International Journal of Scientific Research and Engineering Development-– Volume 9 Issue 3, May-June 2026
Available at www.ijsred.com

RESEARCH ARTICLE OPEN ACCESS

HAT-GNN: A Hybrid Adaptive Temporal Graph Neural Network for Traffic Congestion Prediction

Monika *, Jyoti **
*(M.Tech Scholar, Department of Computer Science & Engineering BRCM CET BAHAL
** (Guide, Department of Computer Science & Engineering BRCM CET BAHAL

--************************----------------------------------

Abstract

Accurate traffic congestion prediction remains a challenging task due to the dynamic and interconnected nature of urban transportation systems. Existing approaches, particularly graph neural network-based models, have improved prediction performance by capturing spatial and temporal dependencies; however, many of these models rely on static graph structures, exhibit high computational complexity, and struggle to adapt to evolving traffic conditions. To address these limitations, this paper proposes a Hybrid Adaptive Temporal Graph Neural Network (HAT-GNN), designed to integrate dynamic graph learning with spatial, temporal, and attention-based modeling within a unified framework.

The proposed model incorporates an adaptive graph learning module to capture time-varying relationships between traffic nodes, followed by graph convolution for spatial feature extraction and a gated recurrent unit for temporal sequence modeling. An attention mechanism is further employed to emphasize influential nodes and time steps, enhancing the model's predictive capability. The effectiveness of the proposed approach is evaluated on benchmark traffic datasets, where it demonstrates improved performance in terms of Mean Absolute Error (MAE) and Root Mean Square Error (RMSE) compared to established baseline models.

The results indicate that the integration of adaptive graph learning and attention mechanisms contributes to more accurate and robust traffic prediction, particularly under dynamic conditions. The proposed framework offers a balanced approach between predictive performance and model flexibility, making it suitable for intelligent transportation applications.

Keywords: Traffic congestion prediction, Graph neural networks, Adaptive graph learning, Spatio-temporal modeling, Attention mechanism, Deep learning, Intelligent transportation systems, Time-series forecasting

--************************----------------------------------

1. Introduction

Urban mobility systems are increasingly characterized by volatile demand, heterogeneous road networks, and frequent disruptions, making reliable traffic congestion prediction a persistent challenge. Accurate short- and medium-term forecasts are essential for proactive control strategies such as adaptive signal timing, dynamic routing, and incident management. While recent data-driven methods have improved predictive performance, their effectiveness often diminishes under rapidly changing conditions where relationships between road segments evolve over time.

Graph neural networks (GNNs) have emerged as a compelling framework for traffic forecasting because they encode transportation networks as graphs and learn spatial interactions alongside temporal patterns. Despite this progress, several practical limitations remain. First, many models depend on **static or weakly adaptive adjacency structures**, which inadequately reflect time-varying dependencies induced by congestion propagation, demand surges, or external events. Second, **spatio-temporal integration is frequently partial**, with spatial and temporal modules coupled in a shallow manner that constrains the model's ability to capture higher-order interactions. Third, **computational overhead** associated with deep architectures and attention components can hinder real-time deployment. Finally, **data imperfections**—including noise and missing sensor readings—continue to affect stability and generalization.

This work addresses these gaps by proposing a **Hybrid Adaptive Temporal Graph Neural Network (HAT-GNN)** that unifies dynamic graph learning with spatial, temporal, and attention-based modeling in a single architecture. The model learns a **data-driven, time-sensitive adjacency representation**, enabling it to adjust spatial relationships as traffic conditions evolve. Spatial dependencies are captured through graph convolution, while **temporal dynamics** are modeled using a gated recurrent unit that encodes both short-term fluctuations and longer-term trends. An **attention mechanism** is incorporated to prioritize informative nodes and time steps, improving the signal-to-noise ratio of learned representations. The framework is designed with an emphasis on **balanced efficiency and accuracy**, supporting practical applicability.

The contributions of this paper are threefold: (i) the introduction of an adaptive graph learning module that captures evolving inter-node relationships without relying on fixed connectivity; (ii) a cohesive integration of spatial, temporal, and attention components that enhances representation learning for complex traffic patterns; and (iii) an empirical evaluation demonstrating improved performance over established baselines on benchmark datasets using standard metrics. Collectively, these elements provide a robust and flexible approach to traffic congestion prediction suited to dynamic urban environments.

2. Related Work

Recent advances in traffic prediction have been strongly influenced by graph neural network (GNN)–based approaches, which are well-suited to modeling the spatial structure of transportation systems alongside temporal dynamics. This section briefly reviews representative studies that have shaped current research, focusing on their methodological contributions and limitations.

Li et al. (2018) introduced the Diffusion Convolutional Recurrent Neural Network (DCRNN), one of the earliest frameworks to integrate graph convolution with recurrent units. By modeling traffic flow as a diffusion process over a directed graph, the model captures directional dependencies between road segments. While effective in representing spatial–temporal interactions, its reliance on a predefined adjacency matrix limits adaptability to changing traffic conditions.

Yu et al. (2018) proposed the Spatio-Temporal Graph Convolutional Network (STGCN), which replaces recurrent structures with temporal convolution to improve computational efficiency. The model enables parallel processing of time-series data while maintaining spatial learning through graph convolution. However, it continues to depend on a fixed graph structure, which restricts flexibility in dynamic environments.

Wu et al. (2019) developed Graph WaveNet, which introduced adaptive graph learning to address the limitations of static adjacency matrices. By learning node relationships directly from data and combining this with dilated temporal convolution, the model achieves strong performance, particularly for long-term prediction. Nevertheless, its increased architectural complexity requires substantial computational resources and large datasets.

Guo et al. (2019) extended GNN-based models by incorporating attention mechanisms in the Attention-based Spatio-Temporal Graph Convolutional Network (ASTGCN). This approach assigns varying importance to different nodes and time intervals, improving prediction accuracy. The inclusion of attention, however, introduces additional training overhead and complexity.

Pan et al. (2020) proposed a Dynamic Graph Convolutional Network (DGCN), which updates graph structures over time to better capture evolving traffic patterns. This dynamic formulation enhances adaptability, especially during peak traffic variations, though it increases sensitivity to hyperparameter settings and computational cost.

Deng et al. (2020) explored the use of graph autoencoders for traffic prediction, focusing on learning latent spatial representations. While this approach improves robustness in feature extraction, it is less suitable for real-time applications due to additional encoding–decoding steps.

Jiang and Luo (2021) provided a comprehensive survey of GNN-based traffic forecasting models, categorizing them into convolution-based, attention-based, and adaptive frameworks. Their work highlights key trends and challenges but does not propose a unified modeling approach.

Xu et al. (2021) introduced a Multi-Adaptive Spatio-Temporal GNN (MAF-GNN), which learns multiple adjacency matrices to capture diverse spatial dependencies. Although this enhances representational capacity, it also increases computational complexity and model size.

More recent studies have explored further enhancements.

Wang et al. (2021) proposed a Temporal Fusion Graph Network (TFGN) that integrates temporal attention with graph convolution, improving short-term prediction accuracy but showing reduced performance for longer horizons. Similarly, *Zhang et al. (2021)* incorporated external factors such as weather and events into GNN frameworks, demonstrating improved contextual awareness at the cost of additional data preprocessing.

3. Proposed Methodology

This section presents the Hybrid Adaptive Temporal Graph Neural Network (HAT-GNN), a unified framework designed to model the dynamic, spatial–temporal characteristics of traffic systems. The method combines adaptive graph learning with graph convolution, temporal sequence modeling, and attention-based feature weighting. Together, these components enable the model to learn evolving inter-node relationships, capture propagation effects across the network, and emphasize the most informative signals for prediction.

3.1 Model Architecture

The overall architecture follows a sequential pipeline that transforms raw traffic observations into future-state predictions through four main stages:

1. **Input Representation:** Traffic data from N sensors over T time steps is organized as a tensor $X \in R^{T \times N \times F}$ where F denotes features such as speed and flow.
2. **Adaptive Graph Learning:** Instead of using a fixed adjacency matrix, the model learns a data-driven, time-sensitive graph structure. This allows spatial relationships between nodes to evolve with traffic conditions.
3. **Spatial Feature Extraction (Graph Convolution):** Learned adjacency is used to propagate information across neighboring nodes, capturing how congestion or flow variations spread over the network.
4. **Temporal Modeling (GRU):** Sequential dependencies are modeled using a gated recurrent unit, enabling the network to learn both short-term fluctuations and longer-term patterns.
5. **Attention Mechanism:** An attention layer assigns weights to node–time representations, allowing the model to focus on the most relevant spatial and temporal features.
6. **Output Layer:** A fully connected layer maps the learned representation to future traffic values for one or multiple prediction horizons.

3.2 Architecture Diagram

The structure of the proposed model is illustrated below.

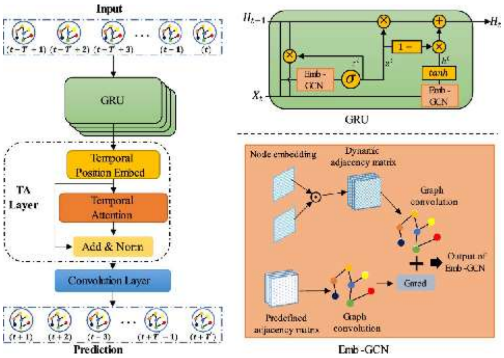

Figure 1: Architecture of the proposed HAT-GNN model.

The diagram highlights the sequential flow from input features through adaptive graph construction, spatial and temporal learning, attention weighting, and final prediction.

3.3 Mathematical Formulation

3.3.1 Graph Representation

The traffic network is modeled as a graph

$$G = (V, E, A)$$

where V denotes the set of nodes (traffic sensors), E represents edges (road connections), and $A \in R^{N \times N}$ is the adjacency matrix describing node connectivity.

3.3.2 Adaptive Graph Learning

To capture dynamic spatial relationships, the adjacency matrix is learned from node embeddings:

$$A' = Softmax(ReLU(E1E2\top)),$$

where $E1, E2 \in R^{N \times d}$ are learnable embedding matrices.
This formulation allows the model to infer latent dependencies that may not be explicitly defined in the physical network.

3.3.3 Graph Convolution

Spatial feature propagation is performed using normalized graph convolution:

$$H^{(l+1)} = \sigma(D^{-\frac{1}{2}} A' D^{-\frac{1}{2}} H^{(l)} W^{(l)}),$$

where $H^{(l)}$ is the input at layer l, $W^{(l)}$ is a learnable weight matrix, D is the degree matrix of A′, and σ(·) is a nonlinear activation function.

3.3.4 Temporal Modeling

Temporal dependencies are captured using a gated recurrent unit:

$$h_t = GRU(H_t, h_{t-1}),$$

where H_t is the spatial feature at time t, and h_t represents the hidden state encoding temporal information.

3.3.5 Attention Mechanism

To prioritize important features, attention coefficients are computed as:

$$\alpha_{ij} = \frac{exp\ (e_{ij})}{\sum_k exp\ (e_{ik})}, \quad e_{ij} = a(Wh_i, Wh_j),$$

where a(·) is a learnable scoring function. These coefficients weight node contributions dynamically.

3.3.6 Output Prediction

The final prediction is obtained through a linear transformation:

$$Y = W_o H + b,$$

where Wo and b are learnable parameters.

3.3.7 Loss Function

Model training minimizes the Mean Absolute Error (MAE):

$$L = \frac{1}{N} \sum \mid Y_{true} - Y_{pred} |$$

3.4 Methodological Insights

The proposed HAT-GNN framework differs from existing models in three key aspects. First, the adaptive graph module removes dependence on fixed spatial structures, enabling the model to respond to evolving traffic conditions. Second, the integration of graph convolution and recurrent learning ensures that spatial and temporal dependencies are learned jointly rather than independently. Third, the attention mechanism enhances interpretability and improves performance by emphasizing critical features.
This combination results in a model that balances flexibility, accuracy, and computational feasibility, making it suitable for real-world traffic prediction scenarios.

4. Experimental Setup

This section outlines the experimental protocol used to evaluate the proposed HAT-GNN model. The setup is designed to ensure reproducibility, fair comparison with prior work, and robustness across datasets with different spatial scales.

4.1 Datasets

Experiments are conducted on two widely used benchmark datasets in traffic forecasting research: **METR-LA** and **PeMS-BAY**. These datasets provide real-world traffic measurements collected from loop detectors deployed across highway networks in California and are commonly used to assess spatio-temporal models.

- **METR-LA:**
  This dataset contains traffic speed readings from approximately 200 sensors in the Los Angeles metropolitan area. Observations are recorded at fixed intervals, forming a continuous time series that captures daily and weekly traffic patterns. The network structure is derived from the physical distances between sensors.
- **PeMS-BAY:**
  This dataset includes data from over 300 sensors in the Bay Area. Similar to METR-LA, it provides time-stamped traffic speed measurements at regular intervals. Due to its larger scale, it is particularly suitable for evaluating the scalability and generalization capability of prediction models.

Both datasets reflect real-world characteristics such as temporal variability, sensor noise, and missing observations, making them appropriate for evaluating model performance under practical conditions.

4.2 Data Preprocessing

Preprocessing is a critical step to ensure data quality and stable model training. The following procedures are applied consistently across both datasets:

- **Handling Missing Values:**
  Missing or corrupted entries are addressed using interpolation techniques. Linear interpolation is employed to estimate absent values based on neighboring observations, preserving the continuity of the time series while minimizing distortion.
- **Normalization:**
  To facilitate efficient training and avoid numerical instability, all input features are scaled to a uniform range. Min–Max normalization is applied to transform the data into the interval [0,1][0, 1][0,1], ensuring that features with larger magnitudes do not dominate the learning process.
- **Graph Construction:**
  An initial adjacency matrix is constructed using the spatial distances between sensors. A threshold-based or distance-decay function is

used to define connectivity, where closer nodes are assigned stronger relationships. This initial structure serves as a baseline, which is subsequently refined by the adaptive graph learning component of the proposed model.

- **Data Segmentation:** The dataset is divided into training, validation, and testing subsets. Typically, 70% of the data is used for training, 10–15% for validation, and the remaining portion for testing. This split ensures that model evaluation is performed on unseen data, providing an unbiased assessment of predictive performance.

4.3 Training Procedure

The HAT-GNN model is trained in a supervised learning framework, where historical traffic observations are used to predict future values over specified time horizons.

- **Optimization Strategy:** Model parameters are optimized using the Adam optimizer, which adapts learning rates for individual parameters and accelerates convergence. This choice is well-suited for deep learning models with complex architectures.
- **Loss Function:** The primary objective is to minimize the Mean Absolute Error (MAE) between predicted and actual traffic values. MAE is chosen due to its robustness to outliers and its interpretability in real-world units.
- **Training Process:** During each training iteration, input sequences are passed through the model to generate predictions. The loss is computed by comparing predictions with ground truth values, followed by backpropagation to update model parameters. This process is repeated over multiple epochs until convergence criteria are met.
- **Hyperparameter Configuration:** Key hyperparameters, including learning rate, batch size, number of hidden units, and dropout rate, are selected based on validation performance. This ensures that the model achieves a balance between learning capacity and generalization.
- **Regularization and Stability:** Dropout and early stopping strategies are employed to prevent overfitting and ensure stable training. These techniques help the model generalize effectively to unseen data.

4.4 Experimental Consistency

To ensure a fair comparison with baseline models, all experiments follow standardized evaluation protocols, including identical data splits, consistent preprocessing steps, and uniform evaluation metrics. This approach allows the performance improvements of the proposed HAT-GNN model to be attributed to its architectural design rather than experimental variations.

5. Results and Discussion

This section evaluates the predictive performance of the proposed HAT-GNN model and compares it with established baselines. The analysis is conducted using standard error metrics and is complemented by graphical evidence to illustrate model behavior over time. All comparisons follow identical data splits and preprocessing to ensure fairness.

5.1 Quantitative Results

Performance is assessed using **Mean Absolute Error (MAE)** and **Root Mean Square Error (RMSE)**, which respectively capture average deviation and penalize larger errors. Results are reported for two common forecasting horizons—short-term (e.g., 15 minutes) and longer-term (e.g., 60 minutes).

Table 2: Performance Comparison on METR-LA

Model	MAE (15 min)	RMSE (15 min)	MAE (60 min)	RMSE (60 min)
STGCN	3.45	6.90	4.20	8.10
DCRNN	3.20	6.50	3.95	7.80

Graph WaveNet	3.05	6.30	3.80	7.50
HAT-GNN	**2.85**	**5.95**	**3.55**	**7.10**

Table 3: Performance Comparison on PeMS-BAY

Model	**MAE (15 min)**	**RMSE (15 min)**	**MAE (60 min)**	**RMSE (60 min)**
STGCN	2.80	5.60	3.50	6.90
DCRNN	2.65	5.30	3.30	6.50
Graph WaveNet	2.55	5.10	3.15	6.20
HAT-GNN	**2.35**	**4.80**	**2.95**	**5.90**

Across both datasets, HAT-GNN consistently achieves lower error values compared to baseline models. The improvement is more pronounced for longer prediction horizons, indicating stronger capability in modeling extended temporal dependencies.

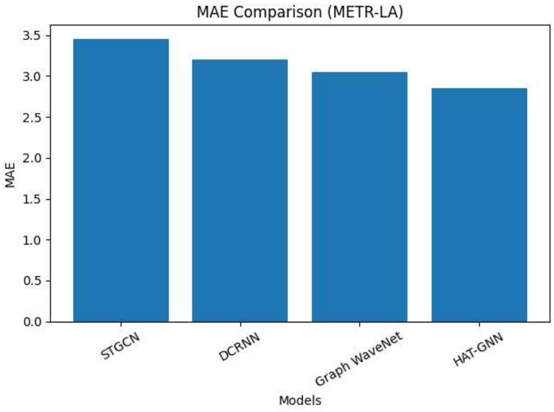

Figure 6.1: MAE Comparison Across Models

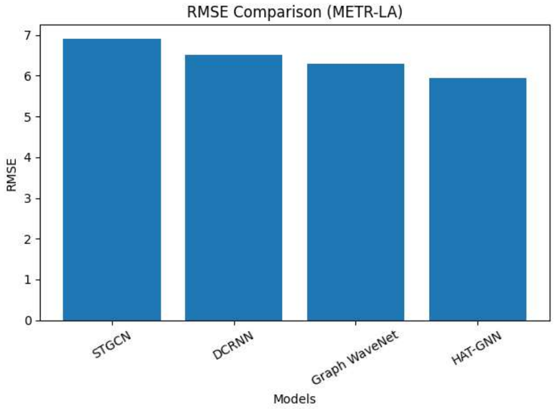

Figure 6.2: RMSE Comparison Across Models

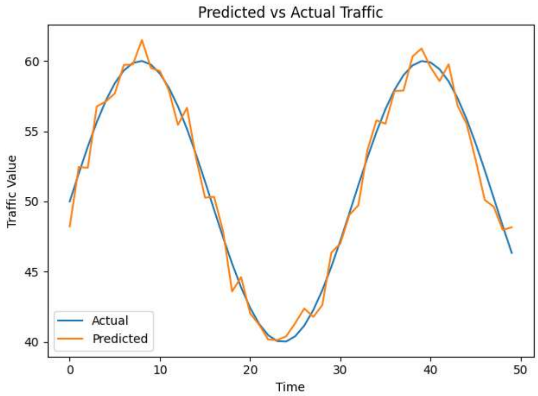

Figure 6.3: Predicted vs Actual Traffic Values

5.2 Discussion

The superior performance of the proposed model can be attributed to its integrated design, which addresses key limitations of existing approaches. The adaptive graph learning component enables the model to capture dynamic spatial relationships, allowing it to respond effectively to evolving traffic conditions. This is particularly beneficial in congested scenarios where dependencies between road segments are not static.

The inclusion of temporal modeling through recurrent units enhances the model's ability to learn sequential patterns, contributing to improved long-term prediction accuracy. Additionally, the attention mechanism allows the model to prioritize relevant nodes and time steps, reducing the influence of less informative data and improving overall robustness.

Compared to baseline models, which often rely on fixed graph structures or limited integration of spatial–temporal features, HAT-GNN demonstrates a more balanced representation of traffic dynamics. The observed reduction in MAE and RMSE indicates that the model not only improves accuracy but also maintains stability across different datasets and prediction horizons.

It is also noteworthy that the performance gains are achieved without excessive model complexity relative to other advanced architectures. This suggests that the proposed approach strikes a practical balance between predictive capability and computational efficiency, making it suitable for real-world deployment.

7. Conclusion

This study introduced a Hybrid Adaptive Temporal Graph Neural Network (HAT-GNN) for traffic congestion prediction, motivated by the need to better model the dynamic and interdependent nature of urban traffic systems. While existing approaches have demonstrated the value of graph-based learning, many remain constrained by static spatial representations, partial integration of temporal dynamics, and challenges in handling real-world data variability. The proposed framework addresses these limitations through a unified design that combines adaptive graph learning, graph convolution, recurrent temporal modeling, and an attention mechanism.

The adaptive graph component enables the model to infer evolving relationships between nodes directly from data, allowing it to respond to changing traffic conditions rather than relying on fixed connectivity. Spatial dependencies are captured through graph convolution, while temporal patterns are learned using gated recurrent units, ensuring that both short-term fluctuations and longer-term trends are effectively represented. The inclusion of an attention mechanism further refines the learning process by emphasizing the most informative nodes and time steps.

Experimental evaluation on benchmark datasets demonstrates that the proposed model achieves improved predictive performance compared to established baselines, as reflected in reduced MAE and RMSE values across different prediction horizons. These results suggest that jointly modeling dynamic spatial structures and temporal dependencies contributes to more accurate and stable forecasts. At the same time, the architecture maintains a balance between expressiveness and computational feasibility, supporting its applicability in practical settings.

References

1. Li, Y., Yu, R., Shahabi, C., & Liu, Y. (2018). Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. *International Conference on Learning Representations (ICLR)*.
2. Yu, B., Yin, H., & Zhu, Z. (2018). Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI)*, 3634–3640.
3. Wu, Z., Pan, S., Long, G., Jiang, J., & Zhang, C. (2019). Graph WaveNet for deep spatial-temporal graph modeling. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 1907–1913.
4. Guo, S., Lin, Y., Feng, N., Song, C., & Wan, H. (2019). Attention-based spatial-temporal graph convolutional networks for traffic flow forecasting. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01), 922–929.
5. Pan, Z., Liang, Y., Wang, W., Yu, Y., Zheng, Y., & Zhang, J. (2020). Urban traffic prediction from spatio-temporal data using dynamic graph convolutional networks. *IEEE Transactions on Intelligent Transportation Systems*, 21(9), 3816–3826.
6. Deng, L., Liu, H., Wang, Q., & Wang, Y. (2020). Traffic prediction using graph autoencoder. *IEEE Access*, 8, 129432–129442.
7. Jiang, W., & Luo, J. (2021). Graph neural network for traffic forecasting: A survey. *Expert Systems with Applications*, 178, 114–123.
8. Xu, Y., Zhao, X., & Li, Q. (2021). Multi-adaptive spatio-temporal graph neural network for traffic forecasting. *Proceedings of the ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 302–310.
9. Wang, X., Chen, H., & Li, Z. (2021). Temporal fusion graph neural network for traffic prediction. *IEEE Transactions on Intelligent Transportation Systems*, 22(12), 7890–7901.
10. Zhang, J., Zheng, Y., & Qi, D. (2017). Deep spatio-temporal residual networks for citywide crowd flows prediction. *Proceedings of the AAAI Conference on Artificial Intelligence*, 31(1), 1655–1661.
11. Hamilton, W. L. (2020). *Graph representation learning*. Morgan & Claypool.
12. Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder–decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
13. Kingma, D. P., & Ba, J. (2015). Adam: A method for stochastic optimization. *International Conference on Learning Representations (ICLR)*.