DOI 10.15622/ia.25.2.7

V.-H. Bui, D.-N. Tran, T.-H. Dao, H.Q. Trung, P.V. Kien,
N.V. Thang, D.-T. Tran

IDRRS: IOT INERTIAL DEVICE FOR REAL-TIME ROAD SURFACE CLASSIFICATION AND POSITION ESTIMATION ENHANCEMENT

Viet-Hoan Bui, Duc-Nghia Tran, To-Hieu Dao, Hoang Quang Trung, Pham Vu Kien, Nguyen Van Thang, and Duc-Tan Tran. **IDRRS: IoT Inertial Device for Real-Time Road Surface Classification and Position Estimation Enhancement.**

Abstract. Road surface condition monitoring is essential for enhancing transportation safety and infrastructure maintenance. This study develops an IoT-oriented inertial sensing framework for real-time pavement classification, pothole detection, and enhanced vehicle position estimation. The framework integrates a memory-constrained XGBoost model designed for microcontroller deployment, a velocity-aided GPS interpolation procedure, and an abnormality-index-based pothole detection algorithm. Experimental results on a private dataset and the PVS dataset show classification accuracies of 95.39% and 93.21%, respectively. To examine transferability, the configuration tuned on the private dataset was applied to the PVS dataset without retraining and achieved 92.45% accuracy. Furthermore, the GPS interpolation procedure reduces mean localization errors from 5.571-11.893 m to 1.835–3.563 m across vehicle speeds of 20–50 km/h. An additional contribution of this study is the release of a private dataset capturing vibration signatures from representative road types, supporting further research in road surface classification.

Keywords: Inertial sensor, real-time, microcontroller, machine learning, IoT.

1. Introduction. In recent years, the development of intelligent transportation systems (ITS) [1] and their sub-applications, such as advanced driver assistance systems [2], has significantly increased the demand for diverse and reliable situational data from the traffic environment [3-6]. The core objective of these systems is to enhance road safety, improve traffic management efficiency, optimize the operational performance of vehicles, and ensure passenger comfort [5, 7-9]. To achieve these goals, autonomous and semi-autonomous vehicles require accurate and real-time environmental perception, wherein understanding the characteristics of the road surface is a foundational requirement [10].

Among situational data types, road surface classification (RSC) is considered one of the most critical pieces of information [5, 6]. The ability to accurately distinguish between road surface types – such as asphalt, cobblestone, and dirt – has a direct impact on vehicle dynamics, including grip, braking distance, and vibration patterns [10]. Real-time RSC information is a vital input for adaptive control systems, for example, anti-lock braking systems [11] and traction control systems [12], allowing them to adjust operational parameters to ensure maximum safety [5, 8].

Many methods have been proposed to address the RSC problem [10, 13-15]. Vision-based methods (using cameras), while common, suffer significant

performance degradation due to environmental factors such as changing light conditions, shadows, or adverse weather (rain, snow) [16-19]. 3D scanning methods, such as those using LiDAR, provide high-precision data but require expensive hardware and consume significant computational resources, making them unsuitable for large-scale deployment [5, 7, 8, 20].

Consequently, methods based on inertial sensors (inertial measurement units - IMU), including accelerometers and gyroscopes, have emerged as an effective alternative [5, 13, 20]. These sensors are low-cost, highly durable, and, most importantly, they operate independently of light and weather conditions [5, 8, 13, 14, 20]. Furthermore, they offer a fundamental advantage: instead of passively observing the surface, they directly measure the kinetic response and vibration patterns of the vehicle itself as it interacts with the road [5, 8, 14, 20].

Despite these advantages, achieving robust and scalable RSC from inertial data remains challenging. Vibro-kinetic responses vary significantly across vehicles, suspension systems, driving behaviors, road environments, and sensor mounting positions [4, 7, 13, 14]. These contextual variations introduce strong dependency factors that can limit the generalization capability of machine learning (ML) models. While previous studies have explored various architectures, two primary gaps remain: (i) the lack of evaluation under standardized, context-aware protocol frameworks; (ii) the high computational demand of modern deep learning (DL) models, which limits their deployment on resource-constrained microcontrollers (MCU) [9, 15, 21, 22].

In response to these challenges, this work develops a lightweight and deployable classical machine learning (CML) pipeline. Unlike many previous studies that prioritize accuracy through complex DL models, this study focuses on practical hardware efficiency without compromising generalization. The contributions are fourfold:

– This study introduces the IoT Inertial Device for Real-Time Road Surface Classification and Position Estimation Enhancement (IDRRS). An Extreme Gradient Boosting (XGBoost) classification model is deployed directly onto the device, integrated with a GPS interpolation algorithm. This approach provides road surface condition data and corresponding positioning, achieving classification accuracies of 95.39% on the private dataset and 93.21% on the PVS dataset, respectively.

– Integration of a GPS interpolation algorithm that reduces positioning deviation by 4-5 times compared to conventional GPS.

– A pothole-detection method using an abnormality index computed from triaxial accelerometer signals to localize surface disturbances on reconstructed trajectories.

– A private dataset was constructed based on inertial sensors and positioning data to support feature engineering analysis and the development of predictive models.

2. Related works. The classification of road surface types using inertial sensors and ML is a critical component of ITS, enhancing safety and efficiency. This approach is crucial for enhancing road safety and maintenance by providing real-time data on road conditions [5, 7-10]. Various ML models, including DL and traditional algorithms, have been employed to classify road surfaces based on data from IMU and GPS sensors [10, 13-15]. While DL models offer automated feature discovery, CML methods remain competitive due to their lower computational footprint on embedded systems.

The author [4] established a foundation for modern IMU-based RSC research by introducing the PVS dataset. Comparison between CML and DL methods showed that a convolutional neural network (CNN) architecture achieved 93.17% accuracy, outperforming CML models such as k-nearest neighbors (74.79%). This result was achieved through a 3-experiment split evaluation protocol designed to test model generalization across unseen vehicle, driver, and environmental contexts.

Subsequent research has increasingly focused on improving classification accuracy through more sophisticated DL architectures. A common characteristic of many of these studies is the use of simplified evaluation procedures, such as random 80/20 splits or k-fold cross-validation (CV), rather than the context-aware three-scenario evaluation protocol defined in [4]. These simplified splits allow data from the same vehicle, driver, or scenario to appear in both training and test sets, increasing the risk of context leakage [4] and thus preventing standardized performance comparisons [9].

One early attempt to enhance performance was proposed by the authors in [5], who introduced the SE-ResNet architecture, a residual CNN augmented with Squeeze-and-Excitation (SE) modules. The reported accuracy reached 98.41% using 5-fold CV on the entire dataset. As a result, the comparison with the baseline should be interpreted with caution.

Recurrent neural network (RNN) variants have also been explored to address temporal dependencies in inertial signals. In paper [6] proposed ResBiGRU-SE, a hybrid architecture combining residual connections, bidirectional gated recurrent units (BiGRU), and SE modules. The reported accuracy reached 98.41% through 5-fold CV, making it difficult to assess stability across different vehicles, drivers, and environments.

In paper [8] proposed a hybrid architecture combining CNN and long short-term memory (LSTM) branches to leverage multiple data representations. The model utilizes frequency-domain data for the CNN component and stacked

discrete wavelet transform coefficients for the LSTM branch. On the PVS dataset, this hybrid approach reported an accuracy of 94.78%. However, the evaluation was restricted to a specific 66/34 split (only Experiment 3), providing limited evidence of the model's robustness against independent contextual shifts.

In paper [23] applied Transformer-based architecture to the task, using multi-head attention to model long-range temporal dependencies. The method incorporated a Random Forest (RF)-based feature selection step to reduce input dimensionality. Using an 80/20 split on PVS 1–8, the model achieved a weighted $F1$-score of 97%. When evaluated on an unseen context (PVS 9), performance decreased to 80.41%, illustrating the sensitivity of such models to contextual variation.

In study [24] extended this line of work by proposing a VGG-inspired 1D-CNN architecture achieving 99.3% accuracy with 101.6k parameters. This result was obtained using a simple 80/20 split. Therefore, while the architecture is compact and efficient, its generalization capability under unseen operational conditions remains unverified.

A parallel line of research has focused on improving practicality and computational efficiency. The performance sensitivity observed in several DL models, together with the substantial drop reported by the authors in [23] when evaluated on an unseen context, further underscores the need for approaches that maintain accuracy while operating reliably across diverse conditions and hardware constraints.

Paper [9] are among the few studies that apply the three-scenario evaluation protocol of the author in [4]. Their work introduced two lightweight CNN architectures, SepRNet-1D and SepSERNet-1D, designed with principles inspired by MobileNet V3 through the use of separable convolutions and residual connections. SepRNet-1D achieved an average accuracy of 94.69% with a reported inference time of under 4 ms using a TensorFlow Lite implementation on a CPU-based desktop system (13 GB RAM).

In paper [20] the author revisited tree-based ML models to re-evaluate their potential for RSC. The authors examined RF, Gradient Boosting (GB), and XGBoost, and proposed a Hybrid Filter-Wrapper (HFW) feature selection strategy that combines the efficiency of filter-based ranking with the effectiveness of wrapper-based selection. Their highest reported accuracy was 94.2%, obtained using an XGBoost model with a substantially reduced feature subset, shrinking the original 126 features to 30 through HFW-GB. However, this result was achieved using a 70/30 split, providing limited evidence of the model's robustness against independent contextual shifts.

In summary, while recent advancements in DL and Transformer-based models have increased classification accuracy, these approaches often involve high computational overhead and limited robustness in unseen environments. Furthermore, existing literature frequently overlooks the interplay between classification performance and spatial localization accuracy. The proposed IDRRS device addresses these limitations by integrating a hardware-optimized XGBoost pipeline with a velocity-aided GPS interpolation strategy. This integration achieves competitive classification performance on resource-constrained hardware while improving positioning accuracy in urban settings.

3. Materials and methods.

3.1. Proposed system. The proposed IoT-oriented framework, illustrated in Figure 1, is a compact device designed for real-time RSC and pothole detection on degraded pavements. The hardware architecture, detailed in Table 1, is centered around an ESP32 MCU that serves as the primary processing unit. This MCU executes embedded CML models and anomaly detection algorithms on vibration data acquired from a GY-85 9-axis IMU. Localization is supported by an ATGM336H GPS module, while a SIMCom A7680C modem facilitates data transmission over the LTE network. The sensor node is housed in a 7.3 × 4.5 × 3.4 cm^3 protective casing and is primarily powered by the vehicle's battery system, with an integrated 3.7 V–2000 mAh Li-Po battery serving as a backup to ensure uninterrupted operation during power fluctuations.

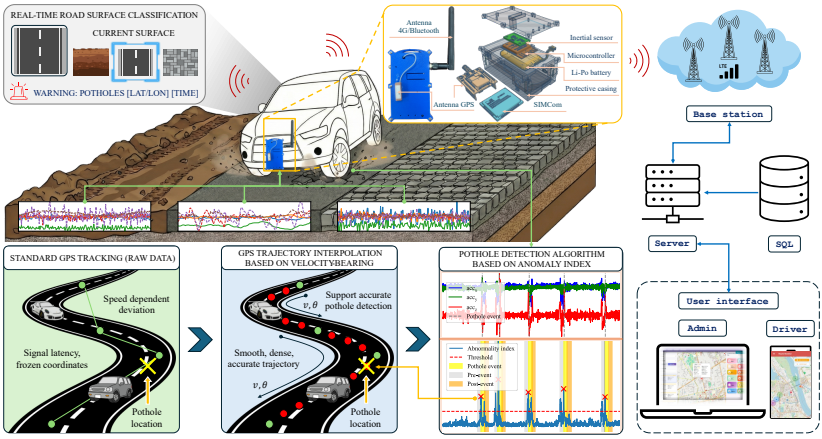

Fig. 1. Architecture of the proposed IoT-based road surface monitoring and alert system

The deployment involves mounting the device on a fleet of vehicles dedicated to road inspection tasks. To capture stable vibration patterns, the sensor node is rigidly mounted at the geometric center of each vehicle, typically near the differential system, with a coordinate convention where the x-axis is oriented leftward, the y-axis forward, and the z-axis downward. To mitigate inertial sensor bias and account for vehicle-specific idle vibrations caused by engine and suspension differences, a calibration procedure is executed whenever the vehicle remains stationary for one second. This process calculates the zero-point offset for the accelerometer and the static bias for the gyroscope. Additionally, the system utilizes the internal thermometer of the ITG3200 for thermal compensation of the inertial readings. The protective casing is also equipped with waterproof gaskets and internal electronic potting to protect the circuitry from humidity and temperature fluctuations.

Table 1. Hardware configuration and experimental platform

Module	**Component**	**Specification**	**Unit**	f_s **(Hz)**
Sensor node	ESP32	520 kB RAM, 448 kB ROM, 4 MB Flash	-	-
	GY-85 9-axis IMU	Accelerometer ADXL345	g (1 g = 9.8 m/s^2)	100
		Gyroscope ITG3200	rad/s	
		Magnetometer HMC5883L	μT	
	GPS ATGM336H	Speed	km/h	1
		Latitude, longitude	Decimal degrees	
	SIMCom A7680C	4G LTE Cat 1	-	-
	Li-Po battery	3.7 V–2000 mAh	-	-
Server HTTP	Compute node	2 × Intel Xeon Gold 6226R (32 cores)	-	-
	Memory	384 GB DDR4 RAM	-	-
	Operating system & Access	AlmaLinux 8, Remote SSH (Visual Studio Code)	-	-

As depicted in the architecture diagram (Figure 1), initial processing occurs locally on the ESP32. Raw sensor signals are segmented and converted into road surface categories while simultaneously being monitored for sudden disturbances using an abnormality-index-based algorithm. Because standard GPS modules often suffer from update delays or signal loss in challenging environments, a position interpolation procedure is integrated to provide continuous and enhanced location estimation for detected anomalies. The processed results are then transmitted via the cellular network to an SQL database on a central server. This centralized repository supports a two-part user interface: an administrator dashboard for traffic authorities to visualize pavement quality and prioritize repairs, and a driver interface that delivers real-time hazard warnings. This integrated flow facilitates a proactive approach to road maintenance and enhances operational safety through timely anomaly reporting.

3.2. Dataset.

3.2.1. Private dataset. The private dataset[1] was collected using a fleet of vehicles dedicated to regional road inspection tasks. Asphalt, cobblestone, and dirt were selected as the target surface categories because they constitute the predominant pavement types in regional transportation networks and represent the primary focus of contemporary road condition monitoring studies [4,20,25].

Data collection followed a controlled experimental protocol where each session was dedicated to a single road surface type under the direct supervision of two researchers. These supervisors observed the conditions and ensured the recorded data corresponded precisely to the targeted pavement category, facilitating ground-truth annotation at the source without the requirement for post-processing software tools.

The acquired signals were stored on a card in plain-text format, with a sampling frequency of 100 Hz for the inertial and magnetometer sensors and 1 Hz for the GPS module. Each data file contained triaxial accelerometer, gyroscope, and magnetometer measurements, alongside GPS latitude, longitude, timestamps, and speed. Figure 2 provides the GPS trajectories and representative photographs of the sampled road surfaces to offer visual context for the experimental environments. In total, the recordings amounted to 903.90 minutes (15.06 hours), corresponding to approximately 196 MB of raw measurements.

Fig. 2. GPS trajectories for dirt (red), cobblestone (blue), and asphalt (green) segments (top) and representative images of the three surface types (bottom)

3.2.2. Public dataset. The public dataset used in this study is the PVS[2] benchmark [4], which complements the private dataset by providing a

[1] https://dx.doi.org/10.21227/x6av-va13

[2] https://github.com/jefmenegazzo/Intelligent-Vehicle-Perception-Based-on-Inertial-Sensing-and-Artificial-Intelligence

structured basis for evaluating reproducibility and generalization under diverse contextual factors. The dataset comprises nine sub-datasets (PVS 1–PVS 9) collected from three vehicles driven by three independent drivers across three road scenarios, with all samples labeled into dirt, cobblestone, and asphalt classes. Each vehicle carries MPU-9250 modules installed at three locations – dashboard (DB), above suspension (AS), and below suspension (BS) – on both left and right sides of the front axle. To ensure valid generalization, this study adheres to the 3-experiment split evaluation protocol proposed by the author in [4], a method designed to test robustness against changes in vehicle type, driver behavior, and environmental context [9].

3.3. Windowing. Continuous sensor signals are segmented into fixed-length windows using a sliding window of 600 samples (6 seconds at 100 Hz) with 50% overlap. This configuration is selected to balance classification performance with real-time operational constraints. Regarding window duration, research on inertial sensor data indicates that recognition accuracy tends to saturate at sizes around 6 seconds, as this interval provides sufficient temporal coverage to capture essential cyclic characteristics of the movement [26]. Specifically, in the context of RSC for MCU deployment, a 6-second window has been shown to effectively balance high accuracy with the limited processing capabilities of low-end hardware [25]. Furthermore, the 50% overlap results in a 300-sample stride, ensuring that new predictions are generated every 3 seconds. This setup complies with established real-time latency requirements, where a prediction delay of no more than 3 seconds is considered acceptable for timely monitoring [4, 25]. The general characteristics of the datasets and the distribution of the classification instances are summarized in Table 2.

Table 2. Summary of dataset characteristics and instances distribution

Dataset	Class	Instances (number of window segments)	Total instances	Split protocol
Private	Dirt road Cobblestone road Asphalt road	3258 2157 3624	9039	10×10-fold NCV
Public PVS	Dirt road Cobblestone road Asphalt road	1996 2096 3088	7180	3-scenario split [4, 9]

For the private dataset, the segmentation process yielded 9039 windows, which are utilized in a 10×10-fold nested cross-validation (NCV) procedure. Regarding the PVS dataset, a total of 7180 unique windows are generated (Table 3). Performance evaluation on the PVS dataset follows a protocol involving three independent experiments (Exp. 1-3) to ensure generalization across different vehicles, drivers, and scenarios [4,9]. Within this three-scenario

framework, the aggregate evaluation comprises 21540 processed instances, consisting of 14146 training windows and 7394 testing windows (Table 4).

Table 3. Distribution of segmented windows per PVS dataset

Dataset	Vehicle	Driver	Scenario	Number of windows		
				Dirt	Cobblestone	Asphalt
PVS 1	Volkswagen Saveiro	1	1	172	412	374
PVS 2	Volkswagen Saveiro	1	2	298	138	392
PVS 3	Volkswagen Saveiro	1	3	190	176	336
PVS 4	Fiat Bravo	2	1	160	384	336
PVS 5	Fiat Bravo	2	2	404	120	366
PVS 6	Fiat Bravo	2	3	160	210	268
PVS 7	Fiat Palio	3	1	158	362	334
PVS 8	Fiat Palio	3	2	300	126	396
PVS 9	Fiat Palio	3	3	154	168	286
Total				1996	2096	3088

Table 4. Train/test split distribution of segmented windows for the three evaluation experiments on the PVS dataset

	Dirt road		Cobblestone road		Asphalt road	
	Train	Test	Train	Test	Train	Test
Exp. 1	994	1002	1712	384	1934	1154
Exp. 2	1272	724	1382	714	2118	970
Exp. 3	1244	752	1438	658	2052	1036

3.4. Feature engineering. A feature engineering stage was applied to transform each segmented window into a compact representation for ML input. As the design targets real-time operation on resource-constrained hardware, only low-complexity time-domain features were considered, while frequency-domain descriptors were excluded due to their high computational complexity [21]. For each inertial channel, a comprehensive set of 13 statistical descriptors was selected to characterize signals across four dimensions: central tendency (mean, median), dispersion (standard deviation–SD, root mean square–RMS, interquartile range–IQR, maximum, minimum, and range), distribution shape (skewness, kurtosis), and temporal dynamics (Hjorth mobility, Hjorth complexity, and autocorrelation). Additionally, the mean and SD of the speed signal were included, as vibration intensity correlates with vehicle velocity [4]. This specific ensemble was designed to provide a multidimensional characterization of vibration patterns while maintaining compatibility with the computational constraints of embedded inference [20, 21].

3.5. Constrained hyperparameter optimization. To optimize the models for deployment on memory-limited MCUs, two primary hyperparameters – `n_estimators` and `max_depth` – were systematically

tuned within a search space of 1–50 and 1–20, respectively. These specific parameters were selected because they exert a dominant influence on both the classification performance and the final memory footprint of tree-based ensembles [21, 27, 28]. Other hyperparameters were maintained at their default settings as defined in the ML libraries, with the `random_state` fixed at 42 to ensure experimental reproducibility. All four tree-based classifier families (XGBoost, RF, Light Gradient Boosting Machine–LGBM, and Extra Trees–ET) were evaluated under this same constrained search space. For each candidate configuration, the model was fitted on the entire dataset and exported to a C/C++ header file using the `micromlgen` or `m2cgen` library [22].

The private dataset was collected using a dedicated inspection fleet operating under homogeneous mechanical conditions and consistent driving routines. Because the data do not exhibit strong driver–vehicle–scenario dependency factors, group-aware or experiment-aware resampling is unnecessary. Hyperparameter optimization (HPO) on this dataset therefore proceeds in two stages. First, a 10-fold flat cross-validation (FCV) is applied to enumerate the constrained search space, evaluate the relative ranking of the candidate configurations, and filter out models whose exported headers exceed the memory constraint. Second, only the deployable configurations identified during FCV are assessed through a full 10×10-fold NCV procedure. FCV determines the final deployable hyperparameter configuration for each classifier family, whereas NCV yields an unbiased estimate of the generalization performance of the model trained with that configuration.

Although formal HPO is conducted only on the private dataset, the PVS dataset features a structurally different evaluation protocol consisting of three experiment-wise splits. To ensure that the proposed pipeline remains reliable under this context-aware partitioning, the same constrained search space is applied within each PVS training split solely for the purpose of examining the stability of the overall pipeline under heterogeneous operational conditions. No information from the PVS test sets is accessed during this process, and the application of the search space on PVS functions as a robustness check to verify that the proposed methodology maintains consistent behavior across distinct data-partitioning structures.

3.6. Velocity-aided GPS interpolation. During data acquisition, GPS receivers often operate at 1 Hz or occasionally fail to provide updates due to satellite occlusion and multipath effects. This data sparsity creates a resolution gap in the recorded trajectory, degrading the temporal alignment with the 100 Hz IMU stream. Consequently, these artifacts reduce the reliability of the positional data associated with road surface types and localized surface anomalies.

To mitigate these issues, a velocity-aided interpolation procedure is applied to reconstruct geographically consistent intermediate positions. Unlike generic polynomial interpolation, this method utilizes the physical motion parameters of the vehicle. The procedure assumes short-term continuous motion and estimates the travelled distance S from the instantaneous speed v and elapsed time t between two valid GPS fixes:

$$S = v \times t. \tag{1}$$

A great-circle navigation model is employed with an Earth radius $R = 6371000$ m. Given two valid observations (ϕ_1, λ_1) and (ϕ_2, λ_2), the bearing angle θ is computed as:

$$\theta = \operatorname{atan2}\Big(\sin(\Delta\lambda)\cos(\phi_2),\ \cos(\phi_1)\sin(\phi_2) - \sin(\phi_1)\cos(\phi_2)\cos(\Delta\lambda)\Big), \tag{2}$$

with $\Delta\lambda = \lambda_2 - \lambda_1$. The interpolated geographic position $(\mathrm{lat}_2, \mathrm{lon}_2)$ at distance S from $(\mathrm{lat}_1, \mathrm{lon}_1)$ along direction θ is determined as:

$$\begin{cases} \mathrm{lat}_2 & = \arcsin\big(\sin(\mathrm{lat}_1)\cos\ \frac{S}{R} + \cos(\mathrm{lat}_1)\sin\ \frac{S}{R}\cos\theta\big) \\ \mathrm{lon}_2 & = \mathrm{lon}_1 + \arctan 2\big(\sin\theta\ \sin\ \frac{S}{R}\cos(\mathrm{lat}_1),\ \cos\ \frac{S}{R} - \sin(\mathrm{lat}_1)\sin(\mathrm{lat}_2)\big) \end{cases}. \tag{3}$$

The interpolation algorithm scans GPS samples and detects intervals where coordinates remain unchanged within an angular tolerance $\varepsilon = 4.5 \times 10^{-6}$ degrees (approximately 0.5 m). Intermediate points are inserted whenever the instantaneous speed is positive. To ensure synchronization with the inertial sensor data, interpolated positions are generated at $F_s = 100$ Hz. The bearing angle is estimated from the nearest available GPS observations, and since θ is recalculated whenever new observations become available, the reconstructed trajectory follows curved paths rather than imposing piecewise-linear interpolation. Figure 3 illustrates the reconstruction principle.

Algorithm 1 summarizes the complete procedure. This method produces a temporally dense trajectory that reflects t he v ehicle motion more closely than raw GPS logs. By bridging the resolution gap between low-frequency GPS updates and high-frequency IMU sampling, this method establishes the spatial reference required to map road surface conditions.

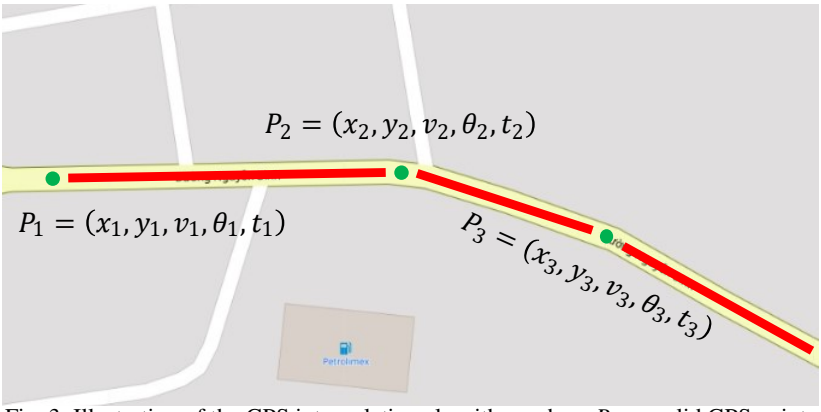

Fig. 3. Illustration of the GPS interpolation algorithm, where P_i are valid GPS points and intermediate points are reconstructed using velocity v and bearing angle θ

Algorithm 1: Velocity-aided GPS interpolation

Input: GPS samples (lat, lon, v, t); interpolation frequency F_s.
Output: Interpolated sequence $\mathscr{L}$.

$step \leftarrow 1/F_s$;
$\mathscr{L} \leftarrow \varnothing$;
$i \leftarrow 1$;
while $i \leqslant N$ **do**
  Find $j > i$ such that GPS$(j) \neq$GPS(i) within tolerance
  Append sample i to $\mathscr{L}$
  if $j > N$ *or* $v_i \leqslant 0$ *or* $j = i+1$ **then**
    $i \leftarrow i+1$; continue
  end
  Choose θ from nearest valid GPS points
  $\Delta T \leftarrow t_j - t_i$
  for *each* Δt *in* $\{step, 2\,step, \ldots, < \Delta T\}$ **do**
    $S \leftarrow v_i \Delta t$;
    $(\hat{lat}, \hat{lon}) \leftarrow$ dest_point$(lat_i, lon_i, \theta, S)$
    Append $(\hat{lat}, \hat{lon}, t_i + \Delta t)$ to $\mathscr{L}$
  end
  Append sample j;
  $i \leftarrow j$;
end
return $\mathscr{L}$

3.7. Pothole detection using abnormality index. Restoring spatial continuity is essential for locating localized road surface anomalies. At an urban vehicle speed of 40 km/h, a standard 1 Hz GPS receiver provides a position update only every 11 m, creating a resolution gap that complicates the localization of potholes. By incorporating the 100 Hz reconstructed path, the system achieves the spatial granularity necessary to map detected potholes onto the estimated vehicle trajectory.

A lightweight abnormality-index method is applied to identify these events from triaxial accelerometer measurements [21]. For each axis, a rolling mean window of size $W = 100$ samples (corresponding to 1 second at $f_s = 100$ Hz) is utilized to filter out low-frequency trends. Let $a_x(n)$, $a_y(n)$, $a_z(n)$ denote the raw accelerations at sample n, and $\bar{a}_x(n)$, $\bar{a}_y(n)$, $\bar{a}_z(n)$ be their corresponding rolling means. The abnormality index (AbI) is defined as:

$$AbI(n) = \sqrt{\left[a_x(n) - \bar{a}_x(n)\right]^2 + \left[a_y(n) - \bar{a}_y(n)\right]^2 + \left[a_z(n) - \bar{a}_z(n)\right]^2}. \quad (4)$$

A detection threshold (Th) is established based on the statistical distribution of the index:

$$Th = \mu_{\text{AbI}} + k \cdot \sigma_{\text{AbI}}, \quad (5)$$

where k is a scalar that controls the balance between sensitivity and robustness. Through experimental analysis, k is set to 2. To prevent redundant detections, samples satisfying $AbI(n) > Th$ are grouped into a single pothole event if they occur within a 1.5-second interval. This interval accounts for the sequential impacts of the front and rear axles on the same surface anomaly, which manifest as double peaks.

Figure 4 illustrates the pothole detection process, including raw accelerometer signals, the abnormality index, and the localization of events along the interpolated trajectory. By leveraging the high-resolution spatial reference from the interpolation algorithm, the system ensures that detected potholes are assigned coordinates consistent with the physical impacts recorded by the IMU. This integration provides the detailed geographic information required for road maintenance and hazard warning.

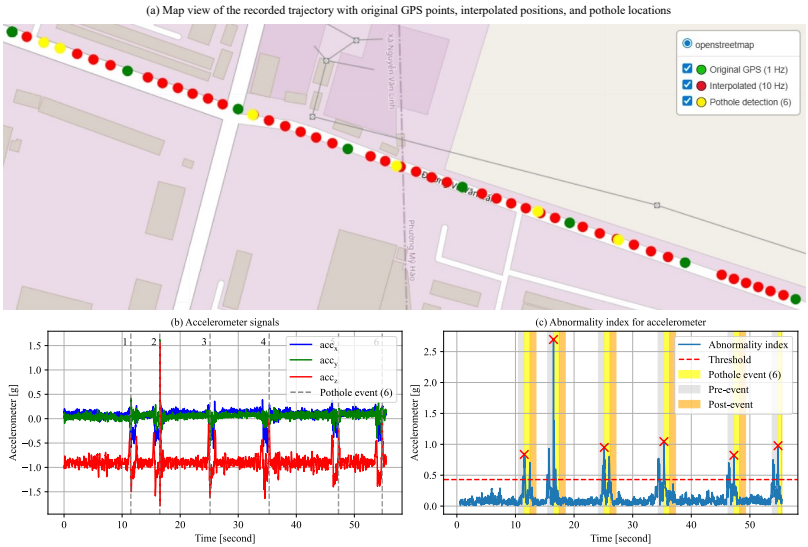

Fig. 4. Detection of pothole events using the abnormality index aligned with interpolated GPS positions

4. Result and discussion.

4.1. Evaluation on the private dataset. Table 5 summarizes the performance and Flash footprint of the four CML model families under the constrained HPO search space on the private dataset. The performance is reported as the mean accuracy $\pm$ SD across the 10-fold NCV, accompanied by the 95% confidence interval (CI). Among the evaluated methods, XGBoost attained the highest accuracy at 95.39 $\pm$ 0.82%, with a 95% CI of [94.80% – 95.97%], demonstrating strong stability across different data splits. LGBM achieved a comparable accuracy of 95.08 $\pm$ 0.56%, while RF and ET yielded lower performance at 92.45% and 91.72%, respectively.

Table 5. Performance and size of optimized CML models on the private dataset

Model	Sensor	Configuration	HPO	Accuracy	95% CI	Flash (kB)
XGBoost	Center	$n_{\text{est}} \in [1, 50]$ $m_{\text{depth}} \in [1, 20]$	{42, 6}	95.39 $\pm$ 0.82	[94.80 - 95.97]	975.79
RF			{13, 12}	92.45 $\pm$ 0.67	[91.97 - 92.94]	1018.04
LGBM			{49, 13}	95.08 $\pm$ 0.56	[94.67 - 95.48]	856.88
ET			{13, 12}	91.72 $\pm$ 0.64	[91.26 - 92.19]	1015.13

To determine whether the performance difference between the two leading models, XGBoost and LGBM, is statistically significant, McNemar's test is conducted (Table 6(a)). The contingency table for McNemar's test

accounts for the counts of correct and incorrect predictions between the two classifiers. For the private dataset, the test yields a χ^2 statistic of 4.1897 with a p-value of 0.0407. Since $p < 0.05$, the performance difference between XGBoost and LGBM is considered statistically significant.

Table 6. Statistical significance comparison between models using McNemar's test

Dataset	Model	LGBM true	LGBM false	χ^2 statistic	p-value
(a) Private (sample size $N = 9039$ windows)	XGB true	8521	101	4.1897	0.0407*
	XGB false	73	344		
(b) Public PVS (sample size $N = 7394$ windows)	XGB true	6744	148	9.3279	0.0023**
	XGB false	99	403		

*$p < 0.05$, **$p < 0.01$ indicates statistical significance levels.

The constraints imposed on the HPO search space are strictly governed by the physical memory limitations of the target ESP32 MCU. According to official specifications from Espressif, the ESP32 features 520 kB RAM, 448 kB ROM, and 4 MB Flash memory (Table 1). In practical implementation, a significant portion of the Flash memory is reserved for system partitions, including the bootloader, Wi-Fi/Bluetooth stacks, and Over-the-Air update buffers. Consequently, the application partition is typically limited to approximately 1.2 MB in standard partition tables. Analysis of the HPO results shows that model configurations exceeding the predefined limits result in header files larger than 1 MB (Figure 5). Such footprints are unsuitable for deployment as they risk exceeding the available application storage. By limiting these parameters, the optimized XGBoost model occupies 975.79 kB, effectively utilizing the Flash budget without compromising system operational integrity.

A detailed characterization of the classification behavior is shown in the per-class report in Table 7 and the confusion matrix in Figure 6(a). For dirt-road segments, XGBoost correctly identifies 3072 out of 3258 windows, corresponding to a recall of 94.29%. Most misclassified cases are assigned to the cobblestone class (184 windows), with very limited confusion involving asphalt (2 windows). Cobblestone represents the most challenging class; its vibration patterns exhibit irregular structures with partial similarity to those of dirt roads. As a result, XGBoost correctly recognizes 1963 out of 2157 cobblestone windows (recall 91.01%), with misclassifications predominantly shifting toward dirt. Asphalt displays the most distinct signature among the three classes and is classified with high reliability: 3587 out of 3624 asphalt windows are correctly predicted, yielding a recall of 98.98% and an $F1$-score of 98.60%.

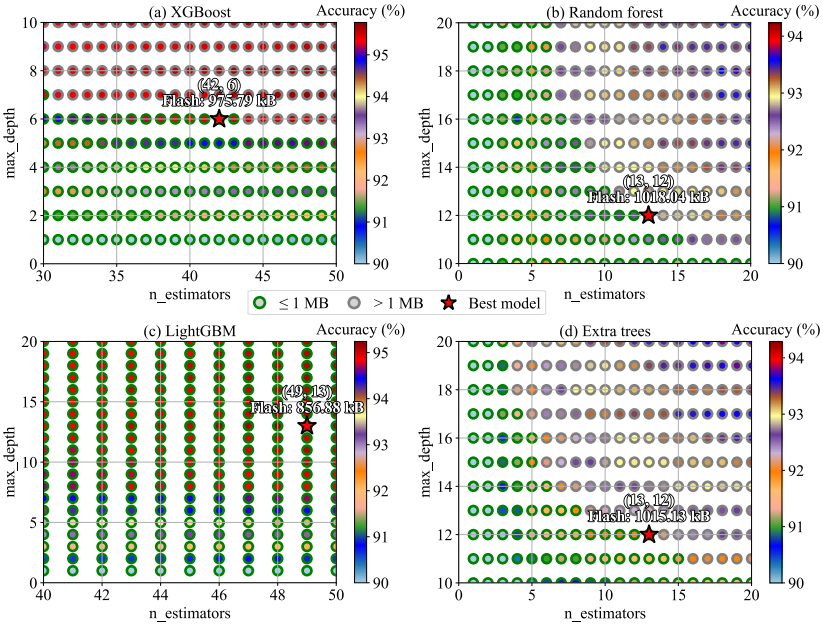

Fig. 5. The correlation between hyperparameter configuration and Flash memory footprint during parameter tuning

Table 7. Classification report of CML models on the private dataset

Model	Class	Precision	Recall	$F1$-score
XGBoost	Dirt road	95.55	94.29	94.92
	Cobblestone road	90.38	91.01	90.69
	Asphalt road	98.22	98.98	98.60
RF	Dirt road	91.75	90.85	91.30
	Cobblestone road	84.32	86.28	85.29
	Asphalt road	98.06	97.57	97.81
LGBM	Dirt road	94.88	94.38	94.63
	Cobblestone road	90.22	89.85	90.03
	Asphalt road	98.11	98.81	98.46
ET	Dirt road	90.31	90.36	90.33
	Cobblestone road	84.00	83.73	83.86
	Asphalt road	97.58	97.71	97.64

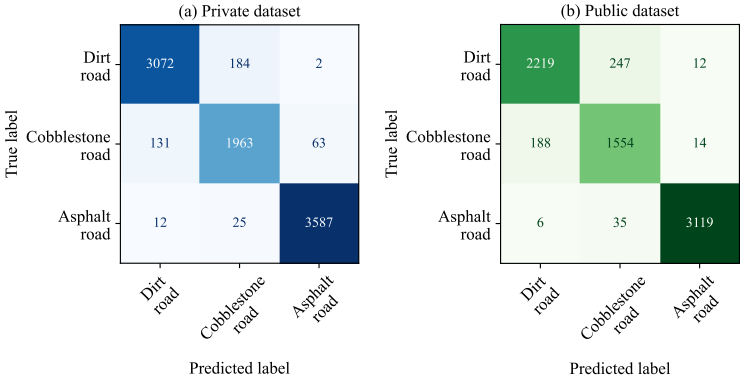

Fig. 6. Confusion matrices of the XGBoost models on: a) the private dataset (pooled 10×10-fold NCV predictions); b) the PVS dataset (combined test sets Exp. 1–3)

4.2. Evaluation on the public dataset.

4.2.1. Impact of sensor placement. A central objective in evaluating the PVS dataset is to examine how different IMU mounting positions influence classification performance. Because the dataset includes three placements – DB, AS, and BS – it enables a systematic assessment of how vibrational richness varies across the vehicle structure. Table 8 reports the average accuracy of the four CML model families across seven sensor configurations.

Table 8. Average accuracy of CML models under different sensor placement configurations on the PVS dataset

Sensor location	Accuracy (%)			
	XGBoost	**RF**	**LGBM**	**ET**
	Single placement			
DB	89.06	82.68	88.78	82.38
AS	90.29	87.62	89.86	84.11
BS	90.90	87.29	91.27	83.97
	Dual placement			
DB, AS	91.21	87.83	89.88	83.51
DB, BS	92.69	87.29	91.72	84.49
AS, BS	**93.21**	**90.73**	**92.56**	**86.50**
	All placements			
DB, AS, BS	93.19	89.68	91.28	85.11

In the single-placement setting, sensors closer to the tire–road interface consistently provide higher accuracy than the cabin-mounted dashboard sensor. Both AS and BS reach accuracy levels around 90%, whereas DB attains 89.06%

with XGBoost. This behavior aligns with mechanical intuition: the suspension system attenuates road-induced vibrations before they propagate to the cabin, reducing the discriminative content available at the dashboard location [4].

For dual-placement configurations, combining AS and BS yields the highest overall performance for all model families, reaching 93.21% with XGBoost. This configuration benefits from capturing two complementary vibration pathways – pre-suspension and post-suspension – thereby preserving a richer representation of road-surface irregularities. Configurations involving the dashboard sensor (DB, AS) or (DB, BS) result in slightly lower accuracy, indicating that the DB channel contributes limited additional information in this context.

The addition of the dashboard sensor to form the three-sensor configuration (DB, AS, BS) does not yield further improvement relative to the dual (AS, BS) setup; XGBoost accuracy decreases marginally from 93.21% to 93.19%. This outcome suggests that the dashboard channel introduces low-informative redundancy that can slightly interfere with decision boundaries established from the more discriminative AS and BS features.

4.2.2. Performance of the models. Having established the optimal sensor configuration, the next step involved evaluating the end-to-end performance of the embedded-oriented pipeline under the three PVS experiments. All evaluations in this subsection relied exclusively on the (AS, BS) dual-sensor configuration.

Table 9 summarizes the performance and Flash footprint of the four CML model families under the constrained HPO search space on the PVS dataset.

Table 9. Performance and size of optimized CML models on the PVS dataset

Model	Sensor	Configuration	HPO	Accuracy (%)	95% CI	Flash (kB)
XGBoost	AS, BS	$n_{\text{est}} \in [1, 50]$ $m_{\text{depth}} \in [1, 20]$	{50, 3}	93.21 ± 0.16	[92.80 - 93.61]	212.88
RF			{20, 9}	90.73 ± 0.52	[89.45 - 92.01]	912.08
LGBM			{45, 5}	92.56 ± 0.49	[91.34 - 93.77]	441.42
ET			{9, 12}	86.50 ± 2.21	[81.00 - 92.00]	915.39

The accuracy is reported as the mean $\pm$ SD across the three PVS experiments, accompanied by the 95% CI. Among the evaluated methods, XGBoost attained the highest overall accuracy, reaching 93.21 ± 0.16% with a 95% CI of [92.80% – 93.61%]. This narrow interval and low SD demonstrated the high stability and generalization capability of the model across different vehicles, drivers, and scenarios. Notably, this performance level aligned with the 93.17% accuracy of the CNN baseline originally reported by the author in [4] while offering a significantly more efficient Flash footprint of 212.88 kB. The remaining models exhibited varying degrees of effectiveness. LGBM

achieved a comparable accuracy of $92.56 \pm 0.49\%$, while RF and ET yielded lower accuracies of 90.73% and 86.50%, respectively.

The statistical advantage of XGBoost is also observed on the PVS dataset (Table 6(b)). The McNemar's test results in a χ^2 statistic of 9.3279 and a p-value of 0.0023. This result confirms the statistical significance of the performance difference ($p < 0.01$), suggesting that the model remains robust despite the higher heterogeneity of the PVS dataset.

To provide a finer-grained view of generalization under the PVS evaluation protocol, the per-class behavior of the four CML models was examined using the aggregated test sets from Exp. 1–3. Table 10 presents precision, recall, and $F1$-score for each road category, and Figure 6(b) visualizes the confusion matrix of the best-performing model (XGBoost), constructed from all 7394 test windows.

Table 10. Classification report of CML models on the PVS dataset

Model	**Class**	**Precision**	**Recall**	$F1$**-score**
XGBoost	Dirt road	91.96	89.55	90.74
	Cobblestone road	84.64	88.50	86.53
	Asphalt road	99.17	98.70	98.94
RF	Dirt road	89.56	84.50	86.96
	Cobblestone road	78.87	85.65	82.12
	Asphalt road	98.76	98.42	98.59
LGBM	Dirt road	91.22	88.10	89.63
	Cobblestone road	82.80	87.70	85.18
	Asphalt road	99.33	98.73	99.03
ET	Dirt road	84.43	76.84	80.46
	Cobblestone road	70.02	79.67	74.53
	Asphalt road	98.57	97.97	98.27

Among all evaluated models, asphalt consistently emerges as the easiest surface type to classify. For XGBoost, 3119 of the 3160 asphalt windows are correctly identified, corresponding to a recall of 98.70% and an $F1$-score of 98.94%. Dirt-road segments display moderately lower, though still robust, generalization. XGBoost attains a recall of 89.55%, correctly classifying 2219 out of 2478 dirt windows. Cobblestone represents the most challenging class for all models. XGBoost reaches an $F1$-score of 86.53% and a recall of 88.50%, correctly identifying 1554 out of 1756 cobblestone windows. The confusion matrix shows that most errors involve confusion with dirt, a phenomenon also noted in prior studies [4, 9].

4.2.3. Robustness across vehicles and scenarios. Table 11 presents the performance metrics under Leave-One-Vehicle-Out (LOVO) and

Leave-One-Scenario-Out (LOSO) CV protocols. These evaluations examine the model sensitivity to variations in vehicles and environmental conditions. Under the LOVO protocol, the model achieves an average accuracy of 89.63% and a macro $F1$-score of 87.18%. The accuracy drop observed for the Fiat Palio (82.49%) indicates that differences in vehicle mechanical responses can influence the recorded vibration data. For the LOSO protocol, the average accuracy and macro $F1$-score are 89.26% and 86.19%, respectively. Scenario 3 represents the most challenging subset with an accuracy of 86.24%, reflecting the influence of varying road profiles or driving environments. The fluctuations in performance across different scenarios and vehicles suggest that domain shifts affect the stability of the classification.

Table 11. Model robustness under LOVO and LOSO CV on the PVS dataset

Training	**Testing**	**Accuracy** (%)	**Macro $F1$-score** (%)
	Leave-One-Vehicle-Out CV		
Volkswagen Saveiro, Fiat Bravo	Fiat Palio	82.49	77.06
Volkswagen Saveiro, Fiat Palio	Fiat Bravo	93.06	92.35
Fiat Bravo, Fiat Palio	Volkswagen Saveiro	93.33	92.12
Average		89.63	87.18
	Leave-One-Scenario-Out CV		
Scenario 1, 2	Scenario 3	86.24	83.23
Scenario 1, 3	Scenario 2	93.39	90.76
Scenario 2, 3	Scenario 1	88.15	84.58
Average		89.26	86.19

In the PVS dataset, Leave-One-Vehicle-Out CV is equivalent to Leave-One-Driver-Out CV.

4.3. Model interpretability and ablation study.

4.3.1. Feature importance and speed sensitivity. The interpretability of the optimized XGBoost model (Fig. 7) is evaluated using SHAP (SHapley Additive exPlanations) values to quantify the contribution of each feature to the classification outcomes. For the asphalt class, low values of `IQR_Acc_Y` and higher `Mean_Speed` are the most influential indicators. For cobblestone roads, high values of `SD_Acc_Y` and lower `Mean_Speed` are significant predictors. In contrast, the classification of dirt roads relies heavily on `IQR_Acc_Y` and low values of `HM_Acc_Y`. These results show that the feature-engineering pipeline effectively captures the physical differences in road-vehicle interactions.

The impact of vehicle speed on model predictions is illustrated in the SHAP dependence plots (Fig. 8). For dirt roads, the model maintains higher confidence over a broader range, up to 40 km/h. Regarding cobblestone roads, the positive contribution is highest at very low speeds (below 25 km/h) and drops significantly as velocity increases. Beyond this threshold, the fluctuating SHAP values suggest that vibration signatures of cobblestone roads can overlap with those of dirt surfaces. Conversely, the asphalt class shows a clear positive correlation with speed. Beyond 50 km/h, the average speed becomes a major

factor for identification, as the surface smoothness allows for higher travel velocities. Therefore, the inclusion of the speed feature provides useful information for surface classification under different driving conditions.

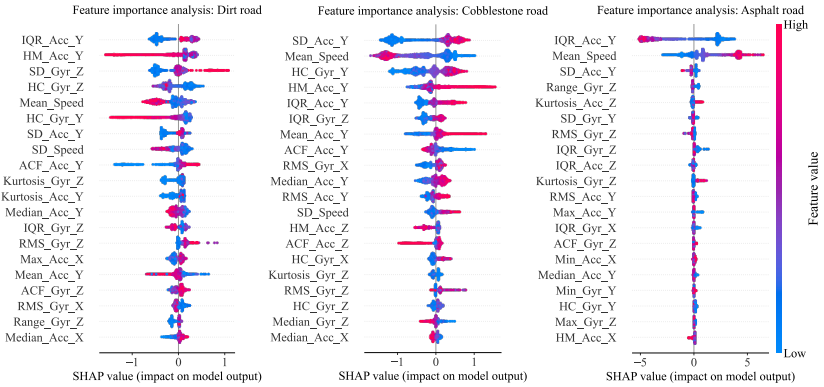

Fig. 7. Feature importance analysis of the XGBoost model using SHAP values

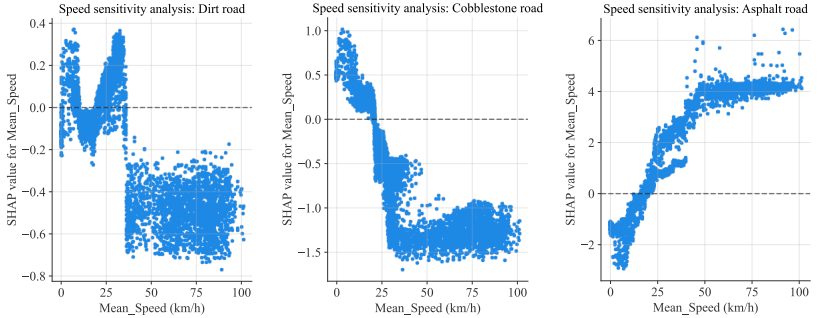

Fig. 8. SHAP dependence plot illustrating the sensitivity of model predictions to vehicle speed variations

4.3.2. Analysis of sensor fusion, feature selection, and window size. The contribution of different system components is evaluated through an ablation study covering sensor configurations, feature sets, and window sizes (Table 12). Dual-sensor fusion (AS, BS) provides the highest accuracy, as single-placement or single-sensor configurations lead to performance drops between 1.83% and 4.24%. Regarding the feature set, removing specific metrics such as range or ACF results in a marginal accuracy increase but leads to wider 95% CIs, suggesting that the proposed 13-feature combination

offers better stability. This feature set also significantly outperforms the combinations proposed in prior studies [4, 20, 25], which show accuracy reductions ranging from 2.54% to 3.32%. Furthermore, excluding the speed feature decreases accuracy by 1.14%. Analysis of window sizes shows that classification performance improves as the duration increases from 1 s to 6 s. Smaller windows fail to capture sufficient vibrational context, while larger windows introduce redundant information and increase identification latency. Therefore, the proposed configuration is effective for real-time RSC, providing a suitable balance between accuracy and operational efficiency.

Table 12. Ablation study investigating the impact of sensor fusion, feature selection, and window sizes on classification accuracy

Ablation category	Configuration variant	Accuracy (%)	Δ (%)	95% CI
Proposed	AS, BS (13 features, 6-second window, 50% overlap)	93.21	-	[92.80 - 93.61]
1. Sensor fusion	Accelerometer only	91.38	-1.83	[90.71 - 92.05]
	Gyroscope only	88.97	-4.24	[86.65 - 91.29]
	Single placement (AS only)	90.29	-2.92	[89.53 - 91.04]
	Single placement (BS only)	90.90	-2.31	[85.91 - 95.89]
2. Feature set	Remove mean	92.43	-0.78	[89.71 - 95.16]
	Remove SD	93.17	-0.04	[92.38 - 93.96]
	Remove RMS	92.86	-0.35	[92.05 - 93.68]
	Remove max	92.82	-0.39	[92.40 - 93.24]
	Remove min	92.94	-0.27	[92.44 - 93.45]
	Remove range	93.25	+0.04	[92.18 - 94.31]
	Remove median	92.99	-0.22	[91.41 - 94.57]
	Remove IQR	92.62	-0.59	[92.06 - 93.18]
	Remove skewness	92.95	-0.26	[91.89 - 94.02]
	Remove kurtosis	92.72	-0.49	[90.37 - 95.07]
	Remove HM	93.17	-0.04	[91.45 - 94.89]
	Remove HC	92.59	-0.62	[91.32 - 93.86]
	Remove ACF	93.23	+0.02	[92.08 - 94.37]
	Without speed feature	92.07	-1.14	[90.83 - 93.30]
	In paper [4] proposed mean, SD, variance	89.89	-3.32	[87.72 - 92.06]
	[20] proposed mean, SD, max, min, range, variance, median	90.67	-2.54	[88.27 - 93.06]
	In paper [25] proposed mean, RMS, SD, variance, median, range	90.24	-2.97	[87.52 - 92.97]
3. Window size	1-second (0% overlap)	87.47	-5.74	[80.91 - 94.03]
	2-second (0% overlap)	89.96	-3.25	[86.42 - 93.49]
	3-second (0% overlap)	90.17	-3.04	[85.25 - 95.09]
	4-second (50% overlap)	91.51	-1.70	[88.25 - 94.78]
	5-second (50% overlap)	92.17	-1.04	[89.65 - 94.69]

Δ *represents the change in accuracy when the model varies parameters.*

4.3.3. Cross-dataset transferability. To examine transferability, the model optimized on the private dataset was applied directly to the PVS dataset without any re-training, hyperparameter adjustment, or experiment-specific modifications. The transferred model attains 92.45% $\pm$ 0.44%, representing a reduction of only 0.76% relative to the PVS-tuned configuration. This limited

degradation, despite the substantial domain shift between datasets, indicates the stability of the proposed 6-second windowing, 13 handcrafted features, and constrained hyperparameter search.

Table 13 compares the performance of the proposed pipeline with existing studies using the PVS dataset. While architectures such as 1D-CNN [24] and Transformers [23] report accuracies exceeding 97%, these results are obtained through 80/20 or k-fold splits, which often do not isolate contextual dependencies [4]. Moreover, the heavy memory footprint and floating-point operations of these models often exceed the hardware limits of low-power MCUs [21]. In contrast, under the 3-scenario protocol, the memory-constrained XGBoost model achieves 93.21%, which is competitive with the 93.17% CNN baseline [4]. The performance is also comparable to modern lightweight architectures such as SepRNet-1D (94.69%) [9], which was evaluated under the same protocol but requires significantly higher computational resources.

Table 13. Comparison of RSC methods on the PVS dataset

Study	Method	Sensor location(s)	Windowing	Evaluation protocol	Accuracy (%)
Proposed	XGBoost + HPO for model size-constrained	4 IMUs (AS+BS, both sides)	6 seconds, 50% overlap	3-scenario split	93.21
[24]	1D-CNN (VGG-based)	Not specified	Not specified	80/20 split	99.3
[23]	Transformer	Not specified	Not specified	80/20 split + test on PVS 9	97.0 ($F1$), 80.41 (PVS 9)
[20]	XGBoost + HFW feature selection	3 IMUs (DB, left wheel, right wheel)	5 seconds, 0% overlap	70/30 split	94.2
[9]	SepRNet-1D	2 IMUs (DB, both sides)	3 seconds, 0% overlap	3-scenario split	94.69
[8]	CNN-LSTM (FFT + Stacked DWT)	2 IMUs (DB, both sides)	3 seconds, 0% overlap	66/34 split (Exp. 3 only)	94.78
[6]	ResBiGRU-SE	2 IMUs (AS+BS, left or right side only)	2 seconds, 50% overlap	5-fold CV	98.41
[5]	SE-ResNet	2 IMUs (AS+BS, left or right side only)	2 seconds, 50% overlap	5-fold CV	98.41
Menegazzo and Wangenheim (2021) [4]	CNN	2 IMUs (BS, both sides)	3 seconds, 0% overlap	3-scenario split	93.17 (Baseline)

The performance gap between the proposed method and modern DL architectures must be interpreted within the context of deployment hardware. While SepRNet-1D [9] was evaluated on a desktop CPU with 13 GB RAM, the current pipeline targets MCU-class hardware. Reliance on high-performance

computing platforms for model inference introduces significant financial and energy burdens [22]. A Cortex-M4 class MCU costs approximately 5-10 USD and can operate on low-power sources, enabling affordable and sustainable artificial intelligence inference [22]. In this context, achieving 93.21% accuracy on the PVS dataset – and 92.45% without any re-tuning – indicates that memory-efficient classical models remain a practical and competitive alternative to lightweight neural architectures.

4.4. Evaluation of GPS interpolation and pothole detection. The spatial accuracy of the proposed framework is evaluated by comparing detected coordinates against a static ground truth. This reference was established by recording GPS coordinates while remaining stationary at each pothole location to obtain a stable position. During test runs at average speeds of 20, 30, and 50 km/h, the localization error is measured as the geodesic distance between this static ground truth and two distinct points: (i) the raw GPS coordinates and (ii) the nearest interpolated coordinates assigned at the moment of impact. This procedure provides a quantitative measure of how closely the reconstructed trajectory aligns with the physical pothole position compared to raw satellite logs.

Table 14 summarizes the error metrics, including the mean error and root-mean-square error (RMSE). At 20 km/h, the raw GPS exhibits a mean error of 5.571 m, which is reduced to 1.835 m through interpolation. This improvement becomes more substantial as the vehicle speed increases and the raw GPS resolution gap widens. Specifically, the mean error decreases from 8.656 m to 2.833 m at 30 km/h, and from 11.893 m to 3.563 m at 50 km/h. Across all evaluated speeds, the proposed method significantly lowers the RMSE, notably decreasing from 11.911 m to 3.571 m at the 50 km/h threshold. These results indicate that the proposed interpolation method reduces the spatial gaps inherent in low-frequency updates, thereby improving the coordinate estimation as travel velocities increase.

Table 14. Comparison of pothole localization errors between raw GPS and the proposed velocity-aided interpolation across various average speeds

Average speed	**20 km/h**		**30 km/h**		**50 km/h**	
Indexes	**Normal**	**Proposed**	**Normal**	**Proposed**	**Normal**	**Proposed**
Min	4.849	1.474	7.841	2.209	11.014	3.278
Max	5.856	2.222	9.924	3.884	13.074	4.090
Mean	5.571	1.835	8.656	2.833	11.893	3.563
SD	0.299	0.220	0.779	0.572	0.679	0.233
RMSE	5.578	1.847	8.687	2.884	11.911	3.571

Normal: without interpolation; Proposed: interpolation method was applied

4.5. Inference latency and power consumption. The suitability of the optimized XGBoost model for resource-constrained hardware is validated through a computational complexity analysis on the ESP32 MCU (80 MHz). During operation, the system acquires 4200 samples within a 6-second window from seven sensor channels at a sampling frequency of 100 Hz. The exported C++ header file, occupying 975 kB of Flash memory, contains 126 decision trees with a maximum depth of 6. Analytical profiling of the 22951 instruction lines shows that the execution of a single prediction requires 16632 clock cycles, resulting in an inference latency of approximately 0.208 ms. This latency represents a significant reduction compared to the 2.64 ms and 3.47 ms reported for the SepRNet-1D and SepSERNet-1D architectures, respectively [9]. Furthermore, the end-to-end processing pipeline, encompassing feature extraction and model inference, incurs a total latency of 25.61 ms. This performance is significantly superior to the total processing time of 110.08 ms reported for the smartphone-based Asfault system [29]. Given the 3-second interval between recognition instances, the total processing time accounts for less than 1% of the available duty cycle, allowing for real-time operation while maintaining significant overhead for tasks such as GPS logging and data transmission.

Power consumption is evaluated based on the operational states of the hardware. The device maintains an idle power of 37-92.5 mW (10-25 mA), which increases to 314.5-388.5 mW (85-105 mA) during data acquisition from the sensors and GPS module. The inference phase requires 407-481 mW (110-130 mA), while LTE data transmission involves peaks between 666 and 814 mW (180-220 mA). Based on an average operating power of approximately 462.5 mW, the 2000 mAh battery provides about 13.6 hours of autonomous operation. This duration is sufficient to maintain system functionality during electrical instabilities until the vehicle power system is addressed. This performance is noteworthy when compared to smartphone-based applications, where the Asfault system demonstrated a battery drain of 13% per hour on a larger 2900 mAh battery, resulting in approximately 7.7 hours of operation [29].

The scalability of the proposed system for large-scale urban environments is supported by the efficiency of the edge-computing pipeline. By performing feature extraction and classification directly on the MCU, only high-level metadata – such as road surface types and coordinates – are transmitted to the central server. This significantly reduces network bandwidth requirements and server-side computational load, even in high-traffic scenarios with numerous active nodes. Furthermore, the low unit cost of the sensor node enables the deployment of extensive crowdsourcing fleets. While the current framework is validated on vehicles, there is significant potential for

extrapolation into multi-agent systems involving swarms of wheeled ground robots [30, 31]. In such configurations, autonomous agents and conventional vehicle fleets can collaborate through collective data fusion to refine spatial maps, ensuring high-resolution monitoring across broad geographical areas.

5. Conclusion. This research developed an IoT-based framework for RSC using a lightweight feature set and memory-constrained ML. The results show that optimized classical models achieve performance comparable to DL architectures while requiring significantly fewer computational resources. This efficiency allows the system to operate in real-time on low-cost MCU.

The framework also integrates GPS interpolation and pothole detection algorithms to improve trajectory reconstruction and hazard identification. Future work will focus on a real-time system to provide road-condition maps and automated alerts. This system will assist authorities in maintenance planning and provide road users with timely information to improve travel safety.

Data usage statement. The complete dataset is available at: 10.21227/x6av-va13. Researchers who wish to access this dataset must cite this paper in their publications and contact hieu.daoto@phenikaa-uni.edu.vn for further information.

Source code. HieuSSALAB/Real-TimeRoadSurface

Reference.

1. Carlos M.R., Aragon M.E., Gonzalez L.C., Escalante H.J., Martınez F. Evaluation of detection approaches for road anomalies based on accelerometer readings – addressing who's who. IEEE Transactions on Intelligent Transportation Systems. 2018. vol. 19. no. 10. pp. 3334–3343. DOI: 10.1109/TITS.2017.2773084.
2. Ganguly B., Dey D., Munshi S. An unsupervised learning approach for road anomaly segmentation using RGB-D sensor for advanced driver assistance system. IEEE Transactions on Intelligent Transportation Systems. 2022. vol. 23. no. 10. pp. 19042–19053. DOI: 10.1109/TITS.2022.3164847.
3. Duong C.C., Nguyen T.T., Duong V.T., Tran D.-N., Chinh T.M., Le A.N., Tran D.-T. Smartphone-based sensing for intelligent inland waterway transportation. International Journal of Interactive Mobile Technologies (iJIM). 2020. vol. 14. no. 18. pp. 195–203. DOI: 10.3991/ijim.v14i18.16449.
4. Menegazzo J., von Wangenheim A. Road surface type classification based on inertial sensors and machine learning: A comparison between classical and deep machine learning approaches for multi-contextual real-world scenarios. Computing. 2021. vol. 103. no. 10. pp. 2143–2170. DOI: 10.1007/s00607-021-00914-0.
5. Hnoohom N., Mekruksavanich S., Jitpattanakul A. A comprehensive evaluation of state-of-the-art deep learning models for road surface type classification. Intelligent Automation & Soft Computing. 2023. vol. 37. no. 2. pp. 1275–1291. DOI: 10.32604/i-asc.2023.038584.
6. Mekruksavanich S., Rojanavasu P., Srisungsittisunti B., Plengvittaya C., Phaphan W., Jitpattanakul A. Enhancing intelligent transportation systems: A deep learning approach

for terrain recognition using vehicular inertial sensors. Lobachevskii Journal of Mathematics. 2024. vol. 45. no. 12. pp. 6324–6342. DOI: 10.1134/S1995080224607628.
7. Sattar S., Li S., Chapman M. Developing a near real-time road surface anomaly detection approach for road surface monitoring. Measurement. 2021. vol. 185.
8. Raslan E., Alrahmawy M.F., Mohammed Y., Tolba A. Evaluation of data representation techniques for vibration based road surface condition classification. Scientific Reports. 2024. vol. 14. no. 1. DOI: 10.1038/s41598-024-61757-1.
9. Manoni L., Orcioni S., Conti M. A lightweight 1D-CNN architecture for accurate and efficient road type classification using vibrational signals. IEEE Access. 2025. vol. 13. pp. 174349–174367. DOI: 10.1109/ACCESS.2025.3617943.
10. Botezatu A.-P., Burlacu A., Orhei C. A review of deep learning advancements in road analysis for autonomous driving. Applied Sciences. 2024. vol. 14. no. 11. DOI: 10.3390/app14114705.
11. Yigit H., Koylu H., Eken S. Estimation of road surface type from brake pressure pulses of ABS. Expert Systems with Applications. 2023. vol. 212. DOI: 10.1016/j.eswa.2022.118726.
12. Biju S., Chammam A., Askar S., Rodrigues P., Jalalnezhad M. Prediction-based controller radial neural network for the traction control system. Journal of Vibration and Control. 2024. DOI: 10.1177/10775463241296911.
13. Martinez-Rıos E.A., Bustamante-Bello M.R., Arce-Saenz L.A. A review of road surface anomaly detection and classification systems based on vibration-based techniques. Applied Sciences. 2022. vol. 12. no. 19. DOI: 10.3390/app12199413.
14. Kim Y.-M., Kim Y.-G., Son S.-Y., Lim S.-Y., Choi B.-Y., Choi D.-H. Review of recent automated pothole-detection methods. Applied Sciences. 2022. vol. 12. no. 11. DOI: 10.3390/app12115320.
15. Manoni L., Orcioni S., Conti M. Recent advancements in deep learning techniques for road condition monitoring: A comprehensive review. IEEE Access. 2024. vol. 12. pp. 154271–154293. DOI: 10.1109/ACCESS.2024.3481649.
16. Coenen T.B., Golroo A. A review on automated pavement distress detection methods. Cogent Engineering. 2017. vol. 4. no. 1. DOI: 10.1080/23311916.2017.1374822.
17. Cao W., Liu Q., He Z. Review of pavement defect detection methods. IEEE Access. 2020. vol. 8. pp. 14531–14544. DOI: 10.1109/ACCESS.2020.2966881.
18. Dib J., Sirlantzis K., Howells G. A review on negative road anomaly detection methods. IEEE Access. 2020. vol. 8. pp. 57298–57316. DOI: 10.1109/ACCESS.2020.2982220.
19. Peraka N.S.P., Biligiri K.P. Pavement asset management systems and technologies: A review. Automation in Construction. 2020. vol. 119. DOI: 10.1016/j.autcon.2020.103336.
20. Cong N.V., Tran D.-N., Long T.T., Thao N.G.M., Tran D.-T. Hybrid feature selection for real-time road surface classification on low-end hardware: A machine learning approach. Results in Engineering. 2025. vol. 27. DOI: 10.1016/j.rineng.2025.105693.
21. Dao T.-H., Tran D.-N., Bui V.-H., Nguyen V.S., Hoa D.K., Thanh P.V., Tran D.-T. RFAR: A real-time firefighter activity recognition system using wearable accelerometer. IEEE Sensors Journal. 2025. vol. 25. no. 17. pp. 33674–33691. DOI: 10.1109/JSEN.2025.3593466.
22. Saha S.S., Sandha S.S., Srivastava M. Machine learning for microcontroller-class hardware: A review. IEEE Sensors Journal. 2022. vol. 22. no. 22. pp. 21362–21390. DOI: 10.1109/JSEN.2022.3210773.
23. Aslam I., Mahfuz S. Transformer-based classification of road conditions using vehicular sensor data. Procedia Computer Science. 2025. vol. 257. pp. 444–451. DOI: 10.1016/j.procs.2025.03.058.

24. Cui J., Zhang H., Wang X., Jing Y., Chou X. Research on road surface recognition algorithm based on vehicle vibration data. Sensors. 2025. vol. 25. no. 18. DOI: 10.3390/s25185642.
25. Thang N.V., Thang P.D., Kien L.M., Thu N.T., Dao T.-H. Development of low-cost road surface classification system using acceleration sensors on motorcycles. Edge Artificial Intelligence: Foundations, Techniques, and Applications. 2025. pp. 567–578. DOI: 10.1002/9781394355037.ch25.
26. Lu D.-N., Nguyen D.-N., Nguyen T.-H., Nguyen H.-N. Vehicle mode and driving activity detection based on analyzing sensor data of smartphones. Sensors. 2018. vol. 18. no. 4. DOI: 10.3390/s18041036.
27. Durap A. A comparative analysis of machine learning algorithms for predicting wave runup. Anthropocene Coasts. 2023. vol. 6. no. 1. DOI: 10.1007/s44218-023-00033-7.
28. Thomas N.S., Kaliraj S. An improved and optimized random forest based approach to predict the software faults. SN Computer Science. 2024. vol. 5. no. 5. DOI: 10.1007/s42979-024-02764-x.
29. Souza V.M., Giusti R., Batista A.J. Asfault: A low-cost system to evaluate pavement conditions in real-time using smartphones and machine learning. Pervasive and Mobile Computing. 2018. vol. 51. pp. 121–137. DOI: 10.1016/j.pmcj.2018.10.008.
30. Khaleghian S., Taheri S. Terrain classification using intelligent tire. Journal of Terramechanics. 2017. vol. 71. pp. 15–24. DOI: 10.1016/j.jterra.2017.01.005.
31. Sebastian B., Ben-Tzvi P. Support vector machine based real-time terrain estimation for tracked robots. Mechatronics. 2019. vol. 62. DOI: 10.1016/j.mechatronics.2019.102260.

Viet-Hoan Bui — Student, Phenikaa School of Engineering, Phenikaa University. Research interests: integrated sensors, machine learning, advanced signal processing for healthcare applications. The number of publications — 4. 21011096@st.phenikaa-uni.edu.vn; Duongnoi, 12116, Hanoi, Vietnam; office phone: +84(869)891-098.

Duc-Nghia Tran — Ph.D., Researcher, Institute of Information Technology, Vietnam Academy of Science and Technology. Research interests: AI, IoT, electron paramagnetic resonance, parameter estimation, data analysis. The number of publications — 51. nghiatd@ioit.ac.vn; 18, Hoang Quoc, 10000, Hanoi, Vietnam; office phone: +84(936)866-335.

To-Hieu Dao — Ph.D., Lecturer, Phenikaa School of Engineering, Phenikaa University. Research interests: machine learning, indoor localization, IMU sensor fusion, signal processing, vibration-based sensing, embedded systems, real-time sensing applications. The number of publications — 12. hieu.daoto@phenikaa-uni.edu.vn; Duongnoi, 12116, Hanoi, Vietnam; office phone: +84(389)959-524.

Hoang Quang Trung — Ph.D., Lecturer, Phenikaa School of Engineering, Phenikaa University. Research interests: signal and data processing for telecommunication systems, internet of things (IoT). The number of publications — 10. trung.hoangquang@phenikaa-uni.edu.vn; Duongnoi, 12116, Hanoi, Vietnam; office phone: +84(389)959-524.

Pham Vu Kien — Student, Phenikaa School of Engineering, Phenikaa University. Research interests: IoT and intelligent sensing solutions for smart infrastructure and transportation systems. 20010645@st.phenikaa-uni.edu.vn; Duongnoi, 12116, Hanoi, Vietnam; office phone: +84(389)959-524.

Nguyen Van Thang — Ph.D., Lecturer, VNU University of Engineering and Technology. Research interests: IoT and sensor applications, biomedical electronics, signal processing. The number of publications — 20. nvthangdtvt@vnu.edu.vn; 144, Xuan Thuy, Hanoi, Vietnam; office phone: +84(389)959-524.

Duc-Tan Tran — Ph.D., Dr.Sci., Professor, Lecturer, Vice Dean of the Faculty, Phenikaa School of Engineering, Phenikaa University. Research interests: representation, processing,

analysis, and communication of information embedded in signals and datasets. The number of publications — 150. tan.tranduc@phenikaa-uni.edu.vn; Duongnoi, 12116, Hanoi, Vietnam; office phone: +84(904)182-389.

Acknowledgements. This research is funded by Phenikaa University, number PU2024-1-A-01. Thanks to the Institute of Information Technology (IoIT-VAST) for supporting this research and allowing us to use the «IoT and Robot intensive laboratory» equipment.

УДК 006.72 **DOI** 10.15622/ia.25.2.7

В.-Х. Буи , Д.-Н. Тран, Т.-Х. Дао, Х.-К. Чунг, Ф.-В. Киен, Н.-В. Тханг, Д.-Т. Тран

IDRRS: ИНЕРЦИАЛЬНОЕ IOT-УСТРОЙСТВО ДЛЯ КЛАССИФИКАЦИИ ДОРОЖНОГО ПОКРЫТИЯ В РЕЖИМЕ РЕАЛЬНОГО ВРЕМЕНИ И ПОВЫШЕНИЯ ТОЧНОСТИ ОЦЕНКИ МЕСТОПОЛОЖЕНИЯ

Буй В.-Х., Тран Д.-Н., Дао Т.-Х., Чунг Х.К., Киен Ф.В., Тханг Н.В., Тран Д.-Т. **IDRRS: инерциальное IoT-устройство для классификации дорожного покрытия в режиме реального времени и повышения точности оценки местоположения.**

Аннотация. Мониторинг состояния дорожного покрытия является важной задачей, направленной на повышение безопасности дорожного движения и оптимизацию обслуживания транспортной инфраструктуры. В настоящей работе разработана инерциальная IoT-система, предназначенная для классификации типов дорожного покрытия в режиме реального времени, обнаружения выбоин и повышения точности оценки местоположения транспортного средства. Предложенная система включает модель XGBoost, адаптированную для развертывания на микроконтроллерах с ограниченным объемом памяти, алгоритм интерполяции GPS с использованием данных о скорости движения, а также метод обнаружения выбоин на основе индекса аномалий. Экспериментальная оценка проведена на собственном наборе данных и общедоступном наборе данных PVS. Достигнутая точность классификации составила 95.39% и 93.21% соответственно. Для анализа переносимости модель, обученная на собственном наборе данных, была применена к набору PVS без дополнительного обучения и обеспечила точность 92.45%, что подтверждает устойчивость предложенного подхода к смене источника данных. Применение процедуры интерполяции GPS позволило снизить среднюю ошибку локализации с 5.571-11.893 м до 1.835-3.563 м при скоростях движения транспортного средства от 20 до 50 км/ч. Дополнительным вкладом работы является публикация собственного набора данных, содержащего вибрационные сигнатуры типовых дорожных покрытий, что способствует дальнейшему развитию методов интеллектуальной классификации дорожного покрытия.

Ключевые слова: инерциальный датчик, режим реального времени, микроконтроллер, машинное обучение, интернет вещей.

Литература

1. Carlos M.R., Aragon M.E., Gonzalez L.C., Escalante H.J., Martınez F. Evaluation of detection approaches for road anomalies based on accelerometer readings – addressing who's who. IEEE Transactions on Intelligent Transportation Systems. 2018. vol. 19. no. 10. pp. 3334–3343. DOI: 10.1109/TITS.2017.2773084.
2. Ganguly B., Dey D., Munshi S. An unsupervised learning approach for road anomaly segmentation using RGB-D sensor for advanced driver assistance system. IEEE Transactions on Intelligent Transportation Systems. 2022. vol. 23. no. 10. pp. 19042–19053. DOI: 10.1109/TITS.2022.3164847.
3. Duong C.C., Nguyen T.T., Duong V.T., Tran D.-N., Chinh T.M., Le A.N., Tran D.-T. Smartphone-based sensing for intelligent inland waterway transportation. International Journal of Interactive Mobile Technologies (iJIM). 2020. vol. 14. no. 18. pp. 195–203. DOI: 10.3991/ijim.v14i18.16449.

4. Menegazzo J., von Wangenheim A. Road surface type classification based on inertial sensors and machine learning: A comparison between classical and deep machine learning approaches for multi-contextual real-world scenarios. Computing. 2021. vol. 103. no. 10. pp. 2143–2170. DOI: 10.1007/s00607-021-00914-0.
5. Hnoohom N., Mekruksavanich S., Jitpattanakul A. A comprehensive evaluation of state-of-the-art deep learning models for road surface type classification. Intelligent Automation & Soft Computing. 2023. vol. 37. no. 2. pp. 1275–1291. DOI: 10.32604/i-asc.2023.038584.
6. Mekruksavanich S., Rojanavasu P., Srisungsittisunti B., Plengvittaya C., Phaphan W., Jitpattanakul A. Enhancing intelligent transportation systems: A deep learning approach for terrain recognition using vehicular inertial sensors. Lobachevskii Journal of Mathematics. 2024. vol. 45. no. 12. pp. 6324–6342. DOI: 10.1134/S1995080224607628.
7. Sattar S., Li S., Chapman M. Developing a near real-time road surface anomaly detection approach for road surface monitoring. Measurement. 2021. vol. 185.
8. Raslan E., Alrahmawy M.F., Mohammed Y., Tolba A. Evaluation of data representation techniques for vibration based road surface condition classification. Scientific Reports. 2024. vol. 14. no. 1. DOI: 10.1038/s41598-024-61757-1.
9. Manoni L., Orcioni S., Conti M. A lightweight 1D-CNN architecture for accurate and efficient road type classification using vibrational signals. IEEE Access. 2025. vol. 13. pp. 174349–174367. DOI: 10.1109/ACCESS.2025.3617943.
10. Botezatu A.-P., Burlacu A., Orhei C. A review of deep learning advancements in road analysis for autonomous driving. Applied Sciences. 2024. vol. 14. no. 11. DOI: 10.3390/app14114705.
11. Yigit H., Koylu H., Eken S. Estimation of road surface type from brake pressure pulses of ABS. Expert Systems with Applications. 2023. vol. 212. DOI: 10.1016/j.eswa.2022.118726.
12. Biju S., Chammam A., Askar S., Rodrigues P., Jalalnezhad M. Prediction-based controller radial neural network for the traction control system. Journal of Vibration and Control. 2024. DOI: 10.1177/10775463241296911.
13. Martinez-Rıos E.A., Bustamante-Bello M.R., Arce-Saenz L.A. A review of road surface anomaly detection and classification systems based on vibration-based techniques. Applied Sciences. 2022. vol. 12. no. 19. DOI: 10.3390/app12199413.
14. Kim Y.-M., Kim Y.-G., Son S.-Y., Lim S.-Y., Choi B.-Y., Choi D.-H. Review of recent automated pothole-detection methods. Applied Sciences. 2022. vol. 12. no. 11. DOI: 10.3390/app12115320.
15. Manoni L., Orcioni S., Conti M. Recent advancements in deep learning techniques for road condition monitoring: A comprehensive review. IEEE Access. 2024. vol. 12. pp. 154271–154293. DOI: 10.1109/ACCESS.2024.3481649.
16. Coenen T.B., Golroo A. A review on automated pavement distress detection methods. Cogent Engineering. 2017. vol. 4. no. 1. DOI: 10.1080/23311916.2017.1374822.
17. Cao W., Liu Q., He Z. Review of pavement defect detection methods. IEEE Access. 2020. vol. 8. pp. 14531–14544. DOI: 10.1109/ACCESS.2020.2966881.
18. Dib J., Sirlantzis K., Howells G. A review on negative road anomaly detection methods. IEEE Access. 2020. vol. 8. pp. 57298–57316. DOI: 10.1109/ACCESS.2020.2982220.
19. Peraka N.S.P., Biligiri K.P. Pavement asset management systems and technologies: A review. Automation in Construction. 2020. vol. 119. DOI: 10.1016/j.autcon.2020.103336.
20. Cong N.V., Tran D.-N., Long T.T., Thao N.G.M., Tran D.-T. Hybrid feature selection for real-time road surface classification on low-end hardware: A machine learning approach. Results in Engineering. 2025. vol. 27. DOI: 10.1016/j.rineng.2025.105693.

21. Dao T.-H., Tran D.-N., Bui V.-H., Nguyen V.S., Hoa D.K., Thanh P.V., Tran D.-T. RFAR: A real-time firefighter activity recognition system using wearable accelerometer. IEEE Sensors Journal. 2025. vol. 25. no. 17. pp. 33674–33691. DOI: 10.1109/JSEN.2025.3593466.
22. Saha S.S., Sandha S.S., Srivastava M. Machine learning for microcontroller-class hardware: A review. IEEE Sensors Journal. 2022. vol. 22. no. 22. pp. 21362–21390. DOI: 10.1109/JSEN.2022.3210773.
23. Aslam I., Mahfuz S. Transformer-based classification of road conditions using vehicular sensor data. Procedia Computer Science. 2025. vol. 257. pp. 444–451. DOI: 10.1016/j.procs.2025.03.058.
24. Cui J., Zhang H., Wang X., Jing Y., Chou X. Research on road surface recognition algorithm based on vehicle vibration data. Sensors. 2025. vol. 25. no. 18. DOI: 10.3390/s25185642.
25. Thang N.V., Thang P.D., Kien L.M., Thu N.T., Dao T.-H. Development of low-cost road surface classification system using acceleration sensors on motorcycles. Edge Artificial Intelligence: Foundations, Techniques, and Applications. 2025. pp. 567–578. DOI: 10.1002/9781394355037.ch25.
26. Lu D.-N., Nguyen D.-N., Nguyen T.-H., Nguyen H.-N. Vehicle mode and driving activity detection based on analyzing sensor data of smartphones. Sensors. 2018. vol. 18. no. 4. DOI: 10.3390/s18041036.
27. Durap A. A comparative analysis of machine learning algorithms for predicting wave runup. Anthropocene Coasts. 2023. vol. 6. no. 1. DOI: 10.1007/s44218-023-00033-7.
28. Thomas N.S., Kaliraj S. An improved and optimized random forest based approach to predict the software faults. SN Computer Science. 2024. vol. 5. no. 5. DOI: 10.1007/s42979-024-02764-x.
29. Souza V.M., Giusti R., Batista A.J. Asfault: A low-cost system to evaluate pavement conditions in real-time using smartphones and machine learning. Pervasive and Mobile Computing. 2018. vol. 51. pp. 121–137. DOI: 10.1016/j.pmcj.2018.10.008.
30. Khaleghian S., Taheri S. Terrain classification using intelligent tire. Journal of Terramechanics. 2017. vol. 71. pp. 15–24. DOI: 10.1016/j.jterra.2017.01.005.
31. Sebastian B., Ben-Tzvi P. Support vector machine based real-time terrain estimation for tracked robots. Mechatronics. 2019. vol. 62. DOI: 10.1016/j.mechatronics.2019.102260.

Буй Вьет-Хоан — студент, инженерная школа, Университет Феникаа. Область научных интересов: интегрированные сенсорные системы, методы машинного обучения, современные методы обработки сигналов для применения в сфере здравоохранения. Число научных публикаций — 4. 21011096@st.phenikaa-uni.edu.vn; Дуонгной, 12116, Ханой, Вьетнам; р.т.: +84(869)891-098.

Тран Дук-Нгиа — Ph.D., научный сотрудник, институт информационных технологий, Вьетнамская академия наук и технологий. Область научных интересов: искусственный интеллект, интернет вещей (IoT), электронный парамагнитный резонанс, оценивание параметров, анализ данных. Число научных публикаций — 51. nghiatd@ioit.ac.vn; Хоанг Куок, 18, 10000, Ханой, Вьетнам; р.т.: +84(936)866-335.

Дао Ту-Хиеу — Ph.D., преподаватель, инженерная школа, Университет Феникаа. Область научных интересов: машинное обучение, определение местоположения внутри помещений, объединение данных с инерциальных измерительных блоков (IMU), обработка сигналов, датчики на основе вибрации, встроенные системы, приложения для сбора данных в реальном времени. Число научных публикаций — 12. hieu.daoto@phenikaa-uni.edu.vn; Дуонгной, 12116, Ханой, Вьетнам; р.т.: +84(389)959-524.

Чунг Хоанг Куанг — Ph.D., преподаватель, инженерная школа, Университет Феникаа. Область научных интересов: обработка сигналов и данных для телекоммуникационных систем, интернет вещей (IoT). Число научных публикаций — 10. trung.hoangquang@phenikaa-uni.edu.vn; Дуонгной, 12116, Ханой, Вьетнам; р.т.: +84(389)959-524.

Киен Фам Ву — студент, инженерная школа, Университет Феникаа. Область научных интересов: IoT, интеллектуальные сенсорные системы, поддерживающие умную инфраструктуру, транспортные системы. 20010645@st.phenikaa-uni.edu.vn; Дуонгной, 12116, Ханой, Вьетнам; р.т.: +84(389)959-524.

Тханг Нгуен Ван — Ph.D., преподаватель, Университет инженерии и технологий Вьетнамского национального университета. Область научных интересов: приложения интернета вещей (IoT) и сенсорных систем, биомедицинская электроника, обработка сигналов. Число научных публикаций — 20. nvthangdtvt@vnu.edu.vn; Суан Туй, 144, Ханой, Вьетнам; р.т.: +84(389)959-524.

Тран Дук-Тан — Ph.D., Dr.Sci., профессор, преподаватель, заместитель декана факультета, инженерная школа, Университет Феникаа. Область научных интересов: представление, обработка, анализ и передача информации, содержащейся в сигналах и наборах данных. Число научных публикаций — 150. tan.tranduc@phenikaa-uni.edu.vn; Дуонгной, 12116, Ханой, Вьетнам; р.т.: +84(904)182-389.

Поддержка исследований. Работа выполнена при финансовой поддержке Университета Феникаа (проект № PU2024-1-A-01). Авторы выражают благодарность Институту информационных технологий (ИИТ ВАНТ) за поддержку исследования и предоставленную возможность использовать оборудование лаборатории «Интенсивные робототехнические и IoT-системы».