Short-Term Traffic Congestion Prediction Using Public Transportation Data

Robert-Lucian Kis
Department of Electrical Engineering and Information Technology George Emil Palade University of Medicine, Pharmacy, Science, and Technology of Targu Mures, Romania
kis.robert-lucian.23@stud.umfst.ro

Adrian-Silviu Roman
Department of Electrical Engineering and Information Technology George Emil Palade University of Medicine, Pharmacy, Science, and Technology of Targu Mures, Romania
adrian.roman@umfst.ro

Teri Lenard
Department of Electrical Engineering and Information Technology George Emil Palade University of Medicine, Pharmacy, Science, and Technology of Targu Mures, Romania
teri.lenard@umfst.ro

***Abstract*—Urban traffic congestion imposes significant costs on commuters, the local economy, and the environment, with its prediction remaining a challenging problem due to the complex spatiotemporal dynamics of traffic flow. In this paper, we present a congestion prediction approach for short-term forecasting that leverages location data from public bus data in the city of Târgu Mureș, Romania. The proposed pipeline integrates Automatic Vehicle Location (AVL) data collection, map matching, Travel Time Index (TTI) estimation, and machine learning-based forecasting to deliver 15-minute-ahead congestion predictions. Ten prediction architectures of increasing complexity were evaluated, ranging from naive baselines to recurrent and graph-based deep learning models. The results demonstrate that temporal traffic history is essential for accurate forecasting, and that Graph WaveNet achieves the lowest Mean Absolute Error (MAE = 0.546), outperforming models that rely on the fixed road graph structure. The Spatial-Temporal Identity (STID) model achieves superior results on RMSE, MAPE, and R^2, representing an efficient alternative. The findings confirm that publicly available AVL data can serve as a reliable and cost effective source for urban traffic monitoring in cities lacking dedicated sensor infrastructure.**

***Index Terms*—traffic congestion prediction, public transportation data, automated vehicle location, applied machine learning**

I. Introduction

Urban traffic congestion is a challenge in modern cities that directly affects the costs of public and private commuting, the surrounding environment health, and the local economy. The ability to predict traffic congestion and issue warnings enables road users to make informed routing decisions, allowing in parallel traffic management authorities to respond proactively rather than reactively [1]. Short-term traffic forecasting is particularly useful for real-time applications, but it remains a difficult problem due to the complex dynamics of urban traffic flow.

Traditionally, traffic congestion prediction relies on data from sources such as roadside infrastructure (e.g., radar detectors, video recorders), that requires significant initial investment and continuous maintenance, while offering coverage only at fixed installation points [2], [3]. Moreover, sensors are exposed to degradation and installation damage, further limiting their reliability [4]. The use of video cameras and other sensors-based systems also raises privacy concerns, as continuously collected data may involve personally identifiable information of pedestrians and drivers [5].

An alternative is to use public vehicles equipped with Global Positioning System (GPS) systems that log and report location data to a centralized server. This approach is more cost effective and offers broader spatial coverage compared to fixed roadside infrastructure [6]. Public transportation vehicles represent a valuable source of such data, as they are fitted with Automatic Vehicle Location (AVL) systems that record vehicle location at fixed time intervals [7]. Since such vehicles operate on fixed routes across the road network, they function as mobile probes of the traffic stream, providing reliable and low-cost data for traffic state estimation and Travel Time Estimation (TTE) [3], [8].

Despite the growing interest in AVL-based traffic monitoring, several research gaps remain. Most existing studies [1], [9], [10] rely on data from large, sensor-rich cities, leaving medium-sized urban areas underexplored. Additionally, the suitability of state-of-the-art deep learning models for such scenarios has not been systematically evaluated.

To address these gaps, this paper presents a short-term traffic congestion prediction system for the city of Târgu Mureș (Romania), built exclusively on GPS data collected from the municipal public transit fleet via a publicly available source. The main contribution of this work is three-fold:

1) A complete congestion prediction pipeline, from raw AVL data collection and map matching to Travel Time Index (TTI) estimation and short-term forecasting, demonstrating that traffic prediction is achievable without proprietary sensor infrastructure or specialized datasets;
2) A comparative evaluation of several prediction models of increasing complexity (i.e., Multilayer Perceptron (MLP), Gated Recurrent Unit (GRU), Long Short-Term Memory (LSTM), Graph Neural Network

979-8-3195-0700-6/26/$31.00 © 2026 IEEE

(GNN) with GATv2, Spatio-Temporal Graph Neural Network (STGNN) with T-GCN, Spatial-Temporal Identity (STID), and Graph WaveNet) on a sparse road graph; and
3) Evidence that Graph WaveNet achieves the lowest Mean Absolute Error (MAE) among all evaluated architectures, and that models with learned or implicit spatiotemporal representations outperform those relying on fixed graph structures in sparse network settings.

The remainder of this paper is organized as follows. Section II presents the background and the related work on traffic data collection, road network representation, deep learning approaches on traffic prediction and traffic congestion prediction. Section III describes the proposed approach for short-term traffic congestion prediction. The paper continues with Section IV where the experimental results are discussed. Finally, Section V concludes the paper, presents the limitations and directions for future work.

II. Background and Related Work

A. Traffic Data Collection

Accurate traffic congestion prediction depends on the availability of reliable traffic data. Traditional roadside infrastructure has long been the backbone of urban traffic monitoring [2], [3]. Additionally, traffic congestion research has relied on data collected from probe vehicles such as taxis and smartphone-based sensors [11]–[13].

In many cities, public transit buses are equipped with GPS devices, which provide AVL data at fixed time intervals [7]. The data recorded by these units installed in buses are a potential and reliable source of travel time and speed data. Thus, these buses can be regarded as useful probe vehicles for inferring traffic conditions, providing continuous and real-time information. This data is both reliable and cost-effective compared to other data collection methods. AVL data from buses have proven useful in traffic studies, particularly for TTE and profiling [8], as well as for traffic state estimation [7].

B. Road Network Representation

Modeling urban traffic as a graph-structured problem has become standard practice in the Intelligent Transportation Systems (ITS) literature [14], [15]. In a graph representation of a road network, nodes correspond to intersections and edges to road segments, allowing spatial relationships between locations to be expressed explicitly. A popular approach in this direction are solutions such as OSMnx [16], which automates the retrieval and construction of directed road network graphs from OpenStreetMap data, enabling the modeling of real-world urban street networks. The graph formulation is essential for traffic analysis, as it preserves one-way constraints and turn restrictions. In this work, the road network of Târgu Mureş is modeled as a directed graph using OSMnx, with intersections as nodes and road segments as edges, onto which the bus routes and associated GPS observations are mapped.

C. Machine Learning for Traffic Prediction

Traffic prediction methods have evolved from classical statistical approaches toward deep learning architectures. Early methods relied on statistical time-series models such as ARIMA and its variants, which assume linearity and stationarity, limiting their effectiveness in capturing the complex dynamics of urban traffic [1]. Machine learning approaches including Support Vector Machines (SVMs) and Random Forests (RFs) improved upon these limitations, but required extensive manual feature engineering and struggled with long-range temporal dependencies [1].

Recurrent Neural Network (RNN) architectures advanced the field by enabling models to learn from sequential traffic histories. LSTM networks [17] address the vanishing gradient problem inherented from standard recurrent networks, allowing the model to retain relevant information over extended time horizons. The GRU [18] offers a lighter alternative to LSTM with comparable performance on traffic prediction tasks. Both architectures have been widely used for short-term traffic speed and flow prediction, by capturing the temporal dynamics of congestion accumulation and dissipation [1].

An alternative to recurrent architectures, are graph-based architectures which use learned embeddings that encode the identity of each road segment and time slot. Shao et al. [19] proposed STID, a lightweight model based on simple MLP layers that attaches spatial and temporal identity embeddings directly to the input features. Despite its architectural simplicity, STID has been shown to match or exceed the performance of more complex STGNN architectures across multiple benchmark datasets [19].

GNNs have become a popular approach for spatiotemporal traffic forecasting, as they encode road network topology and enable information propagation between adjacent segments. The T-GCN model [14], a representative STGNN architecture, combines a Graph Convolutional Network (GCN) for spatial modeling with a GRU for temporal modelling, and has become a widely used benchmark for urban traffic prediction. However, T-GCN and related models assume that the road network graph accurately reflects the true spatial dependencies between segments. This assumption may not hold when the graph is sparse or incomplete.

Graph WaveNet [15] addresses this by combining causal convolutions with a self-adaptive adjacency matrix learned end-to-end through node embeddings. It enables the model to discover hidden spatial dependencies without relying on the fixed graph structure. For attention-based neighborhood aggregation, GATv2 [20] extends the original Graph Attention Network (GAT) by enabling fully dynamic attention and has been shown to outperform GAT across multiple benchmarks at no additional parameter cost.

Despite promising results, graph-based spatiotemporal models are sensitive to the density and completeness of the underlying graph. In the current approach, the road graph derived from bus routes is sparse and predominantly linear, which limits the performance of models that depend on fixed graph connectivity. This is discussed in more detail in Section IV.

A systematic review by Attioui and Lahby [1] traced the evolution of traffic congestion forecasting, identifying data-related issues (including limited sensor coverage, sparse observations, and noisy inputs) as the most frequently cited limitation. This is consistent with the challenges encountered in the present work, where traffic observations are restricted to road segments covered by bus routes. Additionally, the review notes that advanced deep learning architectures achieve higher prediction accuracy at the cost of increased computational requirements. This reinforces the importance of selecting models that balance performance with practical implementation restrictions, as in data-constrained urban environments.

III. Proposed Approach

The proposed system for short-term traffic congestion prediction consists of the following steps: (i) data collection, (ii) map matching, and (iii) congestion modeling and prediction.

A. Data Collection

The first step consists of collecting location data from AVL systems and storing it in a centralized database. Each record contains the vehicle identifier, the active route, GPS coordinates in the form of latitude and longitude, and a timestamp. The raw data is cleaned to remove noise and correct positioning errors introduced by the GPS equipment.

B. Map Matching

Raw GPS coordinates do not carry information about the underlying road structure. Map matching is a necessary preprocessing step that associates each recorded position with a specific edge of the road network graph. This process allows vehicle trajectories to be expressed in terms of the directed graph, where nodes represent intersections and edges represent road segments. Without map matching, it would not be possible to aggregate observations by segment or to compute per-segment traffic indicators.

C. Congestion Modeling

To perform congestion modeling, the matched observations are aggregated into 15-minute time windows. A *transit* is defined as the sequence of records generated by a single vehicle traversing a specific road segment within one time window. The transit time and distance are used to compute the average speed for each segment and window. Validation filters are applied to ensure data quality: records corresponding to dwell time at stops and anomalous speeds caused by GPS errors are excluded.

Based on the validated speed observations, two traffic indicators are defined. The *free-flow speed* represents the vehicle's speed observed under uncongested conditions, and is estimated for each segment by analyzing speed recordings during off-peak hours (00:00–05:00), using the 85th percentile of observed speeds. This percentile is widely used in traffic engineering as a robust estimate of the speed at which vehicles travel in the absence of congestion. Based on this reference speed, the TTI (Eq. 1) is computed for each segment and time window as the ratio of actual travel time to ideal (free-flow) travel time:

$$\text{TTI} = \frac{TT_{\text{actual}}}{TT_{\text{ideal}}}, \tag{1}$$

where TT_{actual} is the observed travel time for the segment in the given time window, and TT_{ideal} is the estimated travel time under free-flow conditions. A TTI value of 1.0 indicates free-flow conditions, while higher values indicate increasing levels of congestion.

The objective of the prediction models is to estimate the TTI for each road segment 15 minutes ahead, formulated as a supervised regression task. The input to each model consists of an 8-step historical window (covering the preceding 2 hours at 15-minute resolution), along with contextual features including time of day, day of week, meteorological conditions, and static road segment properties such as length and number of lanes. The target variable is the continuous TTI value, capped in the range $[1.0, 10.0]$, where the lower bound reflects free-flow conditions and the upper bound represents an extreme congestion threshold. For the user interface, TTI values are mapped to discrete congestion levels (e.g., free flow for TTI < 1.3, severe congestion for TTI ≥ 5.0). However, all models are trained and evaluated on the continuous scalar value using standard regression loss functions.

D. Prediction Models

To quantify the contribution of different types of information (e.g., historical, contextual, and topological) multiple model architectures were implemented and evaluated.

As baseline models, two naive baselines establish a minimum performance threshold. The *Historical Average* predicts the mean TTI for a given hour and day of week based on historical data. The *Last Value* model persists the most recent observed TTI as the prediction. Any model that fails to outperform these baselines cannot be considered useful in practice. The MLP was considered with two variants. *MLP-context* uses only temporal (time of day, day of week), meteorological, and static topological features, without direct access to recent TTI observations. *MLP-traffic* additionally incorporates real-time TTI observations, allowing the model to exploit the current traffic state as a predictor.

RNNs such as GRU [18] and LSTM [17] architectures were implemented to capture temporal dynamics by modeling sequences of the 8 most recent TTI observations for each road segment. The T-GCN model [14], used here as a representative STGNN architecture, fuses graph convolutions for spatial modeling with GRU cells for temporal modeling, thus learning spatiotemporal dependencies on the road network graph.

A GNN model using the GATv2 attention operator [20] is evaluated to incorporate spatial dependencies between road segments. Unlike T-GCN, this model does not include a temporal component, and therefore captures only the spatial structure of the network at each time step independently. The STID model [19] presents an architecture based exclusively on MLP networks with learned embedding layers for nodes and time. Despite its simple MLP-based architecture, STID has been shown to match or exceed the performance of more complex spatiotemporal graph models. Graph WaveNet [15] captures long-term temporal dependencies through dilated causal convolutions. Thus, Graph WaveNet discovers the true traffic topology directly from observations, making it particularly suited to settings where the physical road graph does not fully reflect the actual spatial dependencies between segments.

IV. Experimental Results

A. Data collection and pre-processing

Bus location data was retrieved every 10 seconds from the official public transit data of Târgu Mureș[1] via a Google Cloud server. Each record was recorded in JSON format and contained the vehicle identifier, active route, geographic coordinates (latitude and longitude), and a timestamp (see Table I). The raw data was further cleaned to remove noise and correct positioning errors introduced by GPS equipment.

TABLE I
Sample of public transportation data collection at 10s intervals.

vehicle_id	route	lat	lon	timestamp
MS100	18	46.51142	24.51215	2026-01-30 12:04:59
MS100	18	46.51142	24.51215	2026-01-30 12:05:09
MS100	18	46.51142	24.51215	2026-01-30 12:05:19
MS100	18	46.51142	24.51215	2026-01-30 12:05:29
MS100	18	46.51147	24.51216	2026-01-30 12:05:39
MS100	18	46.51150	24.51217	2026-01-30 12:05:49
MS100	18	46.51150	24.51217	2026-01-30 12:05:59

Due to the volume of collected data, the raw JSON files were first flattened and converted into tabular CSV format. During this conversion, a series of validation filters were applied: (i) only records from mid-February to the end of April 2026 were retained, (ii) operational routes were validated using regular expressions, excluding technical routes such as depot trips, (iii) records with null, missing, or corrupted coordinates were discarded, and (iv) duplicate records were removed.

Given the influence of weather conditions on traffic flow, meteorological data was retrieved for each hourly interval of the monitored period from the Open-Meteo archive API, using the geographic coordinates of Târgu Mureș. The extracted features include air temperature, precipitation volume, and the World Meteorological Organization (WMO) weather code.

Raw GPS coordinates were associated with specific edges of the road network graph before any traffic indicator can be computed per segment. This was achieved through the following procedure. First, *spatial filtering* was applied where GPS points were intersected with buffer zones generated around the official General Transit Feed Specification (GTFS) routes and stops, removing outliers and distinguishing records corresponding to vehicles in transit from those at a stop (dwell time). Second, for computational efficiency, coordinates are *spatially aggregated* into 10-meter grids. For each grid cell, the nearest edge in the road network graph is identified, and the association is applied to all GPS points within that cell. Third, *direction correction* was performed to infer the direction of travel of either from the geometry of the official GTFS routes or by computing the bearing between consecutive GPS points. This ensured that each observation was correctly assigned to the directed edge that reflects the vehicle's actual direction of movement.

The processed dataset contains traffic observations aggregated at the road segment level in 15-minute time windows. Each observation includes the following features:

- *Traffic features*: mean TTI and mean speed for the segment in the time window;
- *Temporal features*: time of day and day of week;
- *Meteorological features*: temperature, precipitation, and WMO weather code;
- *Road segment properties*: segment length and number of lanes.

Additionally, expected congestion patterns were revealed during TTI calculation as can be seen in Fig. 1.

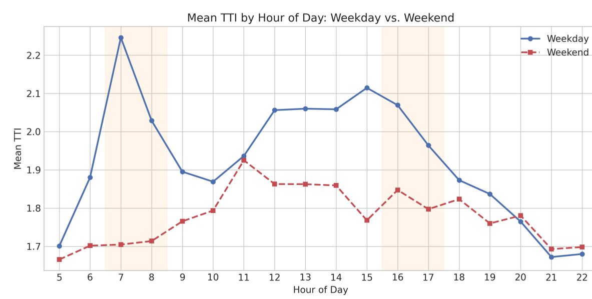

Fig. 1. Mean TTI: Workday vs. Weekend

B. Evaluation Metrics

Model performance was evaluated using four standard regression metrics: MAE, Root Mean Squared Error (RMSE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R^2). MAE measures the average absolute deviation between predicted and actual TTI values and is treated as the primary performance metric. RMSE penalizes larger errors more heavily, MAPE expresses prediction error as a percentage of the true value, and R^2 quantifies the proportion of variance in the target variable explained by the model.

C. Results

The main objective of the considered models in this system was the supervised prediction of the TTI for a specific road segment. The prediction horizon was set to 15 minutes ahead, using a 2-hour historical input window, representing a sequence of 8 consecutive time steps of 15 minutes each.

To ensure an objective evaluation and prevent data leakage, the dataset was divided chronologically: 67% for training, 16% for validation, and 17% for testing. The selection of the best model was based on the MAE metric, as it interprets

[1]https://new.transportlocal.ro/live/

errors linearly and maintains relevance in high-congestion scenarios. Additionally, RMSE, MAPE , and R^2 were also monitored. Since different architectures produce different output structures (e.g., MLP evaluates each record independently, while T-GCN evaluates an entire graph simultaneously), the testing methodology included a protocol which extracted the exact intersection of observations evaluated by all models. This ensures that comparative metrics were computed on an identical dataset.

The results in Table II confirm the theoretical assumptions of the study. The poor performance of the *MLP-context* model ($MAE = 0.728$), which underperformed even the naive *Last-Value* model, demonstrates that relying solely on calendar and topological data is insufficient. This confirms that recent traffic observations are an essential input for accurate short-term forecasting. With the introduction of sequential historical data (2 hours) through recurrent models such as GRU ($MAE = 0.580$), prediction error decreased significantly. On the other hand, the spatial model (GNN, $MAE = 0.632$) delivered better results, confirming that road topology alone, without temporal dynamics, tends to smooth values and cannot accurately anticipate congestion peaks.

TABLE II
MODEL PERFORMANCE EVALUATION.

Model	MAE	RMSE	MAPE (%)	R^2
Last Value	0.709	1.503	36.8	0.168
Historical Average	0.644	1.311	28.4	0.367
MLP-context	0.728	1.382	29.3	0.297
MLP-traffic	0.626	1.272	26.9	0.404
GRU	0.580	1.195	**25.1**	0.460
LSTM	0.576	1.191	25.2	0.464
GATv2	0.632	1.308	28.7	0.370
T-GCN(STGNN)	0.686	1.420	28.4	0.272
STID	0.571	**1.176**	**25.1**	**0.501**
Graph WaveNet	**0.546**	1.213	27.2	0.470

The best overall performance (Fig. 2) on the primary metric was achieved by Graph WaveNet ($MAE = 0.546$), whose ability to model long-term temporal dependencies and discover hidden spatial correlations proved to be particularly effective in a sparse road graph. However, STID achieved superior results on RMSE, MAPE, and R^2 (Table II), indicating that it explains a greater proportion of the variance in TTI and is more robust to larger prediction errors. The findings suggest that the choice between Graph WaveNet and STID may depend on the operational priority: Graph WaveNet is preferable when minimising average prediction error is critical, while STID offers better overall fit and is a computationally lighter alternative. For the deployed warning system, Graph WaveNet was selected based on its MAE performance, as this metric most directly reflects the average deviation in predicted congestion levels experienced by road users.

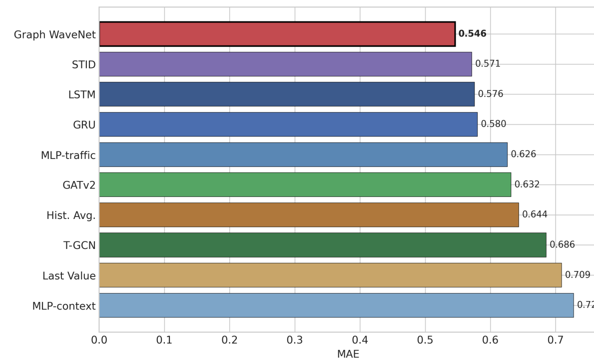

Fig. 2. Model ranking by MAE.

D. Implementation

The system was implemented as a publicly available application[2] that provides real-time, road-level congestion forecasts for Târgu Mureș, Romania, built exclusively on GPS telemetry broadcast by the municipal public transit fleet. The pipeline processes raw location records through map matching, segment aggregation, and TTI estimation to produce a 15-minute time series enriched with temporal and meteorological context, as described in Section III.

The prediction backend exposes a REST API, which serves TTI forecasts 15 minutes ahead for each monitored road segment. A Streamlit[3]-based dashboard presents these predictions and renders them on an interactive map of the city, allowing users to inspect the predicted congestion level for any monitored segment at a selected future time. An example of the interface is shown in Fig. 3.

V. CONCLUSION

This paper presents a short-term traffic congestion prediction system built on AVL data collected from the municipal public transit fleet of Târgu Mureș, Romania. The proposed pipeline demonstrates that accurate short-term congestion forecasting is achievable using only publicly available transit data, without dedicated sensor infrastructure.

A comparative evaluation of multiple prediction architectures demonstrated that incorporating temporal traffic history is essential for accurate congestion forecasting. Graph WaveNet achieved the lowest MAE, while STID outperformed Graph WaveNet on RMSE, MAPE, and R^2. This suggests that lightweight embedding-based models are a viable and efficient alternative when network coverage is sparse. These results confirm that publicly available AVL data can serve as a reliable and cost-effective source for urban traffic monitoring and short-term congestion prediction.

This study has several limitations. The primary limitation is that the road graph was constructed exclusively from segments covered by bus routes, leaving residential and peripheral streets unobserved. This restricts the coverage of the system.

[2]https://github.com/lucikis/Traffic-congestion-prediction
[3]https://streamlit.io/

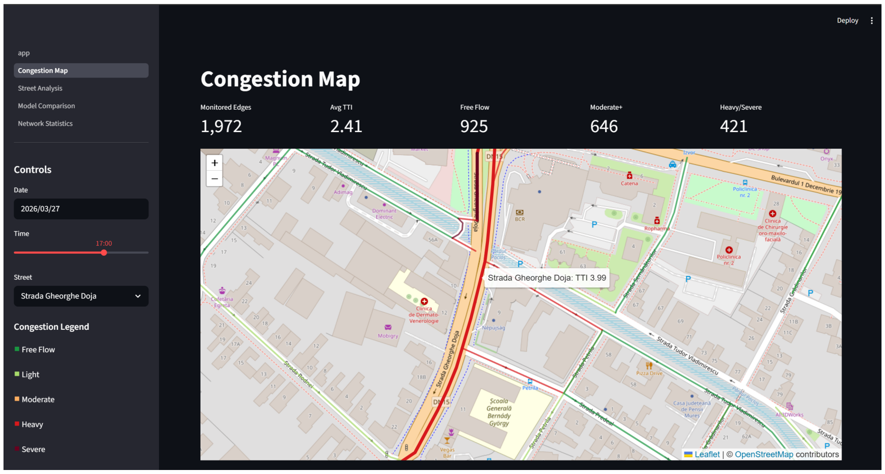

Fig. 3. Congestion map interface displaying predicted TTI values for the Gheorghe Doja Street at 17:00. The dashboard summarises city-wide statistics at the selected time: Monitored edges, Avg. TTI, Free flow, Moderate, Heavy/Severe.

Additionally, the dataset spans approximately two and a half months, preventing the model from capturing seasonal traffic patterns. Moreover, the system was validated on a single city, and its generalisability to other medium-sized urban areas has not been assessed. A further limitation concerns the absence of external validation, as no independent ground-truth references were used.

As future work, we consider extending data collection to allow the model to capture annual traffic cycles, including holiday periods and school breaks. Additionally, adding data sources such as cellular probe data or taxi GPS traces would provide more dense data beyond the bus route coverage. A third direction is the evaluation of transfer learning approaches to assess whether models trained on Târgu Mureș can be adapted to other cities with limited sensor infrastructure.

Acknowledgment

The authors used Claude (Anthropic) to assist with language editing and grammatical refinement of the manuscript. The scientific content, methodology, results, and conclusions are entirely the work of the authors.

References

[1] M. Attioui and M. Lahby, "A systematic literature review of traffic congestion forecasting: From machine learning techniques to large language models," *Vehicles*, vol. 7, no. 4, p. 142, 2025.

[2] K. Micko, P. Papcun, and I. Zolotova, "Review of iot sensor systems used for monitoring the road infrastructure," *Sensors*, vol. 23, no. 9, p. 4469, 2023.

[3] G. Del Serrone, G. Cantisani, and P. Peluso, "Speed data collection methods: a review," *Transportation research procedia*, vol. 69, pp. 512–519, 2023.

[4] K. Čulík, V. Štefancová, and K. Hrudkay, "Application of wireless magnetic sensors in the urban environment and their accuracy verification," *Sensors*, vol. 23, no. 12, p. 5740, 2023.

[5] M. Tahir, Y. Qiao, N. Kanwal, B. Lee, and M. N. Asghar, "Privacy preserved video summarization of road traffic events for iot smart cities," *Cryptography*, vol. 7, no. 1, p. 7, 2023.

[6] B. Ipekyuz, P. K. Sevinen, and H. T. Yaman, "Performance evaluation of fused floating car data (fcd) and bluetooth (bt) data speed estimation on urban arterials," *Arabian Journal for Science and Engineering*, vol. 50, no. 20, pp. 17 091–17 108, 2025.

[7] R. H. Mulangi and V. Kulkarni, "Visualization and assessment of the effect of roadworks on traffic congestion using avl data of public transit," *Journal of Geovisualization and Spatial Analysis*, vol. 6, no. 2, p. 28, 2022.

[8] X. Zhang, L. Lauber, H. Liu, J. Shi, M. Xie, and Y. Pan, "Travel time prediction of urban public transportation based on detection of single routes," *Plos one*, vol. 17, no. 1, p. e0262535, 2022.

[9] Y. Li, R. Yu, C. Shahabi, and Y. Liu, "Diffusion convolutional recurrent neural network: Data-driven traffic forecasting," *arXiv preprint arXiv:1707.01926*, 2017.

[10] X. Liu, Y. Xia, Y. Liang, J. Hu, Y. Wang, L. Bai *et al.*, "Largest: A benchmark dataset for large-scale traffic forecasting," *Advances in Neural Information Processing Systems*, vol. 36, pp. 75 354–75 371, 2023.

[11] J. Stipancic, L. Miranda-Moreno, A. Labbe, and N. Saunier, "Measuring and visualizing space–time congestion patterns in an urban road network using large-scale smartphone-collected gps data," *Transportation letters*, vol. 11, no. 7, pp. 391–401, 2019.

[12] Y. He, B. Hofer, Y. Sheng, Y. Yin, and H. Lin, "Processes and events in the center: a taxi trajectory-based approach to detecting traffic congestion and analyzing its causes," *International Journal of Digital Earth*, vol. 16, no. 1, pp. 509–531, 2023.

[13] A.-S. Roman, R. Bolboacă, T. Lenard, and P. Haller, "Apu-trajgen+: Gru-based adaptive privacy and utility preserving trajectory generation," *IEEE Access*, 2025.

[14] L. Zhao, Y. Song, C. Zhang, Y. Liu, P. Wang, T. Lin *et al.*, "T-gcn: A temporal graph convolutional network for traffic prediction," *IEEE transactions on intelligent transportation systems*, vol. 21, no. 9, pp. 3848–3858, 2019.

[15] Z. Wu, S. Pan, G. Long, J. Jiang, and C. Zhang, "Graph wavenet for deep spatial-temporal graph modeling," *arXiv preprint arXiv:1906.00121*, 2019.

[16] G. Boeing, "Osmnx: New methods for acquiring, constructing, analyzing, and visualizing complex street networks," *Computers, environment and urban systems*, vol. 65, pp. 126–139, 2017.

[17] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.

[18] K. Cho, B. Van Merriënboer, Ç. Gulçehre, D. Bahdanau, F. Bougares, H. Schwenk *et al.*, "Learning phrase representations using rnn encoder–decoder for statistical machine translation," in *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, 2014, pp. 1724–1734.

[19] Z. Shao, Z. Zhang, F. Wang, W. Wei, and Y. Xu, "Spatial-temporal identity: A simple yet effective baseline for multivariate time series forecasting," *Proceedings of the 31st ACM International Conference on Information & Knowledge Management (CIKM '22)*, pp. 4454–4458, 2022. [Online]. Available: https://doi.org/10.1145/3511808.3557702

[20] S. Brody, U. Alon, and E. Yahav, "How attentive are graph attention networks?" *arXiv preprint arXiv:2105.14491*, 2021.