13th Symposium of the European Association for Research in Transportation (hEART2025)
10-12 June 2025, Munich, Germany

Predictive Modeling of Pedestrian Violations Using Ensemble Learning: A Case Study at an Urban Intersection in Athens, Greece

Stella Roussou[1*], Apostolos Ziakopoulos[1], George Yannis[1]

[1] National Technical University of Athens, Department of Transportation Planning and Engineering, Athens, Greece

Introduction

Road safety is a global challenge, with **millions of deaths and injuries caused by road crashes each year**. These injuries are an outcome of various factors interacting together, such as human behavior, infrastructure design, environmental conditions, and socioeconomic disparities.

Densely populated urban environments, where pedestrians, cyclists, and motorized vehicles coexist, present unique challenges for traffic safety management. This study focuses on a busy intersection in Athens, Greece—**Panepistimiou and Vasilissis Sofias**—where limited monitoring and diverse road user behavior present difficulties in managing pedestrian safety.

Objectives

The objectives of this paper are to:

- **Compare the performance of Random Forest and XGBoost** in predicting illegal pedestrian crossings using real-world video data from a major Athens intersection.
- **Integrate YOLOv8, ResNet-50, and Kalman filtering for detecting**, tracking, and analyzing pedestrian behavior across different traffic light phases.

Methodology

- **Object Detection and Feature Extraction**

The YOLOv8 object detection model formulates detection as a single regression task.

- For an input image extracted as a frame from the video, **YOLOv8** (Ultralytics, 2024) outputs a set of bounding boxes $Bi=(xi,yi,wi,hi,ci)$ *(Eq.1)*.
- The detected objects were inputs of the **ResNet-50** for **feature extraction**.
- For **pedestrian tracking**, the system relies on the application of the **Kalman filter**, which balances the uncertainties in the prediction and the measurement.
- Once objects are detected and tracked, **the system analyzes pedestrian behavior in relation to traffic signal phases**.

$$S_{ped}(t) = \begin{cases} illegal & \text{if } S_{light}(t) = red \\ legal & \text{if } S_{light}(t) = green \\ unknown & \text{if } S_{light}(t) = intergreen \end{cases} \qquad (Eq.2)$$

- **Random Forest**

The Random Forest (RF) classifier builds multiple decision trees via bootstrap aggregation (bagging) and outputs class probabilities through majority voting. The model included **hyperparameter tuning** via grid search for:

- a number of trees,
- maximum depth, and
- minimum samples per leaf.

Class imbalance was addressed using adjusted class weights to increase the sensitivity of the model to rare violation events.

- **XGBoost**

The XGBoost, a gradient boosting method in which trees are built sequentially, where each learner **minimizes the residuals** of its predecessors. The loss function incorporates regularization terms to penalize model complexity using L1 (Lasso) and L2 (Ridge) norms. The model underwent extensive **hyperparameter tuning** (e.g., learning rate, max depth, etc.) using grid search.

Algorithmic Results

The dataset for this study was collected by strategically setting up smartphones at the key point of Panepistimiou Street and Vasilissis Sofia's junction, a central and heavily trafficked location in Athens, Greece. More than 8 hours of video footage were recorded across peak and off-peak hours to capture varied pedestrian behaviors and traffic conditions.

Figure 1: *Panepistimiou and Vasilissis Sofias Urban Junction in Athens*

As shown in Figure 1, The video footage used for this study was generated with **cameras from smartphones attached to tripods** at key pedestrian crossing points in Athens, thus providing a flexible and cost-effective solution for data collection.

To analyze the video data, a customized pedestrian violation detection algorithm was developed. This algorithm identifies the active signal phase—red or green—using pre-trained models and region-based image classification.

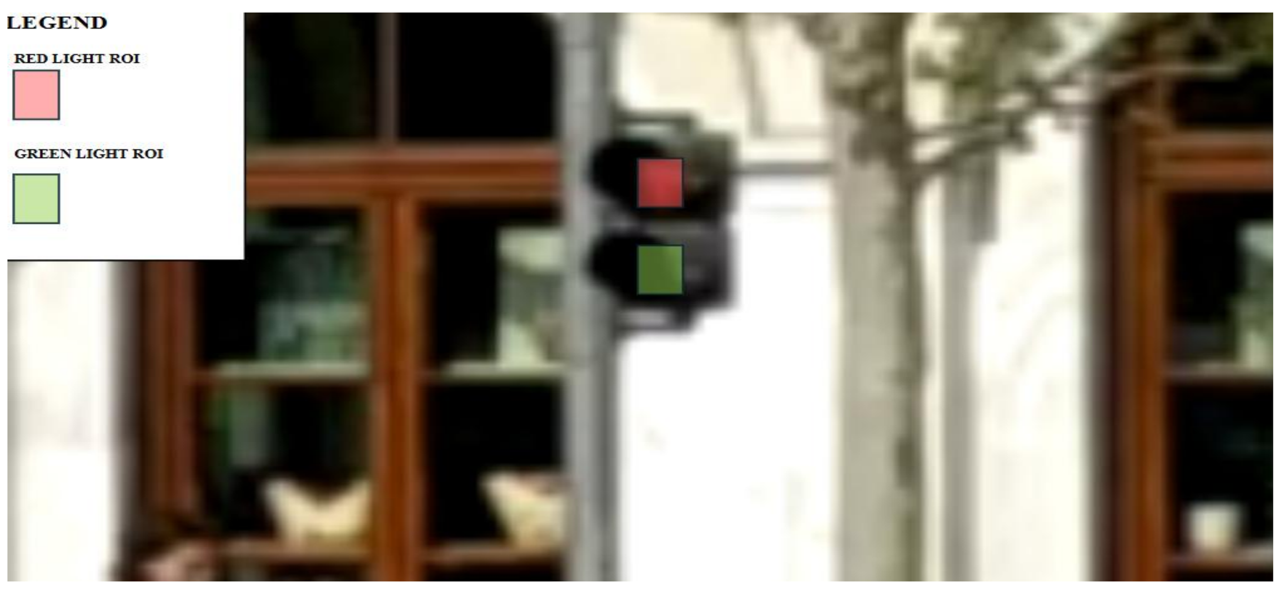

Figure 2: *Traffic Light Detection*

The Region of Interest (ROI) was carefully defined around the pedestrian crossings to isolate relevant activity zones. A crossing was **flagged as illegal** if a pedestrian entered the ROI during a red or intergreen signal phase.

Figure 3: *Region of Interest (ROI)*

This automated process produced a structured dataset of pedestrian behavior under varying signal conditions, which was then used to train and evaluate machine learning models (Random Forest and XGBoost).

Machine Learning Results

Two ensemble classifiers—Random Forest (RF) and XGBoost—were **trained on pedestrian behavior data** to detect **illegal crossings**. Both models were evaluated using precision, recall, and F1-score for both majority and minority classes.

Table 1: *Classification Report for Random Forest Model*

Metric	Class 0 (No Violations)	Class 1 (Violations)	Macro Average	Weighted Average
Precision	0.96	0.55	0.76	0.92
Recall	0.91	0.68	0.79	0.88
F1-Score	0.93	0.61	0.77	0.89
Support	139,386	11,386	-	150,772

The **Random Forest** model achieved an overall accuracy of 88%. Its performance in detecting violations (**minority class) was lower**, with **55% precision**, 68% recall, and an F1-score of 61%.

Table 2: *Classification Report for XGBoost Model*

Metric	Class 0 (No Violations)	Class 1 (Violations)	Macro Average	Weighted Average
Precision	0.96	0.72	0.84	0.93
Recall	0.95	0.88	0.91	0.93
F1-Score	0.96	0.79	0.87	0.93
Support	139,386	11,386	-	150,772

By contrast, **XGBoost** significantly **outperformed** RF, reaching an accuracy of 93%. For the **minority class**, it achieved **72% precision**, 88% recall, and an F1-score of 79%.

RF was optimized with:

- 200 trees,
- depth at 15, and
- minimum samples per split = 10.

XGBoost optimal settings:

- 100 estimators,
- depth at 10,
- learning rate at 0.1, and
- subsample = 0.8.

Conclusions

- Random Forest and XGBoost models were trained to predict **illegal pedestrian crossings** at a busy intersection in Athens.
- The models used data from **pedestrian-vehicle interaction** with features extracted from the built algorithm which used **YOLOv8 for the object detection**, the ResNet-50 for feature extraction and Kalman filtering for trajectory prediction.
- XGBoost outperformed Random Forest in Precision, Recall, Overall Accuracy, and effectively handled class imbalance.
- Findings support the use of machine learning for:
  - Monitoring pedestrian violations.
  - Enhancing traffic light compliance.
- Future research directions:
  - Incorporate contextual data (weather, road layout, pedestrian demographics.
  - Validate models with larger and more diverse datasets in field conditions.

Acknowledgments

The research was carried out within the research project “PHOEBE - Predictive Approaches for Safer Urban Environment”, which has received funding from the European Union’s Horizon Europe research and innovation programme under grant agreement No 101076963.

Contact Information:
Stella Roussou, PhD Candidate, Research Associate NTUA
Department of Transportation Planning and Engineering
Email: s_roussou@mail.ntua.gr
Website: https://www.nrso.ntua.gr/p/s_roussou/