Interpretable Machine Learning for Municipal Road Safety: A Spatiotemporal Analysis of Crash Severity in Athens (2016–2020)

Paper 34

Eleni Maria Theodoraki[a*], Paraskevi Koliou[a], George Yannis[a]

[a]National Technical University of Athens, Iroon Polytechniou 5-9, 15772 Athens, Greece

*Corresponding author: e_theodoraki@msil.ntua.gr

Introduction

Introduction

The Urban Road Safety Challenge

- Road crashes cause ~1.35 million deaths and 50 million injuries globally per year (WHO, 2018)
- Road fatalities are projected to become the 6th leading cause of death worldwide by 2030
- Athens, Greece: intense traffic, dense urban infrastructure, high exposure risk

The Gap: Traditional vs. Modern Approaches

- Traditional statistical models (logit, probit) struggle with nonlinear, high-dimensional crash dynamics
- Machine learning improves predictive power but lacks transparency, the “black box” problem
- Crash severity exhibits strong spatial heterogeneity across municipalities

This Study Proposes

- Interpretable ensemble ML (Random Forest + XGBoost) for crash severity prediction
- SHAP-based explainability to identify key predictors
- Municipal-level spatial risk mapping using SmartMaps geospatial data
- Study area: Athens municipalities, 2016–2020

Literature Review

Literature Review

Urban Crash Severity

- Conventional models fail on nonlinear relationships (Hamdan & Sipos, 2025; Dong et al., 2022)
- Shift to flexible ML approaches for complex, high-dimensional crash data (Kashifi et al., 2023)

Ensemble Machine Learning

- XGBoost: high accuracy and speed, 86.73% for pedestrian crashes (Hamdan & Sipos, 2025)
- Random Forest: resistant to overfitting, 97.4% accuracy on 3.5M records (Hamdan & Sipos, 2025)

Interpretability & SHAP

- SHAP: reliable feature attribution based on cooperative game theory (Aziz et al., 2024)
- Partial Dependence Plots: reveal monotonic/nonlinear global model behaviour (Li, 2022)

Spatial Heterogeneity

- Crash risk is spatial, land use, infrastructure, traffic vary across municipalities (Guo et al., 2024)
- GIS variables among strongest predictors in crash outcome models (Li, 2022; Xiao & Duan, 2025)

Research Gap

- Few studies combine interpretable ML with fine-grained municipal geospatial data → this study fills that gap for Athens

Methodology

Data Framework & Preprocessing

Data Source

- SmartMaps geospatial infrastructure + Athens city traffic records, 2016–2020
- 119 crash records after excluding <2% incomplete severity entries
- Severity distribution: Slight 94.6% | Serious 2.8% | Fatal 2.6%
- Stratified 80/20 train-test split based on severity level

Feature Categories

- Crash: severity level, crash type, number of vehicles, vulnerable road users
- Temporal: year, month, day of week, hour of day
- Environmental: lighting, weather, road conditions
- Spatial: latitude, longitude, municipality ID, GIS variables (road density, intersections, land use)

Preprocessing

- Categorical: target encoding or one-hot encoding depending on cardinality
- Numerical: standardisation or min-max scaling; missing values imputed at municipality level
- Spatial variables left unaltered for SHAP analysis and mapping

Analytical Framework

Descriptive Analytics

- Annual and categorical distributions summarised across 18 fatality dimensions
- Mean, std, min/max, and coefficient of variation computed to identify key risk groups and temporal fluctuations

Statistical Modelling

- T-test (Welch): assessed mean severity scores pre- vs post-2020 to detect pandemic-related shifts
- Linear Regression: $Severity_y = \beta 0 + \beta 1 \cdot Year + \varepsilon$, slope β1 shows direction and rate of change; R^2 shows temporal predictability

Machine Learning Models

- Random Forest: 500 trees, bootstrap, equal class weights, resistant to overfitting
- XGBoost: gradient boosted trees, early stopping, Bayesian hyperparameter optimisation
- Evaluation: R^2, MAE, RMSE, CV-R^2 (5-fold cross-validation)

SHAP Interpretability & Spatial Risk Index

- TreeSHAP: global summary plots, dependence plots, interaction values; local municipality profiles
- Spatial Risk Index = equal-weight composite of (1) mean predicted severity, (2) mean SHAP spatial contribution, (3) empirical KSI rate, normalised 0–1 and mapped via SmartMaps

Machine Learning Models & SHAP

Random Forest

- 500 decision trees, bootstrap sampling, equal class weights
- Resistant to overfitting; strong baseline for tabular crash data

XGBoost

- Gradient Boosted Decision Trees with early stopping to prevent overfitting
- Hyperparameter tuning via Bayesian optimisation (learning rate, max depth, subsampling)

Evaluation Metrics

- R2: explained variance | MAE: mean absolute error | RMSE: root mean square error
- CV-R2: cross-validated R2 from 5-fold CV (out-of-sample generalisability)

SHAP Interpretability (TreeSHAP)

- Global: summary plots show dominant predictors; dependence plots reveal nonlinear effects
- Local: municipality-level profiles identify features that raise or lower severity per area

Spatial Risk Index

- Mean predicted severity + mean SHAP spatial contribution + empirical KSI rate
- Normalised 0-1, equal weights, mapped via SmartMaps layers

Descriptive Analytics

Temporal Trends

- Stable crash and fatality patterns 2016–2019
- Sharp drop in 2020 due to COVID-19 mobility restrictions
- Slight injuries dominate; fatalities stable across all years

Spatial Distribution

- Highest crash volumes in central Athens municipalities
- Non-motorway sections: far more crashes and injuries than motorways

Key Insight

- High crash volume does NOT equal high severity
- Western municipalities: fewer crashes but disproportionately high severity
- Severity-focused analysis is essential alongside volume monitoring

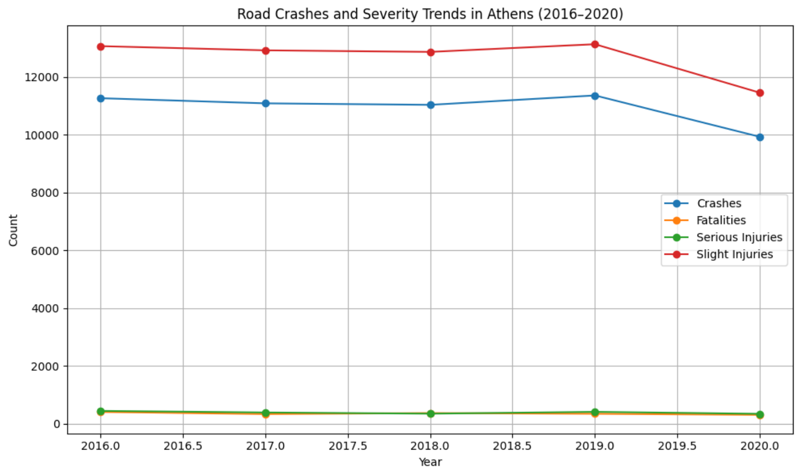

Fig. 1. Road Crashes and Severity Trends in Athens (2016–2020)

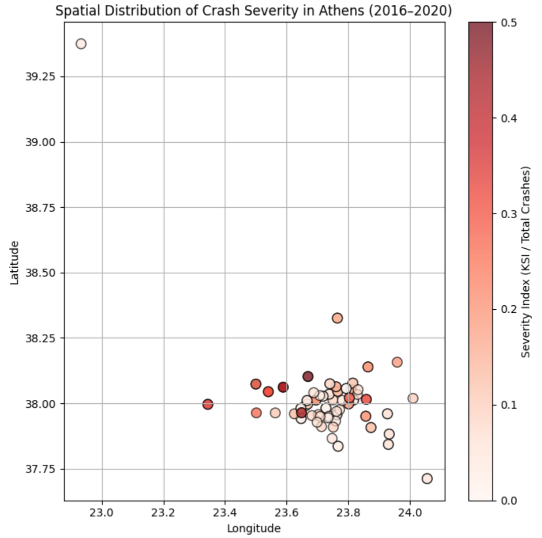

Fig. 2. Crash Density Map of Athens (2016–2020

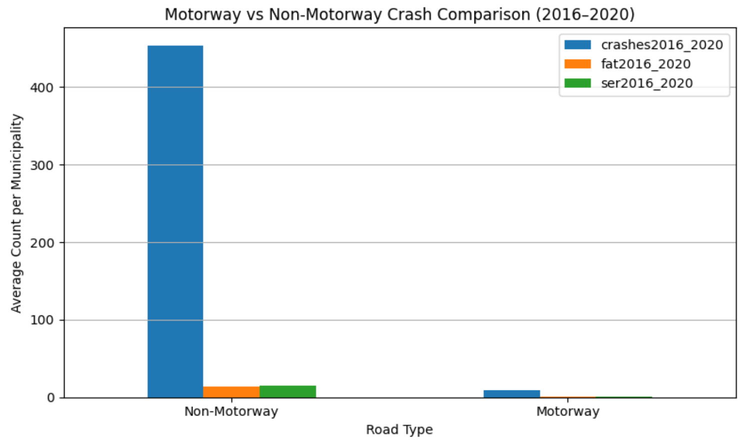

Figure 3: Motorway vs Non-Motorway Crash Comparison

Results

Spatial Patterns

Spatial Risk Map Strong spatial variability in crash severity across Athens

Western Municipalities High Severity

- Aspropyrgos, Elefsina, Megara, Salamina: high severity index, low crash counts
- Specific local conditions (road type, speeds) drive severe outcomes

Central Municipalities

- High crash frequency but lower KSI ratios
- Inverse volume-severity correlation across the city

Planning Implication

- East-west severity gradient confirmed → longitude later validated as top SHAP predictor
- Spatial risk maps = actionable decision-support tool for municipal interventions

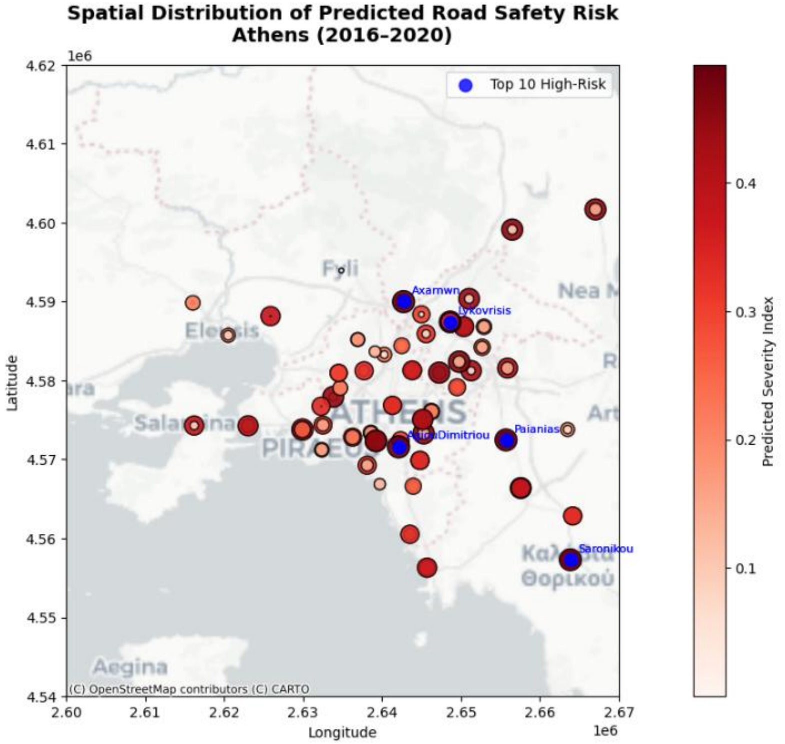

Fig. 4. Spatial Distribution of Historical Crash Severity

Model Performance

XGBoost outperforms Random Forest

- XGBoost: R2=0.302 | MAE=0.0626 | RMSE=0.1096 | CV-R2=0.5499
- Random Forest: R2=0.296 | MAE=0.0620 | RMSE=0.1101 | CV-R2=0.4238

Interpretation

- R2 ~0.30: consistent with benchmarks for municipal aggregate crash models
- Unexplained variance expected: driver behaviour, real-time traffic, enforcement unobserved
- CV-R2=0.55 confirms strong generalisation beyond training data
- Predicted vs actual (Fig. 5) shows good fit across severity levels

Table 1. Model Performance Metrics

Model	R^2	MAE	RMSE	CV-R^2
Random Forest	0.296	0.06197	0.11008	0.4238
XGBoost	0.302	0.06261	0.10961	0.5499

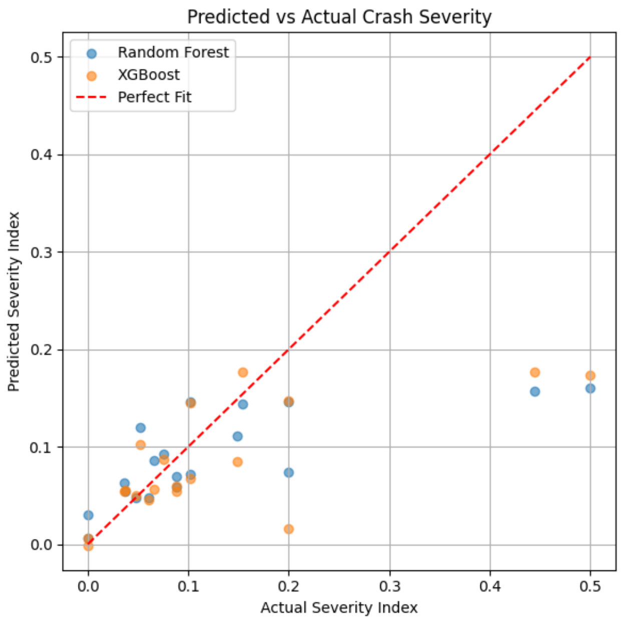

Fig. 5. Predicted vs Actual Crash Severity

SHAP Feature Importance

Top SHAP Features — XGBoost (Table 2)

- caplon (longitude) 0.0353: strongest predictor
- mean_fatalities 0.0246 | mean_crashes 0.0192 | mean_slight 0.0160
- motorway 0.0146 | ksi_total 0.0104

Key Findings

- Longitude = most influential → east-west spatial gradient dominates severity
- Casualty intensity features capture municipal exposure patterns
- Motorway presence shapes distinct severity profiles

Partial Dependence Plots

- Caplon: severity escalates non-linearly in western municipalities (lon < 23.4)
- Mean fatalities: threshold effect, severity spikes above 1 avg. fatality/municipality
- Mean crashes: diminishing returns as crash volume grows

Table 2. SHAP Feature Importance (top 9 features)

Feature	Mean
caplon (longitude)	0.0353
mean_fatalities	0.0246
mean_crashes	0.0192
mean_slight	0.0160
motorway	0.0146
ksi_total	0.0104
mean_serious	0.0049
lon_zscore	0.0042
crash_change_rate	0.0037
fatality_change_rate	0.0033
caplat (latitude)	0.0022
lat_zscore	0.00037

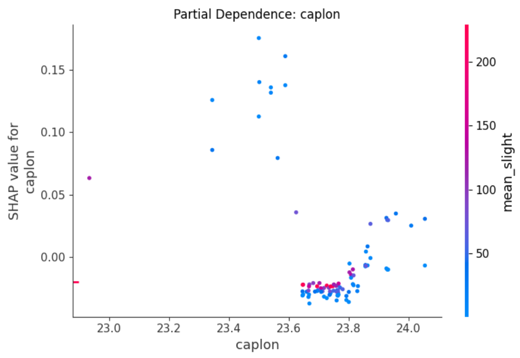

Fig. 6. Partial dependence: caplon

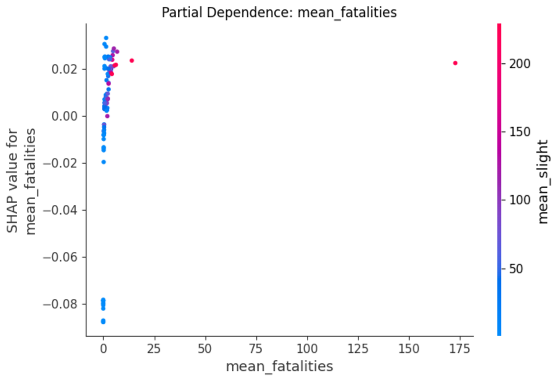

Fig. 7. Partial dependence: mean fatalities

Conclusions & Future Work

Conclusions & Future Work

Main Findings

- XGBoost outperforms RF; CV-R2=0.55 confirms strong generalisation for spatial risk mapping
- Longitude is the strongest predictor, east-west severity gradient confirmed across Athens
- Western municipalities show disproportionately high severity relative to crash volume
- SHAP exposes nonlinear thresholds invisible to traditional models
- Spatial risk maps provide actionable decision-support tools for urban safety planning

Limitations

- Municipal-level aggregation, road-segment data would improve precision
- Missing: traffic volume, socioeconomic factors, enforcement data
- Temporal dynamics (COVID-19, modal shifts) not yet captured

Future Work

- Integrate multimodal data: traffic volume, land use, population density
- Spatiotemporal models tracking annual severity evolution
- Real-time decision interface for municipal road risk monitoring

Interpretable Machine Learning for Municipal Road Safety: A Spatiotemporal Analysis of Crash Severity in Athens (2016–2020)

Paper 34

Eleni Maria Theodoraki[a*], Paraskevi Koliou[a], George Yannis[a]

[a]National Technical University of Athens, Iroon Polytechniou 5-9, 15772 Athens, Greece

*Corresponding author: e_theodoraki@msil.ntua.gr