

An IoT–AI–ML Integrated Smart Embedded System for Real-Time Safety Monitoring, Driver Behaviour Classification and Predictive Maintenance of Heavy Commercial Vehicles

Brajesh Kumar Sen¹; Dr. R. S. Verma²

¹(Research Scholar) Department of Physics. Maharaja Chhatrasal Bundelkhand University Chhatarpur Madhya Pradesh

²(Assistant Professor Physics) Govt Girls Autonomous PG College of Excellence Sagar Madhya Pradesh

Publication Date: 2026/09/26

Abstract: Heavy commercial vehicles (HCVs) are responsible for a disproportionately high number of fatal road traffic crashes worldwide, due to long driving hours, mechanical fatigue and dangerous driving behaviour. In this paper, a smart embedded system for real-time monitoring, tracking and safety management of HCVs is designed, developed and experimentally validated using IoT, AI and ML. The proposed framework is composed of an embedded acquisition layer based on GPS-GSM-sensors and a data preprocessing and analytics pipeline based on Python that can ingest real-time vehicle, driver, environmental and road-condition data. Parameters most strongly associated with hazardous driving and accident-prone conditions are identified by exploratory and statistical analyses (Pearson correlation, one-way ANOVA, chi-square tests). Six supervised machine learning algorithms (Logistic Regression, Decision Tree, Random Forest, Gradient Boosting, K-Nearest Neighbours and Support Vector Machine) are trained and comparatively evaluated for real-time safety-risk classification, with Gradient Boosting yielding the best weighted F1-score of 0.878 and ROC-AUC of 0.935. The weighted F1-score of 0.995 is achieved for overspeeding, harsh braking, rapid acceleration, and unsafe following distance detection by a parallel AI-based driver-behaviour classification framework with Random Forest. A predictive-maintenance module, which is based on a failure-classification model (ROC-AUC = 0.872) and a Random-Forest-based Remaining-Useful-Life regressor (MAE = 9.92 days, $R^2 = 0.897$), further predicts abnormal operating conditions prior to mechanical failure. All three model outputs are combined into a single real-time warning and decision-support mechanism that provides NORMAL, CAUTION, or CRITICAL ALERT notifications, which is tested on a held-out test set of 1,500 records. Fifteen figures and eighteen tables of comparative metrics are presented to validate the experimental results, which demonstrate that the proposed IoT-AI-ML integrated architecture is technically feasible and scalable for enhancing HCV driver safety, vehicle security, traffic management, and accident prevention.

Keywords: *Internet of Things; Heavy Commercial Vehicles; Machine Learning; Driver Behaviour Classification; Predictive Maintenance; Real-Time Safety Monitoring; GPS-GSM Telematics; Decision Support System.*

How to Cite: Brajesh Kumar Sen; Dr. R. S. Verma (2026) An IoT–AI–ML Integrated Smart Embedded System for Real-Time Safety Monitoring, Driver Behaviour Classification and Predictive Maintenance of Heavy Commercial Vehicles. *International Journal of Innovative Science and Research Technology*, 11(9), 1584-1601. <https://doi.org/10.38124/ijisrt/26sep770>

I. INTRODUCTION

Heavy commercial vehicles (HCVs) such as trucks, trailers, tankers and multi-axle carriers are the backbone of freight logistics in most economies, but are also disproportionately involved in severe and fatal road-traffic crashes, because of long operating hours, high mass to braking distance ratios and driver fatigue. Traditional fleet-management practices involve periodic manual inspection and post-incident reporting, and are reactive, not predictive. With the advent of low cost embedded sensors, Global Positioning System (GPS) modules, Global System for

Mobile Communications (GSM) connectivity and onboard diagnostic (CAN-bus) interfaces, continuous, real-time vehicle telematics is now technically and economically possible [1,4]. The first fleet-monitoring systems were mainly concerned with tracking the location of the vehicles [1] and more recent IoT-based architectures include GPS, sensor fusion, and cloud analytics to allow proactive safety management instead of reactive [4,5]. Meanwhile, the development of vehicular ad hoc network (VANET) and vehicle-to-infrastructure communication protocols has increased the number of contextual information (road condition, weather, inter-vehicle spacing, etc.) that can be

used in safety-critical decision making [3]. At the same time, the development of Python-based data-science and machine-learning ecosystems (pandas, NumPy, scikit-learn) has reduced the complexity of creating end-to-end analytics pipelines that can process high-frequency, multi-modal sensor streams and provide real-time predictions. Telematics data has been used to train machine learning (ML) models that have shown good performance in driving-risk classification, aggressive-behaviour detection, and predictive maintenance, which can be superior to rule-based threshold systems as they can capture non-linear interactions between speed, braking, acceleration, following distance, and environmental covariates. But most of the previous studies focus on these three issues separately: safety-risk prediction, driver-behaviour classification and predictive maintenance. They must be integrated into a single real-time decision support layer that can simultaneously notify the driver, fleet manager and maintenance scheduler to be a truly intelligent embedded safety system. Although individual IoT tracking systems, driver-behaviour monitoring applications and predictive-maintenance tools exist, heavy-commercial-vehicle fleets still lack a single, low-cost, embedded framework that.

- Collects multi-sensor real-time data,
- Statistically determines which parameters are most correlated to hazardous driving and accident-prone conditions,
- Applies and comparatively validates multiple machine-learning algorithms for safety-risk and driver-behaviour classification,
- Predicts mechanical failure before it happens and
- Integrates all of the above into a single real-time warning and decision-support mechanism.

This fragmentation makes integration more expensive, slows down the adoption by small and medium fleet operators and leaves a gap between the state of the art in individual sub-problems and a deployable end-to-end safety system that is measurable.

• *Building on the Problem Stated Above, this Study Pursues Eight Specific Objectives:*

- ✓ To design and develop an IoT based smart embedded system for real time monitoring, tracking and safety management of heavy commercial vehicles using GPS, GSM, sensors and appropriate vehicle communication technologies.
- ✓ To build a data acquisition, data preprocessing and data analytics framework using Python to process real-time data from the proposed IoT enabled system, which includes data from vehicles, drivers, environment and road conditions.
- ✓ To use Python based statistical and exploratory data analysis to identify significant patterns and parameters related to hazardous driving behaviour and accident prone conditions in heavy commercial vehicles.
- ✓ To create and test machine-learning models to predict and classify vehicle safety risks in real-time from parameters like vehicle speed, acceleration, braking behaviour,

following distance, location, and other relevant sensor data.

- ✓ To create an AI-based driver-behaviour monitoring and classification framework to identify potentially unsafe driving behaviours such as over speeding, sudden braking, rapid acceleration and unsafe vehicle following distance.
- ✓ To create a machine learning based predictive maintenance model for the detection of abnormal operating conditions and prediction of possible mechanical or safety related failures in heavy commercial vehicles.
- ✓ To integrate the AI/ML prediction models with the IoT based smart embedded system to create an intelligent real-time warning and decision support system for enhancing driver safety, vehicle security, traffic management and accident prevention.

II. LITERATURE REVIEW

➤ *IoT-Based Vehicle Tracking and Embedded Telematics Systems*

The research on fleet monitoring started with the development of GPS based location-tracking systems to give dispatchers information about the location of their vehicles and adherence to routes [1]. Later systems added GPS and GSM data links, allowing near real-time location and basic status telemetry to be sent to a central server at a reduced cost and increased capability of the embedded microcontroller platform [4]. Eren et al. proposed a real-time vehicle tracking and accident-detection framework based on the lightweight publish-subscribe protocol MQTT and cloud analytics, which showed that low-bandwidth embedded devices can stream continuous telemetry without overloading cellular data links [5]. The works together form the technical feasibility of the acquisition layer proposed in Objective 1 of this study: an embedded GPS-GSM-sensor node that can stream location, speed and status data at a fixed sampling interval.

➤ *Vehicle Communication Technologies*

In addition to point-to-point GSM telemetry, vehicular ad hoc networks (VANETs) allow direct vehicle-to-vehicle and vehicle-to-infrastructure communication, which can be used for cooperative safety applications like collision warnings and platooning. Al-Sultan et al. give a detailed overview of VANET architectures, routing protocols and security issues related to safety-critical vehicular communication [3]. The present study does not provide full VANET connectivity, but the architecture is designed to be extensible towards vehicle-to-infrastructure alerting, as is the communication-technology component named in Objective 1.

➤ *Data Acquisition, Preprocessing, and Analytics Pipelines*

For heavy commercial vehicle telematics, Zhang, Chen and Wang present a CAN-bus data-acquisition and real-time processing architecture, which requires buffering, timestamp synchronisation and noise filtering before the data from onboard sensors can be analysed reliably [6]. Gupta and Singh suggest an edge-cloud hybrid data-acquisition framework that combines on-device preprocessing with cloud-side batch analytics for fleet-scale telemetry [7] and

Kumar, Verma, and Sharma propose a scalable preprocessing pipeline implemented in Python specifically for noisy, intermittent IoT sensor streams, which includes imputation and outlier-robust scaling [8]. These studies directly inspire the preprocessing methodology adopted in Section 3.4 of this paper, namely, median imputation for sensor/GSM dropouts and percentile-based winsorising of outliers, which is similar to the noise-handling strategies validated in [6,7,8].

➤ *Exploratory and Statistical Analysis of Driving Patterns*

Paefgen et al. performed a large-scale field study of driving behaviour based on smartphone sensing, and showed that braking and acceleration events derived from the accelerometer are measurable proxies for driving risk [9]. Ma, Zhang and Liu conducted exploratory data analysis on commercial-vehicle crash-risk factors based on telematics data and found that the frequency of harsh braking and speed variance are among the top indicators of crash risk [10]. Unni and Kaimal in particular used statistical methods to analyse heavy-commercial-vehicle driving patterns to identify hazardous-condition signatures, which served as a methodological precedent for the correlation and hypothesis-testing approach used in Section 4 of the present study [11].

➤ *Machine Learning for Safety-Risk Classification*

Tochigi, Sugiura, and Nishiyama proposed a real-time driving-risk-assessment framework for heavy commercial vehicles based on machine learning, and found that ensemble tree-based classifiers are superior to linear baselines for multi-class risk labels [12]. Bozkurt, Yilmaz, and Acarman compared Random Forest and XGBoost for real-time driving-risk classification on telematics data, and found that gradient-boosted ensembles outperform the other algorithms they tested in terms of F1 and AUC scores [13]; this is directly confirmed by the results in Section 7 of this paper, which show that Gradient Boosting has the highest weighted F1-score (0.878) and ROC-AUC (0.935) among six algorithms tested for safety-risk classification. Wang, Zhang, and Zhao continued this research by applying deep-learning and ensemble models to the problem of collision-risk prediction in heavy-duty truck operations, and found that ensembling of classifiers achieved incremental improvements in accuracy compared to single classifiers [14].

➤ *AI-Based Driver-Behaviour Monitoring and Classification*

Dong, Mulgund and Ahern suggested an in-vehicle driver-behaviour-monitoring system using multi-sensor fusion, which classifies unsafe driving manoeuvres (e.g. harsh braking and swerving) from combined accelerometer, GPS and CAN-bus data [15]. Castignani et al. proposed a driver-behaviour-profiling methodology based on a smartphone and a driving-safety scoring function based on the number of events that occurred [16]. Most relevant to Objective 5 of this study, Eren and Soyly used LightGBM specifically to classify aggressive driving behaviour (overspeeding and harsh manoeuvres) of heavy commercial vehicles, achieving classification accuracy of more than 0.95 on held-out data [17] which is comparable with the 0.995

weighted F1-score achieved by the driver-behaviour classifier based on Random-Forest.

➤ *Predictive Maintenance and Remaining Useful Life Estimation*

Al-Sulaifanie and Ali used IoT and machine-learning algorithms for predictive maintenance of heavy vehicles, showing that features from the engine sensors (temperature, oil pressure, vibration) can be used to accurately predict imminent failure before it occurs as a breakdown [18]. In Section 5.3, Bousdekis, Lepenioti, and Mentzas discuss predictive-maintenance architectures in the context of Industry 4.0 and IoT telematics, and list the classification-versus-regression design choice that this paper takes, imminent-failure classification versus Remaining-Useful-Life (RUL) regression [19]. Empirical-mode-decomposition-based fault-diagnosis techniques for rotating machinery are reviewed by Lei et al., highlighting that vibration signal analysis is one of the fundamental diagnostic modalities for mechanical failure, which is why the vibration level is chosen as the leading feature of the maintenance model in this study (see Figure 12) [20]. In particular, Carvalho, Souza, and Ferreira used an XGBoost regressor to predict Remaining Useful Life for heavy-commercial-truck components, and obtained MAE values ranging from 8 to 15 days, depending on the component type, which is similar to the MAE of 9.92 days obtained by the Random Forest Regressor.

➤ *Integrated Decision Support and System Validation*

Alam et al. [22] suggested an integrated IoT-AI edge-computing framework for real-time vehicular safety alerts, as they believe that combining multiple prediction outputs into a single alert taxonomy will decrease the driver's cognitive load as opposed to multiple notification streams. Patel and Jhaveri present an intelligent decision support mechanism for fleet management, which is similar in structure to the three-tier NORMAL/CAUTION/CRITICAL ALERT scheme used in Section 6 of this paper [23] and is based on cloud-edge telemetry. Yuan, Li, and Wang [24] compare the performance-metric reporting practices of various machine-learning models applied to vehicular telemetry, and suggest that accuracy should not be the only metric reported, but rather accuracy, precision, recall, F1-score, and ROC-AUC should be reported together, which is done throughout Sections 5 and 7 of this paper. Lastly, Fawcett's basic treatment of ROC analysis gives the theoretical underpinning for the ROC-AUC metric and ROC-curve visualisations (Figures 6 and 11) that are used to assess the discrimination ability of classifiers in this study [25].

➤ *Research Gap*

The literature reviewed shows well-developed and independently validated solutions for IoT-based tracking [1,4,5], vehicular communication [3], data-acquisition pipelines [6,7,8], exploratory driving-pattern analysis [9,10,11], ML-based safety-risk classification [12,13,14] and AI-based driver-behaviour classification [15,16,17] and predictive maintenance [18,19,20,21]. Few studies, however, combine all these components: acquisition, preprocessing, statistical pattern analysis, dual ML classification (risk and behaviour), predictive maintenance and a fused real-time

decision-support layer, in a single, comparatively validated pipeline, specific to heavy commercial vehicles.

III. PROPOSED SYSTEM ARCHITECTURE AND METHODOLOGY

➤ System Overview

The proposed system is a four-layer architecture: (1) IoT embedded acquisition layer mounted on each heavy commercial vehicle, which consists of a GPS module, GSM/cellular modem, and an array of driving-dynamics, engine-health, and environmental sensors; (2) data-transmission layer that streams timestamped records to a central server at a fixed sampling interval; (3) Python-based preprocessing and analytics layer that cleans, structures, and statistically analyses the incoming stream; and (4) AI/ML decision-support layer that classifies safety risk, driver behaviour, and maintenance condition and fuses these outputs into a single real-time alert. This architecture is directly related to the Objectives 1-7, with Section 3 corresponding to the layers 1-3 (Objectives 1-2), Sections 4-5 to the analytics and modelling aspects (Objectives 3-6), and Section 6 to the fused decision-support layer (Objective 7).

➤ IoT-Based Embedded System Design

The embedded acquisition node is based on three groups of sensors. The location subsystem is based on a GPS receiver to obtain latitude/longitude fixes and a GSM module for cellular uplink, as was done in previous fleet-monitoring and IoT safety-management systems and has been proven to work well in the tracking-and-transmission paradigm [1,4,5]. The driving-dynamics subsystem includes speed (from wheel-speed and GPS-derived velocity), tri-axial acceleration and braking intensity (from an onboard accelerometer), following distance (from a forward-facing range sensor), lane deviation, and harsh-event g-force. The vehicle-health

subsystem monitors the engine and drivetrain using CAN-bus style sensors for engine RPM, engine temperature, oil pressure, vibration level, brake-pad wear, tyre pressure, battery voltage and fuel level, as described in the CAN-bus acquisition architecture in [6]. Ambient temperature, humidity, visibility and rainfall are recorded by an environmental subsystem, enabling the fusion of road-condition and weather context with driving-dynamics data. The sampling of all sensor groups is done synchronously and the data is transmitted as a single structured record every 5 minutes or on-demand when a harsh-event threshold is reached, following the MQTT-style event-driven transmission strategy of [5].

➤ Dataset Description

A structured dataset of 6,000 timestamped records was created, covering 40 heavy commercial vehicles and 60 drivers, and comprising 36 acquired and derived fields per record (Table 1) to operationalise and experimentally validate the proposed architecture. Each record is a five-minute snapshot of telemetry data, including GPS location, driving-dynamics, engine-health and environmental sensor data, as well as three ground-truth analytical targets derived from the sensor data according to the risk-scoring logic validated in [10,11,12]: a three-class safety-risk level (Low/Medium/High), a five-class driver-behaviour label, and a binary imminent-maintenance-failure flag with an associated continuous Remaining-Useful-Life (RUL) value in days. Overall, 82.33% of records were classified as Low risk, 16.12% as Medium risk and 1.55% as High risk, which is the expected class imbalance of real world driving telemetry, where hazardous events are comparatively rare. The average speed recorded in the data set was 51.54 km/h, 3.3% of the records were classified as overspeeding the road speed limit and 15.98% of the records were classified as imminent-maintenance-failure.

Table 1 Structure of the Acquired and Derived IoT Telemetry Dataset (6,000 Records, 40 Vehicles, 60 Drivers).

Attribute Group	Representative Fields	Count
Identifiers & Timestamp	record_id, timestamp, vehicle_id, driver_id	4
GPS Location	latitude, longitude	2
Driving Dynamics	speed_kmph, acceleration_ms2, braking_intensity, following_distance_m, lane_deviation_cm, harsh_event_gforce, overspeed_flag	7
Environmental Context	road_type, road_condition, weather, ambient_temp_c, humidity_pct, visibility_m, rainfall_mm	7
Vehicle Health / Engine	engine_rpm, engine_temp_c, oil_pressure_psi, vibration_level_mm_s, brake_pad_wear_pct, tyre_pressure_psi, battery_voltage_v, fuel_level_pct, vehicle_age_years, odometer_km	10
Derived Analytical Targets	safety_risk_level, driver_behaviour, maintenance_failure, maintenance_failure_risk, remaining_useful_life_days, risk_score_raw	6

➤ Data Acquisition, Preprocessing, and Analytics Framework

In line with the noise-handling strategies described in [6,7,8] the raw telemetry stream was processed using a three-stage Python pipeline using pandas and NumPy. First, missing-value handling was performed to correct for intermittent dropouts in the sensor and GSM-link, with the latitude, tyre-pressure and humidity fields having about 1% missing values, which were imputed using column-wise

median substitution, which does not introduce any directional bias but preserves the underlying distribution. Second, four dynamic sensors with high variance – speed, engine temperature, vibration level and following distance – were winsorised (clipped at the 1st and 99th percentiles) using outlier control, which capped the extreme sensor-glitch values while preserving genuine high-risk observations. The results of outlier-capping are summarised in Table 2. Third, one-hot encoded the categorical road-type, road-condition,

and weather fields and z-score standardised (scikit-learn StandardScaler) all continuous features before training the models, thus preventing gradient-based and distance-based

classifiers from being biased towards large-magnitude features like odometer reading.

Table 2 Outliers Capped (1st-99th Percentile Winsorising) by Sensor Parameter.

Sensor Parameter	Records Capped
speed_kmph	120
engine_temp_c	120
vibration_level_mm_s	119
following_distance_m	60

The cleaned data was then split into training and held-out test sets (75%/25% stratified split) for each of the three modelling tasks described in Section 5, maintaining the class-imbalance in the safety-risk and driver-behaviour labels across both sets.

IV. EXPLORATORY AND STATISTICAL DATA ANALYSIS

This section deals with Objective 3 by using descriptive statistics, correlation analysis and formal hypothesis testing to determine which sensor parameters and contextual

conditions are most strongly related to hazardous driving and accident prone operation, following the exploratory approach used in [9,10,11].

➤ Descriptive Statistics

Table 3 presents the central tendency, dispersion and shape (skewness, kurtosis) of the eleven key dynamic and derived parameters following the preprocessing. Braking intensity and harsh-event g-force are highly positively skewed, as would be expected if harsh braking and sudden manoeuvres are relatively uncommon tail events in normal driving.

Table 3 Descriptive Statistics of Key Dynamic and Derived Sensor Parameters (Post-Preprocessing).

Parameter	Mean	Std	Min	25%	50%	75%	Max	Skew
speed_kmph	51.54	20.39	5.92	36.64	51.06	66.42	97.80	0.04
acceleration_ms2	0.31	0.91	-3.16	-0.30	0.31	0.92	3.37	-0.02
braking_intensity	0.25	0.23	0.00	0.07	0.17	0.35	1.00	1.38
following_distance_m	34.98	17.38	2.00	22.80	34.90	46.70	76.30	0.10
lane_deviation_cm	7.97	8.10	0.00	2.30	5.40	11.00	60.00	1.97
harsh_event_gforce	0.32	0.24	0.00	0.13	0.27	0.46	1.55	0.97
engine_temp_c	99.95	8.24	82.30	94.10	99.40	105.40	120.00	0.21
vibration_level_mm_s	8.94	2.87	3.68	6.74	8.57	10.84	15.95	0.39
oil_pressure_psi	33.13	6.67	10.40	28.70	33.40	37.90	53.00	-0.20
brake_pad_wear_pct	65.46	26.55	0.00	44.30	66.20	91.10	100.00	-0.23
risk_score_raw	1.35	1.81	-3.66	0.06	1.11	2.41	11.86	0.70

➤ Correlation Analysis

The Pearson correlation matrix of the principal sensor and derived-risk parameters is shown in figure 1. Two clusters emerge. Moderate correlation is found between the driving-dynamics variables (braking intensity, following distance, harsh-event g-force) and the raw risk score, confirming the relevance of these variables as safety-risk

predictors. Second, the engine-health variables (engine temperature, vibration level, oil pressure, brake-pad wear) are highly correlated ($|r|$ in the range 0.64-0.88) and all of them are highly correlated with the maintenance-failure-risk target ($r = 0.78-0.86$), which confirms that the wear-driven sensor drift is the dominant signal for the predictive-maintenance task.




Fig 1 Pearson Correlation Matrix of Sensor Parameters with Safety-Risk Score and Maintenance-Failure Risk.

Table 4 shows the driving-dynamics parameters sorted by absolute correlation with the raw risk score, where braking intensity, following distance, and harsh-event g-force are the three most significant driving-dynamics parameters that

indicate crash risk, consistent with the crash-risk-factor results of [10] and the risk-assessment feature importance reported in [12,13].

Table 4 Parameters Ranked by Absolute Pearson Correlation with the Computed Driving-Risk Score.

Parameter	Correlation with Risk Score
braking_intensity	0.360
following_distance_m	-0.242
harsh_event_gforce	0.202
lane_deviation_cm	0.125
acceleration_ms2	0.058
speed_kmph	0.035
oil_pressure_psi	0.016
engine_temp_c	-0.012
vibration_level_mm_s	-0.011
maintenance_failure_risk	-0.010
brake_pad_wear_pct	-0.005
battery_voltage_v	0.002

➤ Statistical Significance Testing

Two hypothesis tests were performed to ensure that the observed associations are not likely to be due to chance. A one-way analysis of variance (ANOVA) was used to determine if there is a significant difference between the mean speed for the three safety-risk levels (Low/Medium/High). The F value is 8.235 and the p value is 0.0002683 (Table 5) which is below the $\alpha = 0.05$ significance level, indicating that there is a significant difference between

the speed of the vehicles in the different risk categories. A chi-square test of independence between road condition and safety-risk level was statistically significant ($\chi^2 = 135.038$, $df = 8$, $p = 2.55e-25$, Table 6), suggesting that there is a relationship between road surface condition and the risk of a hazardous driving event. These results provide empirical evidence for the use of both speed-derived and road-condition features in the machine-learning models.

Table 5 One-Way ANOVA — Vehicle Speed Across Safety-Risk Levels.

Test	F-statistic	p-value	Significant ($\alpha=0.05$)
One-way ANOVA: speed_kmph across safety_risk_level (Low/Medium/High)	8.235	0.0002683	Yes

Table 6 Chi-Square Test of Independence — Road Condition vs. Safety-Risk Level.

Test	Chi2-statistic	d.f.	p-value	Significant ($\alpha=0.05$)
Chi-square: road_condition vs safety_risk_level	135.038	8	2.55e-25	Yes

➤ Hazardous-Driving and Accident-Prone Pattern Identification

Box plots are used to show the distribution of speed for the three risk levels in Figure 2. The median speed is monotonically increasing from Low to High risk, and the interquartile range is larger for the High-risk category, suggesting that the most hazardous records have a larger

range of speeds. The frequency distribution of the five driver-behaviour categories is shown in Figure 3, where the large majority of records are in the Normal/Safe driving category, Unsafe Following Distance and Harsh Braking are the next most frequent hazardous categories, and the behaviour-frequency patterns are broadly similar to those reported by [16,17].




Fig 2 Distribution of Vehicle Speed Across Low, Medium, and High Safety-Risk Levels.




Fig 3 Frequency Distribution of AI-Classified Driver-Behaviour Categories.

Figure 4 breaks down the safety-risk composition by road condition. The proportion of Medium- and High-risk records is significantly higher for the Pothole and Under Construction road segments than for the dry road surfaces, and the proportion of records for the Foggy and Wet road

segments is intermediate — consistent with the statistically significant chi-square association reported in Table 6 and with the road-hazard association findings of [11]. The percentage breakdowns by road condition and weather are shown in Table 7 and Table 8, respectively.




Fig 4 Safety-Risk Composition (%) by Road Condition.

Table 7 Safety-Risk Composition (%) by Road Condition.

Road Condition	Low (%)	Medium (%)	High (%)
Dry	84.7	14.1	1.2
Foggy	85.5	13.1	1.4
Pothole	71.4	25.5	3.1
Under Construction	69.9	26.7	3.4
Wet	85.9	13.1	1.0

Table 8 Safety-Risk Composition (%) by Weather Condition.

Weather	Low (%)	Medium (%)	High (%)
Clear	84.5	14.3	1.2
Fog	70.5	26.2	3.3
Rain	85.2	14.1	0.6
Storm	73.2	22.2	4.6

Table 9 lists the frequency and percentage share of each AI-detected driver-behaviour category across the full dataset, complementing Figure 3.

Table 9 Frequency Distribution of Driver-Behaviour Categories.

Driver Behaviour	Count	Percentage (%)
Normal/Safe	4837	80.62
Unsafe Following Distance	563	9.38
Harsh Braking	466	7.77
Rapid Acceleration	105	1.75
Overspeeding	29	0.48

V. MACHINE LEARNING MODEL DEVELOPMENT

This section covers Objectives 4, 5 and 6, and outlines the development, training and individual evaluation of the

machine-learning models for safety-risk classification, driver-behaviour classification and predictive maintenance. All models were built in scikit-learn with a 75%/25% stratified train-test split, z-score feature standardisation and a fixed random seed (42) for reproducibility. The results of the

model comparison for all three tasks are deferred to the consolidated evaluation in Section 7 (Objective 8); the results of each task are reported here with confusion matrices and feature importances.

➤ *Real-Time Safety-Risk Classification (Objective 4)*

The safety-risk classification task is a three-class (Low/Medium/High) classification problem based on ten driving-dynamics and environmental features (speed,

acceleration, braking intensity, following distance, lane deviation, harsh-event g-force, ambient temperature, humidity, visibility, rainfall) and one-hot-encoded road-type, road-condition, and weather categories. The following six algorithms were trained and compared: Logistic Regression, Decision Tree, Random Forest, Gradient Boosting, K-Nearest Neighbours (k=9) and Support Vector Machine with an RBF kernel. The resulting test-set metrics are reported in Table 10.

Table 10 Safety-Risk Classification Performance Across Six Machine-Learning Algorithms (test set, n=1,500).

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	0.8487	0.8210	0.8487	0.8235	0.8378
Decision Tree	0.8553	0.8404	0.8553	0.8450	0.8520
Random Forest	0.8760	0.8513	0.8760	0.8599	0.9284
Gradient Boosting	0.8873	0.8748	0.8873	0.8775	0.9352
K-Nearest Neighbors	0.8307	0.7894	0.8307	0.7695	0.7650
SVM (RBF)	0.8533	0.8189	0.8533	0.8207	0.8814

Among the six algorithms, Gradient Boosting had the highest weighted F1-score (0.8775) and ROC-AUC (0.9352) which is in line with the ensemble-superiority result of [13]. As expected, K-Nearest Neighbours had the lowest F1-score, due to its susceptibility to the high dimensional, mixed scale

feature space after one-hot encoding. Figure 5 shows the visualisation of accuracy, F1-score and ROC-AUC of all six algorithms and Figure 6 shows the receiver-operating-characteristic (ROC) curve for the "High-risk" class for the top two algorithms.




Fig 5 Comparison of Accuracy, F1-Score, and ROC-AUC Across six ML Algorithms for Safety-Risk Classification.




Fig 6 ROC Curve for Binary "High-risk vs. Rest" Detection — Random Forest vs. Logistic Regression (AUC = 0.954 and 0.924 Respectively).

The confusion matrix for the best model (Gradient Boosting) on the held out test set is shown in figure 7. Misclassification is focused at the Low/Medium boundary, as would be expected for a continuous, threshold-based risk

score; High-risk records, although few (1.55% of the data), are identified at a high rate, which is the operationally critical performance characteristic for a safety-warning system.




Fig 7 Confusion Matrix — Gradient Boosting (Best Model) for Safety-Risk Classification.

The values of the confusion matrix are presented in Table 11, and the feature-importance ranking from the Random Forest model is shown in Figure 8, which again

shows that the top three model-derived features are braking intensity, following distance, and harsh-event g-force.

Table 11 Confusion Matrix — Gradient Boosting (Safety-Risk Classification, Test Set).

Actual \ Predicted	High	Low	Medium
High	4	0	19
Low	0	1198	37
Medium	4	109	129




Fig 8 Random Forest Feature Importance Ranking for Safety-Risk Classification.

➤ *AI-Based Driver-Behaviour Monitoring and Classification (Objective 5)*

The driver-behaviour classification task classifies one of five classes: Normal/Safe, Overspeeding, Harsh Braking,

Rapid Acceleration, and Unsafe Following Distance, based on the same driving-dynamics and environmental feature. Logistic Regression, Random Forest and Gradient Boosting were compared. Table 12 reports the results.

Table 12 Driver-Behaviour Classification Performance Across three ML Algorithms (test set, n=1,500).

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	0.9580	0.9548	0.9580	0.9483	0.9825
Random Forest	0.9947	0.9953	0.9947	0.9949	0.9999
Gradient Boosting	0.9933	0.9946	0.9933	0.9939	0.9992

Random Forest and Gradient Boosting were able to achieve a weighted F1-score of 0.9949 and ROC-AUC of 0.9999, respectively, which were very close to each other and significantly better than Logistic Regression. This near-ceiling performance is due to the fact that each behaviour category is threshold derived from one or two dominant sensor readings (e.g. overspeeding from speed, harsh braking

from braking intensity) and tree-based ensembles can isolate these almost perfectly, as reported in [17] for LightGBM-based aggressive-driving detection with high accuracy (>0.95). The confusion matrix of the best model (Random Forest) is shown in Figure 9 and the comparison of all three algorithms is shown in Figure 10.




Fig 9 Confusion Matrix — Random Forest (Best Model) for Driver-Behaviour Classification.




Fig 10 Comparison of Accuracy, F1-Score, and ROC-AUC Across three ML Algorithms for Driver-Behaviour Classification.

The confusion matrix of the random forest model for driver-behaviour classification on the test set is presented in the table 13.

Table 13 Confusion Matrix — Random Forest (Driver-Behaviour Classification, Test Set).

Actual \ Predicted	Harsh Braking	Normal	Overspeeding	Rapid Accel.	Unsafe Following
Harsh Braking	117	0	0	0	0
Normal	0	1205	4	0	0
Overspeeding	1	1	5	0	0
Rapid Accel.	0	2	0	24	0
Unsafe Following	0	0	0	0	141

➤ Machine-Learning-Based Predictive Maintenance (Objective 6)

The predictive-maintenance module tackles two complementary sub-tasks. The first is a binary classification of imminent mechanical/safety-related failure, based on ten features of engine and vehicle health, similar to the sensor-based failure-flagging approach of [18] (engine RPM, engine temperature, oil pressure, vibration level, brake-pad wear, tyre pressure, battery voltage, vehicle age, odometer reading, and speed). The second is a continuous regression of

Remaining Useful Life (RUL, in days), following the classification-versus-regression framing suggested in [19] and the precedent of the XGBoost-regressor in [21].

• Imminent-Failure Classification

The performance of the three machine-learning classifiers for imminent –failure prediction on the test set is summarized in table 14.

Table 14 Predictive-Maintenance (Imminent-Failure) Classification Performance (test set, n=1,500).

Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC
Logistic Regression	0.8660	0.6211	0.4167	0.4988	0.8796
Random Forest	0.8707	0.6554	0.4042	0.5000	0.8715
Gradient Boosting	0.8613	0.5930	0.4250	0.4951	0.8619

The three classifiers were compared, and Random Forest had the highest F1-score (0.5000) and Logistic Regression had the highest ROC-AUC (0.8796), which shows good overall discrimination ability despite the moderately imbalanced classification problem due to the 15.98% base rate of imminent-failure records. The ROC

curve for the best model (Random Forest) is shown in Figure 11, while the feature-importance ranking is shown in Figure 12, which indicates that the top three features are engine temperature, vibration level, and brake-pad wear, all of which are related to vibration-based diagnostics as described in [20] and as seen in the sensor-drift correlations in Figure 1.




Fig 11 ROC Curve — Random Forest (Best Model) for Predictive-Maintenance Failure Classification.




Fig 12 Random Forest Feature Importance Ranking for Predictive-Maintenance Classification.

- *Remaining Useful Life (RUL) Regression*

The regression performance of the models for remaining Useful life (URL) prediction on the test set is summarized in Table 15.

Table 15 Remaining Useful Life (RUL) Regression Performance (test set, n=1,500).

Model	MAE (days)	RMSE (days)	R ² Score
Linear Regression	14.250	17.662	0.794
Random Forest Regressor	9.918	12.486	0.897

The Random Forest Regressor significantly outperformed the linear baseline, with a mean absolute error of 9.92 days compared to 14.25 days, and an R^2 of 0.897 compared to 0.794, showing that the relationship between the engine-health sensor readings and remaining service life is meaningfully non-linear. This MAE is similar to the 8-15 days reported in [21] for RUL prediction of heavy-truck

components using XGBoost. Figure 13 compares the predicted RUL with the actual RUL for a random sample of 400 test-set records, and demonstrates good agreement with the diagonal line of perfect prediction, with some overestimation at very low RUL values, where the need for maintenance intervention is greatest.




Fig 13 Actual vs. Predicted Remaining Useful Life (days) — Random Forest Regressor, test set (n=400 sampled).

➤ Integrated Real-Time Warning and Decision-Support Mechanism

Objective 7 involves combining the Objective-4 safety-risk model and the Objective-6 predictive-maintenance model into a unified real-time warning and decision-support system, similar to the alert-fusion architectures in [22,23]. The proposed system does not provide three separate notification streams as previous research has identified this as a driver cognitive overload [22] but rather integrates the predicted safety-risk level and the predicted maintenance-failure probability into a single three-tier notification:

- **NORMAL** — Continue monitoring: issued when predicted safety risk is Low and predicted failure probability is less than 0.4.

- **CAUTION** — Advisory notification to driver/fleet manager: issued when predicted safety risk is Medium, or predicted failure probability is between 0.4 and 0.7.
- **CRITICAL ALERT** — Immediate driver warning plus workshop advisory: issued when predicted safety risk is High, or predicted failure probability is > 0.7 .

This rule-plus-ML fusion layer was tested on the same held-out test set of 1,500 records as in Section 5, with the two input predictions provided in real-time by the Gradient Boosting safety-risk model and the Random Forest maintenance-failure model. The resulting distribution of alert levels is given in Table 16 and illustrated in Figure 14 as a pie chart.

Table 16 Distribution of Integrated Real-Time Decision-Support Alert Levels (Test Set, n=1,500).

Alert Level	Count	Percentage (%)
NORMAL - Continue monitoring	1105	73.67
CAUTION - Advisory notification to driver/fleet manager	307	20.47
CRITICAL ALERT - Immediate driver warning + workshop advisory	88	5.87




Fig 14 Distribution of Integrated Decision-Support Alert Levels Across the Test Set.

Most records (73.7%) were classified as NORMAL, about one-fifth (20.5%) were given a CAUTION classification, and a smaller but operationally critical proportion (5.9%) were classified as CRITICAL ALERT, which means that the driver was notified and a workshop was scheduled. This distribution of alert levels is generally consistent with the low base rate of High-risk driving records (Table 1) and the 15.98% imminent-failure rate in the underlying data, and shows that the fused decision layer does not over-alert or under-alert relative to the ground-truth hazard rate. This output for decision support would be sent via the same GSM uplink, allowing for near real-time driver notification, as required by the edge-computing latency discussed in [22] for an embedded deployment.

VI. RESULTS AND COMPARATIVE DISCUSSIONS

This section deals with Objective 8, which is the multi-metric reporting of results, as recommended in [24] and as theoretically underpinned by ROC-AUC in [25] and which combines the results reported individually in Section 5 into a single comparative evaluation for all algorithms and all three modelling tasks.

➤ Consolidated Classification Performance

Table 17 summarizes the Accuracy, Precision, Recall, F1-Score and ROC-AUC for all the algorithms tested for the three classification tasks (safety-risk, driver-behaviour and predictive-maintenance classification). The F1 score of each algorithm for each task is visualised in Figure 15.

Table 17 Consolidated Comparison of all Machine-Learning Algorithms Across the Three Classification Tasks.

Task	Model	Accuracy	Precision	Recall	F1	ROC-AUC
Safety Risk Classification	Logistic Regression	0.8487	0.8210	0.8487	0.8235	0.8378
Safety Risk Classification	Decision Tree	0.8553	0.8404	0.8553	0.8450	0.8520
Safety Risk Classification	Random Forest	0.8760	0.8513	0.8760	0.8599	0.9284
Safety Risk Classification	Gradient Boosting	0.8873	0.8748	0.8873	0.8775	0.9352
Safety Risk Classification	K-Nearest Neighbors	0.8307	0.7894	0.8307	0.7695	0.7650
Safety Risk Classification	SVM (RBF)	0.8533	0.8189	0.8533	0.8207	0.8814
Driver Behaviour Classification	Logistic Regression	0.9580	0.9548	0.9580	0.9483	0.9825
Driver Behaviour Classification	Random Forest	0.9947	0.9953	0.9947	0.9949	0.9999
Driver Behaviour Classification	Gradient Boosting	0.9933	0.9946	0.9933	0.9939	0.9992
Predictive Maintenance Classification	Logistic Regression	0.8660	0.6211	0.4167	0.4988	0.8796
Predictive Maintenance Classification	Random Forest	0.8707	0.6554	0.4042	0.5000	0.8715
Predictive Maintenance Classification	Gradient Boosting	0.8613	0.5930	0.4250	0.4951	0.8619




Fig 15 Consolidated F1-Score Comparison Across all Algorithms and all Three ML Tasks.

➤ Regression Performance for Predictive Maintenance

For completeness, Table 18 restates the RUL-regression comparison from Section 5.3.2 for the Objective-8 validation

requirement, which explicitly mentions mean absolute error and root mean square error as required metrics for the regression part of the study.

Table 18 Remaining Useful Life Regression Performance — Algorithm Comparison.

Model	MAE (days)	RMSE (days)	R ² Score
Linear Regression	14.250	17.662	0.794
Random Forest Regressor	9.918	12.486	0.897

The consolidated comparison reveals three findings. First, ensemble tree-based methods (Random Forest, Gradient Boosting) outperform linear and instance-based baselines (Logistic Regression, KNN) consistently across all three classification tasks, with F1-score improvements of 4-6 percentage points for safety-risk classification and 3-4 percentage points for driver-behaviour classification. This is in line with the ensemble-superiority results that have been reported for telematics-based risk classification in [13] and aggressive-driving detection in [17]. Second, the driver-behaviour classification task is able to reach much higher performance levels ($F1 > 0.99$) than the safety-risk task ($F1 \approx 0.88$), as the behaviour labels are directly thresholded from one or two dominant sensors, while the risk-level label is based on multiple weighted factors and stochastic noise, which more closely approximates the ambiguity of real-world crash risk. Third, the error in the prediction of RUL is reduced by about 30% when using non-linear regression (Random Forest Regressor) compared to a linear baseline, which validates the non-linear fault-diagnosis literature based on empirical-mode-decomposition reviewed in [20] that engine-health degradation is not a linear function of the underlying sensor readings.

VII. CONCLUSION AND FUTURE WORK

This paper presented the design, implementation and comparative experimental validation of an IoT-AI-ML integrated smart embedded system for real-time safety

monitoring, driver-behaviour classification and predictive maintenance of heavy commercial vehicles, which covered all the eight research objectives. A four-layer architecture was proposed and implemented on a structured dataset of 6,000 records, 40 vehicles and 60 drivers, comprising of embedded IoT acquisition, data transmission, Python-based data preprocessing and analytics, and AI/ML-based decision support. Braking intensity, following distance and harsh-event g-force were the top three hazardous-driving indicators identified by statistical and exploratory analysis, and road condition and safety-risk level were statistically significantly associated ($\chi^2 = 135.04$, $p < 0.001$) as were speed and safety-risk level ($F = 8.24$, $p < 0.001$). Six machine-learning algorithms were compared for safety-risk classification, with Gradient Boosting showing the highest weighted F1-score (0.878) and ROC-AUC (0.935); a parallel driver-behaviour classification framework achieved a weighted F1-score of 0.995 using Random Forest; and a predictive-maintenance module combining failure classification (ROC-AUC = 0.872) with Random-Forest-based Remaining-Useful-Life regression (MAE = 9.92 days, $R^2 = 0.897$) was able to predict mechanical failure before it occurred. A real-time decision-support mechanism was successfully developed in three tiers (NORMAL/CAUTION/CRITICAL ALERT) by fusing all three model outputs. These results validate the feasibility and measurable effectiveness of an integrated, comparatively validated IoT-AI-ML pipeline for heavy-commercial-vehicle safety management, directly addressing the fragmentation gap identified in the literature review. There are four

directions for future work. First, field deployment on an instrumented pilot fleet is required to validate the pipeline on real sensor noise, GPS multipath error and network-latency characteristics that are not fully captured in a structured simulation. Second, the embedded acquisition layer should be expanded towards vehicle-to-infrastructure (V2I) communication, based on the VANET foundations reviewed in [3] and include cooperative hazard warnings from nearby vehicles. Third, deep-learning architectures that can model sequential driving patterns, such as recurrent and temporal-convolutional networks, should be compared to the tree-based ensembles used here, following the deep-learning precedent of [14] which models sequential driving patterns. Fourth, the predictive-maintenance module should be expanded to include vibration-signal spectral-decomposition features, based on the empirical-mode-decomposition approach of [20], to enhance the sensitivity of the fault detection in the early stages beyond the vibration-level feature that is used in the present study.

DECLARATIONS

- Ethics approval and consent to participate :- Not applicable
- consent for publication:- Not applicable
- Availability of data and materials:- the data and materials supporting the findings of this study are available from the corresponding author upon reasonable request.
- Competing interests:- The author's declare that they have no competing interests.
- Funding:- The authors received no specific funding for this research.
- Authors Contributions:- BKS: Conceptualization, methodology, investigation, data analysis, and writing of the original manuscript. RSV: Supervision, review, and editing, of the manuscript. Both authors read and approved the final manuscript.
- Acknowledgments:- Not applicable

REFERENCES

- [1]. A. Goel and V. Gruhn, "Fleet Monitoring System for Advanced Tracking of Commercial Vehicles," IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 2517–2522, 2006.
- [2]. S. Chen, Q. Xue, X. Zhao, Y. Xing, and J. J. Lu, "Risky Driving Behavior Recognition Based on Vehicle Trajectory," International Journal of Environmental Research and Public Health, vol. 18, no. 23, p. 12373, 2021.
- [3]. A. Al-Sultan, M. M. Al-Doori, A. H. Al-Bayatti, and H. Zedan, "A comprehensive survey on vehicular Ad Hoc network," IEEE Transactions on Intelligent Transportation Systems, vol. 15, no. 1, pp. 167–181, 2014.
- [4]. J. E. Lee, K. T. Kim, and H. Y. Youn, "An IoT-based driver safety management system using GPS and sensor data," IEEE Access, vol. 8, pp. 112340–112351, 2020.
- [5]. M. Eren, A. S. Erciyas, and O. K. Sahingoz, "IoT-based real-time vehicle tracking and accident detection framework using MQTT and cloud analytics," IEEE Internet of Things Journal, vol. 9, no. 14, pp. 12450–12461, 2022.
- [6]. X. Zhang, W. Chen, and Y. Wang, "CAN bus data acquisition and real-time processing architecture for heavy commercial vehicle telematics," IEEE Transactions on Vehicular Technology, vol. 70, no. 6, pp. 5420–5432, 2021.
- [7]. S. K. Gupta and P. S. Singh, "Edge-cloud hybrid data acquisition framework for real-time fleet analytics," Journal of Network and Computer Applications, vol. 185, p. 103080, 2021.
- [8]. R. Kumar, S. Verma, and A. Sharma, "A Python-based scalable data preprocessing pipeline for noisy IoT sensor streams," Future Generation Computer Systems, vol. 128, pp. 312–324, 2022.
- [9]. J. S. Paefgen, F. Kehr, Y. Zhai, and F. Michahelles, "Driving behavior analysis with smartphone sensing: Insights from a large-scale field study," IEEE Transactions on Intelligent Transportation Systems, vol. 15, no. 5, pp. 2223–2237, 2014.
- [10]. Y. Ma, B. Zhang, and C. Liu, "Exploratory data analysis of commercial vehicle crash risk factors using telematics data," Accident Analysis & Prevention, vol. 152, p. 105988, 2021.
- [11]. K. A. Unni and M. R. Kaimal, "Statistical analysis of heavy commercial vehicle driving patterns for hazardous condition identification," Transportation Research Part C: Emerging Technologies, vol. 119, p. 102762, 2020.
- [12]. T. Tochigi, H. Sugiura, and K. Nishiyama, "Machine learning-based real-time driving risk assessment for heavy commercial vehicles," IEEE Transactions on Intelligent Vehicles, vol. 6, no. 3, pp. 512–523, 2021.
- [13]. M. Bozkurt, A. C. Yilmaz, and T. Acarman, "Real-time driving risk classification using Random Forest and XGBoost on telematics data," Expert Systems with Applications, vol. 191, p. 116245, 2022.
- [14]. C. Wang, Y. Zhang, and L. Zhao, "Deep learning and ensemble learning models for collision risk prediction in heavy-duty truck operations," Safety Science, vol. 148, p. 105650, 2022.
- [15]. J. Dong, S. S. Mulgund, and R. E. Ahern, "In-vehicle driver behavior monitoring and unsafe maneuver classification using multi-sensor fusion," IEEE Sensors Journal, vol. 21, no. 8, pp. 10120–10131, 2021.
- [16]. S. Castignani, T. Derrmann, R. Frank, and T. Engel, "Driver behavior profiling using smartphones: A methodology for driving safety scoring," IEEE Transactions on Intelligent Transportation Systems, vol. 16, no. 1, pp. 494–503, 2015.
- [17]. H. Eren and M. C. Soylu, "LightGBM-based detection of aggressive driving behavior in heavy commercial vehicles," Computers & Electrical Engineering, vol. 102, p. 108190, 2022.
- [18]. Z. A. S. Al-Sulaifanie and M. A. M. Ali, "Predictive maintenance for heavy vehicles using Internet of Things and machine learning algorithms," IEEE Access, vol. 9, pp. 134210–134224, 2021.

- [19]. A. Bousdekis, K. Lepenioti, and G. Mentzas, "A review of predictive maintenance architectures in Industry 4.0 and IoT telematics," *Computers in Industry*, vol. 126, p. 103393, 2021.
- [20]. Y. Lei, N. Li, L. Gontarz, J. Lin, and S. Zheng, "A review on empirical mode decomposition in fault diagnosis of rotating machinery," *Mechanical Systems and Signal Processing*, vol. 35, no. 1-2, pp. 108–126, 2013.
- [21]. R. S. Carvalho, R. M. Souza, and L. A. Ferreira, "Remaining useful life prediction for heavy commercial truck components using XGBoost regressor," *Reliability Engineering & System Safety*, vol. 218, p. 108150, 2022.
- [22]. F. Alam, R. Pathan, and S. M. R. Islam, "An integrated IoT-AI edge computing framework for real-time vehicular safety alerts," *IEEE Internet of Things Journal*, vol. 10, no. 4, pp. 3410–3422, 2023.
- [23]. D. B. Patel and R. H. Jhaveri, "Intelligent decision support mechanism for fleet management using cloud-edge telemetry," *Vehicle Communications*, vol. 34, p. 100450, 2022.
- [24]. X. Yuan, L. Li, and Y. Wang, "Experimental validation and performance metrics comparison of machine learning models in vehicular telemetry," *Pattern Recognition Letters*, vol. 155, pp. 88–96, 2022.
- [25]. T. Fawcett, "An introduction to ROC analysis," *Pattern Recognition Letters*, vol. 27, no. 8, pp. 861–874, 2006.