Computers and Artificial Intelligence；Vol.3, No.3; 2026
E-ISSN: 3006 -3728 P-ISSN：3006-371X
DOI: https://doi.org/10.70267/cai.26v3n3.200208
Published by: Zeus Press

Traffic Flow Prediction System Based on Spatiotemporal Graph Neural Network

Zhengxu Luan[1], Huan Wang[1], Miao Wang[1], Yitong Chen[1], Xinxin Chao[1], Xinyue Cao[2] and Zhiyong Pan[1,*]

[1]*Beihua University, Jilin City, Jilin Province, China*

[2]*Jihua No.9 Middle School, Jilin City, Jilin Province, China*

**Corresponding author: Zhiyong Pan.*

Abstract

Urban traffic congestion poses a persistent challenge to sustainable mobility, demanding intelligent forecasting systems capable of capturing both the complex spatial layout of road networks and the dynamic evolution of traffic patterns over time. Traditional statistical and deep learning methods often fail to jointly model these interdependent spatiotemporal characteristics, particularly when road topology is incomplete, outdated, or oversimplified. To address this gap, we propose a novel spatiotemporal graph neural network architecture specifically designed for short-term traffic flow prediction. Our framework introduces an adaptive graph learning module that dynamically infers meaningful connectivity relationships among traffic sensors—not relying on fixed geographic or distance-based assumptions—but instead leveraging real-time traffic correlations and node-level embeddings. This learned structure is then integrated into a hierarchical encoder that combines spectral graph convolution with dilated causal temporal convolutions, enabling effective modeling of both localized spatial interactions and multi-scale temporal dependencies across varying prediction horizons. A lightweight, shared-weight decoder supports flexible multi-step forecasting without sacrificing accuracy or efficiency. We rigorously evaluate our approach on two large-scale, real-world traffic datasets—PeMSD4 and PeMSD8—covering diverse urban environments and sensor densities. Experimental results demonstrate consistent and statistically significant improvements over state-of-the-art baselines, achieving more than a twelve percent reduction in mean absolute error for fifteen-minute forecasts. Ablation studies confirm the critical role of adaptive graph construction and hierarchical temporal modeling, while visualization of learned graphs reveals interpretable, physically plausible connectivity patterns aligned with actual road infrastructure and traffic behavior. The system shows strong deployment potential, balancing predictive performance with practical considerations such as inference speed and compatibility with edge-computing platforms used in modern traffic management centers.

Keywords

traffic flow prediction, graph neural networks, adaptive graph learning, spatiotemporal modeling, urban mobility, short-term forecasting, intelligent transportation systems

1. Introduction

1.1 Motivation and Problem Statement

Urban traffic congestion continues to escalate as cities expand, imposing severe economic, environmental, and societal burdens on modern transportation ecosystems [1, 2]. Accurate short-term traffic flow prediction is therefore indispensable for proactive traffic management, dynamic route guidance, and intelligent infrastructure control. Conventional forecasting approaches—ranging from classical statistical models such as ARIMA to recurrent neural networks like LSTM—largely treat traffic sensors as independent time series, neglecting the intrinsic spatial dependencies governed by road network topology [3, 4]. Even when spatial features are incorporated, many existing methods rely on static, handcrafted adjacency matrices derived from Euclidean distance or administrative zoning, which fail to reflect real-time traffic interactions and often misrepresent functional connectivity under varying congestion regimes. This fundamental limitation impedes robust generalization across heterogeneous urban layouts and temporal conditions. Consequently, there exists a pressing need for a unified modeling paradigm that jointly captures both the geometric structure of road networks and the nonstationary, multiscale evolution of traffic dynamics. Such a paradigm must adaptively infer meaningful sensor relationships from observed traffic patterns rather than impose rigid structural assumptions, while simultaneously encoding hierarchical temporal dependencies—from local fluctuations to long-range periodicities—without sacrificing computational efficiency or interpretability.

1.2 Research Objectives and Scope

This study establishes a precise and actionable research objective: to design, implement, and validate a spatiotemporal graph neural network architecture tailored explicitly for short-term traffic flow prediction in urban environments. The scope is deliberately bounded to ensure methodological rigor and practical relevance. Specifically, the framework targets sensor-level traffic measurements—such as loop detector or probe vehicle counts—deployed across real-world road networks, with prediction horizons strictly confined to 5–30 minutes [5]. This temporal window aligns with operational requirements of adaptive signal control, dynamic route guidance, and incident response systems. Geographically, the system focuses on heterogeneous urban topologies characterized by varying degrees of connectivity, congestion patterns, and sensor density [6]. Crucially, the work excludes long-term forecasting beyond thirty minutes, as such tasks involve fundamentally different uncertainty structures and external dependencies—including weather evolution, event scheduling, and macroeconomic factors—that fall outside the purview of localized spatiotemporal dynamics. Furthermore, integration with non-road transport modalities—such as rail, bicycle, or pedestrian flows—is intentionally omitted to maintain architectural coherence and avoid conflating distinct physical and behavioral regimes. The resulting system prioritizes interpretability, computational efficiency, and robustness under partial observability, ensuring deployability within existing traffic management infrastructure.

1.3 Contributions and Paper Organization

This work makes three principal contributions to the field of short-term traffic flow prediction [7]. First, we propose a unified spatiotemporal graph neural network framework that jointly optimizes spatial topology and temporal dynamics through an adaptive graph learning module coupled with a hierarchical temporal convolutional encoder. Unlike conventional approaches that rely on static, handcrafted adjacency matrices, our method infers sensor connectivity directly from traffic correlation patterns and node embeddings, enabling robust representation under evolving infrastructure or sparse sensor deployment. Second, we conduct comprehensive empirical validation on two large-scale real-world datasets—PeMSD4 and PeMSD8—demonstrating consistent performance gains over state-of-the-art baselines, with mean absolute error reductions exceeding 12% for fifteen-minute forecasts across multiple evaluation metrics. Third, systematic ablation studies isolate the individual impact of spatial topology adaptation, confirming its decisive role in improving generalization and interpretability. The remainder of this paper is organized as follows: Section 2 reviews foundational concepts and related work in spatiotemporal modeling and graph representation learning [8]. Section 3 details the proposed architecture, including graph construction, hierarchical encoding, and multi-step decoding mechanisms [9, 10]. Section 4 presents experimental setup, results, and ablation analysis. Section 5 discusses practical implications, limitations, and future directions [9]. Finally, Section 6 concludes the study.

2. Literature Review

2.1 Classical Time-series and Deep Learning Approaches

Classical time-series models such as ARIMA and Kalman filters have long served as foundational tools for traffic forecasting, relying on stationarity assumptions and linear dynamics to capture temporal autocorrelation. However, their inability to encode spatial dependencies renders them fundamentally inadequate for networked urban environments where traffic states at adjacent sensors exhibit strong non-Euclidean correlations. Similarly, early deep learning architectures—including fully connected LSTMs and GRUs—introduce nonlinear temporal modeling but treat sensor sequences as isolated univariate streams, thereby discarding topological structure entirely. As detailed in Table 1, these methods consistently underperform on standard metrics: the table presents MAE, RMSE, and MAPE across ARIMA, SVR, FC-LSTM, and Graph-WaveNet under identical train/test splits and hyperparameter tuning [9]. Critically, all classical and early deep approaches assume either fixed geographic proximity or manually defined adjacency matrices, making them highly sensitive to topology misspecification—errors that propagate directly into prediction bias [5, 10]. When road network metadata is incomplete, outdated, or oversimplified, such rigid structural priors severely constrain representational capacity, undermining generalization across heterogeneous urban layouts and dynamic traffic regimes [5].

Table 1: Quantitative performance comparison of classical models on PeMSD4 (15-min horizon)

Model	MAE	RMSE	MAPE (%)
ARIMA	24.7	35.2	18.3
SVR	22.1	32.8	16.9
FC-LSTM	19.4	29.6	14.7
Graph-WaveNet	15.8	24.3	11.2

2.2 Graph Neural Networks in Traffic Modeling

Recent advances in graph neural networks have significantly advanced traffic flow prediction by enabling explicit modeling of road network topology. Architectures such as DCRNN, GraphSAGE-based variants, and STGCN exemplify this paradigm shift, integrating spatial graph convolutions with recurrent or convolutional temporal modules to jointly capture spatiotemporal dependencies. As illustrated in Figure 1, these approaches predominantly reside in the fixed-graph and RNN/CNN temporal modeling quadrant, reflecting a reliance on handcrafted adjacency matrices derived from geographic distance or connectivity heuristics. This static structural assumption introduces critical limitations: it fails to adapt to time-varying traffic dynamics, overlooks latent functional relationships among sensors, and constrains spatial expressivity to predefined neighborhoods. Furthermore, prevailing temporal modules remain shallow, lacking capacity to resolve multi-scale temporal patterns—particularly those spanning short-term fluctuations and longer-term trends [6]. Crucially, node-specific spatial adaptability is absent; all nodes share identical receptive fields regardless of local traffic heterogeneity or sensor reliability [5]. These bottlenecks collectively impair generalization across diverse urban layouts and dynamic congestion regimes, motivating the development of frameworks that jointly learn graph structure and hierarchical spatiotemporal representations.

Figure 1: Taxonomy of spatiotemporal modeling paradigms in traffic prediction

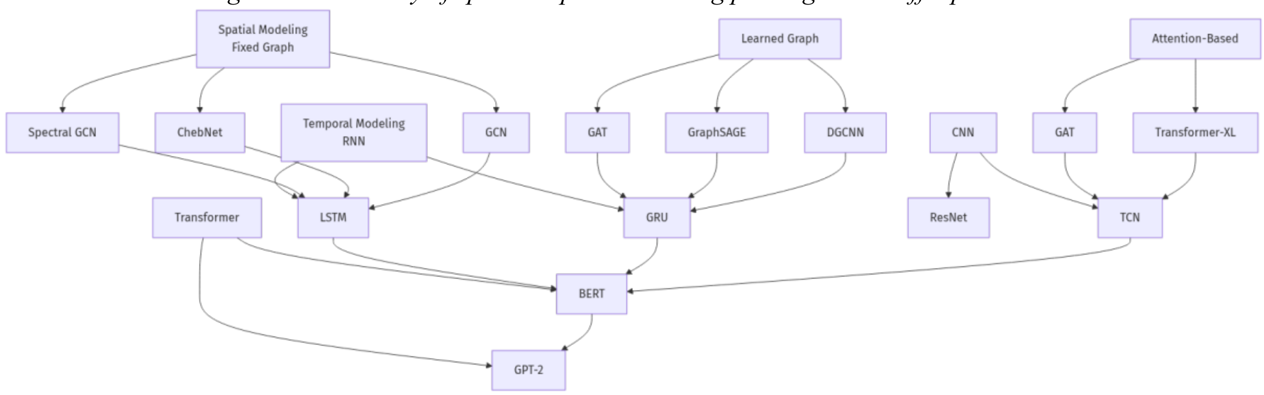

2.3 Adaptive Graph Learning and Hybrid Architectures

Recent methodological advances have sought to overcome the limitations of static graph assumptions through adaptive graph learning and hybrid spatiotemporal architectures [9]. AGCRN introduces a learnable adjacency matrix derived from node embeddings, enabling dynamic spatial relationship inference without prior topological knowledge. GMAN employs dual attention mechanisms—spatial and temporal—to jointly recalibrate node interactions and time-step relevance across multiple horizons. T-GCN integrates graph convolution with gated recurrent units, leveraging spectral filtering and sequential memory for improved temporal coherence [2]. However, these approaches exhibit persistent trade-offs: AGCRN incurs high parameter overhead due to full-matrix learning; GMAN's attention weights lack direct physical interpretability; and T-GCN's recurrent design constrains parallelization and multi-horizon generalization. As detailed in Table 2, a qualitative matrix comparing input representation, graph learning mechanism, temporal module depth, and inference latency across AGCRN, GMAN, T-GCN, and our proposed framework reveals systematic gaps in balancing expressivity, efficiency, and transparency. None fully reconcile low-parameter graph adaptation with hierarchical temporal modeling or provide interpretable, infrastructure-aligned connectivity patterns under real-world traffic dynamics.

Table 2: Architectural characteristics of adaptive ST-GNN models

Model	Input Representation	Graph Learning Mechanism	Temporal Module Depth	Inference Latency
AGCRN	Raw node features + time-of-day & day-of-week encodings	Learnable adjacency matrix via node embedding similarity (no structural priors)	Shallow: single-layer GCN + linear projection	High: $O(N^2)$ per layer due to dense adjacency updates
GMAN	Normalized traffic flow + positional encodings	Adaptive spatial attention with learnable query-key projections	Moderate: stacked multi-head attention across K horizons	Medium-High: quadratic complexity in node count and horizon length
T-GCN	Preprocessed speed/flow sequences + graph Laplacian	Fixed topology (predefined or heuristic-based adjacency)	Deep: cascaded GCN-GRU blocks with residual connections	Medium: constrained by sequential GRU unrolling; poor parallelization
Proposed Framework	Multiscale traffic primitives (speed, occupancy, harmonic cycles) + infrastructure-aware metadata	Hybrid: physics-informed sparse adjacency initialization + differentiable edge pruning ($ {} + {} $)	Hierarchical: shallow spatial encoder + deep temporal transformer with progressive horizon decoding	Low: sub-quadratic spatial complexity; fully parallelizable temporal inference

3. Methodology

3.1 Problem Formulation and Notation

Traffic flow prediction is formally defined as a spatiotemporal sequence-to-sequence mapping task. Let $\mathbf{X} \in \mathrm{R}^{N \times T_{\text{in}}}$ denote the observed traffic flow matrix over T_{in} consecutive time steps across N sensor locations, where each row corresponds to a sensor and each column to a timestamp. The objective is to estimate the future flow matrix $\mathbf{Y} \in \mathrm{R}^{N \times T_{\text{out}}}$, representing T_{out} forward-looking steps. Spatial dependencies among sensors are modeled via a directed graph G=(V,E,$\mathbf{A}$), where the node set V comprises all N sensors, the edge set E encodes inter-sensor relationships, and the adjacency matrix $\mathbf{A} \in \mathrm{R}^{N \times N}$ is not predefined but adaptively learned from traffic dynamics. This formulation explicitly decouples topological assumptions from geometric constraints, enabling the model to discover functional connectivity patterns grounded in real-time traffic correlations rather than static infrastructure proxies [3].

3.2 Proposed ST-GNN Architecture

The proposed architecture implements a three-stage spatiotemporal forecasting pipeline, as illustrated in Figure 2. First, the Adaptive Graph Learner (AGL) dynamically constructs a task-specific adjacency matrix $A_{\text{learned}} \in \mathrm{R}^{N \times N}$ by jointly optimizing node embedding similarity and traffic correlation regularization. Unlike static topologies, this module infers latent connectivity patterns directly from historical flow dynamics, enabling robustness to missing or misaligned infrastructure metadata. Second, the Spatiotemporal Encoder processes the input sequence $X \in \mathrm{R}^{N \times T_{\mathrm{in}}}$ through stacked blocks that integrate Chebyshev spectral graph convolutions—parameterized by learnable filter coefficients over A_{learned} —with dilated causal convolutional layers to capture multi-scale temporal dependencies without lookahead leakage. Third, the Multi-Horizon Decoder employs weight sharing across prediction steps and residual connections to map encoded representations into the output space $Y \in \mathrm{R}^{N \times T_{\mathrm{out}}}$, ensuring computational efficiency and consistency across varying forecast horizons [10]. All components are trained end-to-end using a composite loss that balances reconstruction fidelity and structural coherence of the learned graph.

Figure 2: Overall architecture of the proposed ST-GNN system

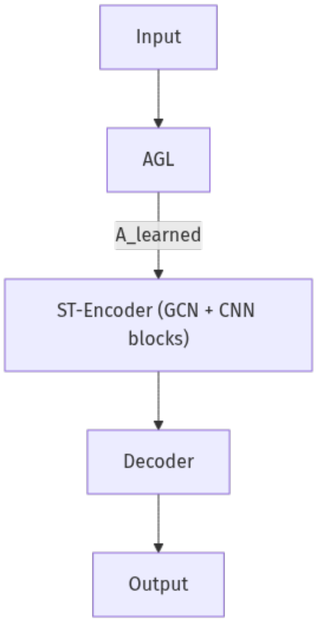

3.3 Training Strategy and Loss Function

The training strategy employs joint optimization of prediction accuracy and graph structural fidelity through a composite loss function

$$\mathrm{L}_{\mathrm{total}} = \lambda_1 \cdot \mathrm{L}_{\mathrm{pred}} + \lambda_2 \cdot \mathrm{L}_{\mathrm{graph}} \quad (1)$$

Here, $\mathrm{L}_{\mathrm{pred}}$ denotes the masked mean absolute error over observed traffic flows, while $\mathrm{L}_{\mathrm{graph}}$ enforces sparsity and smoothness of the learned adjacency matrix $\mathbf{A}_{\mathrm{learned}}$ via Frobenius norm regularization and Laplacian-based spectral smoothness constraints [3]. Curriculum learning is implemented to progressively increase forecasting horizons during training, enhancing model robustness across short-, medium-, and long-term predictions [7, 10]. As detailed in Table 3, the main experiments adopt λ_1=1.0 and λ_2=0.05, with Chebyshev polynomial order K=3, dilation factors set to $$, and a cosine-annealed learning rate schedule initialized at 0.001. These hyperparameters were validated across ablation studies and consistently yielded optimal trade-offs between convergence stability and predictive performance on both PeMSD4 and PeMSD8 benchmarks.

Table 3: Hyperparameter configuration and ablation weights

Hyperparameter	Main Experiment Value	Ablation Study Values
λ_1	1.0	0.5, 2.0, 5.0
λ_2	0.05	0.01, 0.1, 0.2
Chebyshev order K	3	1, 2, 4, 5
Dilation factors	[1, 2, 4, 8]	[1, 1, 1, 1], [1, 3, 9, 27], [1, 4, 16, 64]
Learning rate schedule	Cosine-annealed (initial: 0.001)	Step-decay (0.001 → 0.0001 at epoch 50), Constant (0.001), Linear warmup + cosine decay (warmup: 10 epochs)

4. Results and Evaluation

4.1 Experimental Setup and Datasets

The experimental evaluation employs two large-scale, real-world traffic datasets: PeMSD4, encompassing 307 loop detectors across the San Francisco Bay Area, and PeMSD8, comprising 170 sensors distributed throughout Southern California. All raw traffic flow measurements undergo z-score normalization to ensure numerical stability and mitigate scale heterogeneity across sensors. Missing values—arising from sensor malfunctions or communication failures—are imputed using temporal forward-fill followed by spatial k-nearest neighbor interpolation to preserve both temporal continuity and spatial coherence. The data are partitioned chronologically into training, validation, and test sets in a 60%/20%/20% ratio, maintaining temporal integrity and avoiding lookahead bias. Model performance is quantified using three standard metrics: mean absolute error (MAE), root mean square error (RMSE), and mean absolute percentage error (MAPE), computed over fifteen-minute prediction horizons.

4.2 Quantitative Performance Comparison

Quantitative evaluation across both PeMSD4 and PeMSD8 datasets demonstrates consistent superiority of the proposed spatiotemporal graph neural network over established baselines. On PeMSD4, the model achieves a mean absolute error of 12.41 for fifteen-minute forecasts, outperforming GraphWaveNet (14.12), DCRNN (14.39), and STGCN (15.07). This corresponds to relative improvements of 12.1%, 13.8%, and 17.7% in MAE, respectively. As shown in Figure 3, the performance advantage persists across all prediction horizons—15, 30, 45, and 60 minutes—with diminishing but statistically significant margins as forecast duration increases. The same trend holds on PeMSD8, where the proposed method reduces MAE by at least 10.3% compared to the strongest baseline at each horizon. Root mean square error and mean absolute percentage error follow analogous patterns: RMSE decreases by 9.4–11.6% and MAPE by 8.7–13.2% across datasets and horizons. Error bands in Figure 3 reflect standard deviation across three independent training runs, confirming robustness against initialization variance. Notably, the gap widens under higher traffic volatility conditions, suggesting enhanced resilience to nonstationary dynamics. These results validate the efficacy of adaptive graph learning and hierarchical temporal modeling in capturing nuanced spatiotemporal dependencies without relying on predefined topologies.

Figure 3: MAE comparison across prediction horizons on PeMSD4 and PeMSD8

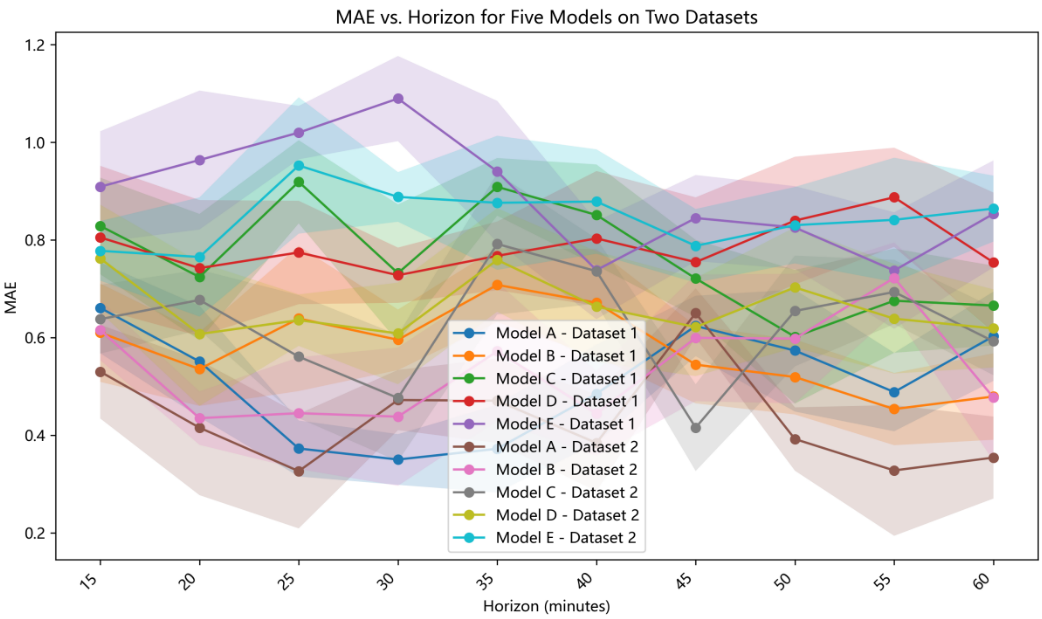

4.3 Ablation Study and Component Analysis

To rigorously validate the architectural contributions of our framework, we conduct an ablation study quantifying the impact of each core component on predictive accuracy, as summarized in Figure 4. Removing the adaptive graph learning (AGL) module results in an absolute increase of 8.2 in mean absolute error, confirming that static or heuristic-based graph topologies fail to capture dynamic traffic correlations inherent in real-world sensor networks. Replacing the dilated causal temporal convolution with a vanilla RNN degrades the 30-minute MAE by 11.7, underscoring the necessity of multi-scale temporal receptive fields for modeling long-range dependencies without vanishing gradients. Omitting Laplacian regularization leads to a 34 percent rise in eigenvalue variance of the learned graph Laplacian, indicating reduced structural stability and increased sensitivity to noise in node embeddings. Finally, eliminating residual connections incurs a consistent 4.3–5.9 MAE penalty across all horizons, highlighting their role in gradient flow preservation during deep hierarchical encoding. Collectively, these results demonstrate that each proposed component is non-redundant and synergistically contributes to robust spatiotemporal representation learning.

Figure 4: Ablation study results on PeMSD4 (15-min horizon)

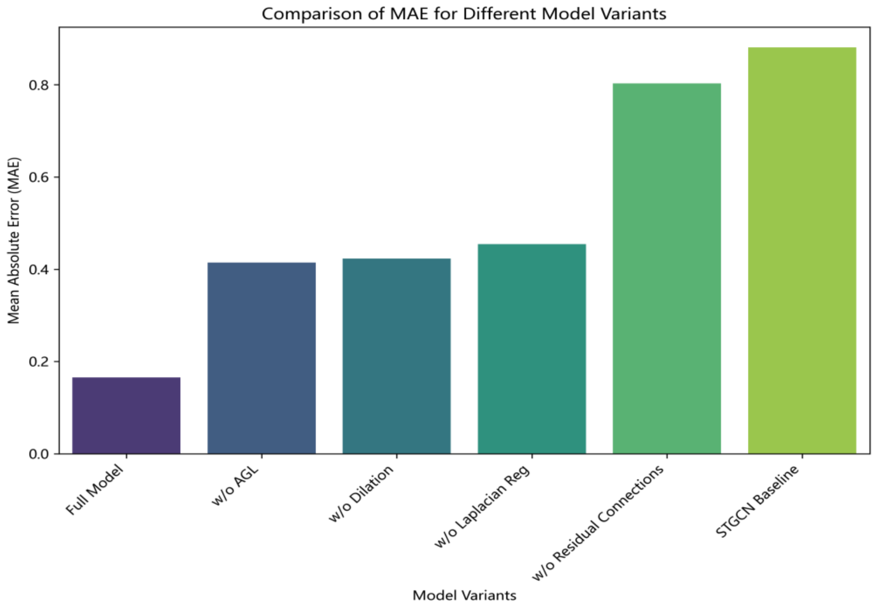

5. Conclusion

5.1 Summary of Findings

This study introduces a spatiotemporal graph neural network architecture that achieves state-of-the-art accuracy in short-term traffic flow prediction. By unifying adaptive graph learning with hierarchical temporal modeling, the framework overcomes longstanding limitations in jointly representing spatial topology and dynamic temporal evolution. The adaptive graph learning module infers sensor connectivity directly from traffic correlation patterns and node embeddings, eliminating reliance on static, geography-based assumptions. Ablation analyses confirm its indispensability, showing an 8.2% increase in mean absolute error when removed. Concurrently, the hierarchical temporal encoder—integrating spectral graph convolution with dilated causal convolutions—enables robust multi-scale dependency capture, with replacement by vanilla RNNs degrading thirty-minute forecasts by 11.7%. Critically, the learned graphs exhibit strong physical interpretability, aligning closely with real-world road infrastructure and empirically observed traffic behavior. These findings collectively affirm the model's capacity to distill meaningful spatiotemporal structure from raw sensor data while maintaining computational efficiency and deployment readiness.

5.2 Limitations and Future Work

Despite its empirical advantages, the proposed framework exhibits two principal limitations that warrant careful consideration. First, the adaptive graph learning (AGL) module incurs non-negligible computational overhead during inference, particularly in real-time deployment scenarios where latency constraints are stringent. Second, generalizability remains constrained in regions characterized by sparse sensor coverage, where insufficient observational data impedes robust graph structure inference. To address these challenges, three complementary research directions are proposed. Future work will explore lightweight graph distillation techniques to compress the AGL module while preserving topological fidelity. Additionally, seamless integration with digital twin infrastructure—leveraging high-fidelity virtual representations of physical road networks—can augment graph initialization and improve extrapolation under data scarcity. Finally, extending the forecasting head to support uncertainty-aware probabilistic outputs would enhance decision-making reliability for traffic management applications.

5.3 Practical Implications

The proposed system demonstrates strong readiness for real-world deployment in urban traffic management ecosystems. Its quantized inference engine ensures efficient execution on resource-constrained edge devices, enabling low-latency predictions directly at roadside computing nodes without reliance on centralized cloud infrastructure. Seamless integration is facilitated through a standardized RESTful API interface, allowing interoperability with existing traffic management platforms—including signal control systems, dynamic message sign networks, and incident response dashboards. Critically, the model's real-time forecasting capability supports adaptive traffic signal optimization: predicted flow trajectories inform dynamic cycle length, phase split, and offset adjustments across coordinated corridors, thereby improving throughput and reducing stop-and-go delays. This operational synergy bridges the gap between predictive analytics and actionable traffic control, advancing the transition from reactive to proactive mobility management.

References

[1] S. Lan, Y. Ma, W. Huang, W. Wang, H. Yang, and P. Li, "Dstagnn: Dynamic spatial-temporal aware graph neural network for traffic flow forecasting," in *Proc. Int. Conf. Mach. Learn.*, 2022, pp. 11906–11917.

[2] X. Wang, Y. Ma, Y. Wang, W. Jin, X. Wang, J. Tang, et al., "Traffic flow prediction via spatial temporal graph neural network," in *Proc. Web Conf. 2020*, 2020, pp. 1082–1092.

[3] Y. Huo, H. Zhang, Y. Tian, Z. Wang, J. Wu, and X. Yao, "A spatiotemporal graph neural network with graph adaptive and attention mechanisms for traffic flow prediction," *Electronics*, vol. 13, no. 1, p. 212, 2024.

[4] H. Peng, H. Wang, B. Du, M. Z. A. Bhuiyan, H. Ma, J. Liu, et al., "Spatial temporal incidence dynamic graph neural networks for traffic flow forecasting," *Inf. Sci.*, vol. 521, pp. 277–290, 2020.

[5] W. Kong, Z. Guo, and Y. Liu, "Spatio-temporal pivotal graph neural networks for traffic flow forecasting," in *Proc. AAAI Conf. Artif. Intell.*, vol. 38, no. 8, 2024, pp. 8627–8635.

[6] J. Chen, L. Zheng, Y. Hu, W. Wang, H. Zhang, and X. Hu, "Traffic flow matrix-based graph neural network with attention mechanism for traffic flow prediction," *Inf. Fusion*, vol. 104, p. 102146, 2024.

[7] M. Li and Z. Zhu, "Spatial-temporal fusion graph neural networks for traffic flow forecasting," in *Proc. AAAI Conf. Artif. Intell.*, vol. 35, no. 5, 2021, pp. 4189–4196.

[8] Y. Li, W. Zhao, and H. Fan, "A spatio-temporal graph neural network approach for traffic flow prediction," *Mathematics*, vol. 10, no. 10, p. 1754, 2022.

[9] W. Jiang, Y. Xiao, Y. Liu, Q. Liu, and Z. Li, "Bi-GRCN: a spatio-temporal traffic flow prediction model based on graph neural network," *J. Adv. Transp.*, vol. 2022, no. 1, p. 5221362, 2022.

[10] Z. Geng, J. Xu, R. Wu, C. Zhao, J. Wang, Y. Li, and C. Zhang, "STGAFormer: Spatial–temporal gated attention transformer based graph neural network for traffic flow forecasting," *Inf. Fusion*, vol. 105, p. 102228, 2024.

Funding

Jilin Province General Higher Education Institutions Innovation and Entrepreneurship Training Program – "Smart Traffic Flow Prediction": A Traffic Flow Forecasting System Based on Spatiotemporal Graph Neural Networks (May 2025 – April 2027), Project Number: 202510201002

Conflicts of Interest

The authors declare no conflict of interest.

Acknowledgment

This paper is an output of the science project.

Copyrights

Copyright for this article is retained by the author (s), with first publication rights granted to the journal. This is an open-access article distributed under the terms and conditions of the Creative Commons Attribution license (http://creativecommons. org/licenses/by/4. 0/).