



Road Safety and Simulation 2026 – RSS2026

A Graph Transformer Approach for Modeling Crash Occurrence at Intersections Using Telematics-Informed Road Networks

Simone Paradiso^{a,b*}, Apostolos Ziakopoulos^a, Haris Sideris^b, George Yannis^a

^aNational Technical University of Athens, Department of Transportation Planning and Engineering, 5 Iroon Polytechniou Street, 15773, Athens, Greece

^bOSeven Single Member Private Company, 27B Chaimanta Street, 15234 Chalandri, Greece

Abstract

Spatial analysis has long been used in road safety to identify high-risk locations and uncover spatial patterns associated with crashes or hazardous behavior. The field has recently expanded thanks to advances in artificial intelligence and the growing availability of telematics data, enabling the development of more spatially informed models. This study leverages smartphone-collected telematics data and historical crash data, mapped onto the road network, along with road network features from OpenStreetMap, to predict crash risk at intersections in the city centre of Athens. The problem is formulated as a binary classification task, where intersections are labelled according to whether at least one crash has occurred. Two modelling approaches were compared: Extreme Gradient Boosting (XGBoost), one of the most commonly used machine learning methods in road safety, and a Graph Neural Network (GNN) with TransformerConv layers, which is well-suited for road network analysis, since it handles data in graph form, thus preserving its structure. While XGBoost captures relationships at the feature level, the GNN leverages the road network topology, improves node representations by incorporating information from neighbouring nodes, and integrates edge features. Both models were trained using optimized hyperparameters to maximize accuracy, and undersampling was applied to address severe class imbalance. The GNN outperformed XGBoost and was selected for further exploratory analysis. Based on model interpretations via Integrated Gradients, speed variability, harsh event intensity, and the number of streets at each intersection were identified as the most influential features, providing actionable insights for targeted interventions. This study demonstrates an approach to combine telematics and crash data within a graph-based AI framework, enabling crash risk estimation and interpretable insights at intersections. The proposed methodology highlights the potential of analyzing a telematics-informed road network with GNNs to support data-driven traffic safety analysis and prioritization of locations for further investigation.

© 2026 The Authors. Published by ELSEVIER B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the scientific committee of the Road Safety and Simulation 2026 – RSS2026

* Corresponding author. Tel.: +30.2107721575.

E-mail address: simone_paradiso@mail.ntua.gr

2352-1465 © 2026 The Authors. Published by ELSEVIER B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0>)

Peer-review under responsibility of the scientific committee of the Road Safety and Simulation 2026 – RSS2026

Keywords: Spatial Road Safety Analysis; Telematics Data; Graph Transformer Networks; Crash Occurrence Prediction; Model Explainability.

1. Introduction

Spatial analysis has been widely adopted in road safety to enable understanding of spatial patterns relevant to crash analysis and road safety monitoring, an essential practice for addressing an issue that causes over one million fatalities worldwide each year (WHO 2023) and thousands in Europe (European Commission 2025). While traditionally studied by employing econometrics and statistical approaches, advances in Artificial Intelligence (AI), particularly deep learning models, along with the growing telematics technologies have opened new avenues for spatial analysis (Ziakopoulos and Yannis 2020).

Tree-based models, a key component of the broader family of machine learning algorithms have been widely used, with algorithms like Random Forest, CatBoost and XGBoost showing strong performance and promise (Silva et al. 2020), exhibiting robustness on tabular and imbalanced datasets (Cunha et al. 2025).

The development of Graph Neural Networks (GNNs) enabled graph-structured information to be incorporated into deep learning models (Hamilton 2020). While they have been employed in spatial road safety modeling before, little work has been conducted by exploiting telematics data, particularly with telematics-informed road networks.

This study utilizes telematics data collected from a smartphone application, reflecting driving behavior, along with road infrastructure features extracted from OpenStreetMap (OSM) and historical crash records provided by police officers attending crash scenes. The aim of the present research is to map telematics data onto the road network and to estimate the likelihood of crash at specific intersections within the road network of the municipality area of Athens, leveraging telematics features as input for a classification model. This problem is indeed formulated as a binary classification task, where predictors include telematics data and road infrastructure data and the target variable indicates the presence or absence of a crash at each location. The analysis was approached through GNNs, leveraging the road network structure to capture spatial dependencies and interactions between intersections.

The study compares a well-known machine learning tool, the XGBoost model, with a GNN model based on TransformerConv layers, which emulate the Transformer framework. While XGBoost is well-established for tabular data, its limitations are addressed by the GNN, which can operate on graph-structured data. In the comparison, the GNN demonstrated superior generalization to unseen intersections through test set validation and was subsequently selected for feature explanation and for identifying the features most important for crash risk. Explainable Artificial Intelligence (XAI) provides a suite of techniques that enable human users to understand and produce more explainable models (Dwivedi et al. 2023). In this study, XAI was applied to assess feature contributions to model predictions, identifying speed and harsh event intensity as the most influential factors.

2. Data Collection and Preparation

A dataset of trips within the city centre of Athens was provided by OSeven Telematics (OSeven 2025), collected via a smartphone application at a 1 Hz frequency, comprising over 2.8 million observations collected during the last four months of 2024. All data were anonymized and processed in compliance with European data protection regulations (GDPR). The dataset includes trip coordinates, per-second speed measurements, and periodic binary indicators (1 or 0), signaling the occurrence of harsh braking, harsh acceleration, speeding, or mobile phone use. It also records the intensity of harsh events on a scale from 1 to 3. These metrics are generated using proprietary machine learning algorithms developed by OSeven, whose accuracy has been validated against OBD data, on-road tests, and established literature benchmarks (Papadimitriou et al. 2019).

A coordinate bounding box defined, based on the range of coordinates in the telematics dataset, was used to extract a graph alongside road features from OpenStreetMap (OSM), a free and editable global map created by volunteers and released under an open-content license (OpenStreetMap 2025), providing the spatial context for aggregating telematics data, allowing features to be summarized appropriately, as illustrated on Table 1.

Each telematics observation was matched to its nearest OSM edge, and metrics were subsequently aggregated at the edge level. Whereas, at node-level characterization, a 50-m radius buffer was defined around each node, and only

telematics observations on edges connected to the corresponding node within the buffer area were considered for aggregation.

Table 1. Node Features.

Feature	Description
street_count	Number of streets connected to the intersection
smoothenedSpeed	Average speed near this intersection
SpeedingFlag_per_trip	Count of speeding events occurred near the intersection per trip
mobile_usage_per_trip	Count of mobile phone usage events near the intersection per trip
harsh_acc_per_trip	Count of harsh acceleration events near the intersection per trip
harsh_brk_per_trip	Count of harsh braking events near the intersection per trip
event_intensity	Average intensity score of harsh events occurring nearby the intersection
speed_std	Speed standard deviation near the intersection

The process ultimately resulted in a network of over 9,615 nodes and 13,656 edges, each characterized by the aggregated telematics data. The node features are shown in Table 1, while the edge features mirror the node features, excluding street_count and including additional attributes such as edge length, movement direction (one- or two-way, represented as a binary variable), and two binary features derived via one-hot encoding from a three-category variable indicating road type (service, urban, or rural).

In this study, historical crash data that contained crashes in Athens, Greece, occurring between 2018 and 2022, were utilized, including road crashes in which at least one involved road user was injured (slightly/seriously) or killed. These data were used to define a binary risk indicator at the node level (i.e., crash occurrence vs. non-occurrence), rather than modeling crash counts, thereby focusing the analysis on spatially persistent risk patterns and partially alleviating concerns arising from the temporal mismatch with the telematics data. Data are collected from the police and codified into the ELSTAT database. Copy files of this database are provided to the Department of Transportation Planning and Engineering of the National Technical University of Athens (NTUA) with personal identification removed, in line with the standing European data protection legislation (GDPR).

The database was then filtered to include only intersection crashes. Since geographic coordinates were unavailable, street names were obtained from other publicly available databases using the geocodes of the two roads forming each intersection. Duplicate entries were removed, as multiple crashes can occur at the same location, and our focus is on crash occurrence. Greek street names were normalized with a custom function and used to query the Google Maps Geocoding API to determine intersection coordinates. As a quality-control step, the resulting coordinates were checked against the spatial extent of the telematics network, and only nodes within the study area were retained. After processing, 572 out of 9,615 telematics nodes had at least one recorded crash.

3. Modelling Approaches

This study investigates the relationship between telematics-derived driving behaviour and crash risk by developing predictive models for crash occurrence. To assess the relative effectiveness of conventional machine learning and advanced deep learning paradigms, an XGBoost model and a GNN model were compared. The XGBoost represents the practical standard model operating on traditional tabular data, while the GNN model reflects the transition to topology-aware learning using graph-structured data.

3.1. Conventional Machine Learning Baseline

XGBoost is a gradient boosting algorithm widely used in the machine learning community. It is an ensemble of decision trees designed for high scalability and efficiency. Like other gradient boosting methods, XGBoost constructs an additive model by iteratively minimizing an objective function, using decision trees as its base learners. The

objective function consists of the term loss function and regularization to control tree complexity and prevent overfitting, ensuring that the model does not grow unnecessarily complex and it generalizes better (Bentéjac et al. 2021). To achieve optimal performance with XGBoost, tuning the model's hyperparameters is required. Specifically, the model's hyperparameters include:

- `n_estimators` – controls how many boosting trees are used.
- `learning_rate` – determines the contribution of each tree to the overall model.
- `max_depth` – limits the depth of each individual tree, controlling model complexity.
- `colsample_bytree` – specifies the fraction of features sampled for each tree, helping reduce overfitting.
- `subsample` – defines the fraction of training data used for each tree, also aiding regularization.

3.2. Graph-Based Deep Learning Model

A GNN is a specialized class of deep learning methods designed to perform inference on data structured as graphs, therefore designed to operate on non-Euclidean data (Liang et al. 2022). A GNN architecture is composed of GNN layers, each of which is a computational unit that transforms node features by aggregating information from neighboring nodes and applying learnable weights and nonlinearities.

Given an input set of node representations $\{h_i \in \mathbb{R}^d | i \in \mathcal{V}\}$ and the set of edges \mathcal{E} , the GNN layer outputs a new set of node representations $\{h_i' \in \mathbb{R}^d | i \in \mathcal{V}\}$, through a combination of a function and an aggregation mechanism, such as averaging neighboring node features or applying attention to weigh neighbors differently. The updated node representation is computed as:

$$h_i' = f_{\phi} \left(h_i, \text{AGGREGATE}(\{h_j | j \in \mathcal{N}_i\}) \right)$$

Where $\mathcal{N}_i = \{j \in \mathcal{V} | (j, i) \in \mathcal{E}\}$ are the neighbors of the node i .

The literature has developed various GNN layers, starting with Graph Convolutional Networks (GCN), which combine simple averaging aggregation with convolution operations (Kipf and Welling 2017). Building on this, Graph Attention Networks (GAT) leverage the attention mechanisms which assigns different weights to neighboring nodes (Veličković et al. 2018). While the GAT layer uses attention to weigh only a node's immediate neighbors, a TransformerConv layer (Shi et al. 2021) applies multi-head self-attention like Transformers, allowing nodes to aggregate information from a broader or even global set of nodes. Moreover, when available, multi-dimensional edge features may be concatenated with node features before computing the attention coefficients, in order for them to affect the attention mechanism.

To achieve optimal performance, the GNN model requires careful hyperparameter tuning, while operating under a different training paradigm than traditional machine learning models. The hyperparameters to be tuned include:

- `hid1` and `hid2` – determine the number of neurons in each GNN layer.
- `dropout` – randomly zeroes some neuron outputs during training to prevent overfitting and improve generalization.
- `lr` – controls the step size for weight updates during training.

3.3. Explainable Artificial Intelligence

To assess the association between each feature and predicted crash risk, a XAI technique termed Integrated Gradients (IG) (Sundararajan et al. 2017) was employed. This technique attributes the model's output to its input features by integrated gradients along a path from the baseline to the actual input. Mathematically, for a given node and a feature i , the Integrated Gradients score is defined as follows:

$$IG_i(x) = (x_i - x_i') \cdot \int_{\alpha=0}^1 \frac{\partial F(x' + \alpha(x - x'))}{\partial x_i} d\alpha$$

Where F is the GNN model, x is the feature vector of a node, x' is the baseline which is often a vector of zeros and α is a scalar parameter that ranges from 0 to 1, controlling the interpolation between the baseline and the actual input. Essentially, this method works by gradually nudging the input from a blank value (when the baseline is zero) to the real value. By observing how the output changes along this path, a score is assigned to each feature, reflecting its contribution.

Additionally, the literature notes that while computing feature importance on the training set highlights what the model learned from the data, computing it on the test set reveals how the model actually uses features to make predictions (Molnar 2020).

4. Results and Discussions

The dataset (9,615 nodes) was highly imbalanced, with far fewer crash-risk nodes (class 1). To reduce bias, all 572 positives were retained and an equal number of negatives were randomly sampled (with a fixed seed for reproducibility), creating a balanced 1,144-node subset (~12% of the data, in a process known as undersampling). A first stratified split was then applied to maintain balanced class representation, dividing the constructed dataset into training (80%), validation (10%), and test (10%) sets. Figure 1 illustrates the data split distribution.




Fig. 1. Stratified Data Split.

4.1. Best Model Selection

At this stage, 5-fold stratified cross-validation was used on the combined training and validation set to obtain the optimal hyperparameters for the XGBoost model. For each split, the training set was used to compute scaling statistics to scale the data of both training and validation set, ensuring no data leakage, applying a min-max strategy to ensure that each feature was within the range 0 to 1.

A total of 324 model configurations were evaluated, formulated with different hyperparameter combinations, and the best-performing model was selected based on cross-validated accuracy. The optimal model achieved a cross-validated accuracy of 71.8%, demonstrating a robust predictive performance on the dataset, and it was tested by using the area under the ROC curve (AUC), providing a comprehensive measure of predictive ability, with results presented in Table 2. Testing was conducted on both the balanced test set and an extended set including unseen data (Figure 1). Since AUC measures ranking ability independently of class proportions, it remains a reliable metric across both balanced and imbalanced evaluation sets.

A similar training strategy was applied when using a Transformer-based GNN for binary node classification. Node features were first scaled using min-max normalization based on scaling parameters computed on the training set. Since the GNN uses edge features only for the aggregation mechanism, they were scaled using statistics from all edges (with binary features unchanged) to ensure consistent aggregation across the graph, and because this normalization is independent of the labels, it does not introduce data leakage.

The model employed in this study consisted of three TransformerConv layers with ReLU activations and dropout to stabilize training, followed by a linear output layer that captures additional relationships among the representations generated by the GNN layers. A total of 54 model configurations of hyperparameters were evaluated, and each hyperparameter combination was trained for up to 150 epochs with early stopping based on the validation set, which was also used to compute the accuracy of the model. The optimal GNN model achieved a cross-validated accuracy of 78.9%, and it was tested by using the area under the ROC curve (AUC), with the results summarized on Table 2.

Table 2. Best Model Performance Comparison on Balanced and Extended Test Sets (AUC).

Model	Balanced Test Set – AUC (best model)	Extended Test Set – AUC (best model)	Hyperparameter Ranges Evaluated	Optimal Hyperparameters
-------	--------------------------------------	--------------------------------------	---------------------------------	-------------------------

XGBoost	0.67	0.72	n_estimators: [50, 100, 200], learning_rate: [0.01, 0.1, 0.2, 0.3], max_depth: [3, 5, 7], colsample_bytree: [0.5, 0.7, 1], subsample: [0.5, 0.7, 1]	n_estimators = 50, learning_rate = 0.01, max_depth = 5, colsample_bytree = 1, subsample = 0.5
GNN	0.79	0.78	hid1: [14, 21], hid2: [21, 28, 35], lr: [0.001, 0.0005, 0.0001], dropout: [0.2, 0.3, 0.4]	hid1 = 14 hid2 = 35, lr = 0.001, dropout = 0.2

4.2. Explainability and Feature Importance

The GNN model outperformed XGBoost, likely due to the preservation of the spatial structures of the data, which naturally translates to better interpretation of road safety correlations. GNN was selected for further analysis to assess feature importance through XAI, to examine associations between features and predicted crash risk. IG scores were computed for each node in the test set using a zero-baseline to evaluate feature importance in terms of the model's generalization performance (Martin et al. 2022), which is critical to deploy the model for crash prediction and understanding how each input feature affects the model's predictions on new unseen data. It is important to note that the min-max normalization aligns the zero-baseline with the semantic meaning of zero in the normalized features, which is the absence of a feature.

These node-level attributions can be aggregated into violin plots for each feature, providing a global view of feature importance. The framework, including model training and IG computation, was repeated with the inclusion of speed standard deviation for the road network spatial entities to capture the variability in traffic flow across the intersections. The results are shown in Figure 2, with the black vertical line indicating the mean and the red vertical line indicating the median for each feature.




Fig. 2. Global Feature Importance With (a) and Without Speed Standard Deviation (b).

Upon a comparison of the violin plots, it is apparent that including the standard deviation of speed does not significantly alter the model's reasoning and it does not alter model's performance.

Most telematics features show relatively low attribution values with respect to predicted crash risk. The presented model identifies the number of streets at an intersection as a key factor for crash occurrence, consistent with previous studies indicating that higher node degrees increase risk due to a greater number of conflict points (Nazir et al. 2025). Similarly, higher intensity of harsh events correlates to crash risk, emphasizing the importance of aggressive driving behaviours characteristics in crash analysis, as highlighted in prior literature (Desai et al. 2021; Ziakopoulos 2024). Finally, greater speed standard deviation reflects how more unstable traffic flow may lead to a higher likelihood of crashing, corroborating findings from specific road safety studies (Wang et al. 2018).

Importance scores (IG_i) can be binned and mapped onto the road network, enabling the AI framework to identify the features' impact on new unseen nodes and highlight specific locations of interest. A threshold can be applied to these scores, with bins visualized in different colours across nodes. In this study, the previously mentioned key features

were considered and a threshold of ± 0.02 was set for each feature at each node, determined through visual inspection of the violin plot distributions. This threshold should be regarded as indicative rather than definitive, as optimal thresholds may vary by dataset and context and require further investigation. Specifically, from the violin plots, it is possible to interpret the feature importance as follows and illustrated in Figure 3:

- $IG_i < -0.02$, feature i reduces crash likelihood at the node.
- $IG_i > +0.02$, feature i increases crash likelihood at the node.
- $-0.02 < IG_i < +0.02$, feature i shows negligible attribution to predicted crash risk at the node.

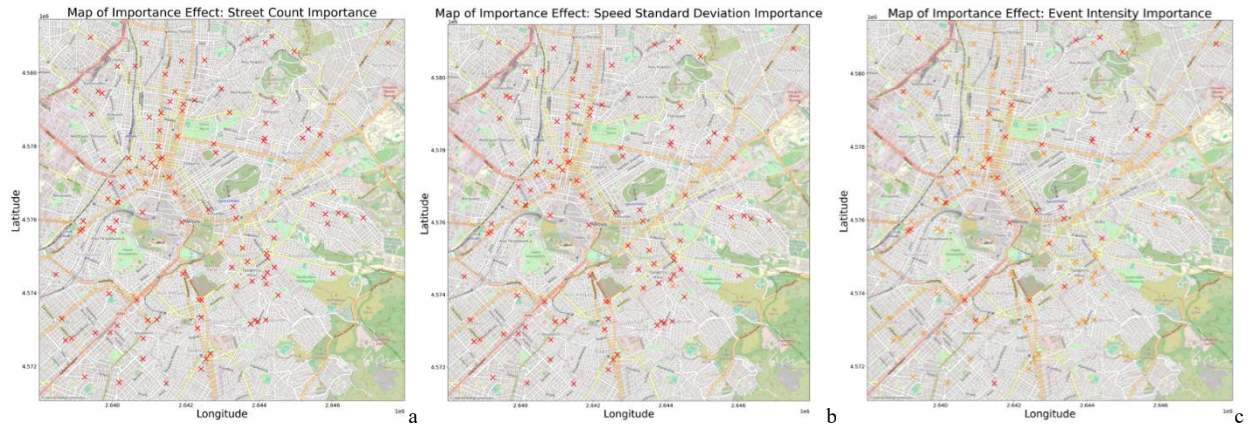


Fig. 3. Spatial Mapping of Feature Impact on Crash Likelihood.

On the map, red nodes indicate that the feature is associated with increased crash likelihood at that location (i.e. $IG > 0.02$), suggesting locations where interventions may be worth exploring. Orange nodes show low feature attribution, suggesting the node may be deprioritized for interventions targeting that feature. Although green nodes are not illustrated in the Figure 3, they would indicate that the feature may be associated with lower crash likelihood.

5. Conclusions

Road safety has been studied within its spatial context using a variety of tools and increasingly evaluated through machine learning and deep learning methods. This work developed a crash occurrence model and compare commonly used machine learning approaches, such as XGBoost, with more advanced methods like GNNs. Results demonstrate that GNN models are highly effective for spatial crash occurrence prediction using graph data, outperforming XGBoost. The Integrated Gradients XAI technique provides a way to assess feature attribution per node in a graph-based model. This is not possible with XGBoost, as it relies on differentiability, whereas XGBoost builds decision trees and therefore requires other techniques, such as SHAP, for feature attribution (Lundberg and Lee 2017). It also enables mapping node feature importance directly onto their spatial locations by binning the importance values.

Despite evaluating 324 hyperparameter combinations for XGBoost versus only 54 for GNN, the GNN still outperformed it, suggesting that further tuning could enhance its performance more. Additional computational resources could also enable a larger study area, extending the evaluation of telematics features beyond the city center. Further research is needed to determine the optimal thresholds for this binning process. Finally, the temporal mismatch between crash and telematics data represents a limitation, and its implications should be further investigated.

Overall, this study proposes an approach that integrates telematics data and police-reported crash records onto a road network, presenting an AI-based methodology for predicting crash risk using both road infrastructure and driving behavior data collected via smartphones. This tool can support road safety analysis and provide actionable feedback on infrastructure, contributing to improved understanding of factors associated with road safety outcomes. Furthermore, by mapping node-level feature importance, the framework can help identify and prioritize potentially risky intersections, allowing authorities to focus safety measures where they are most needed.

Acknowledgements

This research is based on work carried out within the IVORY project. The project has received funding from the European Union’s Horizon Europe research and innovation program under grant agreement No 101119590.

References

- Bentéjac, Candice, Anna Csörgő, and Gonzalo Martínez-Muñoz. 2021. “A Comparative Analysis of XGBoost.” *Artificial Intelligence Review* 54 (3): 1937–67. <https://doi.org/10.1007/s10462-020-09896-5>.
- Cunha, José, José Silvestre Silva, Ricardo Ribeiro, and Paulo Gomes. 2025. “Sustainable Road Safety: Predicting Traffic Accident Severity in Portugal Using Machine Learning.” *Sustainability* 17 (24). <https://doi.org/10.3390/su172411199>.
- Desai, Jairaj, Howell Li, Jijo K. Mathew, Yi-Ting Cheng, Ayman Habib, and Darcy M. Bullock. 2021. “Correlating Hard-Braking Activity with Crash Occurrences on Interstate Construction Projects in Indiana.” *Journal of Big Data Analytics in Transportation* 3 (1): 27–41. <https://doi.org/10.1007/s42421-020-00024-x>.
- Dwivedi, Rudresh, Devam Dave, Het Naik, et al. 2023. “Explainable AI (XAI): Core Ideas, Techniques, and Solutions.” *ACM Comput. Surv.* 55 (9): 194:1-194:33. <https://doi.org/10.1145/3561048>.
- European Commission. 2025. “EU Road Fatalities Drop by 3% in 2024, but Progress Remains Slow.” EU Road Fatalities Drop by 3% in 2024, but Progress Remains Slow, March 18. https://transport.ec.europa.eu/news-events/news/eu-road-fatalities-drop-3-2024-progress-remains-slow-2025-03-18_en.
- Hamilton, William L. 2020. *Graph Representation Learning*. Morgan & Claypool Publishers.
- Kipf, Thomas N., and Max Welling. 2017. “Semi-Supervised Classification with Graph Convolutional Networks.” arXiv:1609.02907. Preprint, arXiv, February 22. <https://doi.org/10.48550/arXiv.1609.02907>.
- Liang, Fan, Cheng Qian, Wei Yu, David Griffith, and Nada Golmie. 2022. “Survey of Graph Neural Networks and Applications.” *Wireless Communications and Mobile Computing* 2022 (1): 9261537. <https://doi.org/10.1155/2022/9261537>.
- Lundberg, Scott M., and Su-In Lee. 2017. “A Unified Approach to Interpreting Model Predictions.” *Advances in Neural Information Processing Systems* 30. <https://proceedings.neurips.cc/paper/2017/hash/8a20a8621978632d76c43dfd28b67767-Abstract.html>.
- Martin, Natalia Garcia, Stefano Malacrino, Marta Wojciechowska, et al. 2022. “A Graph Based Neural Network Approach to Immune Profiling of Multiplexed Tissue Samples.” *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, July, 3063–67. <https://doi.org/10.1109/EMBC48229.2022.9871251>.
- Molnar, Christoph. 2020. *Interpretable Machine Learning*. Lulu.com.
- Nazir, Mehraab, Sai Chand, and Rahul Goel. 2025. “Quantifying the Road Network Structure and Its Impact on Road Traffic Crashes: A Bayesian CAR Modelling Approach.” *IATSS Research* 49 (3): 374–86. <https://doi.org/10.1016/j.iatssr.2025.08.001>.
- OpenStreetMap. 2025. “About OpenStreetMap - OpenStreetMap Wiki.” https://wiki.openstreetmap.org/wiki/About_OpenStreetMap.
- OSeven. 2025. “Oseven.Io.” Oseven.Io. <https://oseven.io/>.
- Papadimitriou, Eleonora, Anastasia Argyropoulou, Dimitrios I. Tselentis, and George Yannis. 2019. “Analysis of Driver Behaviour through Smartphone Data: The Case of Mobile Phone Use While Driving.” *Safety Science* 119 (November): 91–97. <https://doi.org/10.1016/j.ssci.2019.05.059>.
- Shi, Yunsheng, Zhengjie Huang, Shikun Feng, Hui Zhong, Wenjin Wang, and Yu Sun. 2021. “Masked Label Prediction: Unified Message Passing Model for Semi-Supervised Classification.” arXiv:2009.03509. Preprint, arXiv, May 10. <https://doi.org/10.48550/arXiv.2009.03509>.
- Silva, Philippe Barbosa, Michelle Andrade, and Sara Ferreira. 2020. “Machine Learning Applied to Road Safety Modeling: A Systematic Literature Review.” *Journal of Traffic and Transportation Engineering (English Edition)* 7 (6): 775–90. <https://doi.org/10.1016/j.jtte.2020.07.004>.
- Sundararajan, Mukund, Ankur Taly, and Qiqi Yan. 2017. “Axiomatic Attribution for Deep Networks.” arXiv:1703.01365. Preprint, arXiv, June 13. <https://doi.org/10.48550/arXiv.1703.01365>.
- Veličković, Petar, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. “Graph Attention Networks.” arXiv:1710.10903. Preprint, arXiv, February 4. <https://doi.org/10.48550/arXiv.1710.10903>.
- Wang, Xuesong, Qingya Zhou, Mohammed Quddus, Tianxiang Fan, and Shou'en Fang. 2018. “Speed, Speed Variation and Crash Relationships for Urban Arterials.” *Accident Analysis & Prevention* 113 (April): 236–43. <https://doi.org/10.1016/j.aap.2018.01.032>.
- WHO. 2023. “Global Status Report on Road Safety 2023.” <https://www.who.int/publications/i/item/9789240086517>.
- Ziakopoulos, Apostolos. 2024. “Analysis of Harsh Braking and Harsh Acceleration Occurrence via Explainable Imbalanced Machine Learning Using High-Resolution Smartphone Telematics and Traffic Data.” *Accident Analysis & Prevention* 207 (November): 107743. <https://doi.org/10.1016/j.aap.2024.107743>.
- Ziakopoulos, Apostolos, and George Yannis. 2020. “A Review of Spatial Approaches in Road Safety.” *Accident Analysis & Prevention* 135 (February): 105323. <https://doi.org/10.1016/j.aap.2019.105323>.