Self-Supervised Detection of Harsh Cornering Events at Scale using Smartphone Sensor Data

Aristotelis Tsoutsanis[1,2*], Apostolos Ziakopoulos[1], George Yannis[1]

[1] Department of Transportation Planning and Engineering – National Technical University of Athens, Athens, Greece
[2] OSeven Telematics, Athens, Greece
*a_tsoutsanis@mail.ntua.gr

Abstract.
This study presents a self-supervised framework for detecting harsh cornering events using smartphone telematics under real-world deployment conditions. GPS-derived trajectory features, orientation-invariant inertial magnitudes, and OpenStreetMap validation are combined to robustly identify candidate turns despite arbitrary device placement and sensor noise. High-confidence pseudo-labels are generated via ensemble anomaly detection, enabling supervised classifiers to learn without manual annotation. Among four evaluated models, a multi-layer perceptron achieved the best performance, and its latent embeddings revealed clear separation between harsh and normal cornering. The results demonstrate that consumer-grade smartphones can reliably capture lateral risky maneuvers at scale, extending prior work focused mainly on longitudinal events such as hard braking. This approach offers a scalable foundation for applications in driver risk assessment, insurance telematics, fleet management, and urban safety planning.

Keywords: Smartphone telematics, harsh cornering detection, self-supervised learning, sensor fusion, driver behavior.

1 Introduction

Road safety research increasingly utilizes crash surrogates to define crash proneness and associate it with risk [1, 2], particularly as actual crashes are statistically rare and difficult to assess in real-time. Smartphone-based telematics have emerged as a scalable alternative for monitoring driver behavior through sensors like accelerometers, gyroscopes, and GPS [3, 4]. While prior studies have successfully detected longitudinal high G-force events such as hard braking [5, 6, 7] and explored gyroscopic data for rotational motion [8], harsh cornering remains underexplored. Cornering dynamics are complex, involving interactions between lateral acceleration, yaw rate, speed, and road geometry [9]. Despite contributing to only 3% of crashes, harsh cornering is linked to nearly one-third of passenger fatalities via vehicle instability and rollover risk [10].

A primary challenge in detecting these events is the arbitrary orientation of smartphones in real-world deployment, which leads to inconsistent sensor axes. Furthermore, the high cost of ground-truth labeling necessitates self-supervised approaches that leverage orientation-invariant features from unlabeled naturalistic data. Addressing this gap, we propose a self-supervised framework using smartphone telematics that combines GPS trajectory features, inertial magnitudes, and OpenStreetMap validation. By employing anomaly detection for pseudo-label generation and subsequent representation learning, we demonstrate that harsh lateral maneuvers can be identified without reliance on manual annotation.

2 Data Collection and Preprocessing

This section outlines the procedures followed to prepare the data used for detecting harsh cornering events from smartphone-based telematics. 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.

2.1 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, 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.

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.

2.2 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 sum-

marized in Equation (4). These features together capture both the short-term intensity and the spatial extent of vehicle cornering behavior.

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

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

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

$$max_{turn_{angle}} = max_{1 \le i \le 4} = \left|turn_{angle_i}\right| \tag{4}$$

2.3 Map-Based Validation Using OpenStreetMap

Candidate turning events are identified when the heading change exceeds 70° within a 4s window, filtering out minor deviations but still susceptible to noise, lane changes, or abrupt corrections. To reduce false positives, we validate events using Open-StreetMap (OSM): genuine turns must occur within 15 m of an intersection, defined as a node where multiple road segments meet. Events without a corresponding intersection are discarded.

For efficiency, the entire Attica road network is pre-loaded and indexed, enabling fast spatial queries across diverse urban, suburban, and rural geometries. This map-based filtering enforces a topological constraint, ensuring that harsh turns align with actual road structure. The resulting dataset (Fig. 1) comprises 1,900 validated events, predominantly involving 70°–180° turns at speeds below 30 km/h, forming the basis for subsequent clustering and anomaly detection.

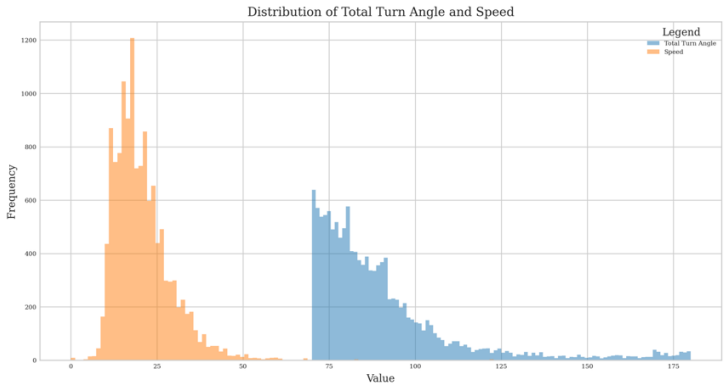

Fig. 1. The histogram shows the distribution of total turn angles (blue) and vehicle speeds (orange) across all turning events in the dataset.

2.4 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 deployments can be placed arbitrarily inside the vehi-

cle—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):

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

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

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.

3 Methodology

3.1 Feature Construction

For each GPS-identified turn, we extract a 3-second multivariate time-series window centered on the timestamp of maximum yaw rate (Fig. 2). From synchronized GPS, accelerometer, gyroscope, and device orientation data , we compute per-corner statistics, including speed, linear and angular magnitudes, yaw rate, temporal derivatives (jerk), and duration. This transforms each detected corner into a fixed-length feature vector suitable for downstream modeling.

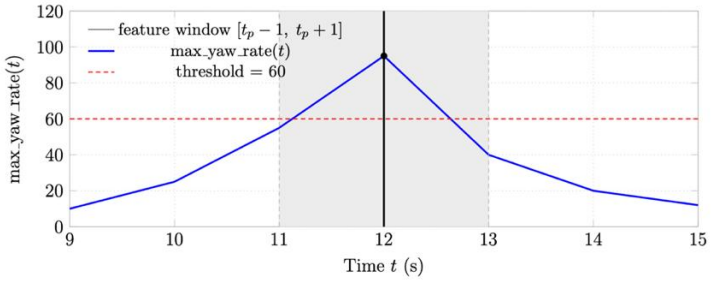

Fig. 2. Six-second segment of max yaw rate(t) centered at t_p= 12 s. The fixed feature window is shaded.

3.2 Anomaly Detection for Pseudo-Label Generation

Lacking ground-truth labels, we employ an unsupervised anomaly seeding strategy, assuming genuinely harsh maneuvers are rare statistical outliers in the high-dimensional feature space. After z-score normalization, features are evaluated by an ensemble of three complementary anomaly detectors: Isolation Forest, Local Outlier Factor (LOF), and One-Class SVM. The continuous anomaly scores from each meth-

od are rank normalized to [0,1] and averaged to produce a single ensemble score per event. Events in the top 5% of this distribution are designated as positive seeds (harsh cornering). To ensure a balanced dataset, an equal number of normal seeds are randomly sampled from the remaining events. These seeded pseudo-labels provide a high-precision training signal for supervised modeling.

3.3 Supervised Modeling on Pseudo-Labels

To evaluate whether the seeded pseudo-labels can support predictive modeling, we train a set of supervised classifiers on the subset of corners assigned to either positive or normal seeds. All events with a non-negative pseudo-label are retained, resulting in a binary classification problem where the positive class corresponds to high-confidence harsh cornering. The labeled subset is randomly partitioned into training (80%) and validation (20%) splits using stratified sampling to preserve the seed-class proportions. All features are standardized using z-score normalization estimated on the training split to prevent information leakage.
Four representative classifiers are benchmarked to capture a range of decision boundaries:

1. Logistic Regression provides a simple linear baseline and is trained with class-weighted loss to mitigate imbalance.
2. Radial-Basis Function (RBF) SVM models non-linear decision boundaries using an RBF kernel and balanced class weights.
3. Random Forest leverages an ensemble of decision trees with balanced subsampling and a large number of estimators to increase robustness.
4. Multi-Layer Perceptron (MLP) with a two-layer feed-forward architecture. The network maps the input features to a 16-dimensional latent representation through a hidden layer of 64 units with ReLU activation and dropout regularization, followed by a single-unit output layer. The model is trained with the AdamW optimizer and early stopping based on validation loss, using a binary cross-entropy objective, which measures the discrepancy between the predicted probability of harsh cornering and the true pseudo-label.

3.4 Representation Learning and Clustering

Beyond classification, we explore the latent structure of cornering behavior by constructing two types of embeddings:

1. MLP Embedding: the 16-dimensional output of the MLP's penultimate layer, representing a task-specific latent space learned from pseudo-labels.
2. Meta-Embedding: a richer representation obtained by concatenating the standardized feature vector with the probabilistic outputs of all four classifiers and subsequently applying Principal Component Analysis (PCA) to reduce dimensionality to eight components.

To assess whether these embeddings reveal natural groupings of cornering events, we apply k-means clustering ($k = 2$) to each embedding and evaluate cluster quality with

the silhouette score. Seeded pseudo-labels are used solely to name the resulting clusters (harsh or normal) for interpretability and are not provided during clustering.

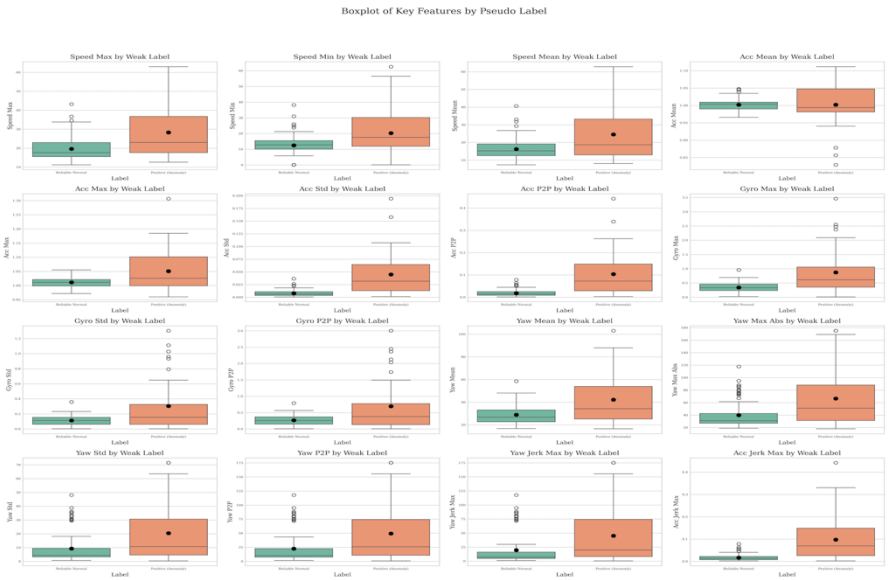

Fig. 3. The distribution of the input features in boxplots.

4 Results

Table 1 reports precision–recall AUC (PR-AUC) and receiver-operating-characteristic AUC (ROC-AUC) on the held-out validation split for all classifiers.

Table 1. Performance of supervised classifiers trained on seeded pseudo-labels. Values show precision-recall AUC (PR-AUC) and ROC-AUC on the held-out validation set.

Model	PR-AUC	ROC-AUC
Logistic Regression	0.9182	0.9086
SVM	**0.9510**	**0.9363**
Random Forest	0.8924	0.8920
MLP	0.8813	0.8726

4.1 Feature Distributions Across Pseudo-Labels

To further examine the discriminative power of the engineered features, Fig. 3 presents boxplots of key feature distributions for positive and normal seeds. Features such as maximum yaw rate, acceleration jerk, and peak-to-peak gyroscope magnitude

exhibit clear separation between the two classes, supporting their relevance for harsh cornering detection.

4.2 Latent Embedding Analysis

Clusters derived from the MLP achieve a silhouette score of Sil_{MLP} = 0.759, indicating clearer separation compared to the meta embedding (Sil_{meta} = 0.332). This shows that the task-specific latent space learned by the MLP captures more discriminative cornering dynamics.

5 Discussion

The results of this study demonstrate that smartphone-based telematics, when combined with unsupervised anomaly detection and orientation-invariant features, can effectively distinguish harsh lateral maneuvers from normal driving at scale. The success of the Multi-Layer Perceptron (MLP) model is rooted in the clear statistical divergence observed across the inertial and kinematic features presented in Fig. 3.
A detailed examination of the feature distributions reveals that "harsh" events are not defined by a single metric, but by a holistic shift in the vehicle's state.
The speed metrics (Speed Max, Min, and Mean) are significantly higher for the "Positive (Anomaly)" group. Notably, the elevated Speed Min suggests that drivers involved in harsh events fail to decelerate sufficiently before entering the turn. This sustained momentum necessitates higher lateral forces to maintain the vehicle's trajectory, leading to the observed spikes in Acc Max and Acc P2P. The most discriminative features are found in the yaw dynamics. Harsh turns exhibit substantially higher Yaw Max Abs and Yaw Std. Quantitatively, the high Yaw Jerk Max in the anomaly group indicates "busy" steering—likely a series of rapid corrections made by the driver to maintain control under high lateral load. In contrast, the "Reliable Normal" group shows a tight distribution with low variance, representing smooth, single-input steering. The significantly higher Acc Std and Gyro Std for anomalies indicate that harsh turns are physically noisy. This suggests that beyond simple high-speed cornering, these events involve secondary vibrations, physical signatures that the MLP model successfully mapped in its latent space to separate harsh versus normal structures.

6 Conclusions

This study presented a self-supervised framework for detecting harsh cornering events from smartphone telematics data. By combining GPS-derived features, orientation-invariant inertial signals, and map-based validation, we constructed a robust pipeline capable of operating under the noisy and variable conditions of real-world smartphone deployment. Using anomaly detection to generate pseudo-labels, we trained supervised classifiers that achieved high predictive performance, with the multi-layer perceptron offering the most promising results.

The findings demonstrate that harsh cornering can be reliably identified without the need for manually annotated data, addressing a long-standing barrier in scaling behavioral telematics research. Beyond methodological contributions, the proposed approach has practical implications for road safety, risk-based insurance, fleet management, and transportation planning.
Future work will extend this framework by validating its effectiveness across diverse regions and vehicle types, investigating temporal patterns of risky driving, and integrating multimodal data sources such as video or vehicular CAN-bus signals.

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., Pu, L., Sayed, T., Shelby, S.G.: Surrogate safety assessment model and validation: Final report. (2008).
2. Nikolaou, D., Ziakopoulos, A., Yannis, G.: A review of surrogate safety measures uses in historical crash investigations. Sustainability 15(9) (2023).
3. Mantouka, E., Barmpounakis, M., Vlahogianni, E., Golias, J.: Smartphone sensing for understanding driving behavior: Current practice and challenges. International Journal of Transportation Science and Technology 10 (2020).
4. Boylan, J., Meyer, D., Chen, W.S.: A systematic review of the use of in-vehicle telematics in monitoring driving behaviours. Accident Analysis Prevention 199, 107519 (2024). https://www.sciencedirect.com/science/article/pii/S0001457524000642
5. Simons-Morton, B., Cheon, K., Guo, F., Albert, P.: Trajectories of kinematic risky driving among novice teenagers. Accident; analysis and prevention 51C, 27–32 (2012).
6. Ziakopoulos, A.: Spatial analysis of harsh driving behavior events in urban networks using high-resolution smartphone and geometric data. Accident Analysis Prevention 157, 106189 (2021). https://www.sciencedirect.com/science/article/pii/S0001457521002207
7. Liu, L., Racz, D., Vaillancourt, K., Michelman, J., Barnes, M., Mellem, S., Eastham, P., Green, B., Armstrong, C., Bal, R., O'Banion, S., Guo, F.: Smartphone-based hard-braking event detection at scale for road safety services. (2022).
8. Barthold, C., Subbu, K.P., Dantu, R.: Evaluation of gyroscope-embedded mobile phones. In: 2011 IEEE International Conference on Systems, Man, and Cybernetics, pp. 1632–1638 (2011).
9. Wahlström, J., Skog, I., Händel, P.: Risk assessment of vehicle cornering events in gnss data driven insurance telematics. In: 17th International IEEE Conference on Intelligent Transportation Systems (ITSC), pp. 3132–3137 (2014).
10. Padmanaban, J., Husher, S.: Occupant injury experience in rollover crashes: an in-depth review of NASS/CDS data. Annu Proc Assoc Adv Automot Med 49, 103–118 (2005).