AI-Driven Framework for Real-Time Prediction of Microscopic and Macroscopic Driving Risk Using Holistic Data

Dimitrios I. Tselentis, Thodoris Garefalakis, Dimitrios Nikolaou, Eva Michelaraki, George Yannis

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

Introduction

Accurate prediction of driving risks is crucial for urban road safety. With traffic crashes resulting in significant human and economic losses globally, real-time risk estimation is vital.

Traffic modeling is approached at **three levels**: macroscopic (overall flow), microscopic (individual driver-vehicle behaviors), and mesoscopic (group interactions). Microscopic models, involving car-following and lane-changing, help identify and mitigate risky behaviors like frequent lane changes. Macroscopic models use aggregated data to reveal broader safety trends, such as speed dispersion correlating with collision risks. **Integrating these two scales** offers a more comprehensive understanding of road safety, yet such integration, especially using new technologies like drones, remains underexplored

Objective

This study utilizes **drone-based data** to capture high-resolution driver behavior and traffic patterns on an urban arterial. Leveraging these insights, AI-driven models are developed to **estimate traffic risk probabilities** at both microscopic and macroscopic levels, addressing research gaps and offering a more holistic, real-time perspective on road safety

Data Overview

The pNEUMA dataset, collected in Athens, Greece (2018), consists of high-frequency (25Hz) vehicle trajectory data captured by ten drones over five days. As shown in Figure 1, the drones covered a 1.3 km² area with over 100 kilometers of roadways and nearly 100 intersections, capturing nearly half a million vehicle trajectories. This study focused on Panepistimiou Street, a five-lane urban arterial in Athens (i.e., Block 2,3 and 5).

Figure 1: Blocks covered by each drone of the swarm

The dataset was restructured into 0.04s intervals, and Panepistimiou lanes were mapped as polygons for accurate positioning. **Three dataframes** were developed and merged into a unified dataframe:

1st Dataframe — Vehicle Metrics

Aggregated data for vehicle pairs (following and leading vehicles) and risk metrics including:

- Track ID
- Position
- Type
- Speed
- Longitudinal acceleration,
- Lane-polygon
- Time-to-Collision (TTC)

$$TTC_i = \frac{x_{i-1}(t) - x_i(t) - l_i}{\dot{x}_i(t) - \dot{x}_{i-1}(t)}$$

Where: $\dot{x}$ speed of the vehicle, x the position, l the vehicle length

2nd Dataframe — Lane-specific Traffic Metrics

Traffic metrics for each polygon at every 0.04s timeframe, including:

- Average vehicle speed

$$v_{avg} = \frac{\sum_{i=1}^{n} v_i}{n}$$

- Total number of vehicles
- Density (veh/km)

$$k = \frac{n}{L}$$

- Traffic flow (veh/h)

$$q = k \times v_{avg}$$

Where: v_i speed of the i-th vehicle (km/h), n the total number of vehicles observed, L the length of the road segment (km)

3rd Dataframe — Traffic Events

Risky traffic events, setting thresholds for safety indicators:

- Thresholds used included:
  - TTC: 1.5 seconds
  - Harsh braking: 4.9 m/s²
  - Harsh acceleration: 4.9 m/s²
  - Speeding: +10 km/h over limit
- Lane change detection was performed by flagging vehicles that appeared in different polygons between two consecutive time frames.

Methodology

- **Long Short-Term Memory (LSTM) networks** were employed to predict driving risk by modeling sequential data. Both **uni-directional and bi-directional LSTMs** were explored, where the former processed data in a single temporal direction, and the latter considered both forward and backward temporal sequences to enhance the contextual understanding of risky events.
- The models utilized driving metrics such as **vehicle type, speed, longitudinal acceleration, relative distance, and time-to-collision** for the ego vehicle and its closest neighbors. Target variables included **lane changing, speeding, harsh acceleration, and harsh braking**. Figure 2 depicts the ego vehicle (**green**) and surrounding vehicles (**black**), whose data were used to estimate microscopic risks.

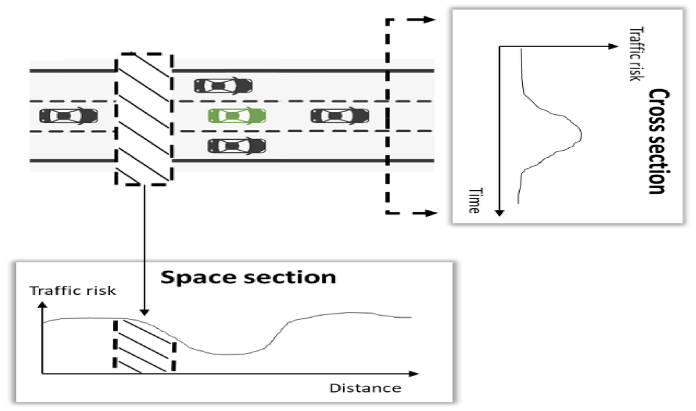

Figure 2: Ego Vehicle and Spatiotemporal Risk Aggregation

- The LSTM models featured a **robust architecture** with at least two layers using 'tanh' activation functions, L2 regularization, and dropout rates over 40% to **prevent overfitting**. The output layer used a sigmoid activation function for binary classification. Models were trained with a batch size of 32 for up to 400 epochs, applying early stopping to halt training when validation loss stopped improving. Configurations were optimized by varying layers, dropout rates, neurons, and optimizers (Adam and SGD).
- The LSTM analyzed driving data through a **10-second observation window (X)**, a **1-second reaction window (Y)** for driver response exclusion, and a **2-second prediction window (Z)** to flag risky events. As shown in Figure 3, the model used sequential data from the observation window (X) to predict events in the prediction window (Z). The input data shape was (250, 35, 69505), representing observations, features, and snapshots.

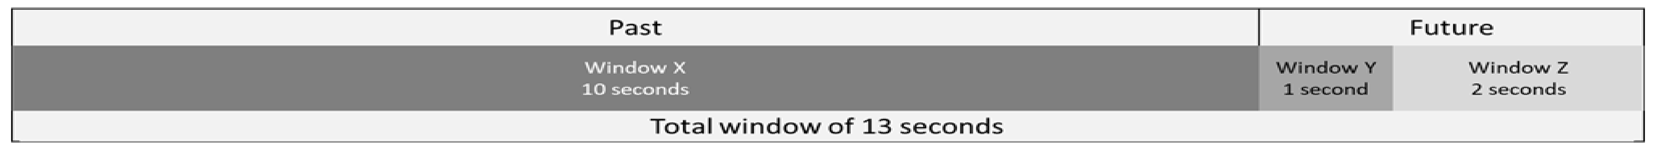

Figure 3: Time-Series Windows for Driving Metrics and Event Prediction

Microscopic Risk Prediction: LSTM Configurations

LSTM models were developed to predict microscopic driving risks, specifically focusing on speeding events. Models included **uni- and bi-directional configurations** with varying layers (2-3), dropout rates (40-50%), neurons (16-128), and optimizers (Adam and SGD). Uni-directional models processed sequences in a single direction, while bi-directional models captured both past and future contexts, offering potentially richer insights and improved accuracy.

Bi-directional models were expected to perform better **by leveraging temporal dependencies** in both directions, which is critical for predicting sequential behaviors such as harsh braking or speed changes based on both past and future contexts. As shown in Table 1, the models were designed with varying complexity to evaluate their effectiveness in capturing these dependencies.

Table 1: Configuration of the LSTM tested for the prediction of speeding events

Model ID	Direction	# LSTM layers	Dropout %	# Neurons in 1st LSTM layer	# Neurons in 2nd LSTM layer	# Neurons in 3rd LSTM layer	Optimizer	Learning rate	Batch size	# epochs
1	Uni	3	40%	64	32	32	Adam	1E-4	64	50
2	Uni	2	50%	32	16	-	Adam	1E-4	64	50
3	Bi	2	50%	32	16	-	Adam	1E-4	64	50
4	Bi	2	50%	64	32	-	Adam	1E-4	64	100
5	Bi	3	50%	32	16	16	Adam	1E-4	64	50
6	Bi	3	50%	64	32	32	Adam	1E-4	64	50
7	Bi	3	40%	128	64	64	Adam	5E-4	64	50
8	Bi	3	40%	64	32	32	SGD (momentum=0.4, decay=0.0001, nesterov=False)	1E-3	64	60
9	Bi	3	40%	32	16	16	SGD (momentum=0.6, decay=0.00001, nesterov=False)	1E-4	64	60
10	Bi	2	40%	64	32	-	SGD (momentum=0.6, decay=0.00001, nesterov=True)	1E-4	64	60
11	Bi	2	40%	64	32	-	SGD (momentum=0.5, decay=0.0005, nesterov=True)	1E-3	64	60
12	Bi	3	40%	128	64	64	SGD (momentum=0.5, decay=0.0005, nesterov=True)	1E-4	256	200
13	Bi	3	40%	32	16	16	SGD (momentum=0.5, decay=0.0005, nesterov=True)	1E-4	32	400

Microscopic Risk Prediction: Performance Metrics

The performance of the LSTM models was evaluated using accuracy, precision, recall, and AUC. These metrics provided a comprehensive assessment of the models' ability to predict driving risk events, ensuring robust evaluation for real-time assessments.

Table 2 summarizes the performance metrics for the tested LSTM models. **Bi-directional models generally outperformed uni-directional models in terms of precision**, likely due to their ability to consider future states.

For example, Model 7, a bi-directional LSTM with three layers and 40% dropout, achieved the **highest precision of 94%** but had a lower recall of 18%, indicating it was highly accurate when identifying speeding events but struggled to detect all occurrences. Model 1, a uni-directional LSTM, had a slightly lower precision of 76% and a slightly higher recall of 22%, showing a trade-off between precision and recall.

Table 2: Performance metrics for the LSTM models

Model ID	Binary accuracy	Precision	Recall	Area Under the Curve (AUC)
1	84%	76%	22%	75%
2	83%	87%	10%	71%
3	84%	89%	17%	75%
4	85%	85%	25%	77%
5	84%	79%	18%	74%
6	86%	86%	27%	77%
7	85%	94%	18%	76%
8	82%	52%	27%	69%
9	81%	37%	4%	64%
10	81%	100%	0%	66%
11	82%	76%	4%	68%
12	82%	72%	2%	65%
13	82%	68%	3%	66%

Macroscopic Driving Risk Probability

Microscopic driving risk probabilities predicted by LSTM models were aggregated to assess risk at the road section level. Risk probabilities of individual vehicles and their interactions within specific road segments were analyzed over **102-second intervals**. This transformed microscopic predictions into a macroscopic view, offering insights into overall risk profiles.

Model 7 initially estimated microscopic risk, focusing on speeding events at the vehicle level, and aggregated this data at the road section level. Future research should expand this approach to include lane changes, harsh acceleration, and braking, considering all road users.

Road segments were divided, and risk probabilities were summed over time to generate segment-specific scores. Aggregated risk data, normalized and visualized in Figure 4, highlight the variation in driving risk for a specific road section across time

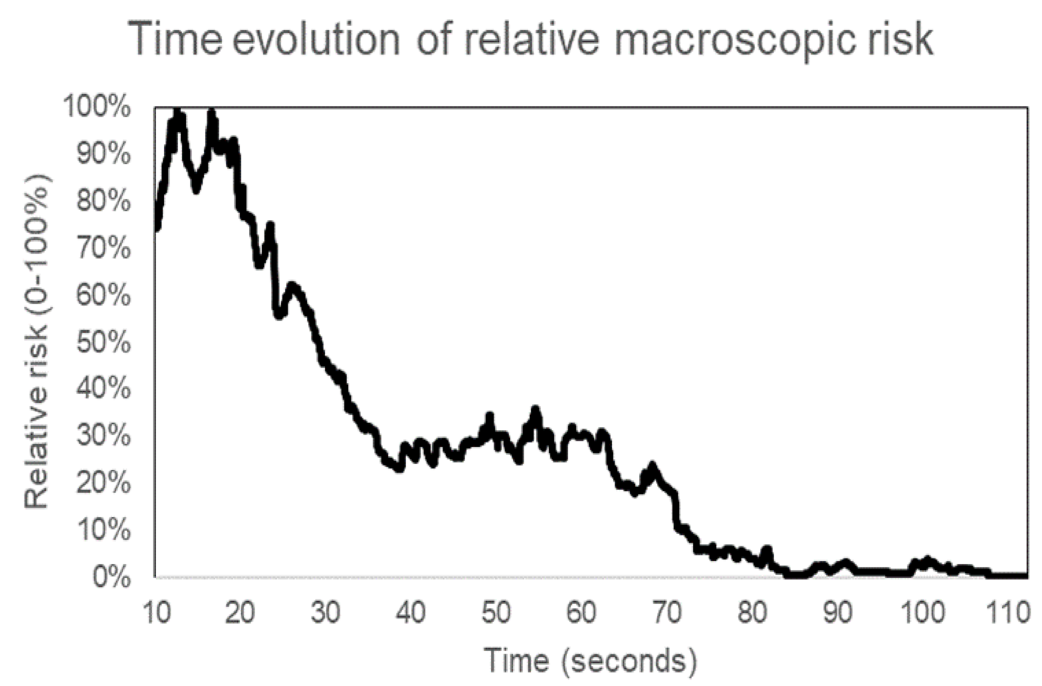

Figure 3: Evolution of aggregated macroscopic risk in time for a specific section of Panepistimiou street for 100 seconds

Conclusions

- **LSTM-based modeling shows strong potential for real-time driving risk prediction**, demonstrating effectiveness at both microscopic and macroscopic levels when combined with drone-collected data.
- **Bi-directional LSTM models offer slightly better precision** than uni-directional models due to their ability to capture future states, but both configurations still struggle with low recall, indicating difficulty in identifying all true risky events.
- **Drone-collected traffic data provides a richer**, more dynamic perspective compared to traditional static or vehicular data, enabling more comprehensive risk assessments.
- **Further refinement and integration of additional data sources** - including traffic, infrastructure, environmental factors, and other road users - can enhance accuracy and generalizability.

Acknowledgments

pneuma The authors would like to acknowledge the use of the pNEUMA dataset in this study – **open-traffic.epfl.ch**

Contact Information:
Eva Michelaraki, PhD , Research Associate NTUA
Department of Transportation Planning and Engineering
Email: evamich@mail.ntua.gr
Website: https://www.nrso.ntua.gr/p/evamich/

104th Annual Meeting of the Transportation Research Board, January 5–9, 2025
TRBAM-25-00817