UNSUPERVISED DETECTION OF HARSH CORNERING BEHAVIOR USING SMARTPHONE-BASED TELEMATICS AND INFRASTRUCTURE DATA

Aristotelis Tsoutsanis, M.Sc. Ph.D Candidate, Corresponding Author
Department of Transportation Planning and Engineering
National Technical University of Athens, Athens, Greece, 15772
OSeven Telematics
Chaimanta 27B, Athens, Greece, 15234
a_tsoutsanis@mail.ntua.gr, atsoutsanis@oseven.io

Apostolos Ziakopoulos, Ph.D
Department of Transportation Planning and Engineering
National Technical University of Athens, Athens, Greece, 15772
apziak@central.ntua.gr

Zoe Vasileiou, M.Sc.
OSeven Telematics
Chaimanta 27B, Athens, Greece, 15234
zvasileiou@oseven.io

George Yannis, Professor
Department of Transportation Planning and Engineering
National Technical University of Athens, Athens, Greece, 15772
geyannis@central.ntua.gr

Word Count: 4480 words + 2 table(s) × 250 = 4980 words

Submission Date: January 8, 2026

ABSTRACT

This study presents an unsupervised approach for detecting harsh cornering behavior using smartphone-based telematics enriched with infrastructure data. Unlike traditional methods that require fixed device orientation or labeled datasets, the proposed framework combines GPS-derived yaw and heading changes, orientation-invariant inertial sensor magnitudes, and spatial validation using OpenStreetMap to isolate turning maneuvers with safety relevance. A preprocessing pipeline identifies candidate turns based on angular thresholds and proximity to mapped intersections. A multivariate feature set is extracted from time-series windows around each turn, capturing translational and rotational dynamics. Outlier detection is performed using DBSCAN, with hyperparameters tuned via k-distance, and complemented by Isolation Forest to enhance robustness. The ensemble of detections shows strong separability when evaluated through supervised classifiers, achieving ROC-AUC scores above 0.98. Visual case studies further confirm the accuracy and interpretability of the identified events. This orientation-agnostic pipeline operates without labeled data, making it highly suitable for real-world deployment in driver behavior analysis, risk profiling, and road safety analytics. The study highlights the feasibility of combining sensor fusion and map-based filtering to detect unsafe driving behaviors at scale.

Keywords: Smartphone telematics, harsh cornering detection, unsupervised learning, sensor fusion, OpenStreetMap, driver behavior, road safety.

INTRODUCTION

Road safety data sources are essential for understanding the causes of crashes and predicting future risks to reduce road traffic accidents. Various factors, including driving speed, road infrastructure, and driver behavior, influence crashes. However, crashes are relatively rare, making it challenging to detect and assess road safety conditions in real-time. To address this challenge, researchers have employed crash surrogates in road safety studies in order to define crash proneness and associate it with crash risk (Gettman et al., Nikolaou et al. (*1, 2*)).

The increasing adoption of smartphone-based telematics offers a scalable and cost-effective alternative for monitoring driver behavior (Mantouka et al. (*3*)). Smartphones, equipped with sensors such as accelerometers, gyroscopes, and GPS, generate rich datasets for analyzing vehicle dynamics. High G-force events, characterized by vehicle acceleration exceeding expected levels, typically result from hard braking or abrupt acceleration maneuvers. These events are widely used by researchers to evaluate driver behavior and identify high-risk drivers (Simons-Morton et al. (*4*)).

Previous studies, such as Liu et al. (*5*), have explored smartphone-based detection of hard-braking events to enhance road safety services, demonstrating the potential of mobile sensors in real-world driving scenarios. Additionally, the use of gyroscopes in mobile applications, as discussed in Barthold et al. (*6*), has shown promise in capturing rotational motion, making it a valuable tool for detecting cornering events.

The accurate classification of vehicle cornering behaviors (e.g., normal vs. harsh turns) is crucial for a wide range of applications, including driver behavior analysis, autonomous driving systems, and vehicle safety enhancement. Identifying and mitigating harsh driving maneuvers is particularly significant, as these actions can lead to severe vehicle instability. For example, while rollover incidents constitute only 3% of all vehicle crashes, they account for nearly one-third of passenger fatalities, highlighting the critical risks associated with extreme driving events (Padmanaban and Husher (*7*)). The ability to detect and analyze these behaviors in real-time could play a pivotal role in reducing fatalities and improving road safety.

While hard braking and rapid acceleration have been extensively studied in the context of smartphone-based telematics, harsh cornering remains comparatively underexplored despite its strong association with vehicle instability and elevated crash risk. Most commercial and academic telematics frameworks focus on longitudinal forces due to their simpler detection using accelerometer data and clearer regulatory implications. However, lateral dynamics during turning are equally critical, particularly in urban or suburban environments where sharp curves and intersections are frequent. Cornering events are inherently more complex to capture, as they depend on rotational motion, trajectory shape, and interaction with road geometry—factors that are difficult to quantify under unknown device orientations. The limited attention given to harsh cornering detection contrasts with its potential importance in crash prevention, driver feedback, and risk-based insurance pricing. This gap highlights the need for robust, orientation-invariant methods capable of identifying aggressive turning behavior in naturalistic driving conditions. Most existing studies that investigate harsh cornering rely on smartphones mounted in fixed orientations within the vehicle, enabling axis-specific thresholding for detecting lateral acceleration or angular velocity (Vlahogianni and Barmpounakis (*8*)). Moreover, a few studies rely on special equipment to align the smartphone to the vehicle's coordinate system (Wang et al., Chen et al. (*9, 10*)). While effective in controlled settings, such methods lack robustness in real-world deployments where the smartphone's position and orientation can vary significantly. To address this limitation, the present study proposes an unsupervised methodology that leverages smartphone sensor and GPS data fused with

infrastructure information from OpenStreetMap. The approach estimates yaw rate over multiple temporal horizons to capture the full cornering profile, normalizes inertial data using vector magnitudes to remain invariant to phone orientation, and filters events based on geospatial validation near road intersections. By employing K-NearestNeighbors to guide hyperparameter tuning and DBSCAN to identify outliers, harsh cornering events are detected without reliance on labeled data or fixed sensor placement. This orientation-agnostic pipeline offers a scalable and flexible solution for real-world driving behavior analysis.

DATA COLLECTION AND PREPROCESSING

This section outlines the procedures followed to prepare the data used for detecting harsh cornering events from smartphone-based telematics. Given the challenges posed by real-world deployment—such as arbitrary device placement, sensor noise, and inconsistent sampling—careful attention was paid to both the selection of meaningful signals and the derivation of robust features. The dataset includes multimodal time-series data recorded during naturalistic driving, and preprocessing steps were designed to extract motion-related cues that can capture turning behavior in a device-agnostic manner.

Data Source

The dataset used in this study originates from OSeven Telematics and consists of anonymized time-series data collected from smartphone sensors at a frequency of 1 Hz. Each record includes GPS-derived location coordinates, speed, and heading (course over ground), as well as readings from embedded inertial sensors such as the accelerometer and gyroscope. The data is captured passively during driving, leveraging consumer-grade smartphones placed arbitrarily inside the vehicle without fixed orientation or mounting. As a result, the dataset reflects realistic deployment conditions, including variability in device orientation, sensor noise, and driving environments. The temporal resolution and multimodal nature of the dataset make it suitable for modeling driving behavior and detecting dynamic events such as harsh cornering.

In total, the dataset after filtering noisy GPS datapoints comprises 4,017 individual trips and 1,758,589 datapoints, representing a large and diverse sample of naturalistic driving behavior across various conditions.

Yaw Rate Estimation

To capture turning dynamics, we extract two features: (a) the maximum yaw rate across short horizons, representing the sharpness of a turn, and (b) the maximum total heading change, representing the overall extent of the maneuver. These are derived from GPS-based heading changes sampled at 1 Hz.

The yaw rate at multiple temporal horizons is computed using Equation (1), which normalizes heading changes over time to account for angular wraparound. The maximum absolute yaw rate across the past four seconds is extracted in Equation (2) to capture the sharpest recent change in heading. Similarly, the total turn angle at each horizon is defined in Equation (3), and the largest absolute heading deviation is summarized in Equation (4). These features together capture both the short-term intensity and the spatial extent of vehicle cornering behavior.

$$\text{yaw\_rate}_n = \frac{((\theta_t - \theta_{t-n} + 180) \bmod 360) - 180}{n \cdot \Delta t}, \quad n = 1,2,3,4 \tag{1}$$

$$\text{max\_yaw\_rate} = \max\left(|\text{yaw\_rate}_1|, |\text{yaw\_rate}_2|, |\text{yaw\_rate}_3|, |\text{yaw\_rate}_4|\right) \tag{2}$$

$$\text{turn\_angle}_n = ((\theta_t - \theta_{t-n} + 180) \bmod 360) - 180, \quad n = 1,2,3,4 \tag{3}$$

$$\text{max\_turn\_angle} = \max\left(|\text{turn\_angle}_1|, |\text{turn\_angle}_2|, |\text{turn\_angle}_3|, |\text{turn\_angle}_4|\right) \tag{4}$$

Map-Based Validation Using OpenStreetMap

After computing the turn angle at each timestamp, we apply a threshold of 70 degrees over the 4 seconds period to identify candidate turning events. This threshold is selected to filter out minor course deviations and emphasize maneuvers that are likely to affect vehicle stability if executed harshly. The resulting candidate events capture high-curvature segments in the trajectory, but they may still include false positives due to GPS noise, lane changes, or sudden heading corrections on straight roads.

To enhance the reliability of turn detection, we introduce a validation step using map data from OpenStreetMap (OSM). The key idea is that genuine cornering events should occur in the vicinity of an actual road intersection. For each candidate turn event, we query the surrounding region within a fixed radius, 15 meters in our case, to check for the presence of an OSM-mapped intersection.

Intersections are extracted using OSM road network data, typically by identifying nodes where multiple road segments meet at an angle, excluding nodes on roundabouts or interchanges. If no such intersection is found near the detected turn point, the event is considered a likely false positive and is discarded from further analysis.

Since the map-based validation is only used to eliminate false positives among the detected cornering events, we avoid building a local road graph for each event—which would be computationally intensive. Instead, we pre-load and index the entire road network of the Attica region in Greece. This area includes a mix of dense urban grids, suburban layouts, and rural roads, offering a realistic and structurally diverse environment. By covering this region in advance, we ensure fast spatial queries for intersection checks while maintaining robustness across varied road geometries.

This spatial filtering step improves precision by enforcing a topological constraint: harsh cornering should correspond not just to sharp heading changes, but to those that align with actual road geometry. As a result, the dataset used in subsequent modeling stages contains a more robust set of turn events grounded in the physical infrastructure of the driving environment.

This ensures that the final dataset focuses on significant cornering maneuvers that are more likely to impact vehicle stability and driver behavior. Figure 1 illustrates the distribution of total turn angles and vehicle speeds across the filtered dataset. As shown, most turns involve angles ranging from 70° to 180°, while vehicle speeds during these maneuvers are typically concentrated below 40 km/h, consistent with expected driving behavior in urban environments.

After validation, the resulting dataset comprises approximately 1,900 high-confidence cornering events in Attica region of Greece, which form the basis for downstream clustering and anomaly detection.

Orientation-Invariant Sensor Features

Following the map-based validation step, the resulting set of candidate turn events was further processed to extract additional features from the inertial sensors. Given that smartphones in real-world

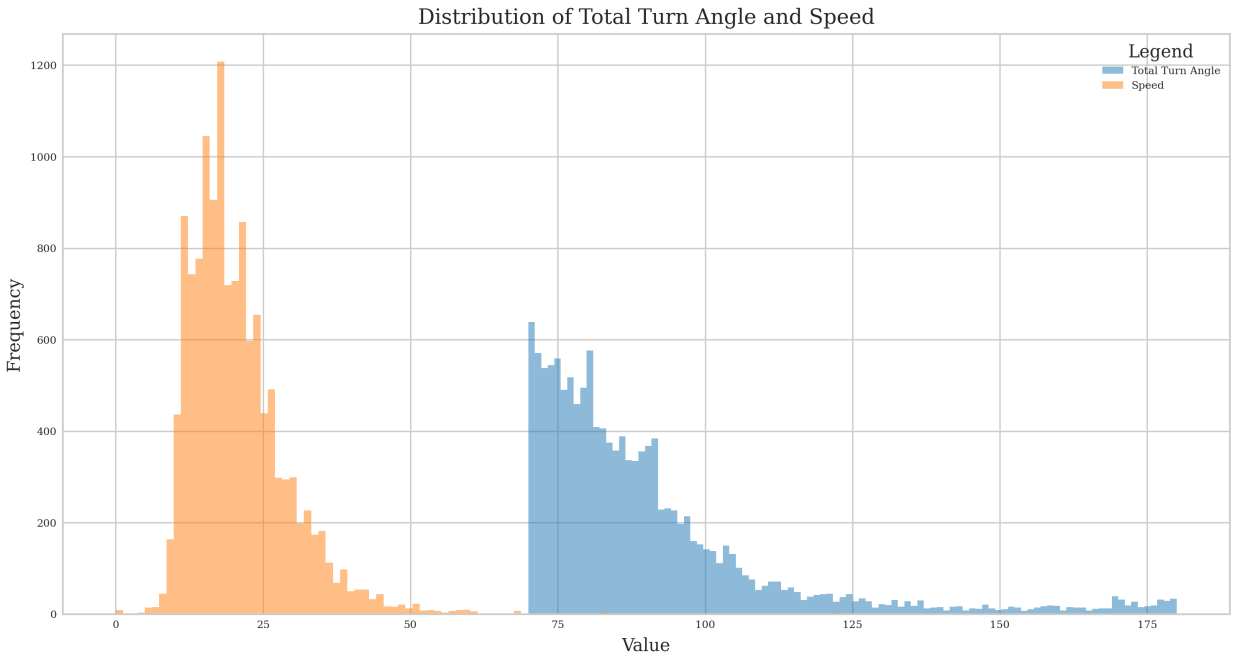

FIGURE 1: The histogram shows the distribution of total turn angles (blue) and vehicle speeds (orange) across all turning events in the dataset.

deployments can be placed arbitrarily inside the vehicle—resulting in unknown and inconsistent orientation—the direct use of individual accelerometer and gyroscope axes (x, y, z) becomes unreliable. To mitigate this, we compute the *magnitude* of the 3-axis accelerometer and gyroscope signals, which is invariant to device orientation and better suited for modeling motion intensity.

The acceleration and gyroscope magnitude are calculated as the Euclidean norm of the three-axis readings at each timestamp t, as shown in Equations (5), (6):

$$\text{accel\_magnitude}_t = \sqrt{a_x(t)^2 + a_y(t)^2 + a_z(t)^2} \tag{5}$$

$$\text{gyro\_magnitude}_t = \sqrt{\omega_x(t)^2 + \omega_y(t)^2 + \omega_z(t)^2} \tag{6}$$

These magnitude features allow us to quantify the overall intensity of linear and angular motion, respectively, without relying on axis-specific orientation. They are particularly useful for identifying sudden vehicle dynamics—such as sharp cornering or abrupt lateral shifts—regardless of how the phone is positioned within the car.

Yaw Rate from Device Orientation Sensors

In addition to estimating yaw rate from GPS-derived heading changes, we also compute a second yaw-related feature using the smartphone's internal orientation sensor, often referred to as `motionYaw`. This value typically comes from a sensor fusion algorithm that combines data from the gyroscope, accelerometer, and magnetometer to provide a stable estimate of the device's yaw (rotation around the vertical axis) relative to the current reference frame.

To make this information comparable across time, we calculate the angular velocity—i.e., the yaw rate—by differentiating the `motionYaw` value over time, taking into account angular

wraparound at 360° boundaries. The formula is similar in structure to the GPS-based yaw rate, but applied to orientation sensor output:

$$\text{yaw\_rate\_motion}_t = \frac{((\psi_t - \psi_{t-1} + 180) \bmod 360) - 180}{\Delta t} \tag{7}$$

where ψ_t denotes the device yaw angle (in degrees) at time t, and Δt is the sampling interval.

Due to the arbitrary phone placement in real-world scenarios, we use this signal as a supplementary feature rather than a sole indicator of cornering.

METHODOLOGY

This section describes the full unsupervised pipeline for detecting harsh cornering events from smartphone-based telematics. It begins with the construction of time-series segments centered on cornering peaks, followed by feature extraction, normalization, and density-based clustering. The goal is to identify outlier turn maneuvers that deviate significantly from typical driving behavior.

Feature Construction

Each cornering event identified by the GPS-based turn angle threshold is further processed to extract a consistent time-series window for feature computation. Specifically, for each record flagged as part of a turn, we identify the *turn peak*—the timestamp where the yaw rate is maximal—and extract a 3-second multivariate time-series window centered on this peak.

This results in a fixed-length sequence of sensor observations spanning 1 second before and 1 second after the peak. Each window includes synchronized values from multiple sensor modalities:

- GPS-derived features: speed, heading
- Accelerometer: 3-axis linear acceleration
- Gyroscope: 3-axis angular velocity
- Device orientation: yaw angle

This segmentation process transforms each turn into a comparable multivariate time-series, which forms the basis for computing robust, orientation-invariant features.

From each time-series segment, we compute a set of 13 statistical and motion-based features that describe the dynamics of the turn. These include speed variations, angular motion intensities, rotational asymmetry, and trajectory properties. Table 1 summarizes the extracted features.

These features collectively encode the intensity, variability, and asymmetry of each cornering maneuver. By capturing both translational and rotational dynamics—including maximum and minimum yaw rates from multiple sources—they offer a rich representation of the vehicle's behavior during turning. The use of magnitudes for inertial sensors ensures robustness to phone orientation, while entry/exit speeds and heading change summarize the broader context of each maneuver.

Unsupervised Detection of Harsh Cornering Events

After feature extraction, we standardize the feature matrix using z-score normalization, ensuring that each feature contributes equally to the distance-based clustering that follows.

We then apply the Density-Based Spatial Clustering of Applications with Noise (DBSCAN) algorithm to identify harsh cornering events (Ester et al. (*11*)). DBSCAN is particularly well-suited for this application because it does not require the number of clusters to be specified in advance, and it explicitly labels as *noise* those points that do not belong to any dense region.

TABLE 1: Description of extracted features for each turn event

Feature	Description
speed_min	Minimum vehicle speed during the turn (km/h)
speed_max	Maximum vehicle speed during the turn (km/h)
entry_speed	Speed at the start of the turn window (km/h)
exit_speed	Speed at the end of the turn window (km/h)
max_speed_change	Maximum speed change observed within the window (km/h)
total_turn_angle_max	Maximum heading change over fixed intervals (degrees)
yaw_rate_max	Max yaw rate computed from GPS heading (°/s)
motion_yaw_rate_deg_max	Max yaw rate from device orientation in degrees/s
motion_yaw_rate_deg_min	Min yaw rate from device orientation in degrees/s
accel_magnitude_min	Minimum acceleration magnitude (m/s^2)
accel_magnitude_max	Maximum acceleration magnitude (m/s^2)
gyro_magnitude_min	Minimum gyroscope magnitude (rad/s)
gyro_magnitude_max	Maximum gyroscope magnitude (rad/s)

In our setting, this "noise" corresponds to abnormal or harsh turning events, under the assumption that most turns follow a typical behavioral pattern. DBSCAN treats these harsh events as statistical outliers in the feature space.

Hyperparameter Tuning via K-Distance Heuristic

DBSCAN requires two hyperparameters: the neighborhood radius ε and the minimum number of neighbors (min_samples) to form a dense cluster. While min_samples is typically set to a fixed value, selecting ε appropriately is crucial for meaningful outlier detection.

To identify an appropriate value for the ε parameter in DBSCAN, we applied the k-distance heuristic, plotting the distance to each point's 50th nearest neighbor after sorting all distances in ascending order as shown in Figure 2. The resulting curve exhibits a characteristic elbow shape, where a steep increase in distance suggests a transition from dense to sparse regions. We selected an ε value of 4.0, which captures the boundary between typical and anomaly turning behaviors. This selection balances precision in detecting dense clusters of regular maneuvers and recall in identifying true outliers representing harsh cornering.

This hybrid approach leverages K-NeirestNeighbors to guide DBSCAN parameter tuning while maintaining the flexibility of density-based clustering to detect rare but significant behaviors.

To assess the internal consistency of clusters formed by DBSCAN, we computed the silhouette score, a common metric for evaluating unsupervised clustering. The silhouette score quantifies how similar each point is to its own cluster compared to other clusters. In our case, DBSCAN achieved a silhouette score of 0.53, indicating moderately well-separated clusters, with outliers forming distinct, sparse regions in the feature space. This confirms the effectiveness of the extracted features and selected ε in capturing cornering behavior patterns.

While unsupervised, this score supports the intuition that regular and harsh cornering events occupy distinct zones in the feature space.

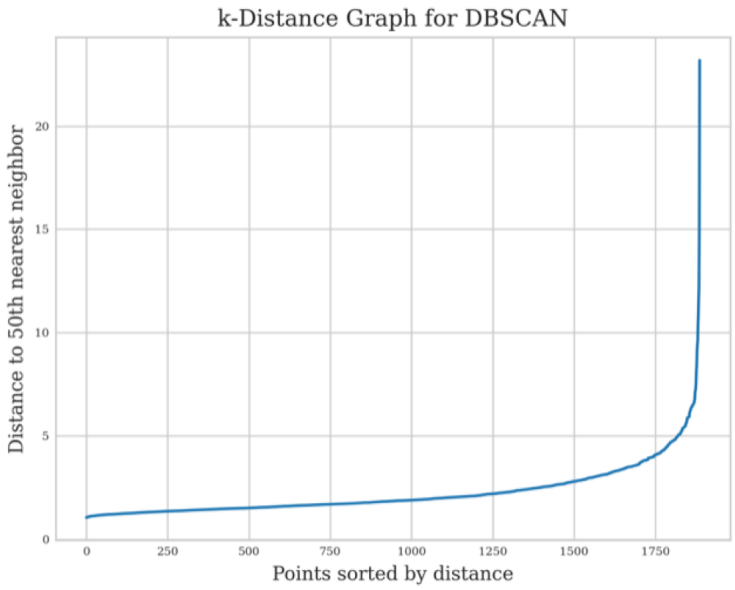

FIGURE 2: k-Distance plot used to determine the optimal ε parameter for DBSCAN. Each point represents the distance to its 50th nearest neighbor, sorted in ascending order

Isolation Forest as a Complementary Detector

In addition to DBSCAN, we evaluated the use of the Isolation Forest algorithm as a model-agnostic unsupervised method for detecting anomalous cornering behaviors. Isolation Forest builds an ensemble of random binary trees, where anomalous points are isolated with fewer splits. We applied the method on the same feature set, using 100 estimators and contamination parameter set to 0.02. Events with the shortest average path lengths were flagged as potential harsh turns.

To further illustrate the characteristics of the harsh events detected by each method, Figure 3 presents a comparison of feature distributions between DBSCAN and Isolation Forest detections. While both methods show elevated values in yaw-related and inertial features compared to normal turns, distinct patterns emerge. DBSCAN-detected events typically exhibit higher turn angles and yaw rates, while Isolation Forest tends to flag events with elevated motion asymmetry and gyroscope magnitude—even at moderate speeds. These differences support the hypothesis that the two methods capture complementary aspects of harsh cornering behavior.

Table 2 summarizes the overlap and divergence between DBSCAN and Isolation Forest in detecting harsh cornering events. Only 50 harsh corner events were detected by both representing nearly 3% of the entire dataset. DBSCAN contributed 22 unique detections, typically involving sharp turns with high yaw rates, while Isolation Forest identified 26 distinct events characterized by subtle but higher speeds. The complementary nature of these results supports the use of an ensemble approach that combines both methods to improve detection robustness.

Validation via Classifier-Based Separability

To assess the internal consistency of the pseudo-labeled harsh cornering events, we evaluated how well the extracted features could separate harsh and normal turns using a supervised learning proxy. Specifically, we trained two binary classifiers—logistic regression and random forest—using the ensemble labels (union of DBSCAN and Isolation Forest detections) as target values.

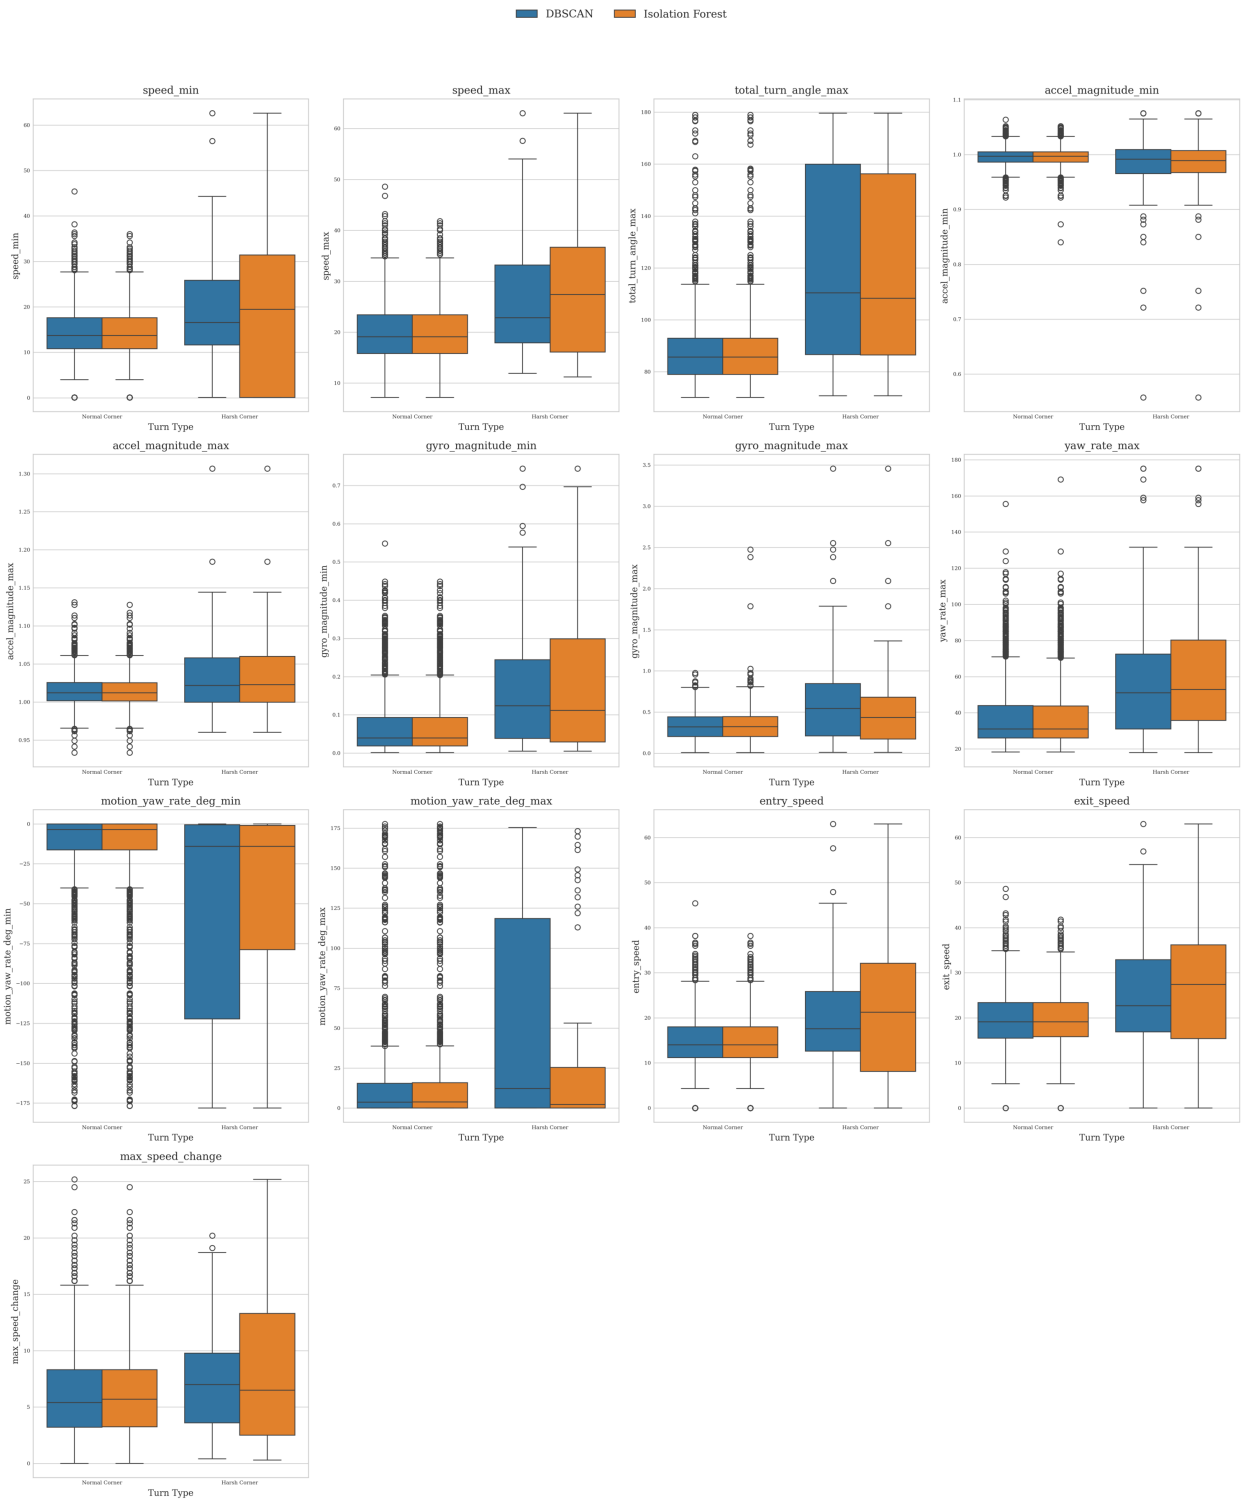

FIGURE 3: Comparison of feature distributions for harsh cornering events detected by DBSCAN (blue) and Isolation Forest (orange). Each boxplot represents one of the 13 extracted features, grouped by detection method.

TABLE 2: Summary of harsh cornering detections by DBSCAN and Isolation Forest.

Method	Harsh Corners Detected	Unique Events
DBSCAN	72	22
Isolation Forest	76	26
Intersection (Both)	50	-

We applied 5-fold cross-validation using ROC-AUC as the evaluation metric. The logistic regression model achieved a mean ROC-AUC of 0.986 ± 0.007, while the random forest classifier reached 0.994 ± 0.003, indicating excellent separability between harsh and normal cornering events in the feature space. These results validate that the extracted features—such as yaw rate, turn angle, acceleration and gyroscope magnitudes—capture the underlying dynamics of risky turning behavior effectively.

To further interpret model behavior, we examined feature importances from the random forest model as shown in Figure 4. The most influential features included maximum gyroscope magnitude, maximum yaw rate and speed confirming their relevance in characterizing high-intensity cornering maneuvers. These findings support the validity of the unsupervised ensemble detection pipeline and highlight the robustness of the feature design in distinguishing abnormal vehicle dynamics.

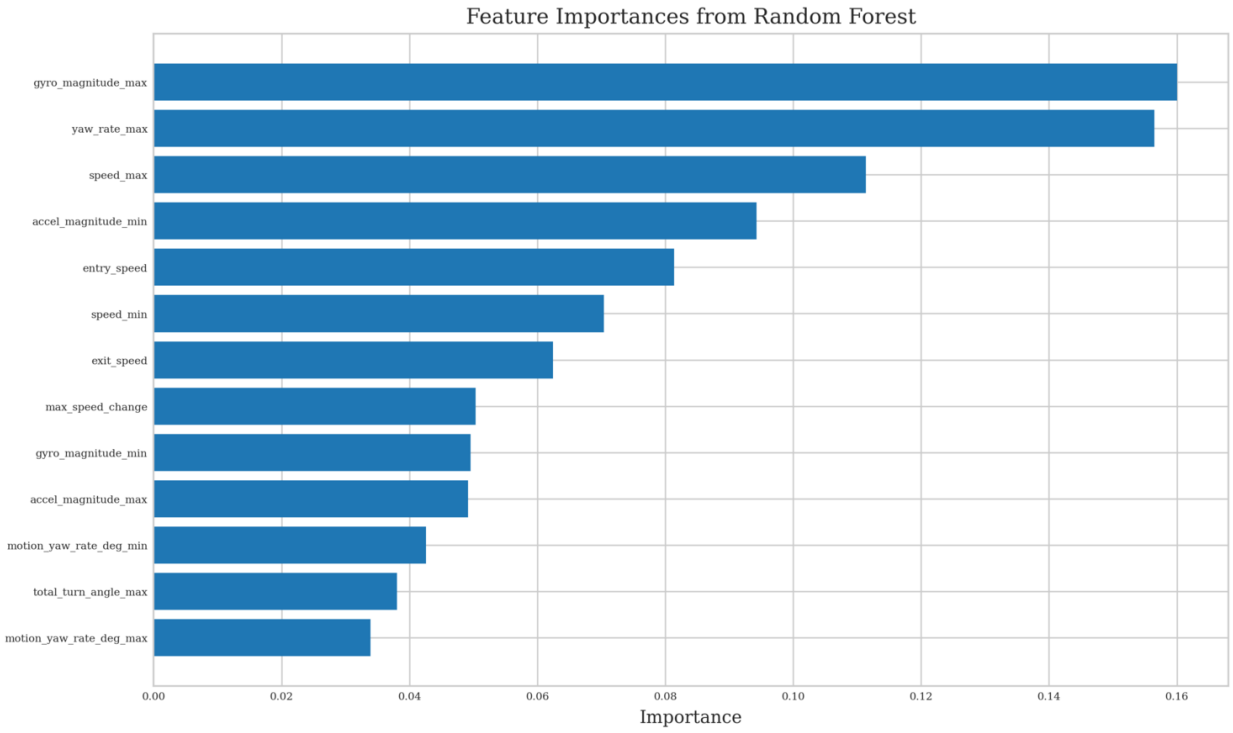

FIGURE 4: Feature importance scores from the random forest classifier trained to separate harsh vs. normal cornering events using ensemble labels.

Qualitative Case Studies and Map-Based Validation

To complement the quantitative validation, we performed qualitative case studies by visualizing predicted harsh cornering events on real-world GPS trajectories. Using the original time-series data, we plotted the location of each detected event on a digital map.

Figure 5 illustrates several representative examples of harsh cornering maneuvers detected by the ensemble model. For each case, we annotated the turn location with always-visible labels indicating vehicle speed and yaw rate. The trajectories clearly show sharp changes in direction or dynamics consistent with risky maneuvers. These visualizations confirm that the detected events correspond to turning behavior, and reinforce the effectiveness of combining orientation-invariant features, map-based filtering, and outlier detection.

These case studies further support the model's real-world interpretability and its potential applicability in driver feedback systems, road safety analytics, and fleet risk profiling.

Figure 5a depicts a harsh cornering event at an urban intersection, where the vehicle enters the turn at a high speed of 42km/h—substantially exceeding typical safe turning speeds in such environments. In Figure 5b, the vehicle executes a moderately fast cornering maneuver, with speeds ranging from 21 to 27km/h and yaw rates up to 27.4°/s, indicative of assertive but not necessarily unsafe driving. Figure 5c captures a sharp turn initiated immediately after departure, where the vehicle accelerates from 0 to 12km/h while reaching a peak yaw rate of 80°/s. This combination of low speed and extremely high angular velocity suggests a potentially unstable or abrupt maneuver. Lastly, Figure 5d presents a harsh 90° intersection turn with sustained high yaw rates, peaking at 47.5°/s. The vehicle performs the turn at speeds between 18 and 23 km/h, reflecting a tight and aggressive maneuver that may compromise stability in a confined urban setting.

DISCUSSION

Summary of Key Findings

This study developed and evaluated a fully unsupervised framework for detecting harsh cornering behavior using smartphone-based telematics and infrastructure data. By extracting yaw and turn dynamics from GPS, aggregating orientation-invariant inertial features, and validating events spatially with OpenStreetMap, the methodology identifies high-risk maneuvers in naturalistic driving conditions. Unsupervised clustering with DBSCAN and Isolation Forest allowed for robust detection without the need for ground truth labels or fixed device mounting. The detected events showed strong separability in feature space, and qualitative map visualizations confirmed their real-world plausibility, demonstrating the strength of the approach in capturing aggressive turning patterns at scale.

Limitations

Despite the promising results, the study has several limitations that inform both the scope of interpretation and future development. One key limitation lies in the temporal resolution of GPS data. Sampling at 1 Hz restricts the granularity of yaw rate estimation and may fail to fully capture the sharpness of brief, high-intensity turns, especially in fast-moving traffic scenarios. Additionally, the lack of labeled ground truth data precludes formal validation of detection precision or recall. Although proxy evaluations using supervised classifiers and qualitative visual inspection support the method's internal consistency, direct confirmation through video annotations or expert reviews would offer stronger validation.

Another constraint arises from the geographical specificity of the dataset. Since the data

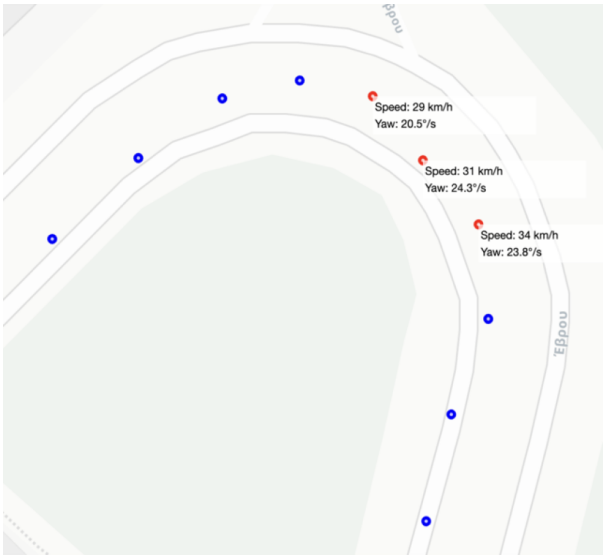

(a) Example of a harsh right turn with a speed above 30km/h in a curved urban road segment.

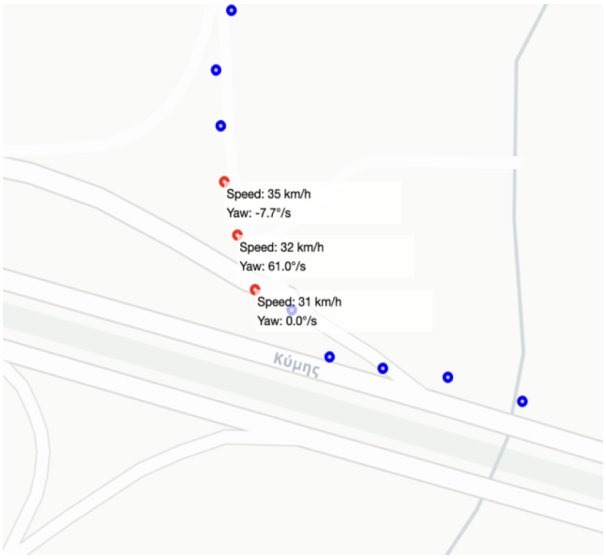

(b) Example of a harsh right turn within an urban road network. At a speed of 32 km/h, the vehicle undergoes a sharp change marked by a yaw rate of 61.0°/s

(c) Example of a harsh cornering maneuver performed at a low-speed urban T-junction. However the vehicle exhibits a sudden yaw rate spike, signaling a highly aggressive turning maneuver.

(d) Example of a harsh 90° intersection left turn with sustained high angular velocity in a narrow road.

FIGURE 5: Representative examples of detected harsh cornering events visualized on real-world map trajectories. Red markers indicate turning peaks, annotated with speed and yaw rate.

predominantly represents urban driving in Greece, the model may not generalize equally well across countries, climates, or vehicle types. Differences in road layout, driving culture, and smartphone sensor characteristics could affect the reliability of feature patterns and detection thresholds.

The reliance on sensor magnitudes to maintain orientation invariance also removes directional context, which could be useful for more granular behavioral profiling or risk assessment. Lastly, the use of OpenStreetMap for validating turns assumes high-quality and up-to-date map data, which may not hold true in all regions, potentially reducing the applicability of the method in poorly mapped areas.

Practical Applications

Despite these limitations, the proposed framework has significant potential for real-world application across several domains. In driver behavior analysis and usage-based insurance, the ability to detect harsh cornering events without requiring sensor calibration or labeled data makes the pipeline well-suited for large-scale deployment on consumer smartphones. It can serve as a basis for driver risk scoring, performance feedback, or targeted coaching interventions. Fleet operators could use the system to monitor aggressive maneuvers and prevent vehicle wear, while also improving safety outcomes across their networks.

Moreover, road safety agencies and urban planners could benefit from spatially aggregating harsh cornering detections to identify risky intersections or infrastructure gaps. Such data-driven insights could support decision-making for traffic calming measures or redesigns. Finally, the mobile-friendly nature of the approach lends itself to real-time use in driver assistance systems or mobility apps, enabling live feedback during trip execution and promoting safer driving habits over time.

Furthermore, the proposed framework can be extended by introducing a severity scoring system for detected harsh cornering events. By aggregating normalized values of yaw rate, turn angle, and gyroscopic magnitude into a composite severity index, each detected event can be ranked on a continuous scale rather than as a binary classification. This enables more nuanced interpretations of driver behavior and supports stratified feedback mechanisms, where different intervention levels are matched to the severity of the maneuver. Severity scoring can also assist insurers in dynamically adjusting driver risk scores based on cumulative event intensity over time, rather than event frequency alone.

Future Research Directions

Several extensions could further improve the methodology and expand its scope. One avenue is the integration of complementary data sources, such as dashcam video or environmental and infrastructure context, to enhance interpretability and aid in severity classification. Employing self-supervised or contrastive learning could unlock richer representations from large unlabeled datasets and improve cross-device generalization. Additionally, introducing directional features—perhaps through soft alignment with vehicle motion or inferred device orientation—could bring back valuable information without losing robustness.

Personalization of detection thresholds based on individual driving baselines may further enhance sensitivity and relevance, particularly when combined with privacy-preserving technologies like federated learning.

Finally, deploying the pipeline in real-time on mobile devices or edge platforms would allow for instant detection and feedback, enabling new applications in connected vehicle systems and behavioral safety platforms. Moreover, aggregating detected events across space and time could support the creation of road safety heatmaps, helping identify high-risk locations and informing infrastructure improvements or targeted traffic interventions.

CONCLUSIONS

This work presents a novel, unsupervised, and orientation-invariant approach to detecting harsh cornering behavior using standard smartphone sensors and open-source map data. The proposed pipeline successfully circumvents the limitations of fixed sensor placement and labeled training data, instead relying on outlier detection in a multi-sensor feature space to isolate aggressive turning events. By validating candidate maneuvers against real road geometry and demonstrating strong internal consistency in the resulting detections, the framework proves suitable for real-world driving behavior analysis.

The method's flexibility, low deployment cost, and robustness to phone placement make it highly applicable to modern telematics systems, driver monitoring solutions, and safety intelligence platforms. While limitations related to sampling rate, contextual generalization, and ground truth remain, the study establishes a strong foundation for unsupervised behavioral detection in naturalistic settings. As smartphone sensing and computational capabilities continue to improve, the proposed approach stands to play a significant role in advancing data-driven road safety and personalized mobility technologies.

ACKNOWLEDGEMENTS

This research has been conducted within the IVORY project. The project has received funding from the European Union's Horizon Europe research and innovation programme under grant agreement No 101119590.

REFERENCES

1. Gettman, D. M., L. Pu, T. Sayed, and S. G. Shelby, Surrogate Safety Assessment Model and Validation: Final Report, 2008.
2. Nikolaou, D., A. Ziakopoulos, and G. Yannis, A Review of Surrogate Safety Measures Uses in Historical Crash Investigations. *Sustainability*, Vol. 15, No. 9, 2023.
3. Mantouka, E., M. Barmpounakis, E. Vlahogianni, and J. Golias, Smartphone Sensing for Understanding Driving Behavior: Current Practice and Challenges. *International Journal of Transportation Science and Technology*, Vol. 10, 2020.
4. Simons-Morton, B., K. Cheon, F. Guo, and P. Albert, Trajectories of Kinematic Risky Driving Among Novice Teenagers. *Accident; analysis and prevention*, Vol. 51C, 2012, pp. 27–32.
5. Liu, L., D. Racz, K. Vaillancourt, J. Michelman, M. Barnes, S. Mellem, P. Eastham, B. Green, C. Armstrong, R. Bal, S. O'Banion, and F. Guo, *Smartphone-based Hard-braking Event Detection at Scale for Road Safety Services*, 2022.
6. Barthold, C., K. Pathapati Subbu, and R. Dantu, Evaluation of gyroscope-embedded mobile phones. In *2011 IEEE International Conference on Systems, Man, and Cybernetics*, 2011, pp. 1632–1638.
7. Padmanaban, J. and S. Husher, Occupant injury experience in rollover crashes: an in-depth review of NASS/CDS data. *Annu Proc Assoc Adv Automot Med*, Vol. 49, 2005, pp. 103–118.
8. Vlahogianni, E. I. and E. N. Barmpounakis, Driving analytics using smartphones: Algorithms, comparisons and challenges. *Transportation Research Part C: Emerging Technologies*, Vol. 79, 2017, pp. 196–206.
9. Wang, Y., Y. Chen, J. Yang, M. Gruteser, R. Martin, H. Liu, L. Liu, and Karataş, Determining Driver Phone Use by Exploiting Smartphone Integrated Sensors. *IEEE Transactions on Mobile Computing*, Vol. 15, 2015, pp. 1–1.
10. Chen, D., K.-T. Cho, S. Han, Z. Jin, and K. Shin, Invisible Sensing of Vehicle Steering with Smartphones, 2015, pp. 1–13.
11. Ester, M., H.-P. Kriegel, J. Sander, and X. Xu, A Density-Based Algorithm for Discovering Clusters in Large Spatial Databases with Noise. In *Knowledge Discovery and Data Mining*, 1996.