

Spatiotemporal Forecasting of Incidents and Congestion with Implications for Sustainable Traffic Control

Tony Kinchen, Ting Bai, Nishanth Venkatesh S., and Andreas A. Malikopoulos

Abstract—Urban traffic anomalies, such as collisions and disruptions, threaten the safety, efficiency, and sustainability of transportation systems. In this paper, we present a simulation-based framework for modeling, detecting, and predicting such anomalies in urban networks. Using the Simulation of Urban MObility (SUMO) platform, we generate reproducible rear-end and intersection crash scenarios with matched baselines, enabling controlled experimentation and comparative evaluation. We record vehicle-level travel time, speed, and emissions for both edge- and network-level analysis. Building on this dataset, we develop a hybrid forecasting architecture that combines bidirectional long short-term memory networks with a diffusion convolutional recurrent neural network to capture temporal dynamics and spatial dependencies. Our simulation studies on the Broadway corridor in New York City demonstrate the framework’s ability to reproduce consistent incident conditions, quantify their effects, and provide accurate multi-horizon traffic forecasts. Our results highlight the value of combining controlled anomaly generation with deep predictive models to support reproducible evaluation and sustainable traffic management.

I. INTRODUCTION

Traffic congestion and unexpected incidents, such as collisions, are recurring phenomena in urban transportation systems. These events disturb normal traffic flow, increase travel times, and exacerbate fuel consumption and emissions, leading to negative impacts on urban mobility [1], [2]. Unlike other disruptions, congestion induces stop-and-go driving patterns, where prolonged idling and frequent acceleration cycles raise carbon emissions (CE) sharply, especially when average freeway speeds fall below 45-50 mph (72-80 km/h) [3]. Mitigating congestion is therefore critical not only for improving traffic efficiency but also for reducing environmental impacts. Despite advances in vehicle platooning [4], [5], [6], optimal route planning [7], [8] and recommendation [9], [10], and intersection signal control [11], [12], traffic congestion remains an inevitable challenge. As urban road networks grow denser and more interconnected, the need for effective approaches to detect and predict traffic anomalies has become increasingly urgent to enhance system resilience and promote sustainable mobility.

Conventional incident detection approaches typically rely on predefined thresholds or rule-based models that compare

observed traffic parameters (e.g., vehicle speed, occupancy, traffic flow) against nominal ranges, triggering alarms when significant deviations are detected [13], [14]. While these methods provide foundational baselines, they fall short in capturing the nonlinear and evolving dynamics of urban traffic systems, particularly in traffic environments where high variability and spatiotemporal dependencies often undermine detection accuracy.

Recent research has explored data-driven techniques for identifying anomalies in traffic systems. Graph-based models and recurrent neural networks have been employed to forecast traffic states and identify irregularities in congestion patterns [15]. In parallel, predictive modeling has advanced with techniques such as sequence-to-sequence learning and diffusion-based recurrent frameworks, which are widely adopted for spatiotemporal forecasting [16]. Attention-based mechanisms further extend these approaches by improving long-horizon prediction accuracy and reducing sensitivity to noise in sensor data. Informer-inspired architectures and Siamese network designs have also shown promise in capturing long-term dependencies and distinguishing abnormal traffic conditions [17], [18]. More recently, informer-based models applied to urban mobility [19] have demonstrated robust performance in anomaly detection tasks, achieving higher recall and F1 (a metric used to evaluate model accuracy) scores than competing baseline methods.

Video-based methods provide another important avenue for traffic anomaly detection. Deep learning pipelines that integrate convolutional neural networks, long short-term memory networks (LSTMs), and fuzzy reasoning have been developed to capture spatiotemporal cues for accidents and near-misses in traffic streams [20]. Rough-fuzzy models and temporal attention mechanisms further extend these capabilities to structured video summarization and anomaly classification [21], [22]. Moreover, edge-server architectures have been introduced to support the detection of abnormal pedestrian behaviors in large-scale closed-circuit television environments, enabling scalable and real-time analysis [23].

In this paper, we propose a simulation-based framework that unifies anomaly generation, metric logging, and predictive modeling. Using the SUMO platform, we construct reproducible collision scenarios and matched baselines for systematic evaluation. The framework records fine-grained per-vehicle metrics, enabling both localized and network-wide analysis, including travel time index comparisons [24]. We then develop a hybrid forecasting architecture combining bidirectional LSTM (BiLSTM) networks with diffusion convolutional recurrent neural

This research was supported in part by NSF under Grants CNS-2401007, CMMI-2348381, IIS-2415478, and in part by MathWorks.

T. Kinchen and N. V. Senthil Kumar are with the Systems Engineering program, Cornell University, Ithaca, NY, USA. Email: {tjk238, ns942}@cornell.edu

T. Bai and A. A. Malikopoulos are with the School of Civil & Environmental Engineering, Cornell University, Ithaca, NY, USA. E-mails: {tingbai, amaliko}@cornell.edu

networks (DCRNN) to capture both temporal dependencies and spatial dynamics [25]. This approach enables real-time prediction of traffic disruptions and reveals how controlled anomalies impact urban mobility. Note that the BiLSTM–DCRNN architecture follows established spatiotemporal forecasting practice and is adopted here as a testbed rather than as a methodological contribution. The main contributions of this paper are as follows: (i) We design a reproducible SUMO-based incident generation pipeline with matched baseline scenarios that enables controlled and fair benchmarking. (ii) We establish a unified metric collection and temporal–spatial alignment protocol for vehicle-, edge-, and network-level analysis. (iii) We develop an interval-based spatiotemporal localization and forecasting formulation that links event-level collision containment with post-incident traffic propagation. Simulation results demonstrate high spatiotemporal containment accuracy for collision localization and accurate multi-horizon forecasts of congestion and emissions, showing the validity of the proposed integrated framework.

II. PROBLEM FORMULATION

We consider a mixed urban traffic network simulated in SUMO with scripted collision events. The road network is characterized as a directed graph $\mathcal{G}(\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is the node-set with each node representing a road intersection, an entry, or an exit point. Edges in $\mathcal{E}$ denote road segments connecting pairs of nodes. Time is discretized as $\mathcal{T} = \{1, \dots, T\}$. For each edge $e \in \mathcal{E}$ and time $t \in \mathcal{T}$, let $f_{e,t} \in \mathbb{R}^2$ denote a feature vector comprising the travel time index (TTI) and CE, which capture the congestion level and environmental cost on that edge, respectively. More specifically, we define the TTI [24] on edge e at time t as $\text{TTI}_{e,t} = \text{TT}_{e,t} / \text{TT}_e^{\text{ff}}$, where $\text{TT}_{e,t}$ is the observed travel time to traverse edge e based on the real traffic flow at time t , and TT_e^{ff} denotes the free-flow travel time without congestion. CE captures all vehicles' emissions on edge e within the time interval $[t-1, t]$, and we describe it using a fitted polynomial function of vehicle speed v (m/s) and acceleration a (m/s²), expressed as $\text{CE}(v, a) = c_0 + c_1 va + c_2 va^2 + c_3 v + c_4 v^2 + c_5 v^3$, where the coefficients $c_0, \dots, c_5$ are determined empirically for each pollutant and vehicle class [26]. These two metrics, TTI and CE, form the basis for modeling network-level traffic states.

With these edge-level definitions of travel time and emissions, we next specify the vehicle classes and admissible events used in our scenarios. The vehicle set $\mathcal{V}$ comprises passenger vehicles (PVs), buses, and automated vehicles (AVs), i.e., $\mathcal{V} = \{\text{PV}, \text{bus}, \text{AV}\}$. Event admissibility follows $\mathcal{R}$, which we define as

$$\mathcal{R} := \{(\text{PV}, \text{PV}, \text{rear}), (\text{PV}, \text{bus}, \text{rear}), (\text{PV}, \text{AV}, \text{rear}), (\text{PV}, \text{PV}, \text{intersection})\}, \quad (1)$$

where the second vehicle in each combination is the leading vehicle in the interaction. For intersections, we treat the scenario under the assumption that PVs exhibit more reckless driving and uncoordinated entry into unsignalized intersections, unlike buses with professional drivers or AVs

with automated stopping features, and therefore restrict events to PV-PV interactions. In this paper, we aim to present a framework that is capable of reproducing consistent incident conditions, quantifying their effects, and providing accurate multi-horizon traffic forecasts. To this end, we employ a spatiotemporal forecasting model, DCRNN, for predicting congestion and emissions, complemented by a BiLSTM module that provides interval estimates of collision time and location.

A. Collision Scripting and Admissibility

We use collision scripting to introduce controlled crashes into the SUMO simulation to reproduce repeatable disruption scenarios. Let $\mathcal{V} = \{\text{PV}, \text{bus}, \text{AV}\}$ denote the set of vehicle classes, $\mathcal{S} = \{\text{rear}, \text{inter}\}$ the set of event types, and $\mathcal{R} \subseteq \mathcal{V} \times \mathcal{V} \times \mathcal{S}$ the admissibility relation. Only events that lie in $\mathcal{R}$ are retained for learning and evaluation. Specifically, we consider (PV, PV, rear), (PV, bus, rear), (PV, AV, rear), and (PV, PV, inter). Two event classes are scripted: rear-end and intersection. We assume rear-end events to occur among PV-PV, PV-bus, and PV-AV, while intersection conflicts are restricted to PV-PV. Events are inserted by halting vehicles at designated locations for a dwell duration and resuming movement thereafter; logs record (x, y, t) , where $x, y \in \mathbb{R}$ are the spatial coordinates of the collision in meters, and $t \in \mathcal{T}$ denotes the simulation time in seconds.

B. Spatiotemporal Forecasting

We aim to forecast edge-level values of TTI and CE over multiple horizons. Let $f_{e,t}$ denote the feature vector for edge $e \in \mathcal{E}$ at time $t \in \mathcal{T}$. The forecasting problem is defined as learning to approximate the true target values $z_{e,t}^{(q)}$, where $z_{e,t}^{(q)}$ represents the scalar corresponding to metric $q \in \{\text{TTI}, \text{CE}\}$ from past observations. For a forecast origin τ , the model ingests the recent lookback window $\{f_{e,\tau-L+1}, \dots, f_{e,\tau}\}$ across all edges, where L is the look-back length and $|\mathcal{E}|$ is the number of edges. The model then produces forecasts $\{\hat{z}_{e,\tau+h}^{(q)}\}_{h=1}^H$, where H is the forecast horizon and each $\hat{z}$ corresponds to a prediction of TTI or CE.

The training objective reflects the interval nature of predictions. Let β be a tunable hyperparameter that controls the trade-off between interval sharpness and coverage. For each target quantity $q \in \{\text{TTI}, \text{CE}\}$, the model outputs lower and upper bounds $(\hat{z}_{\text{low}}^{(q)}, \hat{z}_{\text{high}}^{(q)})$. The loss penalizes under- and over-coverage while encouraging tight intervals, with additional weighting for rare high-magnitude spikes by

$$\mathcal{L} = \frac{1}{N} \sum_{i=1}^N \left[\beta (\hat{z}_{\text{high},i}^{(q)} - \hat{z}_{\text{low},i}^{(q)}) + (\max\{0, \hat{z}_{\text{low},i}^{(q)} - z_i^{(q)}\} + \max\{0, z_i^{(q)} - \hat{z}_{\text{high},i}^{(q)}\}) \cdot w_i^{(q)} \right], \quad (2)$$

where N is the number of training samples, $w_i^{(q)}$ is a target-specific weight that emphasizes rare spike events in either TTI or CE, and $\beta > 0$ penalizes excessively wide prediction intervals. At inference time, intervals are further calibrated using where residuals on a held-out calibration set determine edge-specific padding values to ensure desired coverage.




Fig. 1: Flow chart of the modeling framework.

C. Event Localization with BiLSTM

To localize collision occurrences, our objective is to predict when and where an event will take place in terms of its spatial coordinates (x, y) and time t . For each dimension $z \in \{x, y, t\}$, the model utilizes the upper and lower predictions to define an interval expected to contain the true event value z^* . We train the BiLSTM regressor by minimizing the mean squared error between the predicted bounds and the observed values, encouraging the intervals to be both accurate and appropriately tight. We evaluate model performance using root mean squared error of the endpoints, the probability that true events fall inside the predicted intervals (coverage), the average width of the predicted ranges, and the Dice score, which quantifies overlap between predicted and actual event intervals.

III. MODELING FRAMEWORK

This section presents the modeling framework used to capture both the spatiotemporal occurrence of collision events and their impacts on system-level traffic metrics, including TTI and CE. The framework integrates event localization with network-level forecasting to characterize how localized disruptions propagate through urban traffic networks.

As illustrated in Fig. 1, the framework consists of two complementary components: (i) event localization, which estimates the spatiotemporal coordinates of collision events, and (ii) spatiotemporal forecasting, which predicts the evolution of edge-level traffic states over multiple horizons. Both components rely on recurrent architectures that leverage temporal sequences and graph-structured representations of the road network. Below, we describe the BiLSTM encoder, the diffusion convolution operator, and the diffusion-convolutional recurrent module with fusion layers.

A. BiLSTM Encoder

Bidirectional long short-term memory (BiLSTM) networks are used to capture temporal dependencies by processing input sequences in both forward and backward directions. Given a sequence of features $\{r^{(t)}\}_{t=1}^T$ with $r^{(t)} \in \mathbb{R}^{F_{\text{seq}}}$, the BiLSTM encoder concatenates the final hidden states from both directions to form a representation in $\mathbb{R}^{2d_{\text{LSTM}}}$, where d_{LSTM} is the hidden dimension in one direction. This encoding captures both past and future temporal context and serves as the basis for spatiotemporal event localization. In implementation, the BiLSTM consists of two stacked bidirectional layers with hidden dimension d_{LSTM} in each direction. The final forward and backward hidden states are concatenated and passed through three parallel linear output heads that produce lower

and upper bounds for each of x , y , and t , yielding a six-dimensional interval representation per prediction horizon.

B. Diffusion Convolution

To model directed spatial dependencies on the traffic network $\mathcal{G}(\mathcal{N}, \mathcal{E})$, we adopt the diffusion convolution introduced by Li et al. [25]. The forward and backward row-normalized transition matrices are defined as

$$S_f = D_O^{-1}W, \quad S_b = D_I^{-1}W^T, \quad (3)$$

where W is the weighted adjacency matrix, $D_O = \text{diag}(WM)$, and $D_I = \text{diag}(W^T M)$, with M encoding network connectivity.

For a graph signal $G \in \mathbb{R}^{N \times P}$, the K -step bidirectional diffusion convolution is given by

$$\sum_{k=0}^{K-1} (\theta_{k,1} S_f^k + \theta_{k,2} S_b^k) G_{:,p}, \quad p = 1, \dots, P,$$

where K controls the diffusion depth, $\theta_{k,1}$ and $\theta_{k,2}$ are trainable parameters, and $*_G$ denotes diffusion convolution as a weighted polynomial of S_f and S_b applied to $G_{:,p}$. Here, S_f^k and S_b^k denote matrix powers with $S_f^0 = S_b^0 = I$, and $P = 5$ corresponds to the edge-level traffic features. This formulation enables the model to capture both upstream and downstream spatial influences on each edge and has been shown to be effective for multi-horizon traffic forecasting [27].

C. DCRU Encoder and Fusion Multi-Layer Perception

The diffusion-convolutional gated recurrent unit (DCGRU) extends the standard GRU by replacing dense transformations with diffusion convolutions [25]. At each step $k = 1, \dots, T_g$, the hidden state $h^{(k)}$ is updated using graph signals $\{G^{(k)}\}$ and the previous hidden state $h^{(k-1)}$, enabling the model to jointly capture temporal evolution and spatial propagation of traffic states. The graph-level representation is obtained by pooling the hidden states

$$g = \text{Pool}(\{h^{(k)}\}_{k=1}^{T_g}) \in \mathbb{R}^{d_{\text{DCGRU}}}, \quad (4)$$

where $\text{Pool}(\cdot)$ denotes a temporal aggregation function. The pooled representation is concatenated with other encodings and passed through a feed-forward fusion network

$$h = \phi(W_2 \phi(W_1 u + c_1) + c_2), \quad (5)$$

where $\phi(\cdot)$ is a pointwise nonlinearity. Multi-horizon forecasts are produced via step-specific affine heads

$$\hat{p}^{(s)} = U_s h + d_s, \quad s = 1, \dots, S, \quad (6)$$

with U_s and d_s denoting horizon-specific parameters.

IV. SIMULATION

A. Simulation Setup

As shown in Fig. 2, our study focuses on a segment of the Broadway corridor in New York City that includes multiple urban intersections. To capture realistic traffic heterogeneity, we populate the network with PVs, buses, and AVs in proportions summarized in Table I. We applied a mixed




Fig. 2: SUMO map of NYC Broadway intersections. (a) Rear-end (AV-PV). (b) Rear-end (Bus-PV). (c) Rear-end (PV-PV). (d) Intersection (PV-PV).

TABLE I: Vehicle composition and incident involvement

Vehicle Type	Total Number	Incident Involvement
PV	330	Intersection, Rear-end
Bus	250	Rear-end
AV	80	Rear-end

routing scenario to reflect diverse travel behaviors across vehicles and to generate realistic interaction patterns for the incident simulations. Training, validation, and test datasets are constructed from independent simulation runs with different random seeds, while preserving identical network geometry, traffic demand, and admissibility rules. This design isolates generalization across stochastic traffic realizations while holding structural conditions fixed, reducing the risk of memorizing deterministic simulator behavior.

1) *Collision Incident*: The simulation process begins with the creation of a collision-injected route file. We use a scripted pre-processing step to insert artificial stopping events into selected vehicle trajectories to simulate two common collision types: rear-end and intersection collisions. A rear-end collision models a PV that does not stop in time when approaching a stationary vehicle ahead, such as a bus at a stop, another PV, or an AV. To capture typical rear-end dynamics in a mixed-traffic network, the lagging PV is assigned the maximum speed, representing a speeding vehicle colliding with the stationary one. We generate intersection collisions as two PVs approaching an intersection from perpendicular directions. One vehicle halts after entering the intersection, while the other approaches in a higher direction from the side, mimicking a scenario where a speeding PV collides with a stopped vehicle inside the intersection. We simulate a collision-injected route file using all vehicles, while omitting collision stops in the control variant. We obtain free-flow baselines by simulating each control variant vehicle independently.

2) *Congestion Incident*: To quantify the impact of congestion on network performance, we adopt the TTI as described by Erdelic et al. [24], which serves as a standardized measure of congestion. As introduced earlier in Section II, TTI represents the ratio of actual travel time to an ideal free-flow benchmark, where $TT_{e,t}$ is the observed dwell time on edge e during collision-injected simulation and TT_e^{ff} denotes the travel time under free-flow, no-collision conditions. Values greater than 1 indicate congestion-induced delays.




Fig. 3: Stacked histogram of actual counts inside predicted t range for rear-end collisions by vehicle pair.

An event is considered correctly localized if the true collision coordinates (x^*, y^*, t^*) lie simultaneously within their respective predicted intervals. Figs. 3–6 report marginal containment along individual dimensions for interpretability.

B. Results Analysis

1) *Collision Event Localization*: Figs. 3–6 present the containment performance of the BiLSTM in predicting collision time t and spatial coordinates (x, y) . The stacked histograms report the number of actual events captured within the predicted intervals across different collision types.

For *rear-end collisions*, the model consistently captures most actual events within the predicted ranges. PV-PV collisions dominate the dataset, producing high frequencies at tighter containment levels (see Figs. 3–5). Predictions for PV-AV and PV-Bus collisions are more dispersed but still achieve reasonable coverage, indicating that the model generalizes well to heterogeneous vehicle interactions. Notably, temporal containment t is particularly strong, with most collisions captured within narrow time horizons. In Fig. 4, PV-Bus collisions exhibit wider dispersion along the y-axis compared to PV-PV and PV-AV cases, reflecting greater variability in how buses interact spatially with trailing vehicles.

For *intersection collisions* (PV-PV only), predictions are even sharper. As shown in Fig. 6, the model achieves near-complete containment with narrow predicted intervals across t , x , and y . This performance reflects the deterministic spatiotemporal structure of intersection events and the fact that we conducted experiments on a single intersection geometry at varying times. The reduced variability in roadway context likely contributes to the higher prediction accuracy compared to rear-end collisions. As we can see, the predicted intervals tightly overlap with the ground-truth (x, y, t) collision points, with only minor temporal deviations, demonstrating the model’s ability to provide reliable containment ranges suitable for real-time event flagging.

2) *Network-Level Impacts*: Fig. 7 and Fig. 8 show the DCRNN’s predicted network-level CE emissions and TTI, respectively. These results highlight how the model captures the propagation of anomalies through the network and their sustainability consequences. As shown in Fig. 7, collisions




Fig. 4: Stacked histogram of actual counts inside predicted y range for rear-end collisions by vehicle pair.




Fig. 5: Stacked histogram of actual counts inside predicted x range for rear-end collisions by vehicle pair.

and congestion significantly elevate emissions on affected links. The model accurately reproduces both the magnitude and spatial distribution of emissions, though in certain areas, predicted values overshoot the ground truth, particularly along heavily congested edges. Compared to the free-flow scenario, where emissions remain low and uniform, the experimental case exhibits sharp spikes near rear-end sites, reflecting the impact of stop-and-go traffic and idling.

For TTI, the DCRNN predictions (see Fig. 8) effectively capture localized spikes induced by anomalies. Forecasts align with observed patterns of sharp increases in travel times downstream of collisions, though peak magnitudes are sometimes amplified. In the free-flow scenario, vehicles move smoothly across the network with low TTI values close to baseline conditions, indicating minimal delay. By contrast, the experimental case shows a sharp rise in TTI exceeding 10, especially along central vertical and horizontal corridors. The comparison demonstrates that collisions and disruptions introduce localized bottlenecks that propagate through adjacent segments.

3) Integrated Model Perspective: The above results highlight the complementary strengths of the two modeling approaches. The BiLSTM provides fine-grained localization of collision events, achieving narrow interval predictions in both space and time (see Fig. 6), which is critical for pinpointing disruptions as they occur in real time. The DCRNN, in contrast,




Fig. 6: Predicted ranges in (x, y, t) space for intersection collisions compared to the actual observed collisions, as shown by the red nodes.




Fig. 7: Total CE compared between the experimental (left) and free-flow (right) scenarios.

captures how these localized anomalies propagate across the network, forecasting their effects on congestion and emissions at broader scales (see Figs. 7–8). Together, the models bridge event-level detection and system-level forecasting: the BiLSTM identifies where and when disruptions arise, while the DCRNN predicts how they spread and affect sustainability outcomes. This integration enables both precise detection and system-wide prediction, laying the foundation for anomaly-aware traffic management strategies.

4) Limitations and Future Work: Despite the promising results, the framework has some limitations. The simulation was tested mainly on a representative set of scripted collisions in a single urban network, limiting scenario diversity and variability. While the DCRNN captured the spatial footprint of congestion and emissions, it occasionally over- or underestimated magnitudes that were found to be spikes in the dataset, reflecting challenges in modeling nonlinear traffic-environment feedback. Additionally, the framework currently operates in an offline evaluation mode, without integration into live traffic management systems. These constraints suggest the need for further validation and refinement before deployment in real-world settings.

We will focus future work on extending the framework in several directions. Incorporating additional collision types, multimodal interactions, and more complex urban networks will improve the robustness of the models. Embedding adaptive control strategies such as real-time rerouting or




Fig. 8: The TTI compared between experimental (right) and free-flow (left) scenarios.

signal timing within the prediction loop could demonstrate how anomaly-aware forecasts directly support mitigation and sustainability objectives. Finally, exploring online learning and real-time deployment through traffic sensor streams would enable proactive responses to emerging disruptions. By bridging predictive modeling with control and decision support, the proposed framework has the potential to advance sustainable and resilient urban mobility.

V. CONCLUSION

In this paper, we presented a simulation-based framework that unifies controlled anomaly generation, metric logging, and hybrid deep learning to study the impacts of collisions and congestion on urban traffic networks. We showed that by integrating BiLSTM models with a DCRNN for network-level forecasting, the framework can detect, localize, and predict anomalies while quantifying their effects on both mobility and sustainability. Our results show that the BiLSTM achieves strong containment accuracy, particularly for intersection collisions, while the DCRNN effectively translates localized disruptions into forecasts of travel time index and CE across the network. These contributions establish a reproducible and extensible pipeline for resilient traffic management research. A potential direction of future research should focus on extending the framework with multimodal interactions and larger urban networks to strengthen model robustness. In addition, embedding adaptive control strategies such as real-time rerouting will enable proactive responses to disruptions.

REFERENCES

- [1] A. Choudhary, S. Gokhale, P. Kumar, C. Pradhan, and S. K. Sahu, "Urban traffic congestion: Its causes-consequences-mitigation," *Research Journal of Chemistry and Environment*, vol. 26, no. 12, pp. 164–176, 2022.
- [2] X. Ru, T. Bai, W. Xia, and K. Henrik Johansson, "Adaptive fault-tolerant perimeter control for two-region networks with actuator faults," *IEEE Transactions on Intelligent Transportation Systems*, vol. 27, no. 1, pp. 1170–1181, 2026.
- [3] M. Barth and K. Boriboonsomsin, "Real-world CO₂ impacts of traffic congestion," *Transportation Research Record*, vol. 2058, no. 1, pp. 163–171, 2008.
- [4] A. I. Mahbub, V.-A. Le, and A. A. Malikopoulos, "A safety-prioritized receding horizon control framework for platoon formation in a mixed traffic environment," *Automatica*, vol. 155, p. 111115, 2023.
- [5] T. Bai, A. Johansson, K. H. Johansson, and J. Mårtensson, "Large-scale multi-fleet platoon coordination: A dynamic programming approach," *IEEE Transactions on Intelligent Transportation Systems*, vol. 24, no. 12, pp. 14 427–14 442, 2023.
- [6] —, "Approximate dynamic programming for platoon coordination under hours-of-service regulations," in *61st IEEE Conference on Decision and Control (CDC)*, 2022, pp. 7663–7669.
- [7] V.-A. Le, B. Chalaki, F. N. Tzortoglou, and A. A. Malikopoulos, "Stochastic time-optimal trajectory planning for connected and automated vehicles in mixed-traffic merging scenarios," *IEEE Transactions on Control Systems Technology*, vol. 33, no. 4, pp. 1403–1417, 2025.
- [8] A. A. Malikopoulos and L. Zhao, "Optimal path planning for connected and automated vehicles at urban intersections," in *58th IEEE Conference on Decision and Control (CDC)*, 2019, pp. 1261–1266.
- [9] H. Bang and A. A. Malikopoulos, "Congestion-aware routing, rebalancing, and charging scheduling for electric autonomous mobility-on-demand system," in *American Control Conference (ACC)*, 2022, pp. 3152–3157.
- [10] T. Bai, A. Li, G. Xu, C. G. Cassandras, and A. A. Malikopoulos, "Routing guidance for emerging transportation systems with improved dynamic trip equity," in *2025 IEEE 64th Conference on Decision and Control (CDC)*, 2025, pp. 1168–1173.
- [11] A. A. Malikopoulos and L. Zhao, "A closed-form analytical solution for optimal coordination of connected and automated vehicles," in *2019 American Control Conference (ACC)*. IEEE, 2019, pp. 3599–3604.
- [12] A. A. Malikopoulos, L. E. Beaver, and I. V. Chremos, "Optimal time trajectory and coordination for connected and automated vehicles," *Automatica*, vol. 125, no. 109469, 2021.
- [13] O. ElSahly and A. Abdelfatah, "A systematic review of traffic incident detection algorithms," *Sustainability*, vol. 14, no. 22, p. 14859, 2022.
- [14] Y. J. Stephanedes, A. P. Chassiakos, and P. G. Michalopoulos, "Comparative performance evaluation of incident detection algorithms," *Transportation Research Record*, vol. 1360, pp. 50–57, 1992.
- [15] M. Xu, J. Wu, H. Wang, and M. Cao, "Anomaly detection in road networks using sliding-window tensor factorization," *IEEE Transactions on Intelligent Transportation Systems*, vol. 20, no. 12, pp. 4704–4713, 2019.
- [16] W. Zhao, W. Xu, J. Sun, and H. Zhang, "Prediction-based anomaly detection method for traffic flow," *IEEE Transactions on Intelligent Transportation Systems*, vol. 22, no. 11, pp. 6846–6855, 2021.
- [17] X. Peng, Y. Lin, Q. Cao, Y. Cen, H. Zhuang, and Z. Lin, "Traffic anomaly detection in intelligent transport applications with time series data using informer," in *25th IEEE International Conference on Intelligent Transportation Systems (ITSC)*, 2022, pp. 3337–3342.
- [18] J. Li, X. Wang, J. Liu, X. Li, X. Liu, B. Zheng, and H. Yao, "Deepflow: Abnormal traffic flow detection using siamese networks," *arXiv preprint arXiv:2004.12094*, 2020.
- [19] X. Kang, Y. Li, Y. Wu, H. Duan, and L. Zhang, "Traffic anomaly detection in intelligent transport applications with time series data using informer," *Sensors*, vol. 23, no. 4552, 2023.
- [20] S. W. Khan, Q. Hafeez, M. I. Khalid, R. Alroobaea, S. Hussain, J. Iqbal, J. Almotiri, and S. S. Ullah, "Anomaly detection in traffic surveillance videos using deep learning," *Sensors*, vol. 22, no. 17, p. 6563, 2022.
- [21] A. Pramanik, S. K. Pal, J. Maiti, and P. Mitra, "Traffic anomaly detection and video summarization using spatio-temporal rough fuzzy granulation with z-numbers," *IEEE Transactions on Intelligent Transportation Systems*, vol. 23, no. 12, pp. 24 116–24 128, 2022.
- [22] S. Natha, F. Ahmed, M. Siraj, M. Lagari, M. Altamimi, and A. A. Chandio, "Deep bilstm attention model for spatial and temporal anomaly detection in video surveillance," *Sensors*, vol. 25, no. 1, p. 251, 2025.
- [23] J. Song and J. Nang, "Pedestrian abnormal behavior detection system using edge-server architecture for large-scale CCTV environments," *Applied Sciences*, vol. 14, no. 4615, 2024.
- [24] T. Erdelic, T. Caric, M. Erdelic, L. Tisljaric, A. Turkovic, and N. Jelusic, "Estimating congestion zones and travel time indexes based on floating car data," *Computers, Environment and Urban Systems*, vol. 90, p. 101604, 2021.
- [25] Y. Li, R. Yu, C. Shahabi, and Y. Liu, "Diffusion convolutional recurrent neural network: Data-driven traffic forecasting," in *International Conference on Learning Representations*, 2018.
- [26] D. Krajzewicz, M. Behrisch, P. Wagner, S. Hausberger, and M. Krumnow, "Second generation of pollutant emission models for SUMO," in *Proceedings of the SUMO2014 Conference "Modeling Mobility with Open Data"*, Berlin, Germany, 2014, pp. 1–22.
- [27] F. Sun et al., "Diffusion convolutional recurrent neural network for traffic forecasting: Reproducibility study and insights," *arXiv preprint arXiv:2303.16321*, 2023.