ITC 2/55

Information Technology
and Control
Vol. 55 / No. 2/ 2026
pp. 585-604
DOI 10.5755/j01.itc.55.2.43987

RM-DSTGN: Dynamic Graph Topology Evolution via Rainfall-Modulated Attention for Spatiotemporal Traffic State Estimation

Received 2025/12/31 | Accepted after revision 2026/04/16

HOW TO CITE: Chen, Y. (2026) RM-DSTGN: Dynamic Graph Topology Evolution via Rainfall-Modulated Attention for Spatiotemporal Traffic State Estimation. *Information Technology and Control*, 55(2), 585-604. https://doi.org/10.5755/j01.itc.55.2.43987

RM-DSTGN: Dynamic Graph Topology Evolution via Rainfall-Modulated Attention for Spatiotemporal Traffic State Estimation

Yi Chen

School of Statistics and Mathematics, Shandong University of Finance and Economics, Jinan, 250014, China

Corresponding author: chenyi20050913@163.com

Accurate traffic speed prediction in Intelligent Transportation Systems (ITS) is challenging under adverse weather due to highly non-linear traffic dynamics. Traditional deep learning relies on static graphs and treats weather as a peripheral feature, failing to capture rainfall-induced structural shifts in traffic propagation. To address this, we propose the Rainfall-Modulated Dynamic Spatiotemporal Graph Network (RM-DSTGN). Unlike conventional models, RM-DSTGN internalizes rainfall intensity as a global structural controller that dictates spatial connectivity and temporal reliance. The framework integrates a rainfall-modulated graph generator to dynamically reconfigure spatial dependencies, a dual-branch stability module decoupling macroscopic trends from microscopic speed variances, and a gating mechanism calibrating temporal reliance based on rain severity. Evaluated on real-world data from a critical Hong Kong arterial intersection, RM-DSTGN significantly outperforms state-of-the-art benchmarks, achieving a mean absolute percentage error (MAPE) of 3.62%. Comprehensive ablation and error analyses confirm the model's robustness across varying rainfall intensities, prediction horizons, and look-back windows. Ultimately, RM-DSTGN maintains consistent service levels during extreme weather, validating its potential to enhance urban mobility and safety.

KEYWORDS: Traffic speed prediction; Spatiotemporal graph neural networks; Dynamic graph topology learning; Rainfall modulation

1. Introduction

With the rapid advancement of urbanisation and the corresponding increase in vehicle ownership, the issue of traffic congestion has become a severe problem that should be addressed in modern cities. Such issues not only result in environmental pollution but also impact economic development and diminish the daily quality of life for residents. Intelligent Transportation Systems (ITS) are widely regarded as a crucial infrastructure for mitigating these issues and advancing the development of smart cities. As a vital component of any effective ITS, the capability to make timely traffic predictions lays a solid foundation for its development. Among the various traffic parameters, traffic speed is one of the most important and direct indicators of traffic conditions and congestion levels. Therefore, reliable real-time traffic speed prediction is crucial for advanced traffic applications, such as dynamic route planning, intelligent traffic management, and traffic congestion forecasting.

However, achieving accurate and reliable traffic speed prediction is an intrinsically difficult task, fraught with numerous challenges. The first major hurdle lies in the data itself. Real-world traffic data is collected from a wide array of multi-source and heterogeneous sensors, including fixed-loop detectors, vehicle GPS trajectories, and UAV imagery. Its diversity presents significant challenges in data processing. Raw data streams are often incomplete and unreliable, plagued by issues such as noise, missing data, information redundancy, and significant data sparsity, mainly when relying on partially observed trajectories from probe vehicles or fleets.

Beyond data quality, traffic speed is governed by highly complex, dynamic, and non-linear characteristics. Traffic flow does not merely follow simple, repeatable trends. Its instability is often exacerbated by a multitude of external factors that are difficult to model, such as unexpected events (e.g., accidents and construction work), adverse weather conditions (e.g., rain), and special days (e.g., national holidays). Furthermore, traffic speed is governed by complex spatiotemporal dependence. The state of one road segment is not only dependent on its own history (temporal dependence) but is also intricately linked to the conditions of its neighbouring roads and connected roads (spatial dependence).

To address these issues, traffic speed prediction models have undergone significant evolution. Initial approaches are dominated by traditional statistical methods, such as the historical average (HA) model, autoregressive integrated moving average (ARIMA) [21], and Kalman filters. These models, however, were proven to be limited in their ability to capture the complex, non-linear, and stochastic nature of traffic flow. Consequently, the field has increasingly shifted towards data-driven machine learning approaches. Models such as long short-term memory (LSTM) [13], and gated recurrent units (GRU) [15] have proven effective in learning temporal dependencies. More recently, graph neural networks (GNNs) [7] have been effectively employed to model the non-Euclidean spatial relationships within road networks, leading to the development of sophisticated spatiotemporal models. The evolution has been fueled by increasingly diverse data sources, including traditional loop detectors, connected vehicles (CVs), GPS trajectories, and even unmanned aerial vehicles.

Despite these advancements, two significant challenges persist. First, many state-of-the-art spatiotemporal models are computationally intensive and exhibit high inference latency. Its complexity makes it difficult to deploy on resource-constrained edge devices (e.g., roadside units) for the instantaneous calculations required in real-time applications. The resulting latency is often unacceptable for rapidly changing traffic conditions, as it can prevent the system from promptly managing traffic in response to sudden fluctuations in flow.

Second, the majority of existing research has concentrated on mining endogenous patterns solely from historical traffic data. However, external dynamic factors (e.g., adverse weather) have a significant impact on traffic conditions. Adverse weather, particularly rainfall, directly degrades driving conditions by reducing road surface friction and visibility. It compels drivers to adopt more cautious behaviours (e.g., lowering speed and increasing headway), which collectively diminish road capacity, lead to sharp, non-linear drops in average traffic speed, and can increase the likelihood of accidents. Although a few studies have attempted to incorporate weather data, they often fail to capture the complex structural shifts in traffic propagation caused by these environmental changes.

Even recent advancements in dynamic graph learning, such as DMSTG [8], while having made strides in capturing evolving spatial dependencies, still predominantly rely on implicit correlations inferred from historical traffic states. Crucially, these models typically treat weather conditions merely as peripheral features fused via simple concatenation or attention mechanisms, failing to model how rainfall physically dictates the topological deterioration of the road network. Consequently, they lack the capability to explicitly reconfigure spatial connectivity and isolate the microscopic traffic instability induced by adverse weather, which limits their robustness in extreme conditions.

To address the critical research gap, this paper proposes a novel traffic speed prediction model that explicitly integrates rainfall data, with a primary focus on the impact of rainfall on traffic speed. The central innovation of this work is the re-characterization of rainfall intensity from a passive feature to a global structural controller. By embedding environmental context directly into the spatial topology and temporal gating mechanisms, the model adaptively restructures its spatiotemporal logic in direct response to real-time weather deterioration.

2. Literature Review

To better contextualize the proposed RM-DSTGN model and its unique integration of rainfall data, it is necessary to examine the methodological evolution of traffic forecasting. This section provides a comprehensive overview of the research landscape, tracing the transition from early statistical foundations to contemporary data-driven paradigms, thereby highlighting the gaps that our work intends to bridge.

2.1. Machine Learning Approaches

Traditional machine learning (ML) techniques marked a significant step forward by improving the capability to model non-linear relationships. Unlike purely statistical methods, data-driven algorithms can learn complex mappings between input features and traffic speed. For example, [14] successfully applied a clustered k-nearest neighbors (KNN) algorithm for short-term travel speed prediction. Furthermore, the flexibility of ML allows for the integration of multi-source data; [1] demonstrated this by employing machine learning models to predict traffic conditions using both traffic features and weather information. Such data-driven foundations have also been effectively extended to modern distributed environments for real-time traffic flow prediction [2].

Despite being more robust than their statistical predecessors, classical ML models heavily rely on manual feature engineering. More importantly, they fundamentally struggle to capture the complex, high-dimensional spatiotemporal dependencies inherent in large-scale road networks. This structural limitation inevitably leads to significant performance degradation under highly dynamic traffic scenarios, ultimately catalyzing the shift toward advanced deep learning architectures.

2.2. Deep Learning for Spatiotemporal Traffic Prediction

The rapid growth of transportation big data, combined with advancements in computational power, has been a key factor. It has pushed deep learning (DL) [19] to the forefront of traffic prediction research. Because deep learning models possess the capability to learn hierarchical features, they excel at uncovering complex, non-linear spatiotemporal patterns that traditional methods struggle to detect.

2.2.1. RNNs and their Variants

Recurrent Neural Networks (RNNs) and their advanced variants, such as LSTM networks and GRU, have fundamentally transformed temporal sequence modeling in traffic prediction. By maintaining an internal memory state, these architectures effectively process sequential data to capture long-term dependencies in traffic flow and speed time series. For instance, [13] proposed a hierarchical LSTM-based method for vehicle trajectory prediction that explicitly incorporated interaction information from neighboring vehicles, successfully reducing prediction errors. Expanding on this temporal capability, [18] integrated a Bidirectional LSTM (BiLSTM) within a deep learning neural network to capture complex temporal dynamics under varying weather conditions.

While standalone RNNs excel at capturing temporal dynamics, they inherently struggle to model spatial correlations across a complex road network. Traffic conditions at any given location are influenced not only by their own historical states but also by the dynamic fluctuations of adjacent and even distant road segments. This fundamental limitation spurred the integration of RNNs with spatial-focused deep learning architectures to capture network-wide topological dependencies more comprehensively.

2.2.2. Graph Neural Networks (GNNs) and Spatiotemporal Graph Neural Networks (STGNNs)

GNNs have demonstrated significant potential in effectively capturing spatiotemporal correlations by modeling traffic networks as graphs [7]. In these frameworks, graph convolutions typically capture spatial dependencies, while temporal modules handle dynamic changes. Building upon these foundational concepts, recent studies have further refined STGNN architectures by integrating attention mechanisms and advanced encoding techniques capture inherent traffic dynamic. For instance, [23] introduced spatiotemporal-aware position encoding (STAPE) combined with a multi-head attention mechanism to address multi-period heterogeneity and enhance model adaptability.

Beyond static structures, specialized research has shifted toward dynamic graph learning to capture evolving spatial dependencies. State-of-the-art dynamic models, exemplified by DMSTG [8], focus on generating dynamic adjacency matrices based on the similarity of hidden traffic states. However, this purely data-driven approach often overlooks the physical mechanisms of environmental interference. In contrast, RM-DSTGN transcends this limitation by elevating rainfall from a passive input feature to a global structural controller. By explicitly modulating graph topology and temporal gating based on real-time rainfall intensity, our framework captures distinct regime shifts in traffic dynamics that latent-state models like DMSTG [8] fail to distinguish.

2.2.3. Attention Mechanisms and Transformers

Attention mechanisms and Transformer architectures have become increasingly popular because they capture long-range dependencies better than traditional recurrent or convolutional layers. Rather than treating all inputs equally, they allow models to adaptively weight the most critical spatiotemporal features. For example, the STTF model [20] embeds spatial and temporal information through a dedicated attention module to improve congestion forecasting. [5] also applied a Transformer-based framework, designing a spatial mask to filter out noise while encoding complex spatial interactions.

Incorporating attention into graph structures has driven significant progress as well, most notably with Graph Attention Networks (GAT) [9], which dynamically evaluate the importance of neighboring nodes. Extending this concept, MAT-WGCN [17] couples weighted graph convolutions with multi-head attention, resulting in higher spatiotemporal integration and model robustness. Another recent variation, MTS-GATN [9], extracts features across multiple time scales and augments them with dynamic spatial semantic embeddings. Ultimately, by selectively focusing on the most informative contexts, these attention-driven designs consistently yield more accurate and stable predictions.

2.2.4. Hybrid and Ensemble Deep Learning Models

No single deep learning architecture can capture the full complexity of traffic networks. This limitation has driven a shift toward hybrid and ensemble models that combine the strengths of RNNs, GNNs, and attention mechanisms. [16] demonstrated this by integrating an S-shaped traffic stream model with multi-graph convolutional and LSTM networks. This physics-informed framework improved both performance and generalization, even with limited training data.

Ensemble methods offer another route to prediction stability by aggregating multiple base learners. [3] built a deep stacking model using XGBoost, Random Forest, and Extra Trees, relying on an MLP meta-learner to synthesize the final short-term predictions.

Yet, a critical gap remains. While these combined architectures perform well under normal conditions, their reliability degrades rapidly during environmental anomalies. Standard frameworks generally lack the structural mechanisms to process external disruptions like rainfall—treating weather as minor background noise rather than a force that fundamentally shifts underlying traffic patterns.

2.3. Incorporating Rainfall in Traffic Speed Prediction

Traffic speed is not just a product of road topology and historical patterns; it is highly sensitive to environmental conditions. Rainfall directly alters traffic dynamics by reducing road friction and impairing visibility, forcing drivers to increase headways and lower their speeds. This physical disruption naturally increases the likelihood of congestion.

Standard traffic prediction models often marginalize these weather effects. However, robust forecasting requires explicit meteorological modeling. [24], for example, linked rainfall intensity directly to mandatory speed reductions when optimizing safe speed limits for freeways. This highlights the tight coupling between precipitation and traffic states.

Incorporating weather data into predictive models is difficult due to the non-linear spatiotemporal interactions between meteorological and traffic variables. [18] tackled this by developing a deep learning framework that explicitly feeds multiple weather factors—such as rainfall, temperature, and wind speed—into the prediction pipeline.

Alternatively, some methods correct the input data before training. [4] proposed a pre-tuning methodology that adjusts historical speed data based on weather parameters prior to applying deep learning algorithms. Furthermore, handling the rainfall data itself requires robust feature engineering, as demonstrated by [10] in their machine learning-based rainfall prediction models. These data-processing techniques are directly transferable when building rainfall-aware traffic systems.

2.4. Limitations of Existing Research and Entry Points for Its Study

Significant progress has been made in predicting traffic speeds, particularly with the advent of deep learning and spatiotemporal graph neural networks. Nevertheless, several limitations remain regarding the comprehensive integration of external factors like rainfall. These existing gaps present crucial opportunities for further study and define the starting points for this research.

While many spatiotemporal models effectively capture complex dependencies, they frequently rely on static graph structures that ignore the fluid nature of traffic. The influence between road segments changes significantly based on real-time conditions, incidents, or time of day. Although recent studies have begun to explore dynamic graph construction [6, 8, 12, 22], adaptively modeling these evolving spatial relationships remains challenging—especially under the influence of temporary external factors. Relying on static graphs inevitably leads to suboptimal performance when traffic deviates from historical norms.

Furthermore, while many models now incorporate external factors, the depth of this integration varies widely. Weather data is often treated as a simple appended feature without explicitly modeling its causal impact on traffic physic [1, 18]. This "black-box" approach limits the model's ability to handle unseen extreme conditions. Rainfall, for instance, is frequently reduced to a binary indicator (rain/no-rain) or a single numerical value, discarding vital information about how different precipitation patterns distinctively alter speed. While methods like weather-oriented pre-tuning [4] offer a step toward more specific meteorological modeling, these adaptations often occur outside the core architecture. A comprehensive approach integrated directly into the main prediction structure is still necessary.

The specific role of rainfall remains underexplored compared to other advancements in spatiotemporal modeling. High-performing STGNNs primarily focus on optimizing graph structures or employing attention mechanisms to capture inherent traffic dynamics [23], placing far less emphasis on how environmental conditions fundamentally disrupt these dynamics. Even models that do include weather variables and demonstrate improved accuracy [18] often fail to capture the complex, spatially heterogeneous ways rainfall alters network flow. Therefore, models must dynamically adjust their spatiotemporal reasoning based on real-time rainfall characteristics, rather than treating precipitation as just another static input.

Finally, the interpretability of complex deep learning models remains a persistent challenge, particularly when integrating multiple spatiotemporal dependencies and external factors. Understanding the rationale behind a prediction is crucial for gaining the trust of traffic managers and enabling actionable insights. While studies like Wang et al. [19] have begun exploring counterfactual explanations for deep learning traffic forecasting, significant work is still needed to provide clear, rainfall-specific interpretations of model behavior.

Despite advancements in spatiotemporal traffic prediction, current methodologies exhibit significant limitations when applied to adverse weather conditions.

First, the majority of existing models rely on static graph topologies to represent road networks, assuming that spatial dependencies between nodes remain constant. The assumption is inadequate during rainfall, as environmental factors dynamically alter road friction and visibility, causing structural shifts in traffic propagation patterns that static graphs cannot capture.

Second, the integration of rainfall data is often shallow and simplistic; weather data is typically treated as a peripheral input feature via direct concatenation. This approach fails to model rainfall as a "structural controller" that fundamentally dictates traffic physics.

Finally, traditional deep learning models often struggle with the instability and concept drift caused by extreme weather conditions. They frequently fail to decouple microscopic speed variances from macroscopic trends and lack adaptive mechanisms to reduce reliance on historical periodic patterns when rain renders past data unreliable.

To bridge the identified research gaps, this study offers three primary contributions within the proposed RM-DSTGN framework:

Contribution 1 (C1): Constructing a rainfall-adaptive dynamic graph topology: To overcome the inadequacy of static graphs, we focus on generating dynamic adjacency matrices driven explicitly by real-time rainfall intensity. It allows the model to capture shifting spatial dependencies and dynamically reconfigure network connectivity as environmental conditions deteriorate.

Contribution 2 (C2): Modelling rainfall-induced instability via a dual-branch architecture: Addressing the failure to capture non-linear volatility, we design a dual-branch structure that explicitly separates the learning of macroscopic trends from microscopic speed variances (instability). It ensures that the stochastic impact of rainfall on traffic safety and turbulence is modelled distinctively from standard flow patterns.

Contribution 3 (C3): Dynamically calibrating temporal dependencies with a gating mechanism: To resolve the issue of unreliable historical patterns during anomalies, we introduce a gating mechanism to calibrate the model's temporal reliance dynamically. Its mechanism acts as a "regulating valve," adaptively weighing the importance of historical periodic features against immediate short-term disturbances based on the severity of the rainfall.

By tackling these gaps, this research seeks to contribute to more accurate, robust, and interpretable traffic speed prediction systems, particularly vital for enhancing urban mobility and safety during adverse weather conditions.

3. Methodology

Building on the identified research gaps, this section details the methodology for our adaptive, rainfall-aware prediction system. We first formalize the problem through a dynamic, multi-factor lens to establish the theoretical foundation for the proposed RM-DSTGN model.

3.1. Problem Definition and Preliminaries

Following the research gaps identified in Section 2, particularly the inadequacy of static graph assumptions under adverse weather and the underutilization of rainfall as a structural determinant, we first formalize the traffic prediction problem in a dynamic, multi-factor context.

Definition 1 (Dynamic Traffic Road Network): We define the road network as a graph, where $G = (V, E, A_t)$, $V = \{v_1, v_2, \ldots, v_N\}$ represents the set of N sensors (nodes), and E represents the set of edges (i.e., the road connections between these sensors). Unlike traditional approaches that rely on a time-invariant adjacency matrix A_{static} (as criticized in previous literature), we postulate that the spatial dependencies between nodes are time-variant and modulated by environmental conditions. Thus, we aim to learn a dynamic adjacency matrix $A_t \in R^{N \times N}$ at each time step t to capture shifting traffic propagation patterns.

Definition 2 (Multi-Stream Feature Tensor): To explicitly capture the heterogeneous physical characteristics of traffic dynamics discussed in Section

2.3, we construct a multi-stream input tensor X consisting of three distinct components. First, the Macro-Traffic Stream (X_{trend}) comprises historical speed and flow observations, representing macroscopic traffic trends and periodicity. Second, the Stability Stream (X_{stab}) incorporates the speed variance sequence; as highlighted in our literature review, high variance serves as a critical proxy for traffic turbulence and safety risks induced by rainfall, thereby representing the microscopic instability of the flow. Finally, the Contextual Stream and microscopiccontains the rainfall intensity sequence, denoted as $R = \{R_{t-T+1},...,R_t\}$, where R_t represents the real-time rainfall intensity at time step t, which functions not merely as an input feature but as a global environmental controller regulating the system dynamics.

Problem Statement: Given the historical multi-stream observations over a look-back window of length T, denoted as $X = \{X_{t-T+1},...,X_t\}$, the objective is to learn a mapping function $f(\cdot)$ to predict the future traffic speed $\hat{Y}_{t+1}$.

3.2. Framework Overview: RM-DSTGN

To address the challenges of non-linear rainfall impacts and spatial heterogeneity, we propose the rainfall-modulated dynamic spatiotemporal graph network (RM-DSTGN). As illustrated in Figure 1 the framework departs from the conventional concatenation-based fusion strategies. Instead, it adopts a structure-aware design that explicitly models the causal mechanisms of adverse weather on traffic. The framework comprises three coupled modules.

Specifically, the rainfall-modulated dynamic topology learning mechanism addresses the limitations of static topologies by dynamically reconfiguring spatial dependencies (A_t) driven by real-time rainfall intensity. It is complemented by the dual-branch stability-aware modelling, a parallel architecture that decouples the learning of macroscopic trends. (X_{trend}) and microscopic variabilities (X_{stab}) via a cross-attention mechanism to rectify speed predictions under unstable conditions. Finally, the Rainfall-Gated Temporal Fusion module employs an adaptive gating mechanism to dynamically calibrate the model's reliance between historical periodic patterns and immediate real-time disturbances based on rain severity.

Figure 1

The framework of the proposed model

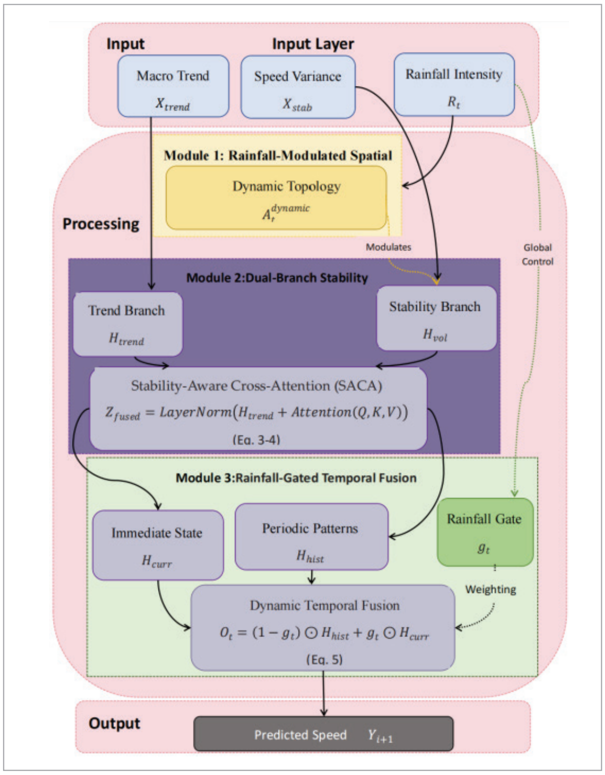

3.3. Rainfall-Modulated Dynamic Topology Learning (Spatial Module)

Standard graph convolutional networks (GCNs) generally rely on pre-defined, static adjacency matrices derived from Euclidean distances. However, the static assumption proves inadequate under adverse weather conditions, where environmental factors fundamentally alter traffic dynamics. For instance, heavy rainfall significantly degrades road friction and visibility, leading to shifts in congestion propagation patterns—such as extending the effective influence range of upstream nodes on downstream traffic. To address its non-stationarity, we propose generating a time-dependent graph structure wherein the prevailing environmental context explicitly modulates nodal connectivity.

Conceptually, we reframe rainfall intensity as a global structural modulator rather than a passive attribute of the node. To operationalize It, we introduce the Rainfall-Modulated Graph Generator. Let $E_{node} \in R^{N \times d}$ denote learnable static node embeddings designed to encode intrinsic topological properties, such as road geometry and functional type. To align the environmental context with these spatial features, the real-time rainfall intensity R_t is projected into the shared latent space via a learnable linear transformation W_r. Subsequently, the dynamic adjacency matrix A_t is derived through a modified self-attention mechanism, formulated as:

$$A_t^{dynamic} = Softmax\left(\frac{(E_{node}W_Q)(E_{node}W_K + \phi(R_t))^T}{\sqrt{d}}\right), \quad (1)$$

where W_Q and W_K are learnable weight matrices, and d is the dimension of node embeddings, and $\phi(\cdot)$ denotes a linear projection layer.

From a physical perspective, it terms functions as an environmental bias integrated into the key vector of the attention mechanism. Under clear weather conditions, the learned topology remains dominated by the static road network structure, thereby capturing routine spatial dependencies. Conversely, during heavy rainfall, the bias term effectively shifts the attention distribution, enabling the model to capture long-range dependencies—such as rapid congestion propagation—that are typically overlooked by static graph representations.

This dynamic matrix is then fed into the GCN layers to aggregate spatial features:

$$H_S^{(l)} = \sigma\left(A_t^{dynamic} H_S^{(l-1)} W^{(l)}\right), \quad (2)$$

where $H_s^{(l)}$ represents the spatial representation at layer l, $W^{(l)}$ is the learnable weight matrix for layer l, and $\sigma(\cdot)$ denotes the non-linear activation function.

The design ensures the spatial aggregation is adaptive to the severity of the weather, effectively mitigating the "static graph limitation" discussed in existing literature.

3.4. Dual-Branch Stability Modelling (Speed Variance Module)

A fundamental challenge in traffic forecasting is the mean-reversion phenomenon, where predictive models tend to converge toward smoothed average speeds, consequently failing to resolve abrupt fluctuations caused by traffic instability. As established in Section 2.3, speed variance serves as a robust proxy for such volatility, encapsulating behaviours such as stop-and-go waves and cautious driving adjustments. To explicitly encode these microscopic dynamics alongside macroscopic trends, we implement a dual-branch architecture.

Structurally, the model processes inputs through two parallel branches to capture distinct traffic characteristics. The Main Branch focuses on trend modelling by taking the macroscopic traffic trend input, denoted as X_{trend}, and processing it through feature extraction layers to generate its hidden representation, H_{trend}. Simultaneously, the Stability Branch targets volatility by processing the speed variance input, X_{stab}, to extract the microscopic disturbance hidden features, denoted as H_{vol}. To fuse these branches, we introduce a stability-aware cross-attention (SACA) module. We utilize the stability hidden features (H_{vol}) as the Query to query the trend hidden features (H_{trend}).

$$Q = H_{vol}W_Q^{attn}, \quad K = H_{trend}W_K^{attn}, \quad V = H_{trend}W_V^{attn} \quad (3)$$

$$Z_{fused} = LayerNorm\left(H_{trend} + Attention(Q, K, V)\right), \quad (4)$$

where W_Q^{attn}, W_K^{attn}, and W_V^{attn} are the learnable projection matrices for the Query, Key, and Value, respectively.

By using variance features to reweight the speed features, the model essentially performs a risk adjustment on the speed prediction. When variance is high (indicating high instability/risk), the attention mechanism highlights these unstable periods, allowing the model to suppress standard historical patterns and adapt quickly to non-linear speed drops.

3.5. Rainfall-Gated Temporal Fusion (Rainfall Intensity Module)

Precipitation introduces significant non-stationarity into traffic dynamics, frequently compromising the validity of historical periodic patterns such as weekly or daily seasonality. As highlighted in Section 2.4, conventional strategies that rely on the direct concatenation of rainfall data are insufficient to capture the concept drift, as they lack the adaptive capacity to distinguish between routine periodicity and weather-induced disruptions.

To regulate information flow, we propose a rainfall gating mechanism that acts as a regulating valve. Its mechanism defines two distinct temporal context vectors: H_{curr}, which captures short-term features extracted from recent time steps (e.g., the last hour) to represent the immediate traffic state, and H_{hist}, which encodes historical periodic features (e.g., the same time last week or day) that typically serve as robust baselines under normal conditions. To mediate between these inputs, a gating signal, g_t, is generated directly from the current rainfall intensity, R_t.

$$g_t = \sigma\left(W_g R_t + b_g\right). \tag{5}$$

The final temporal representation O_t is a dynamically weighted combination:

$$O_t = (1 - g_t) \odot H_{hist} + g_t \odot H_{curr}\ , \tag{6}$$

where $(\cdot)$ denotes the sigmoid activation function, and $\odot$ represents the element-wise multiplication.

From a physical perspective, the gating action interprets the reliability of historical trends based on environmental conditions. In the absence of rain ($R \approx 0$), the gate directs the model to prioritize H_{hist}, effectively leveraging stable daily and weekly periodic patterns similar to standard periodic-based models. Conversely, during heavy rainfall ($R \geq 0$), the gate activates to suppress these historical patterns—which may become invalid due to weather-induced disruptions—and forces the model to rely heavily on H_{curr} to capture immediate traffic deviations.

4. Experiments

Seven comprehensive experiments are presented to evaluate the performance and robustness of the proposed modelling framework, specifically highlighting the three core contributions outlined in Section 1. Relevant comparisons with state-of-the-art benchmarks and ablation studies are carried out to verify the impacts of different factors on prediction accuracy.

First, to validate the effectiveness of Contribution 1 (Constructing a rainfall-adaptive dynamic graph topology) in capturing shifting spatial dependencies and reconfiguring network connectivity, the study employs the first and second experiments (E1 and E2) to analyze the model's spatial performance. Specifically, E1 establishes the framework's predictive performance across diverse spatial locations by demonstrating its advantage in complex and dynamic environments compared to static graph baselines, while E2 assesses the model's reliability in minimizing error rates and managing volatility via the Cumulative Distribution Function (CDF).

Subsequently, the capability of Contribution 2 (Modelling rainfall-induced instability via a dual-branch architecture) to decouple macroscopic trends from microscopic variances is rigorously evaluated through the fourth and fifth experiments (E4 and E5). In this phase, E4 conducts an ablation study to explicitly confirm the structural necessity of the stability module and rainfall intensity inputs, whereas E5 disaggregates performance across varying rainfall intensities to demonstrate the model's robustness during extreme weather events.

Finally, the efficacy of Contribution 3 (Dynamically calibrating temporal dependencies with a gating mechanism) is substantiated by the third, sixth, and seventh experiments (E3, E6, and E7), which examine the model's adaptability to temporal shifts. E3 visually verifies the model's proficiency in tracking dynamic trends and responding to sudden fluctuations; E6 tests temporal robustness over different prediction horizons; and concluding the analysis, E7 identifies the optimal historical look-back window size required to capture the evolution of rainfall-induced shifts effectively.

4.1. Experimental Setup and Data Description

The experiments were conducted on a representative urban arterial intersection in Hong Kong, specifically the junction of Hong Ning Road and Mut Wah Street (Figure 2; Image source: Google Maps). The path follows the westbound traffic flow along Mut Wah Street towards the signalized intersection at Hong Ning Road. The segment represents a critical component of the urban network within a high-density mixed-use district.

Traffic conditions on this path are monitored by roadside sensors equipped with automatic incident detection (AID) capabilities, as illustrated in Figure 2. These sensors provide real-time traffic data,

Figure 2
Study map.

which are aggregated and retrieved from an open data platform developed by the Hong Kong government. As the government officially maintains the traffic speed and volume data from the IT platform, they serve as the ground truth for training the proposed model and evaluating its performance against benchmark methods.

Regarding meteorological data, a weather station is located in proximity to the road segment (indicated by the star icon in Figure 2). The station captures localized weather conditions, including rainfall intensity, which are essential for analyzing the impact of weather on urban traffic flow. The meteorological data are sourced from the Hong Kong government and synchronized with the timestamps of the traffic data. For this study, historical traffic and weather data spanning weekdays from September to November 2024, excluding public holidays, are utilized to ensure the consistency of recurrent traffic patterns.

To further investigate the impact of weather forecasting on prediction accuracy and address the potential unavailability of historical forecast archives, weather forecast scenarios are generated based on observed data from the local weather station. It allows for a comprehensive sensitivity analysis of the model under varying degrees of weather forecast accuracy.

The weekdays in September 2024 are used as training sets, while five specific weekdays in November 2024 are selected for validation of the real-time path travel speed prediction. The mean absolute percentage error (MAPE), mean absolute error (MAE), and root mean squared error (RMSE) are used to evaluate the average performance. The results on MAPE, MAE and RMSE are obtained using:

$$MAPE = \frac{1}{\Delta t}\sum_{i=1}^{\Delta t}\left|\frac{\hat{T}_{t_0+i,p} - T_{t_0+i,p,s_g}}{T_{t_0+i,p,s_g}}\right| \times 100\% \tag{7}$$

$$MAE = \frac{1}{\Delta t}\sum_{i=1}^{\Delta t}\left|\hat{T}_{t_0+i,p} - T_{t_0+i,p,s_g}\right| \tag{8}$$

$$RMSE = \sqrt{\frac{1}{\Delta t}\sum_{i=1}^{\Delta t}\left(\hat{T}_{t_0+i,p} - T_{t_0+i,p,s_g}\right)^2}. \tag{9}$$

To rigorously evaluate the model's robustness and reliability under worst-case scenarios, we introduce the effective maximum absolute percentage error (Effective MaxAPE) and effective maximum absolute error (Effective MaxAE). Unlike standard global maximum metrics, which are highly sensitive to sensor anomalies and instantaneous noise (outliers), the Effective metrics represent the upper bound of prediction errors within the 95% confidence interval. This approach filters out the top 5% of extreme values caused by non-recurrent disturbances, providing a more objective assessment of the model's stability in realistic traffic conditions. These metrics are defined as:

$$Effective\ MaxAE = Q_{0.95}(|y_i - \hat{y}_i|) \quad (10)$$

$$Effective\ MaxAPE = Q_{0.95}\left(\left|\frac{y_i - \hat{y}_i}{y_i}\right|\right) \times 100\%\,, \quad (11)$$

where y_i and $\hat{y}_i$ represent the ground truth and predicted speed values, respectively, and $Q_{0.95}(\cdot)$ denotes the 95th percentile function of the error distribution.

The proposed RM-DSTGN model was implemented using the PyTorch framework. During the training phase, the model parameters were optimized using the Adam optimizer. To achieve optimal convergence, key hyperparameters were dynamically tuned via heuristic search. The search space for the learning rate was set between $[1\times10^{-4}, 5\times10^{-3}]$, and the batch size was selected from {16, 32, 64}, with the optimal batch size ultimately set to 32. The hidden dimensionality for the node embeddings and attention layers was set to 64. The training process was capped at 200 epochs, incorporating an early stopping mechanism with a patience of 15 epochs on the validation set to prevent overfitting.

4.2. Performance of the Proposed Modelling Framework

In the first experiment (E1), the overall prediction performance of the proposed RM-DSTGN is compared with that of five state-of-the-art benchmark models (ATLSTM, GCN, LSTM, SVR, and KNN) to investigate the goodness of fit across different spatial locations. The performance metrics, including MAE, RMSE, MAPE, and Maximum Error, for three distinct sensors are presented and evaluated in Table 1.

It is observed that RM-DSTGN achieves superior performance, particularly in complex and dynamic traffic environments. For instance, on Sensor 1, the proposed framework yields the lowest error rates ($MAPE = 3.62\%$, $MAE = 1.52$, $and\ RMSE = 2.11$), whereas the traditional KNN model exhibits the worst performance ($MAPE = 9.91\%$, $MAE = 4.01$, $and\ RMSE = 5.00$). Notably, in this challenging scenario, the proposed framework significantly boosts prediction accuracy compared to the second-best model (ATLSTM), reducing the RMSE by approximately 7.0%.

When comparing the proposed RM-DSTGN with the second-best deep learning baseline (ATLSTM), our model demonstrates significant superiority under fluctuating conditions. On Sensor 3, the RM-DSTGN reduces the MAPE from 5.66% (ATLSTM) to 4.73%, which represents an accuracy improvement of approximately 16.4% $\left(\frac{5.66\% - 4.73\%}{5.66\%}\right)$. When compared to the worst-performing baseline (KNN) on Sensor 1, the accuracy improvement is as significant as 63.5% $\left(\frac{5.66\% - 4.73\%}{5.66\%}\right)$. It unequivocally demonstrates that the proposed framework, by incorporating rainfall modulation and dynamic topology learning, significantly boosts prediction accuracy compared to both traditional machine learning and standard deep learning approaches in these challenging, high-volatility scenarios.

However, prediction accuracy naturally varies across locations depending on local traffic dynamics. As shown in Table 1, the baseline ATLSTM actually achieves the best performance on Sensor 2 ($MAE = 0.98$, $RMSE = 1.25$, $MAPE = 3.26\%$), slightly outperforming RM-DSTGN. Because Sensor 2 experiences the most stable traffic flow, this result suggests that simpler deep learning architectures are highly effective when weather-induced volatility is absent. While our proposed model remains competitive on this stable segment, its primary advantage emerges in more complex environments. On Sensor 3, a noticeably more challenging location where baseline errors typically rise (e.g., LSTM $MAPE = 8.41\%$), RM-DSTGN effectively holds the error rate down to 4.73%, and it strictly dominates all baselines on the highly dynamic Sensor 1.

Beyond standard accuracy metrics, evaluating a model's reliability under worst-case scenarios is paramount for safety-oriented traffic management systems. To rigorously assess it, we employed the Effective Maximum Absolute Error (Effective MaxAE) and Effective Maximum Absolute Percentage

Error (Effective MaxAPE), as presented in Table 2, which measure performance within the 95% confidence interval to mitigate the influence of extreme sensor noise. In its robustness analysis, the proposed RM-DSTGN demonstrates superior stability, achieving the lowest upper bounds of error among all compared baselines. Specifically, the model recorded an Effective MaxAE of 3.95 km/h and an Effective MaxAPE of 11.22%, indicating that in 95% of the test instances, prediction errors were strictly confined within these minimal limits. By contrast, the second-best model, ATLSTM, exhibited slightly higher instability with values of 4.13 km/h and 12.59%, respectively, while the traditional KNN baseline displayed a significantly wider error margin of 9.26 km/h and 28.72%. These comparative results unequivocally confirm that RM-DSTGN is not only accurate on average but also exceptionally robust, demonstrating a distinct capability to suppress large prediction deviations arising from complex and dynamic traffic conditions.

In the second experiment (E2), the probabilistic distribution of prediction errors is investigated to evaluate the reliability of the proposed framework. The cumulative distribution function (CDF) of the absolute percentage error (APE) is utilized for this purpose, as it provides a comprehensive view of how errors are distributed across different magnitudes. Figure 3 depicts the comparative CDF curves for the Sensor 1. It is unequivocally observed that the RM-DSTGN model exhibits superior performance, as its curve rises the most rapidly among all compared models. Specifically, a significantly larger proportion (94.7%) of predictions produced by RM-DSTGN fall within the low-error range ($APE < 10\%$). Comparatively, the baseline models, particularly KNN and SVR, exhibit significantly slower convergence rates, with probabilities of only 56.2% and 67.8%, respectively, at the same threshold, indicating a lower likelihood of achieving accurate predictions. Furthermore, RM-DSTGN outperforms deep learning baselines such as ATLSTM (90.7%), GCN (85.7%), and LSTM (76.8%), demonstrating a notably shorter tail in its error distribution. It signifies that the proposed framework effectively suppresses extreme outliers and large deviations, ensuring consistent and reliable service levels even under challenging traffic conditions. The ability to maintain such a compact error distribution validates the effectiveness of the proposed spatiotemporal learning mechanism in mitigating the uncertainty of real-time traffic speeds.

Table 2

Evaluation of Model Robustness based on Effective Maximum Errors.

Model	Effective MaxAPE (%)	Effective MaxAE (km/h)
RM-DSTGN	11.22	3.95
ATLSTM	12.59	4.13
GCN	14.65	4.88
LSTM	18.56	5.94
SVR	22.76	7.19
KNN	28.72	9.26

Table 1

Comparison of Overall Prediction Performance Across Different Sensors.

Models	Sensor 1			Sensor 2			Sensor 3		
	MAE (km/h)	RMSE (km/h)	MAPE (%)	MAE (km/h)	RMSE (km/h)	MAPE (%)	MAE (km/h)	RMSE (km/h)	MAPE (%)
RM-DSTGN	1.52	2.11	3.62%	1.03	1.57	3.38%	1.42	2.04	4.73%
ATLSTM	1.78	2.27	4.34%	0.98	1.25	3.26%	1.72	2.22	5.66%
GCN	2.09	2.65	5.11%	1.13	1.43	3.75%	2.00	2.55	6.61%
LSTM	2.55	3.22	6.26%	1.338	1.72	4.49%	2.52	3.18	8.41%
SVR	3.15	3.95	7.77%	1.60	2.01	5.34%	3.03	3.83	10.21%
KNN	4.01	5.00	9.91%	2.06	2.58	6.85%	3.92	4.92	13.19%

In the third experiment (E3), the temporal fitting capability of the proposed RM-DSTGN framework was visually inspected to verify its proficiency in tracking dynamic traffic trends and responding to sudden fluctuations. Regarding different rainfall categories, the category is based on historical rainfall intensity data, and the boundaries for distinguishing light rain, moderate rain and heavy rain are 0.8 and 6.5 mm/h, which is similar to the boundaries set by the previous studies on the effects of rainfall on speed and capacity reduction [11]. Time-series charts, as shown in Figure 4, compare the ground truth traffic speeds with predictions from RM-DSTGN and benchmark models across various dates and sensor locations.

First, the analysis focuses on the model's performance under heavy rainfall conditions observed on November 20, 2024. As illustrated in the results for the Fast Lane of Sensor 1 and the Slow Lane of Sensor 3, traffic speeds experienced precipitous drops due to intense precipitation. The proposed RM-DSTGN (represented by the red curve) demonstrated superior precision, maintaining a low MAPE of 2.02% and 2.05% on the Fast Lane of Sensor 2 and Slow Lane of Sensor 1, respectively, as detailed in Table 3 . This contrasts sharply with baseline models such as KNN and SVR, which exhibited significant lag effects and higher error rates, with MAPEs reaching up to 6.97% and 6.81% during the same interval. These quantitative results effectively prove the model's ability to incorporate real-time rainfall intensity to mitigate prediction delays during extreme weather events.

Second, the model's capability to handle high-frequency fluctuations was evaluated under moderate rainfall on November 19, 2024. On that day, traffic patterns across multiple sensors, specifically the Fast Lane of Sensor 1, the Fast Lane of Sensor 2, and the Slow Lane of Sensor 3, exhibited complex volatility. The RM-DSTGN framework accurately captured these oscillating trends, effectively filtering out noise while retaining the underlying traffic dynamics. The close alignment between the predicted and actual speeds across these diverse locations highlights the model's spatial robustness and adaptability to varying traffic flow characteristics.

Finally, the consistent stability of the proposed framework was further verified through cases on November 13 and November 15, 2024. Whether observing the Fast Lane of Sensor 2 and the Slow Lane of Sensor 1 on November 15, or the Slow Lane of Sensor 1 on November 13, the RM-DSTGN consistently maintained a tight fit with the ground truth throughout the entire day. Comprehensive visual evidence

Figure 3

Comparative CDF curves for the Slow Lane of Sensor 1.

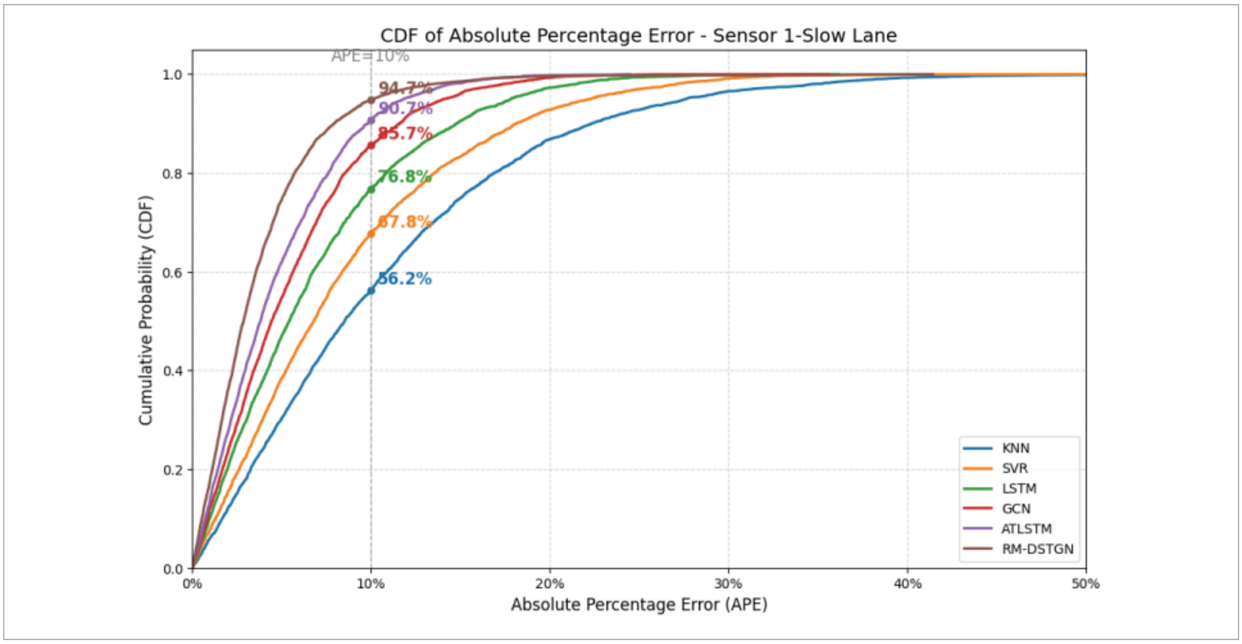

Figure 4

Visual comparison of prediction performance under various weather and traffic conditions.

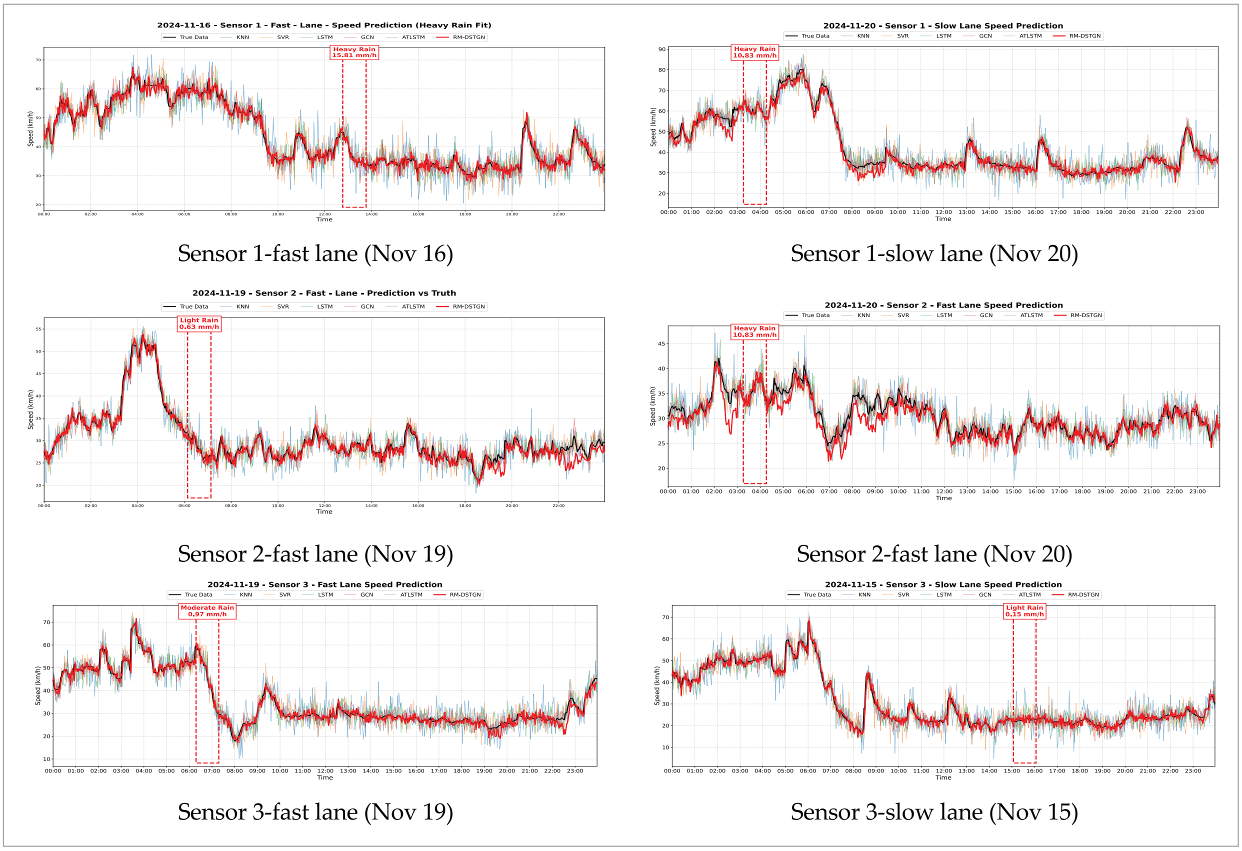

Sensor 1-fast lane (Nov 16)

Sensor 1-slow lane (Nov 20)

Sensor 2-fast lane (Nov 19)

Sensor 2-fast lane (Nov 20)

Sensor 3-fast lane (Nov 19)

Sensor 3-slow lane (Nov 15)

Table 3

Stratified Performance Analysis of RM-DSTGN Across Rainfall Intensity Levels.

Model	Sensor 1-Fast Lane (2024/11/16 12:40-13:50)		Sensor 1 - Slow Lane (2024/11/20 03:15-04:20)		Sensor 2 - Fast Lane (2024/11/20 03:15-04:20)	
	MAPE (%)	RMSE (km/h)	MAPE (%)	RMSE (km/h)	MAPE (%)	RMSE (km/h)
RM-DSTGN	2.43	1.07	2.05	1.44	2.02	1.19
KNN	9.37	4.42	7.17	5.14	6.97	4.74
LSTM	6.47	3.03	4.53	3.52	6.02	3.39
SVR	7.20	3.42	6.12	4.36	6.81	4.02
GCN	4.10	1.97	3.79	2.83	4.94	3.11
ATLSTM	4.33	1.87	3.05	2.36	3.01	1.73

confirms that the proposed framework serves as a reliable tool for real-time traffic prediction, regardless of the specific lane type (fast/slow) or varying degrees of rainfall intensity.

The fourth experiment (E4) serves as an ablation study to systematically validate the contribution of specific model components within the proposed framework. By comparing the absolute percentage error (APE) distribution of the full RM-DSTGN model against three simplified variants that exclude spatial dependencies, speed variance, and rainfall intensity, the experiment confirms the significance of utilizing multi-source data and specific spatial modules for enhancing prediction robustness.

Figure 5 presents the box plots of the APE distribution for Sensor 1-Fast Lane. It is unequivocally observed that the complete RM-DSTGN exhibits superior performance, characterized by the lowest median error and the most compact interquartile range (IQR). It indicates that the whole model offers the highest stability and precision.

In contrast, removing specific components leads to noticeable performance degradation. The w/o Spatial variant (2nd column) exhibits the most extensive spread in error distribution and the highest upper quartile, indicating that ignoring network-wide spatial correlations hinders the model's ability to capture the propagation of traffic congestion, thereby significantly reducing prediction reliability.

Similarly, the w/o Speed Variance variant (3rd column) exhibits a markedly expanded error range, with long whiskers extending to higher error values, compared to the full model. It implies that explicitly modelling speed variance is crucial for understanding traffic flow volatility, especially under unstable traffic conditions.

Furthermore, the w/o Rainfall Intensity variant (4th column) results in a distinct upward shift in the median error and a wider distribution compared to the RM-DSTGN. It confirms that the integration of real-time rainfall data is essential for accurate predictions, as it allows the model to account for the structural impact of adverse weather on driving behaviours and road capacity.

In summary, the ablation study confirms that each component of the RM-DSTGN framework plays a crucial role in ensuring robust and accurate traffic speed prediction.

The fifth experiment (E5) investigates the impact of varying rainfall intensities on the prediction accuracy of RM-DSTGN. Following established studies on speed and capacity reduction, rainfall is categorized into Light, Moderate, and Heavy rain, with boundaries set at 0.8 and 6.5 mm/h. While previous experiments focused on overall spatiotemporal robustness, disaggregating performance by meteorological conditions is crucial for assessing reliability during extreme weather. The results are summarized in Table 4, where the 'Sample percent' column denotes the proportion of time steps for

Figure 5

Robustness Analysis (APE Distribution) - Sensor 1 Fast Lane.

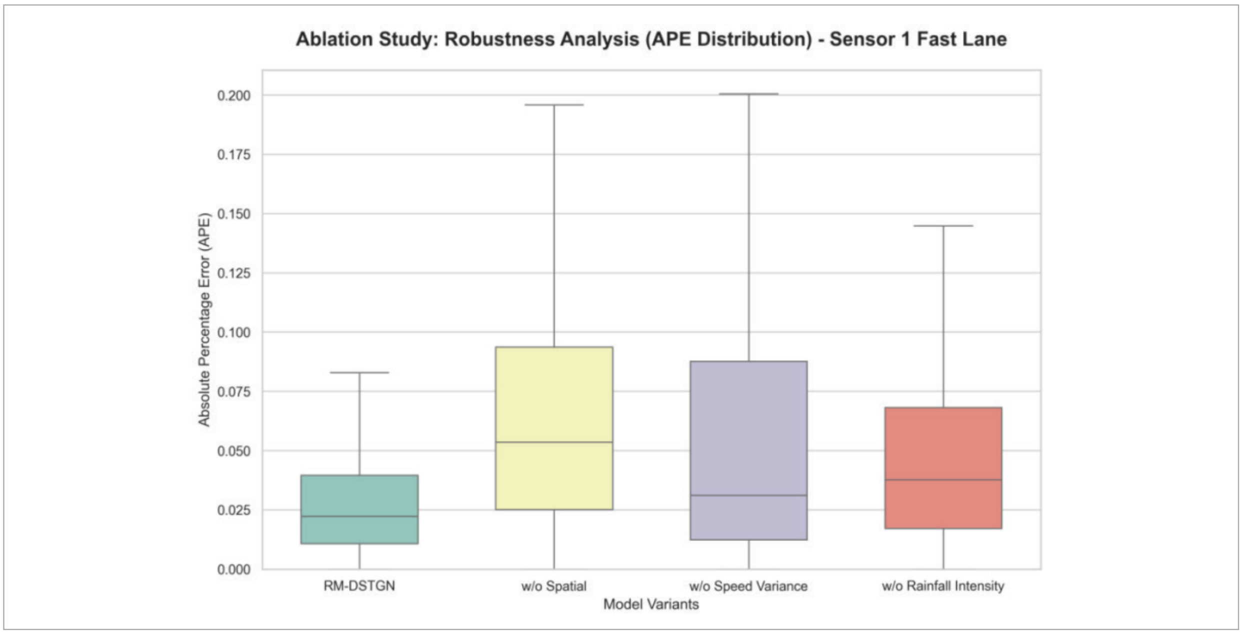

each rainfall category relative to the entire dataset, reflecting the scarcity of high-intensity precipitation events.

It is observed that prediction accuracy correlates with rainfall intensity. Under Light Rain conditions (84.31% of samples), the model maintains high precision with a MAPE of 3.35%, comparable to clear weather performance (E1), which confirms the effective filtering of minor environmental noise. In the Moderate Rain category, errors increase slightly but remain acceptable ($MAPE$ = 6.60%, MAE = 2.22 km/h, $RMSE$ = 2.83 km/h), indicating that the framework successfully captures weather-induced speed reductions without overreacting.

Crucially, RM-DSTGN exhibits significant robustness under Heavy Rain conditions, despite the challenges of data sparsity (1.06% of samples) and chaotic flow. The model successfully constrains the MAE to 7.67 km/h and RMSE to 8.16 km/h. While MAPE naturally rises to 21.27% due to lower traffic speeds, limiting the absolute error to approximately 7–8 km/h demonstrates high practical value. Moreover, the close alignment between RMSE (8.16 km/h) and MAE (7.67 km/h) suggests a lack of extreme outliers, confirming the model actively responds to rare weather events rather than reverting to historical means.

In the sixth experiment (E6), the impact of different prediction horizons (Δt) on the accuracy of the proposed RM-DSTGN model is evaluated. The prediction horizon serves as a critical parameter in real-time traffic management, as authorities require reliable forecasts not only for the immediate future but also for medium-term operational planning. To assess the model's temporal robustness, its performance is tested across five distinct time intervals: 2 minutes, 10 minutes, 20 minutes, 30 minutes, and 60 minutes. The comparative results against benchmark models are summarised in Table 5.

It is generally observed that the prediction error for all models exhibits a monotonically increasing trend as the prediction horizon extends. This phenomenon is attributed to the inherent stochasticity of traffic flow, where uncertainty accumulates over longer time intervals. However, the proposed RM-DSTGN demonstrates superior stability, maintaining a MAPE advantage of 0.83% to 1.76% over the best baseline (ATLSTM) as the horizon extends from 2 to 60 minutes.

Specifically, for short-term prediction (Δt = 2 min), the RM-DSTGN achieves an impressive MAPE of 3.82% and an MAE of 1.32, significantly surpassing the second-best model, ATLSTM ($MAPE$ = 4.65%). It indicates that the model effectively captures the immediate spatiotemporal dependencies and rainfall-induced fluctuations.

As the prediction horizon extends to 30 minutes, while the performance of traditional models, such as KNN and SVR, deteriorates rapidly (MAPE rising to 10.70% and 11.60%, respectively), the RM-DSTGN maintains robust performance with a MAPE of 9.34%. Even at the most extended horizon of 60 minutes, the proposed framework maintains an MAPE of 10.06%, whereas the ATLSTM and GCN models exhibit larger errors of 11.82% and 13.11%, respectively.

The widening performance gap (increasing from a 0.83% MAPE difference at 2 minutes to 1.76% at 60 minutes) suggests that the proposed mechanisms—specifically the dynamic topology learning and rainfall-gated fusion—are particularly effective in mitigating long-term error propagation. By explicitly modelling the evolving impact of rainfall and traffic stability over time, RM-DSTGN provides more reliable support for proactive traffic control strategies compared to state-of-the-art deep learning counterparts.

In the seventh experiment (E7), the impact of the historical look-back window size (T) on the predic-

Table 4

Performance of RM-DSTGN under different rainfall categories.

Rainfall Category	MAE (km/h)	RMSE (km/h)	MAPE (%)	Sample Percent
Light Rain	1.16	1.57	3.35	84.31%
Moderate Rain	2.22	2.83	6.60	14.63%
Heavy Rain	7.67	8.16	21.27	1.06%

Table 5
Comparison of prediction performance under different prediction horizons (Δt).

Δt (intervals)	2min			10min			30min			60min		
Model	MAE (km/h)	RMSE (km/h)	MAPE (%)	MAE (km/h)	RMSE (km/h)	MAPE (%)	MAE (km/h)	RMSE (km/h)	MAPE (%)	MAE (km/h)	RMSE (km/h)	MAPE (%)
ATLSTM	1.59	2.08	4.65	1.80	2.70	4.82	3.23	4.63	8.90	4.09	5.60	11.82
GCN	1.86	2.41	5.44	3.78	5.17	10.63	4.14	5.64	11.90	4.57	6.35	13.11
KNN	3.58	4.58	10.61	2.49	3.63	6.86	3.81	5.42	10.70	4.55	6.42	13.06
LSTM	2.30	2.96	6.77	1.81	2.66	5.00	3.13	4.49	8.66	4.06	5.57	11.53
SVR	2.79	3.60	8.26	2.96	3.65	8.84	4.03	5.30	11.60	4.63	6.202	13.39
RM-DSTGN	1.32	1.89	3.82	1.58	2.28	4.30	3.38	4.70	9.34	3.60	5.06	10.06

tion accuracy of the proposed RM-DSTGN model is systematically investigated. The length of the input sequence is a critical hyperparameter that determines the amount of temporal context available for the model to learn the evolution of traffic states. To identify the optimal temporal receptive field, four distinct window sizes are evaluated: 5 steps (10 minutes), 10 steps (20 minutes), 20 steps (40 minutes), and 30 steps (60 minutes). The quantitative results, including MAE, RMSE, and MAPE, are summarised in Table 6.

It is observed that the prediction accuracy generally improves as the look-back window extends. Specifically, the MAPE decreases monotonically from 6.95% at a 10-minute window (5 steps) to 6.52% at a 60-minute window (30 steps). Similarly, the MAE exhibits a downward trend, reducing from 2.47 to 2.41. The improvement indicates that a longer historical context enables the RM-DSTGN to capture the temporal evolution of traffic trends and the delayed effects of rainfall accumulation more effectively. Unlike short-term fluctuations, the structural shifts in traffic propagation induced by adverse weather often evolve over more extended periods; thus, providing a 60-minute context allows the gating mechanism to differentiate between routine periodicity and weather-induced disruptions more effectively.

It is noted that the RMSE remains relatively stable, fluctuating marginally between 3.48 and 3.50 across the different window sizes. It suggests that while the overall relative accuracy (MAPE) benefits significantly from longer context, the model maintains consistent robustness against large variance outliers regardless of the window size. However, a critical trade-off between accuracy and computational efficiency must be considered. Although extending the look-back window beyond 60 minutes could theoretically yield further marginal reductions in MAPE and improvements in accuracy, it inevitably leads to a substantial increase in computational complexity and training time. The diminishing returns in performance do not justify the elevated resource

Table 6
Comparison of Prediction Accuracy under Different Look-back Windows.

Window size steps	Window size mins	MAE (km/h)	RMSE (km/h)	MAPE (%)
5	10	2.47	3.48	6.95
10	20	2.48	3.58	6.76
20	40	2.42	3.50	6.55
30	60	2.41	3.50	6.52

consumption and inference latency. Consequently, a look-back window of 30 steps (60 minutes) is adopted as the optimal setting for the proposed framework, as it achieves the best balance between maximizing prediction precision and maintaining operational efficiency.

5. Conclusion

A novel rainfall-modulated deep learning framework, RM-DSTGN, is proposed for real-time traffic speed prediction under adverse weather conditions. Traditionally, spatiotemporal models have relied on static graph topologies and incorporated weather data only as supplementary input features, ignoring the dynamic shifts in network correlations induced by environmental factors. However, rainfall fundamentally alters traffic propagation patterns and driver behaviours. To address it, the proposed framework integrates a dynamic topology learning mechanism driven by rainfall intensity, a dual-branch architecture for modelling speed stability, and an adaptive temporal gating mechanism. The proposed modelling framework is validated in comprehensive numerical experiments with data collected from a high-density mixed-use district in Hong Kong.

First, the experimental results demonstrate that the RM-DSTGN significantly outperforms state-of-the-art benchmarks across diverse spatial locations, achieving a minimum MAPE of 3.62% on Sensor 1 (E1). While baseline models perform adequately under highly stable traffic conditions, the proposed framework shows its distinct superiority when traffic is subjected to structural disruptions. Additionally, ablation studies confirm that every component of the framework is necessary; removing the rainfall module significantly degrades performance across all tested sensors (e.g., doubling the median error rate), proving that explicitly modelling the interaction between environmental conditions and traffic flow is essential for robust prediction (E4).

Furthermore, the model exhibits high reliability and stability. Probabilistic error analysis reveals that the framework converges more rapidly to low-error ranges and effectively minimizes extreme outliers compared to baseline models (E2). Even during heavy rainfall, the model demonstrates exceptional sensitivity by tracking sharp speed drops without significant lag (E3). While prediction errors naturally rise under rare, heavy rain conditions due to data sparsity (E5), the model maintains superior temporal robustness, performing well even at a 60-minute prediction horizon, where other models fail to prevent error propagation (E6).

Finally, the study establishes that a 60-minute historical look-back window yields the optimal performance (E7). Its longer historical context allows the model to learn better the gradual structural shifts and delayed effects caused by rainfall accumulation.

A current limitation of this study is that the empirical validation is based on single-corridor data from a specific arterial intersection in Hong Kong. While the framework excels at capturing specific rainfall-induced shifts, its generalizability across diverse urban network topologies (e.g., highly interconnected grid networks or freeways) and varying regional climates remains an open question. Because the dynamic topology generation relies on localized spatial embeddings, deploying this model in a new city would require re-calibrating the static node embeddings (E_{node}) and rainfall thresholds. Future research will investigate domain adaptation and transfer learning techniques to map learned spatiotemporal patterns from data-rich corridors to newly monitored networks, minimizing the need for extensive retraining. Additionally, we aim to explore the integration of other dynamic constraints, such as traffic incidents, to further refine its robustness under multi-hazard scenarios.

Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, author-ship, and/or publication of this article.

Data Sharing Agreement

The datasets used and/or analyzed during the current study are available from the corresponding author on reasonable request.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

References

1. Abduljabbar, R., Dia, H., Liyanage, S. Machine Learning Models for Traffic Prediction on Arterial Roads Using Traffic Features and Weather Information. Applied Sciences, 2024, 14(23), 11047. https://doi.org/10.3390/app142311047
2. Alnami, H., Mahgoub, I., Al-Najada, H., Alalwany, E. A Distributed Machine Learning-Based Scheme for Real-Time Highway Traffic Flow Prediction in Internet of Vehicles. Future Internet, 2025, 17(3). https://doi.org/10.3390/fi17030131
3. Awan, A. A., Majid, A., Riaz, R., Rizvi, S. S., Kwon, S. J. A Novel Deep Stacking-Based Ensemble Approach for Short-Term Traffic Speed Prediction. IEEE Access, 2024, 12, 15222-1523. https://doi.org/10.1109/ACCESS.2024.3357749
4. Bilgin, E., Türkmen, H. I., Güvensan, M. A. A Weather Oriented Pre-Tuning Methodology for Long-Term Traffic Speed Estimation. In 2023 IEEE 24th International Symposium on a World of Wireless, Mobile and Multimedia Networks (WoWMoM), Boston, MA, USA, June 12, 2023, 451-456. https://doi.org/10.1109/WoWMoM57956.2023.00079
5. Chang, A., Ji, Y., Bie, Y. Transformer-Based Short-Term Traffic Forecasting Model Considering Traffic Spatiotemporal Correlation. Frontiers in Neurorobotics, 2025, 19, 1527908. https://doi.org/10.3389/fnbot.2025.1527908
6. Chen, H., Han, H., Chen, Y., Chen, Z., Gao, R., Li, X. A Traffic Speed Prediction Algorithm for Dynamic Spatio-Temporal Graph Convolutional Networks Based on Attention Mechanism. The Journal of Supercomputing, 2025, 81(1), 18. https://doi.org/10.1007/s11227-024-06539-2
7. Chen, J., Yang, L., Yang, Y., Peng, L., Ge, X. Spatio-Temporal Graph Neural Networks for Missing Data Completion in Traffic Prediction. International Journal of Geographical Information Science, 2025, 39(5), 1057-1075. https://doi.org/10.1080/13658816.2024.2381221
8. Diao, Z. L., Wang, X., Zhang, D. F., Xie, G. G., Chen, J. G., Pei, C. H., Meng, X. Y., Xie, K., Zhang, G. X. DMSTG: Dynamic Multiview Spatio-Temporal Networks for Traffic Forecasting. IEEE Transactions on Mobile Computing, 2024, 23(6), 6865-6880. https://doi.org/10.1109/TMC.2023.3328038
9. Fang, J., Wu, Z., Xu, M., Chen, H. Multistep Traffic Speed Prediction from Multiple Time-Scale Spatiotemporal Features Using Graph Attention Network. Applied Intelligence, 2024, 54(15), 7479-7492. https://doi.org/10.1007/s10489-024-05503-0
10. Hassan, M. M., Rony, M. A. T., Khan, M. A. R., Hassan, M. M., Yasmin, F., Nag, A., Zarin, T. H., Bairagi, A. K., Alshathri, S., El-Shafai, W. Machine Learning-Based Rainfall Prediction: Unveiling Insights and Forecasting for Improved Preparedness. IEEE Access, 2023, 11, 132196-132222. https://doi.org/10.1109/ACCESS.2023.3333876
11. Li, A., Lam, W. H., Ma, W., Tam, M. L. A Novel Modelling Framework for Real-Time Prediction of Path Travel Times Using Both Traffic and Weather Data. Transportmetrica A: Transport Science, 2024, 1-34. https://doi.org/10.1080/23249935.2024.2419499
12. Li, L., An, B., Zhang, D., Gan, R., Zhou, Z., Qu, X., Ran, B. Vehicular Speed Prediction Method for Highway Scenarios Based on Spatio-Temporal Graph Convolutional Networks and Potential Field Theory. IEEE Internet of Things Journal, 2025, 12(3), 3330-3349. https://doi.org/10.1109/JIOT.2024.3479221
13. Min, H., Xiong, X., Wang, P., Zhang, Z. A Hierarchical LSTM-Based Vehicle Trajectory Prediction Method Considering Interaction Information. Automotive Innovation, 2024, 7(1), 71-81. https://doi.org/10.1007/s42154-023-00261-0
14. Park, H. S., Park, Y. W., Kwon, O. H., Park, S. H. Applying Clustered KNN Algorithm for Short-Term Travel Speed Prediction and Reduced Speed Detection on Urban Arterial Road Work Zones. Journal of Advanced Transportation, 2022, 2022(1), 1107048. https://doi.org/10.1155/2022/1107048
15. Rehmer, A., Kroll, A. On the Vanishing and Exploding Gradient Problem in Gated Recurrent Units. IFAC-PapersOnLine, 2020, 53(2), 1243-1248. https://doi.org/10.1016/j.ifacol.2020.12.1342
16. Shao, F., Shao, H., Wu, X., Cheng, Q., Lam, W. H. A Physics-Informed Machine Learning Framework for Speed-Flow Prediction: Integrating an S-Shaped Traffic Stream Model with Deep Learning Models. Transportation Research Part C: Emerging Technologies, 2025, 180, 105362. https://doi.org/10.2139/ssrn.5164944

17. Tian, X., Du, L., Zhang, X., Wu, S. MAT-WGCN: Traffic Speed Prediction Using Multi-Head Attention Mechanism and Weighted Adjacency Matrix. Sustainability, 2023, 15(17), 13080. https://doi.org/10.3390/su151713080

18. Wang, C., Huang, S., Zhang, C. Short-Term Traffic Flow Prediction Considering Weather Factors Based on Optimized Deep Learning Neural Networks: Bo-GRA-CNN-BiLSTM. Sustainability, 2025, 17(6), 2576. https://doi.org/10.3390/su17062576

19. Wang, R., Xin, Y., Zhang, Y., Pérez-Cruz, F., Raubal, M. Counterfactual Explanations for Deep Learning-Based Traffic Forecasting. Communications in Transportation Research, 2025, 5, 100176. https://doi.org/10.2139/ssrn.4873045

20. Wang, X., Zeng, R., Zou, F., Liao, L., Huang, F. STTF: An Efficient Transformer Model for Traffic Congestion Prediction. International Journal of Computational Intelligence Systems, 2023, 16(1), 2. https://doi.org/10.1007/s44196-022-00177-3

21. Wang, Y., Yu, C., Hou, J., Chu, S., Zhang, Y., Zhu, Y. ARIMA Model and Few-Shot Learning for Vehicle Speed Time Series Analysis and Prediction. Computational Intelligence and Neuroscience, 2022, 2022(1), 2526821. https://doi.org/10.1155/2022/2526821

22. Yang, W., Wang, Z., Cui, T., Lu, Y., Qu, Z. Traffic Speed Prediction Based on Spatial-Temporal Dynamic and Static Graph Convolutional Recurrent Network. International Journal of Advanced Computer Science & Applications, 2024, 15(12). https://doi.org/10.14569/IJACSA.2024.0151254

23. Xiao, H., Zou, B., Xiao, J., Ku, X., Yuan, L. Utilizing Multihead Attention-Based Graph Convolution Networks for Traffic Speed Prediction. International Journal of Intelligent Systems, 2025, 2025(1), 3471620. https://doi.org/10.1155/int/3471620

24. Yu, L., Luo, Q., He, J. Optimization of Speed Limit Method for Freeways Considering AVs Under Rainy Weather. Journal of Advanced Transportation, 2025, 2025(1), 1720155. https://doi.org/10.1155/atr/1720155