DriveCheck: A Driving Behavior Monitor

Michael Andre Samuel Cuadros-Ccahuana
Faculty of Information Systems Engineering, Universidad Peruana de Ciencias Aplicadas, Lima, Peru
u201713406@upc.edu.pe

Kevin Enrique Servat-Farfan
Faculty of Information Systems Engineering, Universidad Peruana de Ciencias Aplicadas, Lima, Peru
u202010146@upc.edu.pe

Pedro Castaneda
Faculty of Systems Engineering and Electrical Mechanics, Universidad Nacional Toribio Rodriguez de Mendoza, Amazonas, Peru
pedro.castaneda@untrm.edu.pe

Juan Mansilla-Lopez
Faculty of Information Systems Engineering, Universidad Peruana de Ciencias Aplicadas, Lima, Peru
pcsijman@upc.edu.pe (corresponding author)

Alberto Daniel Garcia-Nunez
Universidad Pontificia Bolivariana, Medellin, Antioquia, Colombia
alberto.garcia@upb.edu.co

Received: 23 September 2025 | Revised: 10 November 2025 | Accepted: 17 November 2025

 | DOI: https://doi.org/10.48084/etasr.15030

ABSTRACT

This paper introduces the DriveCheck driving behavior monitor, a web application for the evaluation of risky driving behaviors. The tool uses a random forest supervised learning algorithm to classify normal and risky events based on acceleration, braking, and turning information. Initially, a public Kaggle dataset labeled by driving level was used. Subsequently, a proprietary dataset was collected from 27 real urban trips in Lima, where selected maneuvers (sharp turns, sudden acceleration, and harsh braking) were induced under controlled conditions to calibrate the model. The results showed a risk classification accuracy of 97%, validating the robustness of the model and its applicability in real urban settings (field tests in Lima).

Keywords-safe driving; risk assessment; behavior analysis; road safety

I. INTRODUCTION

A structural problem of vehicle insurance is that rating is often applied uniformly, regardless of each person's driving behavior. This penalizes responsible drivers, who end up subsidizing riskier drivers, and perpetuates inequality in the market. In the USA, for example, the average cost of full-coverage auto insurance has risen to $2,543, a 26% increase over the prior year, whereas wages have not grown at the same pace, a gap affecting safer drivers disproportionately [1, 2]. Empirical evidence also links behaviors such as speeding and driving in high-traffic urban areas with higher accident risk [3]. The Poisson model utilized in [3] that combines driver telematics with contextual factors such as weather can explain claim frequency; however, rare events such as severe crashes and real-time deployment remain challenging [3]. In parallel, usage-adjusted pricing schemes such as Pay-as-You-Drive (PAYD) and Pay-How-You-Drive (PHYD) have been shown to improve fairness by aligning premiums with actual exposure and behavior [4].

From the actuarial perspective, Generalized Linear Models (GLMs) and Generalized Additive Models (GAMs) have been coupled with classical insurance covariates, including driver and vehicle age, credit history, annual mileage, and years without claims, along with territorial clustering, to sharpen risk differentiation. Experiments on a synthetic dataset of 100,000 UBI policies mixing traditional and behavior data showed that annual mileage exhibits a threefold higher relativity for high-distance drivers and that territorial grouping improves prediction [5]. Authors in [6] reviewed driving-style evaluation methods, emphasizing the lack of clear standards for mapping behavioral risks, while documenting practical benefits of

PAYD, Pay-as-You-Save (PAYS), and PHYD implementations based on variables such as speed, acceleration, mileage, and habits. Market dynamics also play a role: evidence from a standardized market indicates that price variation and post-claim learning influence switching, with a 1% price increase raising the probability of changing insurer by 0.2% and a 1% deviation from a floor price raising switching by 2% in a sample of 20,759 customers and 50,553 observations [7].

On the machine learning side, deep architectures have been explored to extract risk-relevant signals from telematics. TabNet improves recall for pricing-related classification on synthetic telematics by using variables such as miles traveled, harsh braking, and sociodemographics, achieving an F1-score up to 0.60 and outperforming GLM and XGBoost in recall under cross-validation with simulated telemetry [8]. The convolutional model in [9] can identify aggressive driving in near real-time using five variables (speed, acceleration, deceleration, frontal distance, and steering wheel turn) captured by OBD-II and GPS, reaching 96.1% accuracy on segments of 1-10 s. The fuzzy-logic system in [10] addresses uncertainty and data scarcity through Mamdani inference on variables such as speed and accelerator/brake usage. It achieved F1 = 0.84, outperforming Random Forest (0.69) and Naive Bayes (0.42), with additional traffic-simulation performance of 82.7% [10]. Hybrid pipelines that fuse tabular telematics with driver imagery reach very high detection accuracy for fatigue, distraction, and aggression when combining models such as Random Forest, ResNet50, and EfficientNetB6 on public datasets, reporting 99.32% accuracy on tabular inputs and 99.87% on images [11]. A new in-vehicle image dataset with 7,286 images across five classes (safe driving, turning, phone use to talk, phone use to write, and other distractions), collected with two smartphones over 30 days and 178 drivers, was presented in [12].

Beyond driver characterization, recent studies model collision risk and decision-making. The comprehensive collision-risk framework of [13] considers both pre-collision probability and post-collision intensity by integrating intentions, speed, acceleration, and inter-vehicle distance within a driving-intention-based trajectory prediction model that combines LSTM networks, social tensors, and Bézier curves. In autonomous-driving decision making, the deep reinforcement learning with a DQN reported in [14] improved the distance traveled by 3.3%, safety by 2.1%, and reduced computation time by 43% compared with DDPG and LSTM baselines in CARLA simulations. In [15], dynamic-risk quantification that jointly considers environment, vehicle motion, and driver behavior attributes obtained weightings of 36.2%, 27.3%, and 36.5% respectively. It also attained a 98.8% correlation with observed risk with an error of 0.007, outperforming time-to-collision and sudden speed-change metrics by about 30% in data collected from 20 drivers with/using simulators and visual sensors [15]. Structural-equation models on 54 UK drivers and more than 14,000 trips over 18 weeks using indicators, such as wiper use, high beams, lane departure warnings, and day of the week, showed a positive correlation between task complexity and risk, with standardized coefficients between 0.32 and 0.53 and higher risk under adverse weather and poor visibility [16]. Improved Intelligent Driver Model simulations, calibrated in VISSIM with real urban trajectories, showed that inverse time-to-collision and related parameters can identify potential, general, and serious risks with accuracies of 96.63%, 97.72%, and 92.71%, respectively [17]. Interpretable indicators such as ηCFDS, derived from 17,002 radar-camera trajectories and modeled with GPR and XGBoost, emerge as the most important predictor of conflict-risk exposure time and separate aggressive and pseudo-shy styles distinguishing them from normal and shy ones, highlighting the value of road interventions [18].

Multisource fusion and large-scale analytics further advance operational viability. The Multi-source Temporal-Attention Fusion Network (MTAFN) in [19] combined physiological state, eye movement, environmental conditions, and vehicular dynamics with static encoders and transfer learning, using GNSS sensors, HD cameras, heart-rate monitors, and eye trackers for 35 drivers and 469,830 risk units. It achieved 94.08% accuracy, 93.22% alert rate, 98.61% prediction accuracy, and 10.33% false alarms, surpassing existing models in real-life scenarios. In [20], a hybrid Bi-LSTM with static and dynamic features trained on 27,057 intercity-bus trips from more than 300 drivers exhibited a weighted F1-score of 0.932 and an F1 of 0.728 for high-risk prediction on long-distance routes in Taiwan, an improvement of up to 9.3 times the logistic regression baseline. The hybrid FMEA plus IF-MARCOS framework in [21] prioritizes risks for young drivers under six criteria (severity, occurrence, detection, cost, time, and scope). Among 17 risks elicited from five experts, risky driving and reverse driving register the highest hybrid RPN values (7,286.4 and 6,912.3), while speeding registers the lowest value (3,021.7). In [22], city-scale telematics was used as a proxy road-safety signal on 4,500 vehicles and 1.3 million crashes in New York and showed moderate Spearman correlations for braking and acceleration with accident rates up to 0.56. Radar validation supports the use of behavior signals to flag high-risk areas beyond collision histories alone [22]. The event-classification pipeline proposed in [23] used smartphone and OBD-II vibration and acceleration data on two Puebla routes with AQ-1 sensors, utilizing Random Forest. It outperformed rule-based methods and enabled real-time deployment in dynamic urban settings, achieving 80% accuracy, 75% recall, 0.77 F1-score, 91.9% overall accuracy, and 0.946 AUC [23]. The smartphone-based crash reporting, designed for low-resource contexts in [24], addresses critical limitations of phone-call-based reporting observed in a Tanzania survey of 193 participants, where 62% of the reports were calls that produced 43% delays and 38% inaccurate locations. A mobile application leveraging GPS, accelerometers, and cameras can mitigate these issues and can be extended with automatic detection via continuous inertial analysis [24]. At the infrastructure level, a scalable IoT architecture for incident reporting (InciComm) was proposed in [25]. The model integrates distance, gyroscope, and fire sensor measurements on a Raspberry Pi, demonstrating provisioning of 200 devices per second and latency below 40 ms in simulated and real tests. It indicates that embedded sensors and microservices can support automated, low-latency monitoring pipelines complementary to behavior-classification research.

The present work proposes a web application that estimates individual driving risk from acceleration and turning data in order to enable usage-adjusted pricing calibrated to each driver. The objective is to deliver deployable, low-cost, and explainable scoring that operates on readily available motion signals and integrates seamlessly into pricing workflows. The contributions of this paper consist of a lightweight pipeline for feature extraction and risk scoring from motion data, a comparative analysis of fairness and accuracy trade-offs against actuarial and machine learning baselines, and a discussion of real-time and operational constraints in UBI pricing. By focusing on simple and ubiquitous signals and on web-native delivery, the approach addresses limitations in real-time applicability, cost, and transparency while aligning premiums more closely with actual driving behavior.

II. SYSTEM DESIGN

A. Technological Architecture

The driving risk event monitoring system was designed with a distributed architecture, accessible through a public web interface. User requests are initially managed by Server 1 (a virtual instance), running on an NGINX server configured as a reverse proxy and load balancer that distributes incoming traffic to the processing nodes. The functional core runs on Servers 2 and 3, where NGINX fronts the application processes (Gunicorn for Flask and Node.js managed by PM2), acting as a reverse proxy. These servers host two main back-end services:

- A service built with Node.js that manages authentication, authorization, JWT issuance, and role-based access control.
- A Python/Flask service responsible for processing telematics data. This API runs a Random Forest-based machine learning model for driving-event classification.

The front end uses modern web technologies and is deployed on Servers 2 and 3. This interface allows users to visualize, in real time, the classified events and generated alerts, and access historical data in an interactive and responsive way.

Information is stored on Server 4, which hosts an SQLite database that centralizes all operational data. Telematics records, classification results, and audit logs are stored there. Because SQLite is designed for a single writer with multiple readers, each service commits only brief transactions, preventing long-held locks and ensuring traceability, data persistence, and reliable decision support.

B. Model

For the automatic detection and classification of driving events, a model based on the Random Forest algorithm was implemented. This model was trained using temporal sequences constructed from telematics data obtained with accelerometer and gyroscope sensors.

1) Preprocessing

Once the data were collected, the most relevant variables for the analysis were selected: accelerations and angular velocities on the X, Y, and Z axes (ax, ay, az and wx, wy, wz). These variables are shown in Figure 2.

The class labels (type) were numerically encoded using label encoding through scikit-learn's LabelEncoder, and the encoder was persisted for use in the inference phase (Table I).

TABLE I. EVENT TYPE AND DESCRIPTION

Type	Description
0	Normal driving
1	Sudden acceleration
2	Sharp left turn
3	Sharp right turn
4	Harsh braking

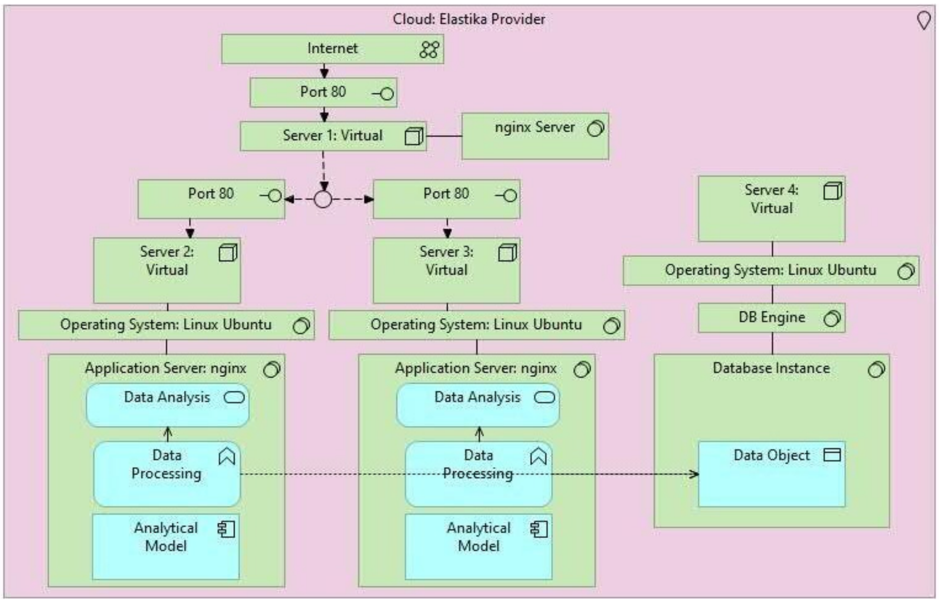

Fig. 1. Architecture of the proposed system.

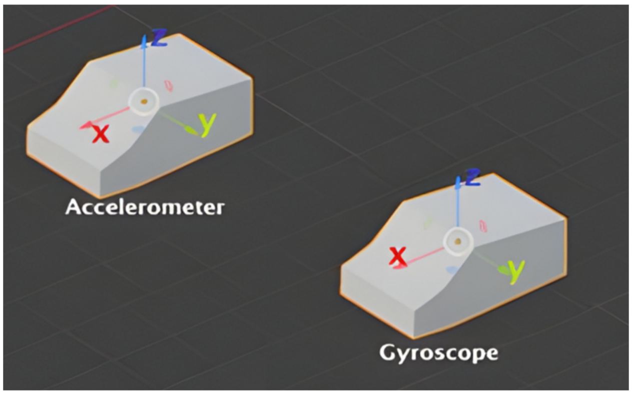

Fig. 2. Orientation of the sensor axes of the accelerometer and gyroscope.

2) Sequence Windowing

The present work used fixed-length sliding windows of 3 s, i.e., six consecutive samples at 2 Hz, yielding 36 features per instance (6 samples × 6 variables). The class of each sequence is the label of its last sample. This policy was applied identically to both sources since both are recorded at 2 Hz, which was selected to balance deployability and signal sufficiency. Multi-second turning, braking, and acceleration patterns remained detectable while reducing on-device and server-side load for real-time use.

3) Training and Validation

Non-overlapping, group-aware splits were utilized for both sources. For the Lima subset, a route-disjoint split was created with 70% of the routes for training, 15% for validation, and 15% for testing. For the public dataset [26], a file/session-level split was created with the same 70/15/15 proportions. Then, the final datasets were built by combining the corresponding parts from both sources: the training, validation, and testing sets include the Lima routes along with the [26] respective files/sessions. No Lima route or [26] file/session appears in more than one split. The classifier is a Random Forest with 100 trees, using scikit-learn defaults. Experiments were run with Python 3.11 and scikit-learn 1.4.2, fixing the random seed at 42 for reproducibility.

4) Evaluation

Model performance was evaluated on the held-out test set using overall accuracy and per-class precision and recall. These metrics are informative under class imbalance and enable an event-wise analysis of detection quality for both non-risky and risky maneuvers (Table III).

5) Model Persistence

Both the trained model and the label encoder were saved to disk using the joblib library. This allows its direct integration with the classification API developed in Flask, enabling real-time inference of driving events within the monitoring system.

6) Dataset

The experiments were carried out with two main sources: The public "Driving Behavior Dataset" [26] (Samsung Galaxy S10 on a Dacia Sandero 1.4 MPI) and our own Lima subset (iPhone 16 on a Kia Seltos), with records of tri-axial accelerometer and gyroscope signals at the same cadence. This study, therefore, applied a single preprocessing policy across sources: fixed-length windows of 3 s (6 samples at 2 Hz) built from consecutive samples without resampling. This alignment avoids domain shift from heterogeneous sampling rates and keeps feature construction identical in both datasets. The current study used raw signals from [26] solely for training/validation/testing under our leakage-safe splits. In addition, all metrics reported in this paper are computed by the authors on those splits. Results, figures, or performance numbers from [26] were not reused.

On the first route, low-speed driving was captured to record only normal (non-risky) events. The journey started at the intersection of Jirón Moquegua and Plaza 2 de Mayo and ended at the intersection of Sacho de Rivera and Avenida Alfonso Ugarte (Figure 3). On the second route, the vehicle was driven at normal speed and only routine maneuvers were performed. This route started at the intersection of Avenida Alejandro Velasco Astete and Avenida Alfredo Benavides and ended at the intersection with Avenida Aviación (Figure 4). On the third route, we again drove at normal speed and recorded only routine maneuvers. This route was driven from Colegio Emblemático Mercedes Cabello de Carbonera to Rímac 15333 (Figure 5).

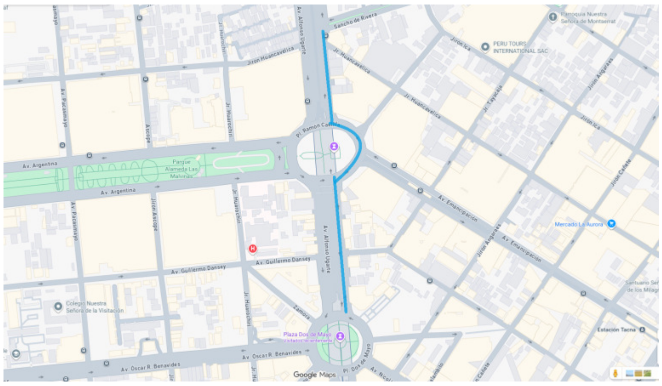

Fig. 3. First driving route.

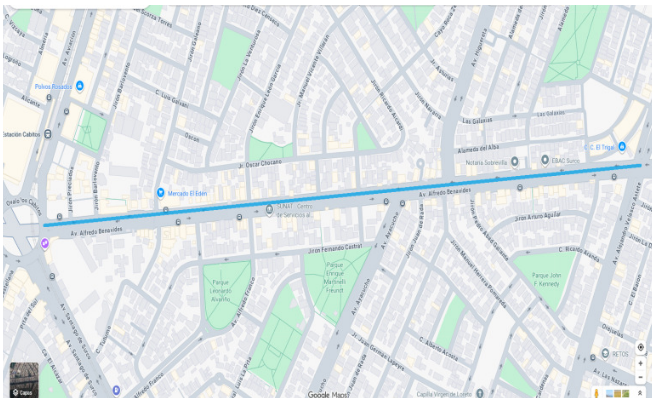

Fig. 4. Second driving route.

Fig. 5. Third driving route.

C. Indicators

The performance of the classification model was evaluated using accuracy and precision. These metrics enable analyzing the system's ability to correctly identify driving events, especially those associated with risky behaviors (Table II).

TABLE II. PERFORMANCE INDICATORS

Accuracy	Proportion of correct predictions over all cases.	$\frac{TP+TN}{TP+TN+FP+FN}$
Recall	Proportion of actual positive events that are correctly detected.	$\frac{TP}{TP+FN}$

TP: True Positives, TN: True Negatives, FP: False Positives, FN: False Negatives

D. Interfaces

The system has a modern web interface, developed with React, Next.js, and Tailwind CSS, which allows users to interact with it from conventional browsers. From this interface, it is possible to view classified events, access alerts generated by the model, and consult the record history in real time through dynamic and intuitive visual components. Access is protected by token-based authentication, which is managed by the user service deployed in Node.js. Once authenticated, users can navigate different sections of the system depending on their role. The interaction with the classified and evaluated data is done transparently through requests to a RESTful API implemented in Flask, which processes new data streams and returns classification and risk assessment results in real time. The interface structure was designed to integrate efficiently with the system's logic, ensuring a smooth, safe, and consistent experience for both technical users and operators or supervisors in the driving environment. The system's main interface allows real-time viewing of recorded events, as well as access to the key functionalities of the application, all through a responsive and easy-to-use design (Figure 6).

Figure 7 illustrates the graphical interface of the individual profile view of the customer. This screen consolidates key information related to recorded driving events, structured into different sections for easy visual analysis.

Fig. 6. Main screen of the application.

Fig. 7. Model processing screen with graphs.

At the top, there is a summary of the customer's data, including name, ID, license plate, status, and dates of registration and update. Below is a bar graph that summarizes the number of events classified on an annual basis, differentiating between risk and non-risk events.

The bottom section provides a detailed analysis for the selected month. It includes the risk level for the month, a breakdown of the types of risky events detected (such as sharp turns or hard braking), the cumulative risk score (represented as the percentage of risky events out of the total), and a ring-chart visualization that shows the ratio of risky to normal events. This presentation facilitates the immediate interpretation of the driver's behavior in specific periods.

III. EXPERIMENTAL RESULTS

To evaluate the performance of the proposed system, this study designed an experimental strategy based on public and real-world telematics data. The Random Forest model was trained using sequences generated from accelerometer and gyroscope signals, captured at a frequency of 2 Hz. A 3 s sliding window was used, and each sequence was labeled with the value corresponding to the last point of the window. The data were divided into 70% for training, 15% for validation, and 15% for testing, ensuring that samples from the same route (Lima) or file/session ([26]) were not mixed across subsets, thus preventing leakage. Figure 8 depicts the confusion matrix on the test set and Table III summarizes the obtained results.

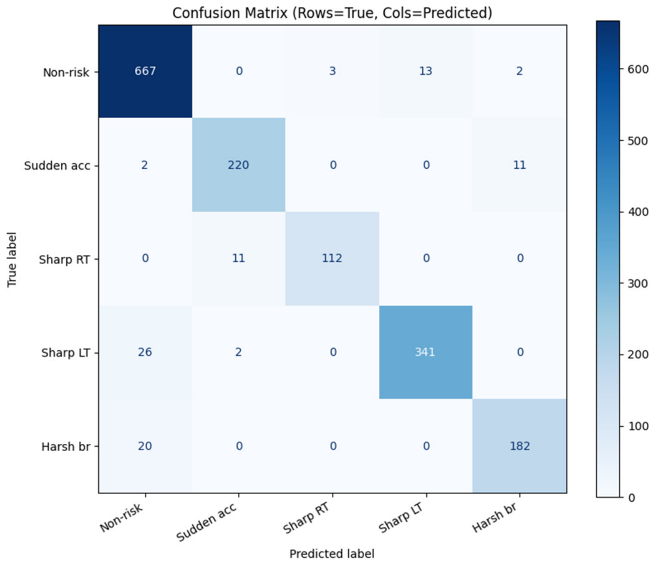

Fig. 8. Confusion matrix on the test set.

TABLE III. MODEL PERFORMANCE PER EVENT

Event Type	Precision (%)	Recall (%)
Non-risky event	93.29%	97.37%
Sudden acceleration	94.42%	94.42%
Sharp right turn	97.39%	91.06%
Sharp left turn	96.33%	92.41%
Harsh braking	93.33%	90.10%
Overall Accuracy	94.42%	

The best results were obtained for non-risky events and sudden accelerations, where the model demonstrated an adequate balance between positive detection and reduction of false positives. Performance was also consistent in complex maneuvers such as sharp turns, with high accuracy and recall rates. The result analysis shows that:

- The model effectively distinguishes between normal and risky maneuvers, with a low margin of error.
- Harsh braking exhibited the lowest recall, likely due to its short duration within the 3 s window at 2 Hz. Future work will evaluate shorter windows (e.g., 2 s), class weighting, and derived features (e.g., jerk) to reduce false negatives for brief braking events.
- The temporal sliding window strategy proved effective in capturing dynamic behavior, without the need for complex derivative features.
- The combined use of collected and publicly available data improved the generalizability of the model. However, it also increased the noise level in the data, which reinforces Random Forest as the chosen solution due to its robustness against nonlinear and noisy patterns.

The overall accuracy of 94.42% indicates reliable labeling of multi-second windows for stable monthly summaries. The balanced performance for normal and sudden acceleration yields robust normal/risky ratios and clear identification of aggressive throttle events for coaching and PHYD adjustments.

IV. DISCUSSION

The proposed model achieves 94.42% overall accuracy with per class recall of at least 90%. Misses concentrate in brief harsh braking episodes, which is consistent with the 3 s window at 2 Hz: short negative longitudinal spikes can be diluted across the window. In contrast, sustained lateral and longitudinal patterns (sharp turns, accelerations) are captured with higher accuracy and recall, enabling stable detection of recurring behaviors.

At an operational level, each event class carries a distinct meaning and action path. Sharp left/right events, driven by lateral acceleration peaks with yaw (wz), serve as proxies for cornering discipline and stability. Repeated detections in the same route segment across trips indicate an intersection conflict risk. The current study identifies segments by aligning trips on start time and elapsed time, and clustering similar IMU patterns (ax/ay/az, wx/wy/wz). These patterns also suggest curb strikes or rollover exposure in vehicles with a high center of gravity, motivating route-level hotspot flags, targeted coaching that says "enter slowly and exit smoothly," and context-aware premium adjustments near sensitive areas such as school zone corners. Harsh braking, a negative longitudinal spike, often reflects tailgating or late response to traffic flow; combined with forward collision alerts (where available), it serves as a near-miss indicator and supports claim triage (timestamped braking before a reported crash) and maintenance risk tracking (brake and tire wear). Sudden acceleration, a positive longitudinal spike, often indicates aggressiveness and potential traction loss, particularly on wet segments, informing eco-driving coaching, traction-aware warnings, and fraud analytics when low-speed claims conflict with high-g traces. Normal driving provides the exposure baseline to normalize risk per time or distance in a way that careful drivers with high mileage are not penalized.

Relative to prior work, deep models, such as TabNet, convolutional and recurrent networks, and multisensor fusion, report strong accuracy but typically assume richer inputs and higher computation costs [8, 9, 19, 20]. The proposed Random Forest pipeline with simple motion signals (accelerometer and gyroscope at 2 Hz) trades a small amount of peak recall for interpretability, low cost, and rapid deployment. This balance remains competitive on accuracy and recall while easing operational adoption.

These characteristics enable three immediate uses in usage-based insurance and operations: monthly risk scores tied to concrete behaviors rather than broad proxies [6], segment level hotspot maps for targeted road or policy interventions, which can be georeferenced when GNSS is available and are consistent with city scale telematics used as a surrogate safety signal [22], and fairer PHYD adjustments that reward consistent normal driving while coaching specific risky patterns (turning, braking, acceleration), avoiding blunt mileage-only penalties [4, 6].

Limitations temper these findings. Temporal resolution (2 Hz, 3 s windows) can underdetect very brief brake taps. The Lima route selection and mixing with a public dataset may introduce distribution shifts. Labels encode the event type but not its severity, which can compress risk gradients. Future work will explore shorter or adaptive windows (for example 2 s) and derived features (jerk, lateral stability indices) to improve harsh braking recall, cost sensitive training or per class thresholds to reduce misses on brief events, simple contextual signals (rain proxies, time of day) and severity scoring to refine pricing and alerting, and streaming inference with debounce logic to trigger near real time coaching without inflating false positives.

V. CONCLUSION

This study presents the development of a comprehensive system, which integrates an efficient distributed architecture, a robust machine learning model, and a web interface focused on user experience, for the monitoring and classification of driving events. The addition of a load-balancing layer using NGINX, along with the use of a reverse proxy, allows for scalable and orderly management of the flow of requests. In addition, services deployed in Node.js for authentication and in Python/Flask for data processing ensure both system protection and real-time execution of the predictive model. The choice of SQLite as the central repository has proved suitable for the experimental environment, providing traceability, persistence, and fast access to telematics data, without introducing bottlenecks in concurrent operations.

From a supervised learning perspective, the trained Random Forest-based model, using sequences created with sliding windows of six records, showed outstanding performance in all classes analyzed. The accuracy and recall obtained values are all above 90%, confirming the proposed model's effectiveness. In addition, the incorporation of the classifier together with the label encoder, packaged with joblib and exposed through a RESTful API, allowed for quick and low-latency implementation, a key aspect in environments where immediate response is required for accident prevention.

As future work, this study proposes expanding the feature set by incorporating variables derived from vehicle dynamics, such as jerk, kinetic energy, or lateral-stability metrics. Likewise, it is proposed to explore hybrid architectures that integrate models such as recurrent neural networks, with the aim of capturing temporal patterns of greater duration and complexity.

Compared to works that rely on high-rate, multi-sensor inputs (e.g., cameras/physiology) and deep models (e.g., TabNet/CNN/LSTM or MTAFN), this study's contribution is a deployable, web-native, and explainable pipeline that attains 94% overall accuracy with per-class recall higher than 90%, using only low-rate (2 Hz) smartphone-IMU signals, sliding-window features, and a Random Forest classifier. This design reduces cost and latency, eases transparency for pricing and operations, and still yields actionable outputs (segment-level hotspots, targeted coaching, and PHYD adjustments). As far as is known, only a few studies report end-to-end system integration (data collection, API inference, interactive dashboards) with this level of simplicity and performance.

ACKNOWLEDGMENTS

The authors are grateful to the Dirección de Investigación of the Universidad Peruana de Ciencias Aplicadas (UPC) for the support provided for this research work through the UPC-EXPOST-2025-2 incentive.

REFERENCES

[1] S. Martin, "The True Cost of Auto Insurance in 2025," *Bankrate*. https://www.bankrate.com/insurance/car/the-true-cost-of-auto-insurance/.

[2] P. Gusner, "7 Best Car Insurance Companies Of December 2025," *Forbes Advisor*, Dec. 05, 2025. https://www.forbes.com/advisor/car-insurance/best-car-insurance-companies/.

[3] J. R. Torra *et al.*, "Weather Conditions and Telematics Panel Data in Monthly Motor Insurance Claim Frequency Models," *Risks*, vol. 11, no. 3, Mar. 2023, https://doi.org/10.3390/risks11030057.

[4] J. Boylan, D. Meyer, and W. S. Chen, "A systematic review of the use of in-vehicle telematics in monitoring driving behaviours," *Accident; Analysis and Prevention*, vol. 199, Dec. 2024, Art. no. 107519, https://doi.org/10.1016/j.aap.2024.107519.

[5] S. Xie, K. Shi, S. Xie, and K. Shi, "Generalised Additive Modelling of Auto Insurance Data with Territory Design: A Rate Regulation Perspective," *Mathematics*, vol. 11, no. 2, Jan. 2023, https://doi.org/10.3390/math11020334.

[6] W. Nai *et al.*, "A Comprehensive Review of Driving Style Evaluation Approaches and Product Designs Applied to Vehicle Usage-Based Insurance," *Sustainability*, vol. 14, no. 13, Jun. 2022, https://doi.org/10.3390/su14137705.

[7] K. Jõeveer and K. Kepp, "What drives drivers? Switching, learning, and the impact of claims in car insurance," *Journal of Behavioral and Experimental Economics*, vol. 103, Apr. 2023, Art. no. 101993, https://doi.org/10.1016/j.socec.2023.101993.

[8] K. McDonnell, F. Murphy, B. Sheehan, L. Masello, and G. Castignani, "Deep learning in insurance: Accuracy and model interpretability using TabNet," *Expert Systems with Applications*, vol. 217, May 2023, Art. no. 119543, https://doi.org/10.1016/j.eswa.2023.119543.

[9] W. A. Al-Hussein *et al.*, "Driver Behavior Profiling and Recognition Using Deep-Learning Methods: In Accordance with Traffic Regulations and Experts Guidelines," *International Journal of Environmental Research and Public Health*, vol. 19, no. 3, Jan. 2022, https://doi.org/10.3390/ijerph19031470.

[10] S. Fernandez *et al.*, "Fuzzy Ontology-Based System for Driver Behavior Classification," *Sensors*, vol. 22, no. 20, Oct. 2022, Art. no. 7954, https://doi.org/10.3390/s22207954.

[11] Md. A. Uddin *et al.*, "Abnormal Driving Behavior Detection: A Machine and Deep Learning Based Hybrid Model," *International Journal of Intelligent Transportation Systems Research*, vol. 23, no. 1, pp. 568–591, Apr. 2025, https://doi.org/10.1007/s13177-025-00471-2.

[12] A. S. Afridi, A. Kafy, Ms. N. N. Moon, and Md. S. Shakil, "A multi-class driver behavior dataset for real-time detection and road safety enhancement," *Data in Brief*, vol. 60, Jun. 2025, Art. no. 111529, https://doi.org/10.1016/j.dib.2025.111529.

[13] S. Jia, J. Xu, S. Wang, and X. Liu, "Connected multi-vehicle crash risk assessment considering probability and intensity," *PLOS ONE*, vol. 20, no. 2, 2025, Art. no. e0313317, https://doi.org/10.1371/journal.pone.0313317.

[14] X. Wang *et al.*, "An Autonomous Vehicle Behavior Decision Method Based on Deep Reinforcement Learning with Hybrid State Space and Driving Risk," *Sensors*, vol. 25, no. 3, Jan. 2025, Art. no. 774, https://doi.org/10.3390/s25030774.

[15] L. Zhu, Z. Zhang, J. Yu, Y. Zhang, and J. Fu, "A real-time synthesized driving risk quantification model based on driver risk perception-response mechanism," *Transportation Research Part C: Emerging Technologies*, vol. 174, May 2025, Art. no. 105073, https://doi.org/10.1016/j.trc.2025.105073.

[16] E. Papazikou, R. Talbot, L. Brown, S. Maynard, and A. Filtness, "Unpacking the relationship between task complexity and driving risk: Insights from a UK on-road trial," *IATSS Research*, vol. 49, no. 2, pp. 127–136, Jul. 2025, https://doi.org/10.1016/j.iatssr.2025.03.003.

[17] S. Qi and A. Zheng, "A vehicle conflict risk identification method based on an improved intelligent driver model," *Applied Sciences*, vol. 15, no. 6, Mar. 2025, Art. no. 3240, https://doi.org/10.3390/app15063240.

[18] X. Ma *et al.*, "Assessing the impact of car-following driving style on traffic conflict risk using asymmetric behavior model and explainable machine learning," *Accident Analysis & Prevention*, vol. 211, Mar. 2025, Art. no. 107904, https://doi.org/10.1016/j.aap.2024.107904.

[19] C. Bai *et al.*, "Multi-source temporal attention fusion network (MTAFN) for driving risk assessment based on naturalistic driving data," *Expert Systems with Applications*, vol. 279, Jun. 2025, Art. no. 127502, https://doi.org/10.1016/j.eswa.2025.127502.

[20] W.-H. Lee and C.-Y. Chang, "Assessing Driving Risk Level: Harnessing Deep Learning Hybrid Model With Intercity Bus Naturalistic Driving Data," *IEEE Access*, vol. 13, pp. 25141–25153, 2025, https://doi.org/10.1109/ACCESS.2025.3538693.

[21] S. Jafarzadeh Ghoushchi, S. Shaffiee Haghshenas, S. Vahabzadeh, S. Shaffiee Haghshenas, V. Astarita, and G. Guido, "Risk assessment of young driver behavior using an extended decision-making approach based on FMEA in uncertain environments," *Neural Computing and Applications*, vol. 37, no. 16, pp. 9565–9602, Jun. 2025, https://doi.org/10.1007/s00521-025-11087-8.

[22] P. Alrassy, A. W. Smyth, and J. Jang, "Driver behavior indices from large-scale fleet telematics data as surrogate safety measures," *Accident Analysis & Prevention*, vol. 179, Jan. 2023, Art. no. 106879, https://doi.org/10.1016/j.aap.2022.106879.

[23] A. Aguilar-González, A. M. Santiago, A. Aguilar-González, and A. M. Santiago, "Road Event Detection and Classification Algorithm Using Vibration and Acceleration Data," *Algorithms*, vol. 18, no. 3, Feb. 2025, Art. no. 127, https://doi.org/10.3390/a18030127.

[24] I. J. Mrema and M. A. Dida, "A Survey of Road Accident Reporting and Driver's Behavior Awareness Systems: The Case of Tanzania," *Engineering, Technology & Applied Science Research*, vol. 10, no. 4, pp. 6009–6015, Aug. 2020, https://doi.org/10.48084/etasr.3449.

[25] S. M. A. Ghaly and M. A. Kadampur, "Scalable Incident Reporting Framework: A Sensor and IoT Research," *Engineering, Technology & Applied Science Research*, vol. 13, no. 3, pp. 10748–10753, Jun. 2023, https://doi.org/10.48084/etasr.5820.

[26] I. Cojocaru, S. Popescu, and C. Mihaescu, "Driving Behavior." kaggle, [Online]. Available: https://www.kaggle.com/datasets/outofskills/driving-behavior.

AUTHOR PROFILES

Michael Andre Samuel Cuadros-Ccahuana is an Information Systems Engineer at the Peruvian University of Applied Sciences (UPC) in Lima, Peru (email: U201713406@upc.edu.pe).

Kevin Enrique Servat-Farfán is an Information Systems Engineer at the Peruvian University of Applied Sciences (UPC) in Lima, Peru (email: U202010146@upc.edu.pe).

Pedro Castañeda is a RENACYT researcher and holds a PhD in Systems Engineering, a master's degree in management and information technology management from UNMSM, and an MBA master's degree from ESAN. He has completed doctoral studies in Public Policy and State Management at the Centro de Altos Estudios Nacionales (CAEN). He leads e-brokerage projects, software development and process improvement, using agile and traditional methodologies. He has the following certifications: Project Management Professional (PMP), Scrum Certified Developer (CSD), IBM Certified Professional in Rational Unified Process, and ORACLE Certifications. Areas of Interest: Artificial Intelligence, Software Productivity, Business Intelligence, Data Analytics, Machine Learning, Software Engineering (email: pedro.castaneda@untrm.edu.pe, ORCID: https://orcid.org/0000-0003-1865-1293).

Juan Mansilla-Lopez received a bachelor's degree in systems engineering from the Universidad de Lima in 1997 and a master's degree in finance from Universidad ESAN in 2011. Since 2022, he has been the coordinator of the Information Systems Engineering program at the Universidad Peruana de Ciencias Aplicadas. His research interests include artificial intelligence, internet of things, finance and stock markets (email: pcsijman@upc.edu.pe, ORCID: https://orcid.org/0000-0003-0039-6044).

Alberto Daniel García-Núñez is a student in Technology and Innovation Management (UPB) and holeds a Master in Information Technology Management (ITESM) (email: alberto.garcia@upb.edu.co, ORCID: https://orcid.org/0000-0002-9402-3785).