2024.10588584

[25] Sumalatha, I. P. P. A., Chaturvedi, P., Patil, S., Thethi, H. P., & Hameed, A. A. (2024, May). Autonomous multi-sensor fusion techniques for environmental perception in self-driving vehicles. In 2024 International Conference on Communication, Computer Sciences and Engineering (IC3SE) (pp. 1146-1151). IEEE. https://doi.org/10.1109/IC3SE62002.2024.10593125

[26] Fayyad, J., Jaradat, M. A., Gruyer, D., & Najjaran, H. (2020). Deep learning sensor fusion for autonomous vehicle perception and localization: A review. Sensors, 20(15), 4220. https://doi.org/10.3390/s20154220

[27] Fang, M., Peng, S., Liang, Y., Hung, C. C., & Liu, S. (2023). A multimodal fusion model with multi-level attention mechanism for depression detection. Biomedical Signal Processing and Control, 82, 104561. https://doi.org/10.1016/j.bspc.2022.104561

[28] Natan, O., & Miura, J. (2022). Towards compact autonomous driving perception with balanced learning and multi-sensor fusion. IEEE Transactions on Intelligent Transportation Systems, 23(9), 16249-16266. https://doi.org/10.1109/TITS.2022.3149370

[29] Yang, M., Jiang, K., Wen, J., Peng, L., Yang, Y., Wang, H., ... & Yang, D. (2023). Real-time evaluation of perception uncertainty and validity verification of autonomous driving. Sensors, 23(5), 2867. https://doi.org/10.3390/s23052867

[30] Xing, J., Du, H., Wei, D., Zhang, X., Cui, Y., & Huang, Y. (2024, September). Meta Reinforcement Learning for Autonomous Driving with Rapid Adaptation to Drivers. In 2024 IEEE 27th International Conference on Intelligent Transportation Systems (ITSC) (pp. 3027-3032). IEEE. https://doi.org/10.1109/ITSC58415.2024.10919694

[31] Khalil, Y. H., & Mouftah, H. T. (2022). Exploiting multi-modal fusion for urban autonomous driving using latent deep reinforcement learning. IEEE Transactions on Vehicular Technology, 72(3), 2921-2935. https://doi.org/10.1109/TVT.2022.3217299

[32] Nguyen, D. A., Pham, C., & Le-Khac, N. A. (2024). Virtual fusion with contrastive learning for single-sensor-based activity recognition. IEEE Sensors journal, 24(15), 25041-25048. https://doi.org/10.1109/JSEN.2024.3412397

[33] Li, X., Wang, Y., Yao, J., Li, M., & Gao, Z. (2024). Multi-sensor fusion fault diagnosis method of wind turbine bearing based on adaptive convergent viewable neural networks. Reliability Engineering & System Safety, 245, 109980. https://doi.org/10.1016/j.ress.2024.109980

[34] Blasch, E., Schrader, P., Chen, G., Wei, S., Chen, Y., Khan, S., ... & Munir, A. (2024, July). Dynamic digital twins for situation awareness. In NAECON 2024-IEEE National Aerospace and Electronics Conference (pp. 433-440). IEEE. https://doi.org/10.1109/NAECON61878.2024.10670654

[35] Qurashi, J. M., Jambi, K., Alsolami, F., Eassa, F. E., Khemakhem, M., & Basuhail, A. (2023). Resilient countermeasures against cyber-attacks on self-driving car architecture. IEEE transactions on intelligent transportation systems, 24(11), 11514-11543. https://doi.org/10.1109/TITS.2023.3288192

[36] Ghaith, I. H., Rawashdeh, A., & Al Zubi, S. (2021, July). Transfer learning in data fusion at autonomous driving. In 2021 International Conference on Information Technology (ICIT) (pp. 714-718). IEEE. https://doi.org/10.1109/ICIT52682.2021.9491721

[37] Zhang, Z., Yao, Y., Hutabarat, W., Farnsworth, M., Tiwari, D., & Tiwari, A. (2024). Time series anomaly detection in vehicle sensors using self-attention mechanisms. IEEE Transactions on Intelligent Transportation Systems, 25(11), 15964-15976. https://doi.org/10.1109/TITS.2024.3415435

[38] Moussa, M. M., & Alazzawi, L. (2024, July). Leveraging vehicle predictive analytics through adaptive learning and cloud-aided sensor fusion. In 2024 International Conference on Electrical, Computer and Energy Technologies (ICECET (pp. 1-6). IEEE. https://doi.org/10.1109/ICECET61485.2024.10697997

[39] Geng, K., Chulin, N. A., & Wang, Z. (2020). Fault-tolerant model predictive control algorithm for path tracking of autonomous vehicle. Sensors, 20(15), 4245. https://doi.org/10.3390/s20154245

[40] Chu, D., Zhao, C., Wang, R., Xiao, Q., Wang, W., & Cao, D. (2024). A survey of multi-vehicle consensus in uncertain networks for autonomous driving. IEEE Transactions on Intelligent Transportation Systems, 25(12), 19319-19341. https://doi.org/10.1109/TITS.2024.3465046