

EFFICIENT ROAD SURFACE CLASSIFICATION ON LOW-COST DEVICES USING VEHICLE VIBRATION DATA

Cong Ngo Van¹, Duc-Nghia Tran², Bui Thi Thu³, Vu Duong Tung⁴, Pham Quang Huy⁵, Manh-Tuyen Vi^{6*}, Duc-Tan Tran^{7*}

Graduate University of Science and Technology, Vietnam Academy of Science and Technology, Hanoi, Vietnam¹

Institute of Information Technology, Vietnam Academy of Science and Technology, Hanoi, Vietnam^{2,3,4}

Faculty of Information Technology, Electric Power University, Hanoi, Vietnam⁵

Faculty of Electrical and Electronic Engineering, Phenikaa University, Hanoi, Vietnam^{6,7}
tuyen.vimanh@phenikaa-uni.edu.vn^{4*}, tan.tranduc@phenikaa-uni.edu.vn^{4*}

Received: 13 August 2025, Revised: 26 November 2025, Accepted: 10 January 2026

*Corresponding Author

ABSTRACT

During road traffic operations, pavement quality directly affects safety, vehicle operating costs, and pavement maintenance activities. Traditional inspection methods are often costly and time-consuming, and they cannot provide continuous data on pavement conditions. This study aims to develop an efficient road-surface classification system capable of real-time operation on low-cost hardware devices. The system uses vibration data collected from vehicles in motion to identify and classify road types with high accuracy and optimized performance. The proposed system employs inertial sensors mounted on vehicles to acquire accelerometer and gyroscope signals and then extracts time-domain statistical features from these signals. To address the main challenge of deploying an effective recognition model in a resource-constrained computing environment, the paper proposes a hybrid feature selection algorithm that combines filter and wrapper methods. This algorithm leverages the fast-processing speed of filter methods and the effective feature selection capability of wrapper methods. The selected feature set is then evaluated using three machine learning models: Random Forest (RF), Gradient Boosting (GBM), and XGBoost. The classification task focuses on three real-world pavement types: smooth asphalt (with less than 10 years of service), degraded asphalt (with more than 15 years of service), and cement concrete pavement. Experimental results show that the proposed feature selection algorithm and classification models achieve high classification performance and fast execution speed. The system attains accuracy higher than 0.95 while reducing computational cost. These findings confirm the feasibility of deploying road-surface classification systems on low-cost devices for real-time pavement monitoring and highlight the importance of appropriate feature selection in balancing system accuracy and performance.

Keywords: Road surface monitoring, Inertial sensors, Machine learning, Real-time monitoring.

1. Introduction

Maintaining the road infrastructure system in an optimal condition is a critical task for transport authorities, as pavement deterioration directly affects traffic safety, vehicle operating costs, and the long-term effectiveness of maintenance strategies (Lekshmiathy et al., 2021; Martinez-Ríos et al., 2022). Surface deterioration can arise from material aging, the impact of climatic conditions, repeated traffic loads, as well as deficiencies in preventive maintenance activities (Varona et al., 2020). In addition, road surface deterioration negatively affects the overall transportation infrastructure. If initial damages are not addressed promptly, they can spread and trigger a chain reaction that impacts the underlying structural layers, eventually leading to complete degradation of the pavement. Common signs of deterioration include cracks, potholes, damaged speed bumps, and rutting. These forms of deterioration increase the risk of unsafe conditions for vehicles and road users (Celaya-Padilla et al., 2018).

At the network level, road infrastructure represents a major public expenditure: Collier et al. (2016) reported that the construction and maintenance of roads in low- and middle-income countries absorb a substantial share of public investment. From a safety perspective, Alhasan et al. (2018) and Huynh et al. (2025) showed that higher levels of surface roughness and rutting are significant predictors of increased crash frequencies and more severe outcomes. Typical signs of

surface damage include cracks, potholes, damaged speed bumps, and wheel-path rutting, all of which require substantial maintenance and rehabilitation budgets if not treated in a timely manner (Celaya-Padilla et al., 2018; Martínez-Ríos et al., 2022). These findings highlight the need for efficient and scalable methods to monitor pavement conditions and to support timely maintenance decisions as a foundation for intelligent transportation systems and data-driven infrastructure management.

Traditionally, road surface condition has been assessed through direct field surveys based on indices such as the Pavement Condition Index (PCI). Although widely used, these surveys are labour-intensive, costly, strongly influenced by human subjectivity, and may pose health and safety risks for field technicians (Martinelli et al., 2022; Shaghilil & Khalafallah, 2018). Moreover, manual inspections are difficult to deploy at high spatial and temporal resolution, which limits their suitability for large-scale, continuous monitoring. These limitations have motivated the development of automated, sensor-based approaches that can support early damage detection and provide timely warning signals.

Recent advances in sensor technologies, non-destructive testing, and data analysis methods have led to a broad range of techniques for road surface condition monitoring and anomaly detection (G. Loprencipe et al., 2017; Varona et al., 2020). Existing approaches can be classified into three main groups:

- The first group relies on vehicle vibration or acceleration data, collected by sensors mounted on vehicles or mobile devices, to infer road roughness and the degree of surface deterioration (Dehnad et al., 2024; Nyirandayisabye et al., 2022).
- The second group leverages computer vision and image/video processing, in which machine learning and deep learning models are applied to automatically detect and classify surface defects or identify surface types (Maeda et al., 2018; Sun et al., 2022).
- The third group uses 2D/3D laser scanning or LiDAR-based geometric measurements to reconstruct detailed surface profiles and structures for quality assessment and mapping (Río-Barral et al., 2022; Guan et al., 2015).

Although these approaches have significantly improved the reliability of pavement condition evaluation, many of them require high-end sensing platforms, substantial computing resources, or complex calibration procedures, which can limit their applicability in large-scale deployments with constrained budgets.

To address these challenges, recent studies have increasingly focused on vibration-based and device-based road surface monitoring using accelerometers, gyroscopes, and smartphone sensors combined with machine learning models (Martínez-Ríos et al., 2022). Review papers have summarised threshold-based, feature-extraction, and deep learning approaches for road surface anomaly detection and classification and have emphasised the need for robust methods that can operate under varying speeds and heterogeneous platforms (Martínez-Ríos et al., 2022; Sattar et al., 2018; Alqaydi et al., 2024). Several works have demonstrated that low-cost inertial sensors can effectively identify road surface anomalies and surface types when combined with supervised learning algorithms, including tree-based models and neural networks (Martinelli et al., 2022; Meocci, 2024; Surblys et al., 2024; Cafiso et al., 2022). Other studies have explored near real-time and crowdsourced road monitoring frameworks, showing that embedded and smartphone-based implementations are feasible but still face limitations related to computational complexity, energy consumption, and model generalisation (Sattar et al., 2021; Khahro et al., 2021).

In the context of scalable and cost-constrained monitoring, vibration-based methods that combine inertial sensors (accelerometers and gyroscopes) with machine learning algorithms on low-cost hardware have emerged as a promising solution (Mihoub et al., 2023). In the context of scalable and cost-constrained monitoring, vibration-based methods that combine inertial sensors (accelerometers and gyroscopes) with machine learning algorithms on low-cost hardware have emerged as a promising solution (Mihoub et al., 2023; Khahro et al., 2021). Inertial sensors can be easily integrated into vehicles, are inexpensive, and can capture motion signals that accurately reflect surface geometry and roughness (Surblys et al., 2024). However, a large portion of existing vibration-based studies focuses on anomaly or pothole detection using thresholding schemes that are sensitive to driving behaviour and require manual tuning of parameters for each deployment (Martínez-Ríos et al., 2022; Sattar et al., 2018). Other machine-learning and deep-learning

approaches often assume powerful processing units or cloud-based platforms and therefore do not explicitly address the constraints of low-end microcontrollers in terms of memory, computation time, and energy consumption (Sattar et al., 2021; Alqaydi et al., 2024; Meocci, 2024). In addition, only a limited number of studies systematically investigate feature selection, time-domain feature design, and data window configuration for inertial signals in order to balance classification performance with computational efficiency on embedded devices, especially in the context of road surface monitoring (Martínez-Ríos et al., 2022; Alqaydi et al., 2024).

This study focuses on evaluating the capability of road surface classification using a single accelerometer and a single gyroscope mounted on a survey vehicle, targeting real-time operation on low-cost embedded hardware within an IoT-based monitoring architecture. Inertial sensors such as accelerometers and gyroscopes are particularly attractive because they can capture motion signals from the object to which they are attached, thereby accurately reflecting geometric irregularities and surface roughness. When these data are processed through feature extraction and combined with appropriate machine learning algorithms, the system can automatically detect and classify damage with high reliability while adapting to various operating conditions. Notably, with low power consumption, low cost, and easy integration into low-performance microcontrollers, these sensors still ensure the capability for real-time processing and classification (Manh-Tuyen Vi et al., 2024). This study aims to design an efficient feature extraction and feature selection pipeline for inertial data, to select lightweight yet accurate machine learning models, and to configure the data window to maintain high classification performance while reducing computational cost and energy consumption. The proposed solution contributes to bridging the gap between high-accuracy algorithms and their practical deployment on resource-constrained hardware devices.

This study focuses on evaluating the capability of road surface classification using accelerometer and gyroscope sensors mounted on a vehicle. The main challenges encountered during the development of the road surface monitoring system are as follows:

- Challenge 1: Limited classification accuracy when using accelerometers and gyroscopes. The system has only one accelerometer and one gyroscope, so its ability to accurately identify different road surface types is constrained.
- Challenge 2: Energy constraints in real-time IoT systems. Applying road surface classification algorithms with high computational complexity leads to significant energy consumption, considerably shortening the device's operating time. This challenge requires optimizing the processing workflow and selecting algorithms that ensure accuracy while minimizing energy consumption.
- Challenge 3: Optimization between algorithms, features, and data windows to balance classification performance, time, and energy usage in the system.

To address these challenges, this paper proposes the design of an IoT-based road surface classification system that operates on vibration data collected from inertial sensors installed on vehicles. The road surface is classified into three categories: asphalt roads less than 10 years old, asphalt roads more than 15 years old, and concrete roads. The system is implemented on a low-cost microcontroller platform and integrates a machine learning model optimised for edge deployment. The proposed approach centres on a hybrid feature selection strategy that combines filter and wrapper mechanisms, implemented through the Recursive Feature Elimination (RFE) technique, to reduce the number of features while maintaining or improving classification accuracy. Three lightweight supervised learning models—RF, GBM, and XGBoost—are then trained and evaluated on the selected features to determine the most suitable option for embedded implementation.

The main contributions of this research are summarized as follows:

- Contribution 1: We develop a road surface classification method based on data from accelerometers and gyroscopes mounted on vehicles, with three road surface types: asphalt roads under 10 years old, asphalt roads over 15 years old, and concrete roads.
- Contribution 2: We compare the performance of three popular machine learning algorithms—RF, GBM, and XGBoost—on the same directly collected dataset.
- Contribution 3: We apply the Recursive Feature Elimination (RFE) technique to select the optimal number of features, reducing model complexity and training time

while maintaining high accuracy. The results show that only about 27 features are sufficient to achieve performance close to that of the full model.

2. Literature Review

The rapid development of sensing and data analysis technologies has strongly driven research on road surface condition detection, and the information obtained can be integrated into various intelligent transportation systems (Menegazzo et al., 2021). Among these, the approach of using inertial sensors mounted on vehicles is considered a highly promising method for assessing road surface conditions (Menegazzo et al., 2021; Rathee et al., 2023).

Threshold-based method.

Research groups using threshold-based methods for detecting road surface anomalies often rely on accelerometer data. Wu et al. (2020) developed a pothole patrol system that combines acceleration threshold values with GPS for digital map localization, along with five noise filtering criteria: speed, high-pass filtering, Z-axis peak, XZ ratio, and the relationship between speed and the Z-axis. Xin et al. (2023) mentioned the trend of acceleration on all three axes approaching zero when a vehicle passes over a pothole and proposed the G-ZERO algorithm, achieving 0.85 accuracy. Sebestyen et al. (2015) analyzed upward acceleration oscillations and achieved about an 0.8 detection rate. In a related line of work, Du et al. (2020) used smartphone accelerometer data and designed a pipeline combining preprocessing, feature extraction, and rule-based detection to recognise abnormal road surface conditions, confirming that low-cost embedded sensors can provide reliable anomaly detection in real traffic.

Although simple and easy to implement, threshold-based methods (Wu et al., 2020; Xin et al., 2023) are easily affected by factors such as sudden braking or sharp turns and require threshold adjustments depending on the road type, device, and vehicle, making large-scale deployment difficult. In addition, most threshold-based systems focus on anomaly detection (e.g., pothole vs. no pothole) rather than multi-class road surface type classification and rarely address computational or energy constraints of low-power IoT devices.

Traditional Machine Learning Methods

Machine learning has been applied to improve road surface classification accuracy by combining inertial sensor and GPS data. Wang et al. (2018) and their earlier 2017 study used acceleration data from the suspension system and GPS speed, applied the Quarter Car (QC) model to compensate for suspension effects, then extracted frequency-domain features using FFT and trained an SVM, achieving an average accuracy of 0.69. Martinelli et al. (2022) used Short-Time Fourier Transform (STFT) to extract time–frequency-domain features and evaluated Decision Tree, SVM, and K-Nearest Neighbours (KNN), with the cubic SVM achieving a PPV of up to 0.97 for potholes/manholes and 0.84 for fatigue cracks.

In the study by Li et al. (2019), acceleration data collected via smartphone were processed using Continuous Wavelet Transform (CWT) with a Daubechies wavelet of order 3, combined with DBSCAN clustering to detect potholes, achieving an accuracy of 0.94. Other smartphone-based studies have also shown that carefully designed features from inertial data can support accurate identification of abnormal surfaces and specific damage types in real-world driving scenarios (e.g., Du et al., 2020).

Despite these advances, many machine learning studies use relatively small or locally collected datasets, which limit the generalizability of their models to different vehicle types, driving behaviours, and pavement structures. Moreover, several works assume processing capabilities comparable to smartphones or single-board computers and do not explicitly analyse the computational or energy cost of deploying the models on low-end microcontrollers.

Deep Learning Methods

Recent advances in deep learning have boosted road surface classification accuracy (Menegazzo et al., 2020). Varona et al. (2020) used CNN and LSTM to classify four different surface types, achieving up to 0.93 accuracy with CNN when classifying potholes, speed bumps, and manholes. Menegazzo et al. (2020) tested CNN, LSTM, and GRU on inertial sensor data

collected from various road types, showing that CNN achieved a training accuracy of 0.93, with asphalt roads reaching 0.986.

While deep learning approaches show strong predictive performance, they typically require large, labeled datasets, high computational resources, and careful hyperparameter tuning. Deploying such models on microcontroller-class devices is challenging due to memory and processing constraints, unless additional model compression, pruning, or specialised frameworks are introduced. Furthermore, deep learning methods often treat feature extraction as a black box, providing limited insight into which signal characteristics are most informative for road surface classification under strict resource budgets.

Low-cost embedded and monitoring systems

Several studies have investigated low-cost hardware architectures and edge-based solutions for pavement monitoring. Martinelli et al. (2022) proposed a low-cost system for real-time screening of pavement conditions based on accelerometers mounted on vehicles, showing that inexpensive sensors can support effective network-level assessments. Loprencipe et al. (2021) validated an inertial-based system consisting of a low-cost three-axis MEMS accelerometer and GPS connected to a Raspberry Pi Zero W, embedded inside a vehicle to indirectly monitor road conditions on urban networks. Bruno et al. (2023) further proposed a low-cost and easy-to-install monitoring system to assess pavement deterioration on urban roads using frequency-weighted vertical acceleration in combination with traditional PCI inspections.

These works demonstrate the feasibility of low-cost, embedded pavement monitoring, but they generally focus on specific sensor configurations and do not systematically explore the trade-offs between feature design, model complexity, processing latency, and energy consumption. In particular, the multi-class classification of road surface types using both accelerometer and gyroscope signals on microcontroller platforms remains underexplored, especially when real-time operation and battery constraints are considered.

Feature Selection and tree-based ensemble methods

Feature selection plays a key role in improving model performance and accuracy (Solorio-Fernández et al., 2016), with three main method groups: filter, wrapper, and hybrid. Filter methods are simple and fast but tend to be less accurate than wrapper methods, which iteratively evaluate feature subsets using a classifier, while hybrid methods combine the advantages of both (Zhang et al., 2019; Yin et al., 2023).

In high-dimensional classification tasks, Li et al. (2018) used the Gini index as a filter criterion, combined it with Gradient Boosted Decision Trees (GBDT), and optimised parameters using Particle Swarm Optimization (PSO), reducing the number of features from 41 to 18 and achieving 0.86 accuracy. Mambwe Sydney and Sun (2020) used XGBoost to select 19 out of 42 features, improving accuracy by 1.9%. Megantara and Ahmad (2020) combined feature ranking with RFE to optimise the number of features and reported performance gains. These studies confirm that combining filter-based ranking with wrapper-style search, such as RFE, can yield compact feature sets with minimal loss in accuracy.

Tree-based ensemble methods, including RF, GBM, and XGBoost, provide a strong algorithmic foundation for such feature selection strategies. RF has been widely recognised as a powerful tree-based ensemble model that offers strong predictive performance, robustness to noisy and high-dimensional data, and built-in measures of variable importance (Hu & Szymczak, 2023). GBM methods build additive ensembles of weak learners via gradient-based optimisation and have been systematically analysed and benchmarked against other tree-based ensembles in recent comparative studies (Bentéjac, Csörgő, & Martínez-Muñoz, 2021), while XGBoost extends this framework with sparsity-aware algorithms, regularisation, and system-level optimisations for high efficiency and scalability (Chen & Guestrin, 2016). In many application domains, these models achieve a good balance between accuracy, interpretability, and computational cost, making them attractive candidates for embedded implementations when paired with appropriate feature selection.

Remarks and Challenges

Overall, machine learning and deep learning methods clearly outperform simple threshold-based techniques in terms of classification accuracy, but their effectiveness still depends strongly on data characteristics, feature design, and the available computational resources (Jeon et al., 2020). For road surface monitoring in resource-constrained IoT environments, three key challenges remain: (1) limited accuracy when using only a single accelerometer and gyroscope; (2) strict energy and real-time constraints on low-cost embedded devices; and (3) the need to jointly optimise algorithms, features, and data windows to balance classification performance, processing speed, and energy consumption. These factors motivate the IoT-based road surface classification solution proposed in this paper, which combines lightweight tree-based models with an RFE-based hybrid feature selection strategy to achieve high effectiveness in real-world conditions under hardware constraints.

3. Research Methods

3.1 IoT-based road surface monitoring system

The proposed IoT-based road surface monitoring system enables early detection of pavement quality problems and the transmission of alerts to management platforms. Figure 1 illustrates the overall architecture, which consists of two main components: the in-vehicle device and the central server. The in-vehicle device simultaneously collects inertial sensor data (acceleration, angular velocity) and positioning data (location, speed, time), together with selected vehicle and environmental information. At the edge, the device performs preprocessing, feature extraction, and road surface classification for each time window. The results are packaged into data frames containing the vehicle ID, timestamp, coordinates, speed, acceleration and angular-velocity components, and the predicted road surface label, and then transmitted to the server. The server receives data from multiple vehicles, stores them in a database, assigns surface-condition labels to a digital map according to position and time, and provides the processed information to monitoring and maintenance decision-support applications.




Fig. 1. IoT-based road surface monitoring system

Figure 2 shows the in-vehicle monitoring device in more detail. The device is powered directly from the car battery and uses an ESP32 microcontroller as the main processing unit, which communicates with the server via Wi-Fi or Bluetooth Low Energy (BLE). Peripheral modules, including the NEO-UBLOX-6M V2 GPS receiver, the MPU-9250 inertial measurement unit (IMU), and a microSD card, are connected to the ESP32 through standard serial interfaces (UART, I²C/SPI). The GPS module operates at 3.3 V, is equipped with a ceramic antenna, EEPROM for configuration storage, and a backup battery, and provides information on position, speed, and acquisition time. The MPU-9250 is a 9-axis IMU (3-axis accelerometer, 3-axis

gyroscope, 3-axis magnetometer) that supports typical measurement ranges of $\pm 2g/\pm 4g/\pm 8g/\pm 16g$ for acceleration and $\pm 250/\pm 500/\pm 1000/\pm 2000$ °/s for angular velocity, making it suitable for capturing vibrations, short impacts, and vehicle body oscillations.




Fig. 2. IoT-based road surface monitoring device

During data acquisition, the MPU-9250 measures raw acceleration and angular velocity along the X, Y, and Z axes at a sampling frequency of 50 Hz per axis. The IMU data are transmitted to the ESP32 via I²C/SPI, and the 32-bit microcontroller performs real-time preprocessing, including low-pass filtering to remove high-frequency noise, gravity compensation on the accelerometer channels when needed, and coordinate normalization. The ESP32 then groups the data into time frames and continuously writes them to the microSD card to avoid data loss in case of connectivity interruptions.

Experiments were conducted on urban and suburban roads in Hanoi, Vietnam, where the tropical monsoon climate causes large variations in temperature and humidity, making pavement particularly prone to degradation. The survey vehicle was a compact sedan operated under normal traffic conditions. The IMU was rigidly mounted inside the vehicle, near the midpoint of the wheelbase, to reduce artefacts due to mechanical looseness. Because raw IMU data are typically noisy, a Kalman filter was incorporated into the processing chain to smooth the signals.

To improve labelling accuracy, the inertial and GPS data were synchronised with video from a dashboard camera installed in the vehicle. The video stream was later reviewed to confirm the road surface type and to align the boundaries of homogeneous segments with the inertial and GPS data, providing visual evidence for abnormal events and increasing the reliability of the training and test datasets. When network connectivity is available, the ESP32 sends compact event summaries (timestamp, vibration index, road surface type, coordinates) via Wi-Fi/BLE to the central server, where the classified segments are aggregated and displayed on a digital map to support network-level road surface monitoring.

Table 1 describes the label classes used in the study. Actual observations combined with dashcam footage were used to synchronize labels with IMU data, ensuring the creation of a clean dataset for model training and evaluation.

Table 1 - Road surface label explanation

Label	Explain
asphalt_10	Asphalt road less than 10 years old
asphalt_15	Asphalt road more than 15 years old
Concrete	Concrete road

3.2 Data Collection

In this study, each survey route corresponding to a road surface type was assigned a unique identifier. The measuring device consisted of an MPU-9250 inertial sensor fixed onto the survey vehicle, combined with an ESP32 microcontroller and an SD card for data recording. The MPU-9250 simultaneously measured three-axis acceleration (A_x, A_y, A_z) and three-axis angular velocity (G_x, G_y, G_z).

Data were sampled at a frequency of 50 Hz for each axis, equivalent to 150 samples per second for the accelerometer module and 150 samples per second for the gyroscope module. Additional information included GPS position (Lat, Lon), odometer readings, altitude (Alt), travel speed (Speed), number of satellites received (Sats), and road surface surface condition (ROADSURFACE) labeled during data collection.

After processing, the ESP32 packaged the data, including acceleration, angular velocity, battery status, sampling time, and, when available, GPS data (synchronized at a 1-second cycle). The SD card was used to store all this data to prevent loss in case of no network connection in the field. When Wi-Fi connectivity was available, the data blocks were transmitted to the server for long-term storage and analysis.

Through multiple sampling sessions at different frequencies, the study found that maintaining a sampling frequency of 50 Hz allowed the system to fully capture the vehicle’s rapid oscillations and short-term vibrations when passing over various road surface irregularities. Table 2 illustrates the accelerometer and gyroscope data collected from the measuring device in this study.

Table 2 - Accelerometer and gyroscope data

Asphalt_10					
Ax	Ay	Az	Gx	Gy	Gz
-0.06	-0.06	-0.97	2.87	-3.48	0
-0.04	0.04	-0.99	2.08	0.98	0.12
-0.04	0.05	-0.98	1.53	-0.98	-0.92
0.06	0.03	-0.94	1.46	0.06	-0.12
-0.06	0.09	-0.99	1.28	-1.16	-0.37
0.08	0.06	-1.05	1.04	0.06	-0.24
-0.05	0.1	-0.95	1.1	-0.49	-0.49
0.01	0.08	-0.96	1.71	0.43	-0.31
0.03	0.06	-0.99	1.77	1.83	-0.37
-0.05	0.11	-0.95	2.14	-2.01	-0.92
-0.04	0.11	-0.96	1.95	-2.26	-0.43
-0.07	0.08	-0.87	1.77	-1.71	-0.67
-0.01	0.09	-0.95	1.53	0.31	-0.37
0.05	0.1	-0.97	0.98	1.34	-0.73
-0.02	0.11	-0.95	1.04	0.55	-0.18

Accelerometer and gyroscope data from the in-vehicle device were synchronized to construct the datasets. In this study, non-overlapping windows provide a significant advantage in terms of computational efficiency. Specifically, the number of windows generated is smaller than when using overlapping windows of the same size, which directly reduces the computational load in feature extraction, model inference, and overall processing time. This is a critical factor for real-time monitoring applications on embedded platforms or low-power devices, which are inherently constrained in memory and processing capability. Datasets with window lengths of 8s, 10s, 12s, and 15s were generated. Tables 3 and 4 present the number of 10-second windows for each road type in the Training and Test sets. With a 60%/40% split between Training and Test data, the Training set contains 3150 asphalt_10 windows, 3460 asphalt_15 windows, and 2340 concrete windows (8950 in total), while the Test set contains 2100 asphalt_10, 2310 asphalt_15, and 1560 concrete windows. The 60%/40% ratio was selected after experimenting with multiple split configurations. This ratio provides sufficient data for model training and hyperparameter tuning, while maintaining a large enough independent Test set to enable a reliable evaluation of system performance.

Table 3 - The number of observations for each road surface type from the Training set

Road surface pattern	Total observations
asphalt_10	3150
asphalt_15	3460
Concrete	2340
Total	8950

Table 4 - The number of observations for each road surface type from the Test set

Road surface pattern	Total observations
asphalt_10	2100
asphalt_15	2310
Concrete	1560
Total	5970

3.3 Feature Selection

Feature selection plays an important role in reducing noise, enhancing the ability to distinguish between road surface types, and optimizing model performance. In this study, the feature set includes statistical measures such as Mean, Standard deviation (STD), Variance, Median, Max, Min, and Range. The selected features have low computational cost and are easy to implement on low-cost microcontrollers.

The features extracted below are calculated separately for each measurement axis of the accelerometer (Ax, Ay, Az) and gyroscope (Gx, Gy, Gz) on fixed-size data windows.

- Mean:

$$\mu = \frac{1}{N} \sum_{i=1}^N x_i \tag{1}$$

- STD:

$$\sigma = \sqrt{\frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2} \tag{2}$$

- Variance:

$$Var = \frac{1}{N} \sum_{i=1}^N (x_i - \mu)^2 \tag{3}$$

- Median:

$$Median(x_i) \tag{4}$$

- Max:

$$Max = \max(x_i) \tag{5}$$

- Min:

$$Min = \min(x_i) \tag{6}$$

- Range:

$$Range = Max - Min \tag{7}$$

Where:

- x_i is the value of the i sample in the data window.
- N is the number of samples in the window.
- μ is the mean value of the window.
- σ is the standard deviation.
- Var is the variance.
- Max and Min are the maximum and minimum values in the data window, respectively.
- $Range$ is the signal amplitude.

3.4 Hybrid Filter-Wrapper algorithm

The proposed Hybrid Filter–Wrapper feature selection algorithm aims to improve road surface classification performance on inertial sensor data while optimizing computational cost for real-time applications on resource-constrained hardware. The process consists of two stages: (i) Filter—ranking and removing features with importance below an optimal threshold (determined through experiments with multiple values), enabling a rapid reduction in feature space size; (ii) Wrapper—applying Recursive Feature Elimination (RFE) on the reduced set, re-evaluating the model’s performance at each step, and stopping when accuracy decreases by more than 2%. This method shows that the number of features and training time can be significantly reduced while still maintaining or improving accuracy. Algorithm 1 illustrates the steps in the filtering stage, and Algorithm 2 presents the steps in the wrapper stage.

Algorithm 1 Filter phase

Input:

- T: training set
- V: validation set
- threshold: feature importance threshold
- $X_{all} = [f_1, f_2, \dots, f_n]$: list of original features

Output:

- $X_{selected} = X_f(k)$: list of selected features

Procedure FILTER_FEATURES(T, V, threshold, X_{all}):

1. (acc, imp) \leftarrow TRAIN_AND_EVALUATE_MODEL(T, V, X_{all})
 2. $X_{selected} \leftarrow []$
 3. For i from 1 to $|X_{all}|$:
 - If $imp[i] \geq$ threshold:
 - Append $X_{all}[i]$ to $X_{selected}$
-

4. Return X_selected

End Procedure

Algorithm 2 Wrapper phase

Input:

- T: training set
- V: validation set
- p: performance monitor metric
- threshold: stopping condition threshold
- init_features = [..., fk, ...]: list of features after the filter phase

Output:

- selected_features = Xf(l): final list of selected features

Procedure WRAPPER_PHASE(T, V, p, threshold, init_features):

1. (Accuracy_i, importance) ← EVALUATE_MODEL(T, V, init_features)
2. While p < threshold:
 - Rank features based on importance
 - Remove the least important feature from init_features
 - (Accuracy_j, importance) ← EVALUATE_MODEL(T, V, init_features)
 - Calculate_p ← Accuracy_i - Accuracy_j
 - If Accuracy_j > accuracy_max:
 - Update accuracy_max
 - Select current subset of features
 - Accuracy_i ← Accuracy_j
3. Return selected_features

End Procedure

3.5 Road surface recognition

To ensure efficient execution on embedded hardware with limited computational resources, three lightweight yet effective classification algorithms were selected: RF, GBM, and XGBoost. These models were implemented using optimized versions of the Scikit-learn and XGBoost libraries suitable for low-latency inference.

For the RF model, the number of trees is limited to 50 estimators with a maximum depth of 5, which reduces memory footprint and speeds up decision-making while maintaining acceptable accuracy. The GBM model is configured with a learning rate of 0.1, maximum depth of 2, and 50 estimators, enabling faster updates and minimizing computation during prediction. The XGBoost model applies a learning rate of 0.1, maximum depth of 3, and 80 estimators with the “gbtree” booster, tuned to balance classification performance and resource constraints. To further assess the generalisation capability of the models, we also performed cross-experiments in which models trained on data from one acquisition session were evaluated on data from other independent sessions collected on different days and routes.

To assess the classification performance, four evaluation metrics were employed: Accuracy, Recall, Precision and F1-score, calculated as follows:

$$Acc = \frac{TP + TN}{TP + FP + TN + FN} \tag{8}$$

$$Recall = \frac{TP}{TP + FN} \tag{9}$$

$$Precision = \frac{TP}{TP + FP} \tag{10}$$

$$F1 = \frac{2 \times Sen \times PPV}{Sen + PPV} \tag{11}$$

Where:

- True Positive (TP) indicates the model correctly predicts a road surface type when it is present.
- False Positive (FP) occurs when the model predicts a road surface type that is not actually present.
- False Negative (FN) occurs when a road surface type is present but the model predicts a different one.
- True Negative (TN) occurs when a road surface type is absent and the model correctly predicts its absence.

4. Results and Discussions

Results

The study evaluated road surface classification performance using different window sizes (8, 10, 12, and 15 seconds) and three machine learning models: RF, GBM, and XGBoost. Based on data collected from accelerometers and gyroscopes installed on vehicles, features such as mean, standard deviation, variance, median, maximum/minimum values, and range were applied. The targeted road surface were asphalt roads less than 10 years old (asphalt_10), asphalt roads over 15 years old (asphalt_15), and concrete roads (Concrete). The evaluation metrics used were Accuracy, Recall, Precision, and F1-score (Equations 8–11).

Table 5 - Accuracy of RF, GBM, and XGBoost

Accuracy			
Window size (s)	RF	GBM	XGBoost
8	0.88	0.93	0.94
10	0.89	0.94	0.95
12	0.89	0.94	0.95
15	0.88	0.93	0.94

Table 6 - Recall of RF, GBM, and XGBoost

Recall			
Window size (s)	RF	GBM	XGBoost
8	0.89	0.93	0.94
10	0.90	0.94	0.95
12	0.90	0.94	0.95
15	0.89	0.93	0.94

Table 5 presents the classification accuracy of three machine learning models—RF, GBM, and XGBoost—when applying different window sizes (8s, 10s, 12s, and 15s). The results show that XGBoost consistently achieves the highest Accuracy in all cases (0.94–0.95), followed by GBM (0.93–0.94), and RF with the lowest performance (0.88–0.89). Table 6 presents the Recall values, reflecting each model’s ability to correctly identify instances belonging to each road surface class. The trend in Table 6 is similar to that in Table 5, with XGBoost again leading (0.94–0.95), GBM following closely (0.93–0.94), and RF performing lower (0.89–0.90).

These results highlight XGBoost’s advantage in capturing nonlinear relationships and feature interactions through its boosting mechanism, enabling the model to achieve better stability and generalization. Furthermore, the performance differences between window sizes show that an 8-second window is too short, causing statistical features to be more susceptible to noise; whereas a 15-second window is too long, potentially mixing multiple driving conditions or road surface types, thereby reducing discriminative ability. In contrast, a 10–12 second window achieves optimal performance by balancing noise reduction and maintaining signal homogeneity. Therefore, it can be concluded that using XGBoost with a 10-second window is the best configuration, ensuring high classification performance while maintaining stability and practicality for deployment in road surface recognition applications.

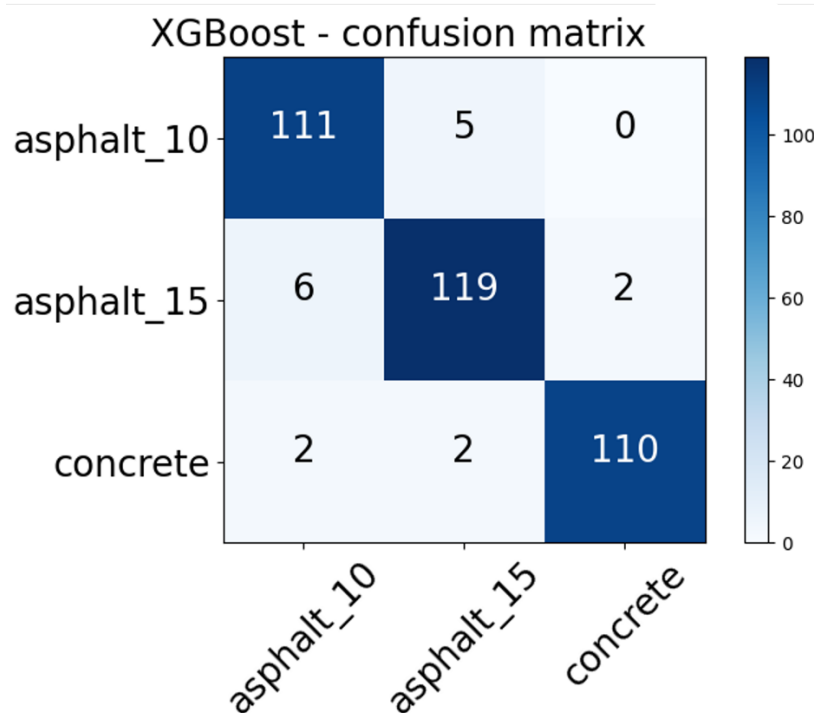


Fig. 3. Confusion matrix of XGBoost on the collected road surface dataset.

Figure 3 shows the confusion matrix of the XGBoost model. The results demonstrate the ability to classify the three road surface types with very high accuracy and clear error distribution. The asphalt_10 class is correctly identified in 111 out of 116 samples, with only 5 misclassified as asphalt_15 and none as concrete. The asphalt_15 class achieves 119 correct identifications out of 127 samples, with most errors being misclassifications as asphalt_10 (6 samples) and very few

as concrete (2 samples), reflecting the high similarity between the two asphalt types, which differ mainly in aging level. The concrete class is correctly identified in 110 out of 114 samples, with only 2 samples misclassified into each asphalt class, indicating that the signal characteristics of concrete are distinctly different from asphalt, enabling the model to achieve near-perfect separation. Overall, most errors occur between the asphalt_10–asphalt_15 pair (11 out of 17 errors), while the boundary between asphalt and concrete is well maintained, confirming that XGBoost achieves high and stable performance, especially for classes with clearly distinct physical characteristics.

Table 7 - Model performance

Road surface pattern	Algorithm performance			
	Accuracy	Recall	Precision	F1-score
asphalt_10	0.96	0.93	0.96	0.94
asphalt_15	0.96	0.94	0.94	0.94
concrete	0.98	0.98	0.96	0.97

The road surface classification accuracy is presented in Table 7. The results show high and balanced performance across all three road surface types, with Accuracy for asphalt_10 and asphalt_15 both reaching 0.96, while concrete is higher at 0.98. For asphalt_10, Precision (0.96) is higher than Recall (0.93), indicating that the model makes confident predictions but still misses some samples, mainly misclassifying them as asphalt_15. Conversely, asphalt_15 has both Recall and Precision at 0.94, reflecting a more symmetrical distribution of classification errors. Concrete is the easiest class to distinguish, with a Recall of 0.98 and Precision of 0.96, demonstrating that the model almost never misses and rarely misclassifies this class, thanks to its distinctly different vibration characteristics. The F1-scores for all three classes are above 0.94, with concrete achieving the highest (0.97), showing that the model maintains consistent effectiveness across classes and only struggles when distinguishing between the two asphalt types with similar characteristics.

From these results, XGBoost with a 10-second window is the optimal choice for the problem in this study. The 10-second window length is long enough to stabilize statistical features and eliminate noise, yet short enough to ensure the signal remains homogeneous in terms of driving conditions and road surface, thereby maintaining high performance while meeting latency requirements for real-world deployment.

Discussions

The experimental results confirm that the proposed configuration, combining XGBoost with a 10-second data window, delivers robust and well-balanced performance for three-class road surface classification. As shown in Tables 5 and 7, the XGBoost model with a 10-second window achieves an overall Accuracy, Recall, Precision, and F1-score of 0.95, while RF and GBM reach peak accuracies of 0.90 and 0.94, respectively, under the same conditions. At the class level (Table 8), the accuracies for asphalt_10, asphalt_15, and concrete are 0.96, 0.96, and 0.98, with F1-scores of 0.95, 0.96, and 0.97, indicating that most misclassifications occur between asphalt_10 and asphalt_15, whereas the separation between asphalt and concrete is almost perfect. This behaviour is consistent with the physical characteristics of the pavements, as both asphalt_10 and asphalt_15 are flexible pavements that differ mainly in aging and deterioration level, whereas concrete exhibits distinct vibration patterns due to its rigid structure.

Compared with threshold-based vibration methods that primarily target pothole or anomaly detection using single-axis accelerations, the proposed system addresses a more challenging multi-class classification problem and removes the need for manual threshold tuning for each deployment. For example, Wu et al. (2020) reported a RF classifier with a precision of 0.885 and a recall of 0.75 for pothole detection using smartphone sensors, with an overall accuracy of about 0.93, while Xin et al. (2023) introduced the G-ZERO algorithm and achieved approximately 0.85 accuracy for pothole detection based on zero-crossing patterns of the acceleration signal. Although these methods are simple and computationally inexpensive, their performance is highly sensitive to driving behaviour and vehicle dynamics and often requires careful calibration of

thresholds for each road segment or sensor configuration. In contrast, the tree-based models used in this study learn decision boundaries directly from multi-axis accelerometer and gyroscope features and generalize well across different segments of the urban and suburban routes considered.

In relation to traditional machine learning approaches for road surface assessment, such as SVM-based frequency-domain methods using suspension accelerations and speed, Wang et al. (2018) reported an average classification accuracy of 0.7 when combining a quarter-car model with FFT features for multiple surface types, while smartphone-based CWT and clustering schemes by Basavaraju et al. (2019) achieved an accuracy of 0.944 in detecting road surface anomalies. The proposed method reaches a comparable or slightly higher accuracy of 0.95 while relying solely on low-cost inertial sensors mounted inside the vehicle and a lightweight embedded platform, and at the same time provides detailed multi-class discrimination between two asphalt aging levels and concrete. Previous studies often assume processing resources equivalent to smartphones or single-board computers and rarely report the computational budget required for real-time inference. By contrast, this work constrains the number of trees and depth of the RF, GBM, and XGBoost models and demonstrates that a compact XGBoost configuration can achieve 0.95 accuracy using only time-domain statistical features, with inference executed on an ESP32-class microcontroller.

Deep learning methods based on CNNs, RNNs, or hybrid CNN–LSTM architectures have reported strong performance for pavement defect detection and surface-type classification. Varona et al. (2020) obtained up to about 0.93 accuracy when using CNNs to classify potholes, speed bumps, and manholes, while Menegazzo & von Wangenheim (2021) reported a CNN model with an average training accuracy of 0.93 and F1-scores up to 0.98 for asphalt segments in a three-class (asphalt, cobblestone, dirt) problem. More recently, Hnoohom et al. (2023) conducted a comprehensive evaluation of state-of-the-art deep learning models for road surface type classification and reported best-case performance around 0.984 accuracy and 0.982 F1-score for multi-class scenarios using CNN-based architectures. However, these approaches typically require large, labeled datasets, substantial memory, and GPU-class processing power, and they often treat feature extraction as a black box, which limits interpretability and complicates deployment on low-cost devices. The results of this study show that, for the considered three-class problem, a carefully designed feature set combined with a tree-based ensemble can reach accuracy levels (0.95) that are competitive with, although slightly lower than, the best deep models reported in the literature, while remaining interpretable and computationally affordable for an ESP32-class microcontroller.

A key contribution of this work lies in the use of the Hybrid Filter–Wrapper feature selection algorithm, which follows the idea of combining a fast, filter-based ranking with a more accurate wrapper stage, similar in spirit to the hybrid strategies proposed by Solorio-Fernández et al. (2016), Zhang et al. (2019), and Yin et al. (2023). In the present study, the original feature set is first reduced by discarding low-importance features and then further refined via Recursive Feature Elimination until the validation accuracy drops by more than a predefined threshold. This process reduces the final feature set to 27 attributes without degrading and, in some configurations, slightly improving the classification performance, while also lowering the computational cost of training and inference and decreasing memory usage, which is crucial for embedded deployment. This behaviour is consistent with previous findings in other domains, where hybrid feature selection improves the trade-off between accuracy and model complexity (Li et al., 2018; Mambwe Sydney & Sun, 2020; Megantara & Ahmad, 2020; Jeon & Oh, 2020). From a system perspective, the results also confirm that the selected 10-second window length is a good compromise between robustness and latency for real-time monitoring. Shorter windows, such as 8 seconds, do not accumulate enough samples to stabilise statistical features and are more sensitive to transient manoeuvres, whereas longer windows (12–15 seconds) tend to mix different driving conditions or surface transitions, which can blur the distinction between classes. With a 10-second window, the end-to-end processing time on the embedded device remains well below the window length, meaning that the classification can be performed in near real time as new windows are completed, which aligns with requirements for continuous pavement monitoring on moving vehicles.

In addition to the aggregate results, cross-session experiments—where models trained on data from one acquisition session are evaluated on data from different sessions showed comparable performance, suggesting that the learned decision boundaries are stable with respect to moderate variations in traffic conditions and environmental factors. This robustness is important for scalable deployment, where the system is expected to operate across multiple days and routes without frequent re-training. Nevertheless, the experiments in this study were conducted on a single vehicle platform and a specific set of urban and suburban roads in Hanoi, and the speed range was limited to typical operating speeds for these environments, which may constrain generalizability to other traffic contexts. Overall, the proposed system bridges a gap in the literature between high-accuracy but resource-intensive deep learning approaches and simple threshold-based methods that are difficult to calibrate and maintain. By combining carefully engineered inertial features, a lightweight XGBoost classifier, and a hybrid feature selection strategy, the study demonstrates that accurate three-class road surface recognition (Accuracy \approx 0.95) can be achieved on low-cost embedded hardware, while offering a favourable trade-off between accuracy, interpretability, and resource consumption compared with existing alternatives.

5. Conclusion

This study proposed and implemented an efficient road surface classification system designed for real-time operation on low-cost embedded hardware using vibration data from a single accelerometer and a single gyroscope mounted on a survey vehicle. Within an IoT-based monitoring architecture, the system combines time-domain statistical features with three lightweight tree-based models—RF, GBM, and XGBoost—and a Hybrid Filter–Wrapper feature selection strategy. Experimental results on three real-world surface types (asphalt under 10 years, asphalt over 15 years, and concrete) show that the XGBoost model with a 10-second window achieves an overall Accuracy of approximately 0.95 and F1-scores above 0.94 for all classes, with concrete being almost perfectly distinguished and only limited confusion between the two asphalt categories. At the same time, the proposed feature selection process reduces the original feature set to 27 attributes without degrading performance, thereby lowering computational cost and enabling near real-time inference on an ESP32-class microcontroller. These findings confirm the feasibility of deploying multi-class road surface recognition on low-cost devices and highlight the importance of carefully designed window configuration and feature selection for resource-constrained embedded systems.

The results of this study offer several practical applications for road authorities and IoT developers. The system provides a low-cost tool to complement traditional inspection methods, enabling higher-frequency and denser pavement monitoring using probe vehicles. At the same time, it allows the sampling rate, window length, feature set, and models to be adjusted for different types of resource-constrained hardware devices.

However, the current evaluation is limited to a single vehicle type, one sensor mounting position, three surface types, and a set of urban/suburban routes in Hanoi, which may restrict the generalizability to other vehicles and environments. In the future, the research will be extended to a wider range of vehicle platforms, pavement types, and operating conditions, investigate multi-sensor fusion, and explore lightweight deep learning or TinyML models, as well as large-scale deployments on vehicle fleets operating in real traffic to assess long-term stability and scalability.

Acknowledgement

Thanks to Institute of Information Technology (VAST) for allowing us to use the “IoT and Robot intensive laboratory” equipment and supporting this research, code NVKN02.01/26-27.

References

- Alhasan, A., Nlenanya, I., Smadi, O., & MacKenzie, C. A. (2018). Impact of pavement surface condition on roadway departure crash risk in Iowa. *Infrastructures*, 3(2), Article 14. <https://doi.org/10.3390/infrastructures3020014>

- Alqaydi, S., Zeiada, W., El Wakil, A., Alnaqbi, A. J., & Azam, A. M. M. (2024). A comprehensive review of smartphone and other device-based techniques for road surface monitoring. *Eng*, 5(4), 3397–3426. <https://doi.org/10.3390/eng5040177>
- Bentéjac, C., Csörgő, A., & Martínez-Muñoz, G. (2021). A comparative analysis of gradient boosting algorithms. *Artificial Intelligence Review*, 54, 1937–1967. <https://doi.org/10.1007/s10462-020-09896-5>
- Bruno, S., De Almeida Filho, F. G. V., De Oliveira, R. H., & Loprencipe, G. (2023). Proposal for a low-cost monitoring system to assess the pavement deterioration in urban roads. *European Transport Research Review*, 15, Article 57. <https://doi.org/10.1186/s12544-023-00679-z>
- Cafiso, S., Di Graziano, A., Marchetta, V., & Pappalardo, G. (2022). Urban road pavements monitoring and assessment using bike and e-scooter as probe vehicles. *Case Studies in Construction Materials*, 16, e00889. <https://doi.org/10.1016/j.cscm.2022.e00889>
- Celaya-Padilla, J. M., Galván-Tejada, C. E., López-Montegudo, F. E., Alonso-González, O., Moreno-Báez, A., Martínez-Torteya, A., Galván-Tejada, J. I., Arceo-Olague, J. G., Luna-García, H., & Gamboa-Rosales, H. (2018). Speed bump detection using accelerometric features: A genetic algorithm approach. *Sensors*, 18(2), Article 443. <https://doi.org/10.3390/s18020443>
- Chen, T., & Guestrin, C. (2016). XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 785–794). Association for Computing Machinery. <https://doi.org/10.1145/2939672.2939785>
- Collier, P., Kirchberger, M., & Söderbom, M. (2016). The cost of road infrastructure in low- and middle-income countries. *The World Bank Economic Review*, 30(3), 522–548. <https://doi.org/10.1093/wber/lhv037>
- Dehnad, M. H., & Yazdi, A. (2024). A review of numerical and experimental studies on hydroplaning of vehicles in motion on road surfaces. *Results in Engineering*, 23, Article 102438. <https://doi.org/10.1016/j.rineng.2024.102438>
- del Río-Barral, P., Soilán, M., González-Collazo, S. M., & Arias, P. (2022). Pavement crack detection and clustering via region-growing algorithm from 3D MLS point clouds. *Remote Sensing*, 14(22), Article 5866. <https://doi.org/10.3390/rs14225866>
- Du, R., Qiu, G., Gao, K., Hu, L., & Liu, L. (2020). Abnormal road surface recognition based on smartphone acceleration sensor. *Sensors*, 20(2), Article 451. <https://doi.org/10.3390/s20020451>
- Guan, H., Li, J., Yu, Y., Chapman, M., Wang, H., Wang, C., & Zhai, R. (2015). Iterative tensor voting for pavement crack extraction using mobile laser scanning data. *IEEE Transactions on Geoscience and Remote Sensing*, 53(3), 1527–1537. <https://doi.org/10.1109/TGRS.2014.2344714>
- Hnoohom, N., Mekruksavanich, S., & Jitpattanakul, A. (2023). A comprehensive evaluation of state-of-the-art deep learning models for road surface type classification. *Intelligent Automation & Soft Computing*, 37(2), 1275–1291. <https://doi.org/10.32604/iasc.2023.038584>
- Hu, J., & Szymczak, S. (2023). A review on longitudinal data analysis with random forest. *Briefings in Bioinformatics*, 24(2), bbad002. <https://doi.org/10.1093/bib/bbad002>
- Huynh, V. N., Truong, L. T., & De Gruyter, C. (2025). Examining the impacts of road pavement roughness and rutting on traffic safety: A macrolevel analysis. *Traffic Injury Prevention*, 26(6), 720–726. <https://doi.org/10.1080/15389588.2024.2448838>
- Jeon, H., & Oh, S. (2020). Hybrid-recursive feature elimination for efficient feature selection. *Applied Sciences*, 10(9), 3211. <https://doi.org/10.3390/app10093211>
- Khahro, S. H., Javed, Y., & Memon, Z. A. (2021). Low cost road health monitoring system: A case of flexible pavements. *Sustainability*, 13(18), Article 10272. <https://doi.org/10.3390/su131810272>
- Lekshmiopathy, J., Velayudhan, S., & Mathew, S. (2021). Effect of combining algorithms in smartphone-based pothole detection. *International Journal of Pavement Research and Technology*, 14(1), 63–72. <https://doi.org/10.1007/s42947-020-0033-0>

- Li, L., Yu, Y., Bai, S., Cheng, J., & Chen, X. (2018). Towards effective network intrusion detection: A hybrid model integrating Gini index and GBDT with PSO. *Journal of Sensors*, 2018, Article 1578314. <https://doi.org/10.1155/2018/1578314>
- Li, X., Huo, D., Goldberg, D. W., Chu, T., Yin, Z., & Hammond, T. (2019). Embracing crowdsensing: An enhanced mobile sensing solution for road anomaly detection. *ISPRS International Journal of Geo-Information*, 8(9), Article 412. <https://doi.org/10.3390/ijgi8090412>
- Loprencipe, G., & Pantuso, A. (2017). A specified procedure for distress identification and assessment for urban road surfaces based on PCI. *Coatings*, 7(5), Article 65. <https://doi.org/10.3390/coatings7050065>
- Loprencipe, G., De Almeida Filho, F. G. V., De Oliveira, R. H., & Bruno, S. (2021). Validation of a low-cost pavement monitoring inertial-based system for urban road networks. *Sensors*, 21(9), Article 3127. <https://doi.org/10.3390/s21093127>
- Maeda, H., Sekimoto, Y., Seto, T., Kashiyama, T., & Omata, H. (2018). Road damage detection and classification using deep neural networks with smartphone images. *Computer-Aided Civil and Infrastructure Engineering*, 33(12), 1127–1141. <https://doi.org/10.1111/mice.12387>
- Mambwe Sydney, K., & Sun, Y. (2020). Performance analysis of intrusion detection systems using a feature selection method on the UNSW-NB15 dataset. *Journal of Big Data*, 7(1), 91. <https://doi.org/10.1186/s40537-020-00379-6>
- Martinelli, A., Meocci, M., Dolfi, M., Branzi, V., Morosi, S., Argenti, F., Berzi, L., & Consumi, T. (2022). Road surface anomaly assessment using low-cost accelerometers: A machine learning approach. *Sensors*, 22(10), Article 3788. <https://doi.org/10.3390/s22103788>
- Martínez-Ríos, E. A., Bustamante-Bello, R., Navarro-Tuch, S., & Perez-Meana, H. (2023). Applications of the generalized Morse wavelets: A review. *IEEE Access*, 11, 667–688. <https://doi.org/10.1109/ACCESS.2022.3232729>
- Megantara, A. A., & Ahmad, T. (2020). Feature importance ranking for increasing performance of intrusion detection system. In *2020 3rd International Conference on Computer and Informatics Engineering (IC2IE)* (pp. 37–42). IEE. <https://doi.org/10.1109/IC2IE50715.2020.9274570>
- Menegazzo, J., & von Wangenheim, A. (2020). Multi-contextual and multi-aspect analysis for road surface type classification through inertial sensors and deep learning. In *Proceedings of the 2020 X Brazilian Symposium on Computing Systems Engineering (SBESC)* (pp. 1–8). IEEE. <https://doi.org/10.1109/SBESC51047.2020.9277846>
- Menegazzo, J., & von Wangenheim, A. (2021). Road surface type classification based on inertial sensors and machine learning. *Computing*, 103(9), 2143–2170. <https://doi.org/10.1007/s00607-021-00914-0>
- Meocci, M. (2024). A vibration-based methodology to monitor road surface: A process to overcome the speed effect. *Sensors*, 24(3), Article 925. <https://doi.org/10.3390/s24030925>
- Mihoub, A., Krichen, M., Alswailim, M., Mahfoudhi, S., & Bel Hadj Salah, R. (2023). Road scanner: A road state scanning approach based on machine learning techniques. *Applied Sciences*, 13(2), 683. <https://doi.org/10.3390/app13020683>
- Nyirandayisabye, R., Li, H., Dong, Q., Hakuzweyezu, T., & Nkinahamira, F. (2022). Automatic pavement damage predictions using various machine learning algorithms: Evaluation and comparison. *Results in Engineering*, 16, Article 100657. <https://doi.org/10.1016/j.rineng.2022.100657>
- Rathee, M., Bačić, B., & Doborjeh, M. (2023). Automated road defect and anomaly detection for traffic safety: A systematic review. *Sensors*, 23(12), Article 5656. <https://doi.org/10.3390/s23125656>
- Sattar, S., Li, S., & Chapman, M. (2018). Road surface monitoring using smartphone sensors: A review. *Sensors*, 18(11), Article 3845. <https://doi.org/10.3390/s18113845>
- Sattar, S., Li, S., & Chapman, M. (2021). Developing a near real-time road surface anomaly detection approach for road surface monitoring. *Measurement*, 185, Article 109990. <https://doi.org/10.1016/j.measurement.2021.109990>

- Sebestyen, G., Muresan, D., & Hangan, A. (2015). Road quality evaluation with mobile devices. In *Proceedings of the 2015 16th International Carpathian Control Conference (ICCC)* (pp. 458–464). IEEE. <https://doi.org/10.1109/CarpathianCC.2015.7145123>
- Shaghlil, N., & Khalafallah, A. (2018). Automating highway infrastructure maintenance using unmanned aerial vehicles. In *Proceedings of the Construction Research Congress 2018* (pp. 486–495). American Society of Civil Engineers. <https://doi.org/10.1061/9780784481295.049>
- Solorio-Fernández, S., Carrasco-Ochoa, J., & Martínez-Trinidad, J. F. (2016). A new hybrid filter-wrapper feature selection method for clustering based on ranking. *Neurocomputing*, 214, 866–880. <https://doi.org/10.1016/j.neucom.2016.07.026>
- Sun, T., Pan, W., Wang, Y., & Liu, Y. (2022). Region of interest constrained negative obstacle detection and tracking with a stereo camera. *IEEE Sensors Journal*, 22(4), 3616–3625. <https://doi.org/10.1109/JSEN.2022.3142024>
- Surblys, V., Kozłowski, E., Matijošius, J., Gołda, P., Laskowska, A., & Kilikevičius, A. (2024). Accelerometer-based pavement classification for vehicle dynamics analysis using neural networks. *Applied Sciences*, 14(21), Article 10027. <https://doi.org/10.3390/app142110027>
- Varona, B., Monteserín, A., & Teyseyre, A. (2020). A deep learning approach to automatic road surface monitoring and pothole detection. *Personal and Ubiquitous Computing*, 24(4), 519–534. <https://doi.org/10.1007/s00779-019-01234-z>
- Vi, M., Tran, D., Thuong, V. T., Linh, N. N., & Tran, D. (2024). Efficient real-time devices based on accelerometer using machine learning for HAR on low-performance microcontrollers. *Computers, Materials & Continua*, 81(1), 1729–1756. <https://doi.org/10.32604/cmc.2024.055511>
- Wang, S., Kodagoda, S., Shi, L., & Dai, X. (2018). Two-stage road terrain identification approach for land vehicles using feature-based and Markov random field algorithm. *IEEE Intelligent Systems*, 33(1), 29–39. <https://doi.org/10.1109/MIS.2017.2581327>
- Wang, S., Kodagoda, S., Shi, L., & Wang, H. (2017). Road-terrain classification for land vehicles: Employing an acceleration-based approach. *IEEE Vehicular Technology Magazine*, 12(3), 34–41. <https://doi.org/10.1109/MVT.2017.2656949>
- Wu, C., Wang, Z., Hu, S., Lepine, J., Na, X., Ainalis, D., & Stettler, M. (2020). An automated machine-learning approach for road pothole detection using smartphone sensor data. *Sensors*, 20(19), Article 5564. <https://doi.org/10.3390/s20195564>
- Xin, H., Ye, Y., Na, X., Hu, H., Wang, G., Wu, C., & Hu, S. (2023). Sustainable road pothole detection: A crowdsourcing-based multi-sensors fusion approach. *Sustainability*, 15(8), Article 6610. <https://doi.org/10.3390/su15086610>
- Yin, Y., Jang-Jaccard, J., Xu, W., Singh, A., Zhu, J., Sabrina, F., & Kwak, J. (2023). Igrf-rfe: A hybrid feature selection method for MLP-based network intrusion detection on UNSW-NB15 dataset. *Journal of Big Data*, 10(1), Article 27. <https://doi.org/10.1186/s40537-023-00694-8>
- Zhang, J., Xiong, Y., & Min, S. (2019). A new hybrid filter/wrapper algorithm for feature selection in classification. *Analytica Chimica Acta*, 1080, 43–54. <https://doi.org/10.1016/j.aca.2019.06.054>