



PDF Download  
3711896.3736904.pdf  
30 December 2025  
Total Citations: 1  
Total Downloads: 1091

Latest updates: <https://dl.acm.org/doi/10.1145/3711896.3736904>

RESEARCH-ARTICLE

## Decoupling Spatio-Temporal Prediction: When Lightweight Large Models Meet Adaptive Hypergraphs

JIAWEN CHEN, Southeast University, Nanjing, Jiangsu, China

QI SHAO, Southeast University, Nanjing, Jiangsu, China

DUXIN CHEN, Southeast University, Nanjing, Jiangsu, China

WENWU YU, Southeast University, Nanjing, Jiangsu, China

Open Access Support provided by:

Southeast University

Published: 03 August 2025

[Citation in BibTeX format](#)

KDD '25: The 31st ACM SIGKDD  
Conference on Knowledge Discovery and  
Data Mining  
August 3 - 7, 2025  
Toronto ON, Canada

Conference Sponsors:  
SIGMOD  
SIGKDD

# Decoupling Spatio-Temporal Prediction: When Lightweight Large Models Meet Adaptive Hypergraphs

Jiawen Chen\*

jiawenchen@seu.edu.cn

School of Mathematics, Southeast University  
Nanjing, Jiangsu, China

Duxin Chen<sup>†</sup>

chendx@seu.edu.cn

School of Mathematics, Southeast University  
Nanjing, Jiangsu, China

Qi Shao\*

mathshaoqi@seu.edu.cn

School of Mathematics, Southeast University  
Nanjing, Jiangsu, China

Wenwu Yu<sup>†</sup>

wwyu@seu.edu.cn

School of Mathematics, Southeast University  
Nanjing, Jiangsu, China

## Abstract

Spatio-temporal prediction is a pivotal task with broad applications in traffic management, climate monitoring, energy scheduling, etc. However, existing methodologies often struggle to balance model expressiveness and computational efficiency, especially when scaling to large real-world datasets. To tackle these challenges, we propose STH-SepNet (Spatio-Temporal Hypergraph Separation Networks), a novel framework that decouples temporal and spatial modeling to enhance both efficiency and precision. Therein, the temporal dimension is modeled using lightweight large language models, which effectively capture low-rank temporal dynamics. Concurrently, the spatial dimension is addressed through an adaptive hypergraph neural network, which dynamically constructs hyperedges to model intricate, higher-order interactions. A carefully designed gating mechanism is integrated to seamlessly fuse temporal and spatial representations. By leveraging the fundamental principles of low-rank temporal dynamics and spatial interactions, STH-SepNet offers a pragmatic and scalable solution for spatio-temporal prediction in real-world applications. Extensive experiments on large-scale real-world datasets across multiple benchmarks demonstrate the effectiveness of STH-SepNet in boosting predictive performance while maintaining computational efficiency. This work may provide a promising lightweight framework for spatio-temporal prediction, aiming to reduce computational demands and while enhancing predictive performance. Our code is available at <https://github.com/SEU-WENJIA/ST-SepNet-Lightweight-LLMs-Meet-Adaptive-Hypergraphs>.

\*Jiawen Chen and Qi Shao contributed equally to this work.

<sup>†</sup>Correspondence to Duxin Chen and Wenwu Yu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*KDD '25, Toronto, ON, Canada.*

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-1454-2/25/08  
<https://doi.org/10.1145/3711896.3736904>

## CCS Concepts

• **Computing methodologies** → **Spatial and physical reasoning; Temporal reasoning.**

## Keywords

Spatio-Temporal Prediction, Graph Neural Networks, Large Language Models, Adaptive Hypergraph Neural Networks

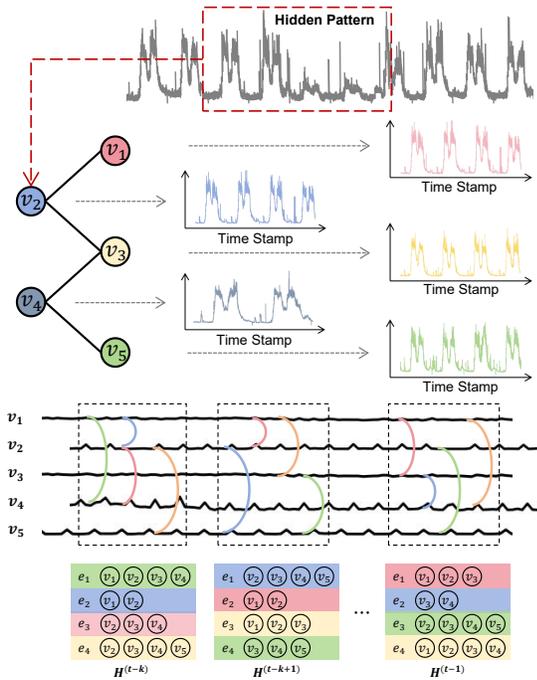
## ACM Reference Format:

Jiawen Chen, Qi Shao, Duxin Chen, and Wenwu Yu. 2025. Decoupling Spatio-Temporal Prediction: When Lightweight Large Models Meet Adaptive Hypergraphs. In *Proceedings of Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.2 (KDD '25)*. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3711896.3736904>

## 1 Introduction

Spatio-temporal prediction serves as a fundamental component of modern data-driven decision-making, with applications in urban traffic forecasting [5, 37], climate modeling [3, 25], and energy grid optimization, etc. Despite its broad significance, the field faces two primary challenges: accurately capturing dynamic spatial dependencies and ensuring computational scalability for large-scale real spatio-temporal datasets [39]. While deep learning has led to notable advancements, existing methods often struggle to achieve a balance between model expressiveness and computational efficiency [21].

Recent advances in graph neural networks (GNNs) and large language models (LLMs) have been extensively explored to address these challenges [18, 33]. GNNs are particularly effective in capturing spatial dependencies through graph-structured representations [42]. However, their dependence on static graph topologies poses a significant constraint, impeding their capacity to accurately model dynamic, higher-order interactions. For instance, in traffic networks, the influence between regions is constantly evolving and influenced by external conditions—factors that are inadequately captured by static adjacency matrices. Meanwhile, LLMs, which distinguish themselves in temporal prediction due to their strong sequence modeling capabilities [22], incur substantial computational expenses when applied to large node sets. Furthermore, their ability to leverage spatial structures remains limited. The dominant approach of jointly modeling spatial and temporal features within a single framework has been shown to exacerbate these challenges,



**Figure 1: (a) Spatio-temporal data exhibit spatial distribution shifts across different nodes. (b) Dynamic adaptive hypergraph captures evolving spatial distribution patterns.**

often leading to overparameterized models that are computationally demanding and difficult to optimize, without yielding proportional performance improvements [17]. This raises a critical question: Can spatial and temporal modeling be decoupled to achieve both efficiency and accuracy?

To tackle these challenges, this work proposes a separation strategy based on two key insights. First, as illustrated in Figure 1, temporal dynamics in spatiotemporal systems often exhibit a low-rank structure, implying that the evolution of system states can be efficiently characterized by a small number of latent factors [1, 26]. This low-rank property facilitates the use of lightweight sequence models, such as distilled versions of LLMs, to capture temporal trends without compromising expressiveness. Second, spatial dependencies in spatiotemporal systems can be viewed as a form of spatial drift, where the influence between nodes shifts over time due to external factors or intrinsic system dynamics. Traditional GNNs, despite their effectiveness in many applications, struggle to capture dynamic drift due to their dependence on static graph structures [32, 37]. To address this limitation, we propose an adaptive hypergraph framework that possesses enhanced representational capabilities for graphs, enabling it to model evolving higher-order interactions. This framework allows hyperedges to dynamically encapsulate shifting relationships among multiple nodes, thereby accurately reflecting the evolving nature of these connections.

Building on these insights, we propose STH-SepNet (Spatio-Temporal Hypergraph Separation Network), a novel framework that specifically integrates lightweight temporal modeling with adaptive spatial modeling. For temporal modeling, compact LLMs

(e.g., BERT [6], GPT-2 [29]) are employed to efficiently capture low-rank temporal dynamics. For spatial modeling, an adaptive hypergraph neural network is introduced to dynamically construct hyperedges, enabling the representation of spatial drift and higher-order interactions. A gating mechanism is further designed to fuse temporal and spatial representations, ensuring seamless integration while maintaining computational efficiency. The main contributions of this work are summarized as follows:

- We propose a lightweight spatio-temporal separation framework (STH-SepNet) for spatio-temporal prediction tasks. The framework integrates textual information and latent spatial dependencies, resulting in significant improvements in predictive performance.
- We design an adaptive hypergraph structure for spatial modeling, which dynamically constructs complex dependency relationships and enhances the extraction of spatial features through effective order modeling.
- We conduct extensive experiments to validate STH-SepNet, demonstrating state-of-the-art performance across multiple benchmarks. The proposed method demonstrates efficient execution on a single A6000 GPU, underscoring its practical applicability for real-world deployment.

## 2 Related works

### 2.1 Large Models for Prediction

Large language models, recognized for their extensive parameter sizes and strong generalization capabilities [9], have been increasingly applied to time-series analysis tasks, including prediction, classification, and imputation. To bridge the gap between numerical data and the text-based processing paradigm of LLMs, researchers have explored novel data formatting techniques. For instance, PromptCast converts numerical sequences into natural language prompts [38], while Gruver et al. encode time-series data as digit strings to enable zero-shot predictions [10]. These approaches demonstrate the potential of LLMs in temporal modeling while also highlighting the need for specialized adaptations to address the unique challenges of time-series data, such as irregular time intervals and long-range dependencies.

Recent efforts have sought to refine tokenization and embedding strategies to improve LLMs’ suitability for forecasting tasks. LLM4TS employs parameter-efficient fine-tuning (PEFT) to adapt pre-trained LLMs for time-series prediction [4], while Zhou et al. propose a unified framework for handling diverse time-series tasks [47]. Additionally, advancements such as reprogramming frameworks [17] and contrastive embedding strategies [20] further align numerical and textual modalities, enhancing LLMs’ capability to process temporal data. However, these methods predominantly focus on temporal modeling while largely overlooking spatial dependencies, which poses a fundamental limitation for spatio-temporal prediction tasks.

The integration of LLMs with transformer-based architectures has further expanded their applicability to domain-specific challenges. Models such as UniST [42] and OpenCity [21] incorporate transformers with graph neural networks to capture spatio-temporal dependencies, while ClimaX [25] and Pangu-Weather [3] illustrate the versatility of transformer-based designs in climate

forecasting. Despite these advancements, existing approaches often struggle to effectively balance spatial and temporal modeling, resulting in increased computational complexity without proportional performance improvements. This underscores the need for novel frameworks that integrate spatio-temporal structures more efficiently while maintaining computational feasibility.

## 2.2 Spatio-Temporal Prediction

Recent advances in spatio-temporal prediction have been driven by the integration of transformer-based models and graph neural networks, aimed at addressing the dual challenges of capturing long-range temporal dependencies and complex spatial interactions. Transformer-based models, such as DLinear [43] and TimesNet [34], have demonstrated strong performance in time-series forecasting by leveraging multi-scale temporal patterns and efficient attention mechanisms. Similarly, PatchTST [27] introduces patch-based attention to enhance local dependency modeling, while iTransformer [24] reconfigures the transformer architecture for multivariate time-series modeling. However, these models often struggle with distribution shifts, such as changes in traffic patterns or external conditions, limiting their robustness in real-world applications.

In the realm of spatial modeling, GNNs and hypergraph structures have emerged as effective tools for capturing complex spatial relationships. STG-NCDE [5] integrates neural controlled differential equations with GNNs to model continuous-time dynamics, while STAEformer [23] incorporates spatial and temporal attention mechanisms within a transformer-based framework. Hypergraph-based models, such as STHGCN [39] and GPT-ST [22], further advance the field by dynamically capturing higher-order dependencies and spatial drift. However, these approaches frequently rely on static or predefined structures, limiting their ability to adapt to evolving spatial relationships and distribution shifts over time.

A key limitation of existing methods is their inability to effectively model distribution shifts, which are inherent in spatio-temporal systems. For example, STID [31] simplifies spatio-temporal prediction but lacks adaptability to dynamic spatial interactions, while FEDformer [46] and Autoformer [36] primarily focus on temporal modeling without accounting for spatial distribution shifts. These gaps underscore the need for adaptive structures capable of dynamically capturing evolving spatial and temporal patterns. To address these challenges, we propose STH-SepNet, a lightweight framework that leverages adaptive hypergraphs to model distribution shifts and lightweight transformers to capture temporal dynamics, achieving state-of-the-art performance.

## 3 Preliminaries

### 3.1 Problem Formulation

Given a graph set denoting the spatial feature  $G = (V, E)$ , where  $V$  and  $E$  represents the set of  $N = |V|$  vertices and the set of edges, respectively. The spatio-temporal prediction problem of multivariate time-series forecasting is defined as follows: suppose the historical observations from  $L$  previous moments  $X_{(t-L+1):t} \in \mathbb{R}^{L \times N \times F}$ , our model STH-SepNet aims to predict the values for the next  $H$  timestamps data  $\hat{X}_{(t+1):(t+H)} \in \mathbb{R}^{H \times N \times F}$ . That is,

$$\hat{X}_{(t+1):(t+H)} = \text{STH-SepNet}_\theta(X_{(t-L+1):t}, \hat{A}, \Phi), \quad (1)$$

where  $\hat{A}$  is the structural adjacency matrix,  $\Phi$  is the prompt of prompt information for input vector, and  $\theta$  the parameter of model.

### 3.2 Adaptive Network Construction

We introduce an adaptive adjacency matrix,  $\tilde{A}_{\text{adp}}$ , as input to the ST-Block, aiming to mitigate similarity between adjacent nodes. Given node features  $E_1, E_2 \in \mathbb{R}^{N \times d}$ , we employ a shared-parameter feed forward neural network (FFN) to generate node embeddings, which are then mapped to  $F_1, F_2 \in \mathbb{R}^{N \times N}$  as follows:

$$F_1 = \tanh(\alpha \text{FFN}(E_1)), \quad (2)$$

$$F_2 = \tanh(\alpha \text{FFN}(E_2)), \quad (3)$$

where  $\alpha$  is a scaling factor that modulates the saturation rate of the activation function. The discrepancy between  $F_1$  and  $F_2$  captures directional relationships between nodes. To introduce non-linearity, we construct an asymmetric adjacency matrix  $A_{\text{adp}} \in \mathbb{R}^{N \times N}$ :

$$A_{\text{adp}} = \text{ReLU}(\tanh(\alpha(F_1^\top F_2 - F_2^\top F_1))). \quad (4)$$

This formulation effectively models asymmetric dependencies by leveraging learned node embeddings in the graph structure.

### 3.3 Incident Matrix

We integrate static spatial topology, e.g., geographic location information in a traffic network, as the static input to build an adjacency matrix  $A$ , defining node similarity via a negative exponential function of pairwise Euclidean distances. The similarity  $A_{ij}$  between nodes  $i$  and  $j$  is defined as:  $A_{ij} = \exp\left(-\frac{d_{ij}^2}{\sigma^2}\right)$ , where  $d_{ij}$  is the distance between nodes  $i$  and  $j$ , and  $\sigma$  is a scaling parameter that regulates the effect of distance on similarity. A fixed threshold is applied to maintain the sparsity of the adjacency matrix.

### 3.4 Adaptive HyperGraph Construction

**DEFINITION 1. (Hypergraph)** A high-order graph  $H(V, E)$  is defined by a set of  $n$  hypernodes  $V = \{v_1, v_2, \dots, v_n\}$  and a set of  $m$  hyperedges  $E = \{e_1, e_2, \dots, e_m\}$ , where  $e_j = (v_1^{(j)}, \dots, v_k^{(j)})$  is an unordered set of nodes on hyperedge  $e_j$ , with  $k = |e_j|$  denotes the number of nodes in the hyperedge.

**THEOREM 1.** For any  $k \geq 2$ , the  $(k-1)$ -hops neighborhood of a node  $v$ , denoted as  $N_{k-1}(v)$ , corresponds to all nodes involved in the  $k$ -order hyperedges in  $H_v^k$ , if and only if the following conditions are satisfied: For each  $w \in N_{k-1}(v)$ , (1) Local Connectivity Condition: there exists at least one path from  $v$  to  $w$  consisting of at most  $k-1$  hyperedges. (2) Hyperedge Coverage Condition: there exists a  $k$ -order hyperedge  $e \in H_v^k$  such that  $w \in e$  and  $e$  contains  $v, w$ , and at most  $k-2$  intermediate nodes. (3) Uniqueness Condition: if there exist multiple  $k$ -order hyperedges containing both  $v$  and  $w$ , then these hyperedges must share the same set of intermediate nodes. Formally:

$$w \in N_{k-1}(v) \iff \{v, F_1, F_2, \dots, u_k, w\} \in H_v^k, \quad (5)$$

where  $F_1, F_2, \dots, u_k$  are intermediary nodes.

Note that please refer to the Appendix A for detailed proof to the above theorem. In constructing a  $(k+1)$ -order hypergraph from  $k$ -hop neighborhoods, each node is interconnected with all nodes within its  $k$ -hops distance, forming one or more hyperedges.

Given a node  $v_i$ , its  $k$ -hops neighborhood  $N_k(v_i)$  includes all nodes reachable within  $k$  edges. Similarly, for node features  $E_3$ , feature representations  $F_3$  are obtained via a feedforward neural network:

$$F_3 = \tanh(\alpha \text{FFN}(E_3)), \quad (6)$$

where  $\alpha$  controls the activation saturation. Higher-order relationships are then constructed using K-Nearest Neighbors (KNN) on feature representations  $F_3 = [f_1, f_2, \dots, f_n]$ . For each node  $v_i$ , its nearest  $k$  neighbors  $N(v_i)$  form a hyperedge  $e_i = \{v_i\} \cup N(v_i)$ , where  $k = \max_j |e_j|$  is the hyperedge order and remains a pre-defined constant for consistency. To determine the hypergraph adaptive adjacency matrix  $H_{\text{adp}} \in \mathbb{R}^{n \times m}$ , where  $n$  is the number of nodes and  $m$  is the number of hyperedges, the adjacency matrix is defined as:  $H_{\text{adp},ij} = 1$ , if  $v_i \in e_j$ , and 0, otherwise in Algorithm A.1. Traditional graph-based methods primarily focus on pairwise interactions between nodes, which can be insufficient for modeling multiple nodes interact simultaneously. By contrast, hypergraphs allow for the representation of higher-order interactions through hyperedges. The adaptive adjacency matrix for hypergraphs can degenerate into traditional graphs, but it leverages this flexibility to capture richer and more complex relationships.

## 4 Methodology

In this work, we propose STH-SepNet, a spatio-temporal forecasting model that integrates a pre-trained LLM with adaptive hypergraphs. As shown in Figure 2, STH-SepNet comprises two key components: lightweight large language models for temporal dynamics and adaptive hypergraphs for spatial dependencies.

### 4.1 Global Trend Module

**Local Aggregation Module.** This module processes the model input node features  $X \in \mathbb{R}^{B \times N \times T \times F}$  by executing an average pooling operation to extract the common features of all nodes within a region, capturing the overall fluctuation trends. That is,

$$X_{\text{pool}} = \text{AvgPool}(X), \quad (7)$$

where  $X_{\text{pool}} \in \mathbb{R}^{B \times T \times F}$ ,  $B$  is the batch size,  $T$  denotes the time steps, and  $F$  equals 1. As shown in Figure 2 ①, the time series embedding module reduces the computational time and memory complexity of the model by aggregating information across adjacent time steps and utilizing temporal patches [27]. Specifically,  $X_{\text{pool}}$  is partitioned into overlapping or non-overlapping blocks  $X_P \in \mathbb{R}^{P \times N_P}$ , where  $P$  is the window length and the number of sliding windows is computed as  $N_P = (T - P)/S + 2$ , with the stride  $S$ . Each patch is treated as a time series token and embedded to obtain  $\hat{X}_P \in \mathbb{R}^{P \times d_m}$ , where  $d_m$  is the hidden dimension of the LLM.

### 4.2 Prompt Adaptation Module

Given that LLMs are primarily trained on extensive text corpora and lack inherent time series knowledge, we propose a cross-modal alignment strategy that transforms time-series data into textual tokens, enabling LLMs to leverage their reasoning capabilities for specialized forecasting tasks. To enhance predictive accuracy, we adopt Pattern-Exploiting Training, which formulates natural language templates as prompts within the embedding space. These prompts integrate three key components: dataset description, task

instructions, and statistical characteristics. More details refer to Appendix B.2. We further prepend this structured information as a prefix prompt, concatenating it with aligned temporal embeddings before feeding it into the LLM. This enables the model to generate valid outputs and adapt to downstream tasks.

**LLM Module.** We utilize a partially frozen pre-trained LLM to capture temporal dependencies in traffic data, fine-tuning its feed-forward layers via LoRA [13]. Pretrained models serves as the backbone for STH-SepNet, comprising stacked transformer decoder modules with  $N$  layers. As depicted in Figure. 2, the input to each layer is represented  $z = \{z^1, z^2, \dots, z^N\}$ , where  $z^1$  consists of concatenated prompt and time-series embeddings For the  $i$ -th layer, the input  $z^i$  undergoes multi-head self-attention (MHSA) followed by layer normalization (LN), producing an intermediate state  $\tilde{z}^i$ , which is further processed by a feed-forward network and another layer normalization step to yield  $z^{i+1}$  for the next layer. It can be summarized as:

$$(Q_i, K_i, V_i) = (W_i^O z^i, W_i^K z^i, W_i^V z^i), \quad (8)$$

$$\text{head}_i = \text{softmax} \left( \frac{Q_i K_i^T}{\sqrt{d}} \right) V_i, \quad (9)$$

$$\text{MHSA}(z^i, z^i, z^i) = W(\text{head}_1 \parallel \dots \parallel \text{head}_h), \quad (10)$$

$$\tilde{z}^i = \text{LN}(z^i + \text{MHSA}(z^i, z^i, z^i)), \quad (11)$$

$$z^{i+1} = \text{LN}(\tilde{z}^i + \text{FFN}(\tilde{z}^i)), \quad (12)$$

where  $z^i$  is the input hidden state at the  $i$ -th layer, and  $\tilde{z}^i$  is the intermediate state after MHSA and layer normalization. Since the LLM module outputs token sequences, we apply a linear layer to align learned patch representations. Additionally, pretrained models such as BERT [6], GPT2, LLAMA [33] and Deepseek [11, 40] can be employed for autoregressive token prediction.

### 4.3 Hypergraph Spatio-temporal Module

To enhance spatio-temporal prediction by incorporating higher-order coupling relationships, STH-SepNet consists of four key components: Mixed Multi-Layer Information Aggregation Module, Adaptive Graph Convolution Network, Hypergraph Convolutional Network, and Temporal Convolutional Network in Figure. 2 ③.

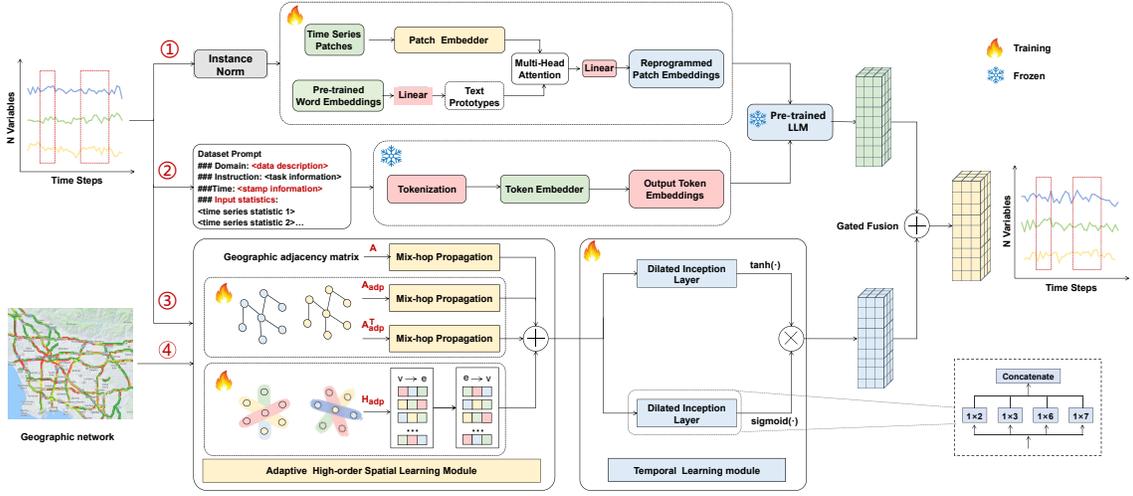
**Mixed Multi-Layer Information Aggregation Module (Mix-Prop).** Given an input  $X^{(0)} \in \mathbb{R}^{N \times C}$ , where  $N$  represents the number of nodes and  $C$  is the feature dimension, the module updates node features through  $k$ -layer propagation. The feature propagation at layer  $k + 1$  is formulated as,

$$X^{(k+1)} = \alpha * X^{(k)} + (1 - \alpha) \hat{A} X^{(k)}, \quad (13)$$

where  $X^{(k)}$  is the node feature matrix at the  $k$ -th layer,  $\hat{A} = D^{-\frac{1}{2}}(A + I)D^{-\frac{1}{2}}$  is the normalized adjacency matrix,  $D$  is the degree matrix,  $I$  is the identity matrix,  $\alpha$  controls the weight of residual connections. To further regulate the flow of information, the MixProp module incorporates a gating mechanism defined by:

$$G^{(k)} = \sigma(W_g X^{(k)}), \quad (14)$$

$$X^{(k+1)} = G^{(k)} \odot X^{(k)} + (1 - G^{(k)}) \odot \hat{A} X^{(k)},$$



**Figure 2: The framework of STH-SepNet.** Given a traffic network  $G = (V, E)$  and time series  $X$  as an example of spatial-temporal datasets. ① Tokenize and embed  $X$  using a customized embedding layer, reprogramming with condensed text prototypes for modality alignment. ② Incorporate dataset descriptions, task instructions, and statistical characteristics as prompt prefixes to guide input transformation. ③ Leverage a Hypergraph Spatio-Temporal module to model complex spatial dependencies and node-level variations via hierarchical representation learning. ④ Incident matrix: real geographic network, if not, adaptive graph or adaptive hypergraph is used. By integrating ①②③④, STH-SepNet generate the forecasts.

where  $W_g \in R^{C \times C}$ ,  $\sigma(\cdot)$  is the sigmoid activation function, and  $\odot$  denotes element-wise multiplication. The MixProp module employs  $k$ -layer propagation to expand the receptive field and capture dependencies between distant nodes.

**Adaptive Graph Convolution Network Module.** To capture both structural and adaptive spatial dependencies, we apply three MixProp-based graph convolutions on the adaptive adjacency matrix  $A_{adp}$  and the real road network  $A$ . These operations extract first-order, transposed first-order, and real relationships, defined as:

$$X_{gconv1} = \text{MixProp}(X, A_{adp}, K, \alpha), \quad (15)$$

$$X_{gconv2} = \text{MixProp}(X, A_{adp}^T, K, \alpha), \quad (16)$$

$$X_{gconv3} = \text{MixProp}(X, A, K, \alpha), \quad (17)$$

where  $X \in R^{B \times N \times T}$  represents the input time-series features,  $K$  is the number of propagation layers, and  $\alpha$  controls residual weighting. The final spatial representation is obtained by fusing these outputs:

$$X_{GCN} = X_{gconv1} + X_{gconv2} + X_{gconv3}. \quad (18)$$

**Adaptive Hypergraph Convolution Network Module.** Given an input feature matrix  $X \in R^{B \times N \times T}$  and an adaptive hypergraph adjacency matrix  $H_{adp}$ , this module employs two information aggregation mechanisms: from node to hyperedge and from hyperedge to node. Initially, a feedforward neural network (FFN) transforms the feature matrix. In the node-to-hyperedge process, each hyperedge  $e_j$  accumulates information from its associated nodes  $N(e_j)$ ,

followed by hyperedge aggregation and linear transformation:

$$X_e^{enc} = \text{FFN}(X), \quad (19)$$

$$X_e^{enc} = \sigma \left( \sum_{i \in N(e)} H_{adp,i} X_i^{enc} W \right), \quad (20)$$

where  $W$  represents the trainable parameter matrix, and  $\sigma(\cdot)$  is the ReLU activation function. Subsequently, features of all hyperedges containing a node  $v_i$  are aggregated back to node representation,

$$X_v^{enc} = \sum_{j \in \mathcal{E}(v_i)} H_{adp,j} X_{e_j}^{enc}, \quad (21)$$

where  $\mathcal{E}(v_i)$  indicates the set of hyperedges for node  $v_i$ . The output of the hypergraph spatial learning module is  $X_{HGCN} = X_v^{enc}$ .

To integrate pairwise and high-order spatial features, we fuse representations from GCNs and HGCNs, ensuring strong scalability in the spatial module:

$$X = \gamma X_{GCN} + (1 - \gamma) X_{HGCN}, \gamma \in [0, 1], \quad (22)$$

where  $\gamma$  is a tunable parameter. When  $\gamma = 1$ , the model degenerates into standard spatial learning module, whereas  $\gamma = 0$  captures only higher-order dependencies.

**Spatial-temporal Convolution Module.** The Spatio-Temporal Convolution Module consists of multiple stacked ST-Blocks, each containing a S-Block for spatial dependencies and a T-Block for temporal dependencies. In the S-Block, each node state  $h_t^{(v)}$  is initialized as the input features  $X \in R^{B \times N \times T \times F}$  and updated by

aggregating features from its neighbors, as follows:

$$m_t^v = \sum_{u \in N(v)} h_{t-1}^u, \quad (23)$$

$$h_t^v = \sigma((1 + \epsilon)h_{t-1}^v + m_t^v), \quad (24)$$

where  $m_t^v$  is the aggregated neighborhood feature at time  $t$ ,  $\sigma$  is an activation function and  $\epsilon$  is a learnable parameter.

The T-Block comprises 1-D dilated convolution layers with a gating mechanism featuring only an output gate. Given the input  $\chi \in \mathbb{R}^{T \times N \times F}$ , the gated output  $h$  is defined as:

$$h = \tanh(q(\chi)) \odot \sigma(q(\chi)), \quad (25)$$

where  $q(\chi)$  is the output of the dilated convolution layers,  $\odot$  denotes the Hadamard product, and  $\sigma$  is the sigmoid activation function.

#### 4.4 Gated Fusion Module

To integrate global trends and node heterogeneity, we fuse outputs from the pre-trained LLM and adaptive high-order spatial module, denoted as  $O_1, O_2 \in \mathbb{R}^{B \times T \times N}$ . A feedforward neural network (FFN) maps the concatenated output vectors to gates of equivalent dimensions. The gating process is formulated as:

$$\text{Gate} = \sigma(\text{FFN}([O_1, O_2])), \quad (26)$$

where  $\sigma$  represents the sigmoid activation function, and  $[\cdot, \cdot]$  denotes the concatenation operation. The ultimate gated fusion process can be expressed as:

$$\tilde{O} = O_1 \odot \text{Gate} + O_2 \odot (1 - \text{Gate}), \quad (27)$$

where Gate signifies the gate map, and  $\tilde{O} \in \mathbb{R}^{B \times T \times N}$  represents the resultant fused output.

#### 4.5 Experiment Settings

**Datasets.** We conduct the experiments on five datasets: BIKE-Inflow, BIKE-Outflow, PEMS03, BJ500 and METR-LA [19], where partition the datasets into train/validation/test sets by the ratio of 7:1:2. Appendix B.1 contains more dataset details.

**Baselines.** As aforementioned, STH-SepNet is designed as a general framework for spatio-temporal prediction tasks. To evaluate its effectiveness, we compare it with several state-of-the-art time series prediction models, including Autoformer [36], Informer [45], FEDformer [46], DLinear [43], TimesNet [35], PatchTST [28], iTransformer [24], TIMELLM [16], AdaMSHyper [30]. Additionally, to demonstrate the robustness of our model, we also include comparisons with spatio-temporal prediction models such as AGCRN [2], MSTGCN [15], MTGNN [37], STGODE [7], STSGCN [32], STGCN [41], GMAN [44], STAEformer [23], STD-MAE [8].

#### 4.6 Main Results

**4.6.1 Effectiveness of STH-SepNet. (RQ1).** Table 1 presents the performance of the proposed STH-SepNet method integrating pretrained model BERT across five datasets. Our method achieves the best Mean Absolute Error (MAE) and Root Mean Squared Error (RMSE) results, which can be attributed to its spatio-temporal separation strategy and adaptive hypergraph structure. Key advantages are detailed below:

**Decoupled Spatio-Temporal Modeling.** STH-SepNet addresses the limitations of joint spatio-temporal modeling by isolating temporal and spatial dependencies. On the BIKE-Outflow dataset, where non-stationary spatial events (e.g., weather-induced station closures) intersect with periodic temporal trends (e.g., rush hours), our method achieves an MAE of 5.33 and RMSE of 14.23, outperforming joint modeling frameworks like TimesNet (MAE: 5.56) and PatchTST (MAE: 5.63). By decoupling temporal dynamics (handled by lightweight modules) and spatial dependencies (modeled via adaptive hypergraphs), we mitigate interference between heterogeneous features, ensuring robust predictions in dynamic scenarios.

**Adaptive Hypergraphs for Dynamic Spatial Drift.** We propose an adaptive hypergraph structure to model evolving spatial dependencies. On the PEMS03 dataset, where traffic accidents or construction events disrupt road network interactions, our method reduces RMSE to 34.17, a 28.8% improvement over dynamic graph-based approaches like STAEformer (RMSE: 48.03). The dynamic hyperedge generation mechanism allows our method to adjust node relationships dynamically, capturing spatial drift. For example, on the BJ500 dataset, our method achieves an MAE of 5.58, surpassing MTGNN (MAE: 5.86) and MSTGCN (MAE: 5.62), demonstrating its ability to adapt to sudden changes in spatial dependencies.

**Scalable Efficiency.** Our approach significantly reduces computational complexity by decoupling the node dimension, transforming spatio-temporal prediction into parallelizable univariate tasks. On the METR-LA dataset (large-scale road network), our method achieves an MAE of 9.42 and RMSE of 16.41, outperforming both graph-based models (MTGNN: MAE: 9.98) and transformer variants (iTransformer: MAE: 10.15). Compared to large language model (LLM)-based baselines like TIMELLM (MAE: 12.36 on METR-LA), our method reduces MAE by 23.8%. This efficiency is achieved through lightweight LLM adaptations for temporal modeling and node-wise decoupling, avoiding the parameter bloat of monolithic LLM frameworks while maintaining high accuracy.

##### 4.6.2 Effectiveness of Adaptive Hypergraph Structure. (RQ2).

Table 2 presents the performance of the proposed method using different graph representation strategies, including a static graph (STH-SepNet-static), an adaptive graph convolutional network (STH-SepNet-GNN), and an adaptive hypergraph structure (STH-SepNet). Additionally, we compare these methods with the baseline large model TimeLLM. The results indicate that STH-SepNet achieves the best performance, and the effectiveness of the adaptive hypergraph can be attributed to the following factors:

**Dynamic Adaptation to Complex Spatial Dependencies.** The core strength of the adaptive hypergraph lies in its dynamic adjustment capability, which enables real-time responses to evolving spatial dependencies. In transportation networks, traditional static graphs or GNNs rely on predefined adjacency relationships and fail to capture sudden disruptions (e.g., traffic accidents) that alter node correlations. On the BIKE-Outflow dataset, STH-SepNet notably reduces prediction errors (MAE: 5.33 vs. 6.34 for the static variant) by dynamically generating hyperedge weights. This adaptability allows the model to flexibly represent shifts in regional traffic flows, such as adjusting inter-regional influence weights under policy-driven restrictions (e.g., traffic bans). Static graphs, constrained by

**Table 1: Performance comparison. Multivariate forecasting results with a prediction horizon of 48 time steps and a fixed lookback window of T=48. Bolded results indicate the best performance. (LLMs: BERT)**

Model	BIKE-Inflow		BIKE-Outflow		PEMS03		BJ500		METR-LA	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
Autoformer (NIPS, 2021)[36]	7.01	17.52	7.19	17.75	44.87	70.84	10.79	16.06	12.47	20.04
Informer (AAAI, 2021)[45]	8.25	20.37	9.21	21.50	33.72	52.15	7.58	11.96	14.50	20.35
FEDformer (PMLR, 2022)[46]	6.28	16.30	6.56	16.67	35.00	50.84	10.77	15.99	12.35	18.79
DLinear (AAAI, 2023) [43]	5.71	15.49	5.82	15.36	45.30	66.81	8.55	13.49	10.90	17.31
TimesNet (ICLR, 2023)[35]	5.54	15.41	<b>5.56</b>	<b>15.18</b>	37.54	62.99	8.67	13.96	10.22	18.29
PatchTST (ICLR, 2023)[28]	<b>5.53</b>	<b>15.39</b>	5.63	15.23	48.42	78.24	8.79	14.28	<b>10.13</b>	<b>18.27</b>
iTransformer (ICLR, 2024)[24]	6.05	16.39	6.15	16.69	43.63	70.61	9.01	14.32	10.15	18.36
TIMELLM (ICLR, 2024)[16]	6.81	16.72	6.93	16.30	<b>32.62</b>	<b>49.77</b>	<b>7.25</b>	<b>11.58</b>	12.36	18.53
AdaMSHyper (NIPS, 2024)[30]	6.72	16.91	7.04	17.14	33.49	50.37	7.41	11.60	12.51	18.60
AGCRN (NIPS, 2020)[2]	6.64	16.14	6.77	16.36	33.14	54.88	6.32	12.81	11.39	23.15
ASTGCN (AAAI, 2019)[12]	6.66	15.87	6.26	14.48	30.65	53.96	6.34	11.34	10.54	22.76
MSTGCN (TNSRE, 2021)[15]	<b>5.91</b>	<b>14.11</b>	6.04	14.24	29.57	47.97	<b>5.62</b>	<b>11.15</b>	10.17	20.24
MTGNN (KDD, 2020)[37]	6.16	14.80	<b>5.93</b>	<b>13.93</b>	<b>29.04</b>	<b>50.32</b>	5.86	10.91	9.98	21.23
STGODE (KDD, 2021)[7]	6.77	15.93	6.82	15.50	33.39	54.16	6.44	12.14	11.48	22.85
STSGCN (AAAI, 2020)[32]	6.73	15.89	6.58	15.36	34.23	58.07	6.40	12.03	11.07	22.79
STGCN (IJCAI, 2018)[41]	7.08	15.72	7.36	16.11	36.02	53.44	6.73	12.62	12.38	22.55
GMAN (AAAI, 2020)[44]	6.73	15.60	6.94	15.84	33.96	53.02	6.41	12.40	11.69	22.37
STAEformer (CIKM, 2024)[23]	5.97	14.57	6.17	14.70	29.62	48.03	5.79	10.42	<b>9.91</b>	<b>21.17</b>
STD-MAE (IJCAI, 2024) [8]	6.13	14.87	6.21	14.37	30.40	48.38	5.92	11.49	10.52	23.11
STH-SepNet (Ours)	<b>5.18</b>	<b>14.40</b>	<b>5.33</b>	<b>14.23</b>	<b>21.03</b>	<b>34.17</b>	<b>5.58</b>	<b>9.77</b>	<b>9.42</b>	<b>16.41</b>

fixed structures, cannot accommodate such changes. This capability is critical in dynamic scenarios (e.g., abrupt traffic flow shifts), enhancing the model’s efficiency in capturing spatial drift.

**Enhancing Lightweight Models to Surpass Large Counterparts.** While large models like TIMELLM leverage massive parameters to capture complex patterns, their performance in spatio-temporal prediction is outperformed by lightweight LLMs integrated with adaptive hypergraphs. For example, on the PEMS03 dataset, STH-SepNet with a BERT backbone achieves an RMSE of 34.17, far superior to TIMELLM’s 50.39. This highlights the division-of-labor advantage of the adaptive hypergraph: it specializes in modeling spatial dynamics, while the lightweight LLM focuses solely on single-node temporal trends, avoiding feature interference inherent in joint modeling. In traffic sensor networks, the hypergraph independently models cascading effects of multi-road congestion, while the LLM predicts traffic trends for individual sensors. This decoupling reduces model complexity and improves task-specific focus, enabling lightweight models to outperform large counterparts even in resource-constrained scenarios.

**Cross-Architecture Consistency and Higher-Order Interaction Modeling.** The adaptive hypergraph’s design is architecture-agnostic, delivering consistent performance across diverse LLM backbones (BERT, GPT, LLAMA, Deepseek). For instance, with the LLAMA7B backbone, STH-SepNet achieves an RMSE of 35.17 on PEMS03, outperforming both the GNN variant (RMSE: 35.24) and static version (RMSE: 49.92). This consistency stems from the hypergraph’s explicit modeling of higher-order spatial interactions. Traditional GNNs capture only pairwise node interactions, whereas

hypergraphs connect multiple nodes via hyperedges, simultaneously modeling multi-region joint influences (e.g., concurrent congestion across multiple roads amplifying impacts on central areas). In large-scale road networks like METR-LA, this capability enables accurate prediction of congestion propagation paths, while static graphs or standard GNNs exhibit significant biases due to their inability to represent multi-node effects. The explicit modeling of higher-order interactions distinguishes adaptive hypergraphs as a uniquely powerful tool for spatial dependency learning.

## 4.7 Ablation Study

We conduct three ablation studies to validate the effectiveness of key components in our proposed framework, exploring the role of LLMs, the mixed-order spatio-temporal convolutional networks and the order of hypergraph on forecast results.

### 4.7.1 LLMs play a critical role in temporal modeling. (RQ3).

To assess the necessity of LLMs in spatio-temporal prediction and their synergistic effects with the adaptive hypergraph structure, we conduct systematic ablation studies. Specifically, we compare a model variant without LLMs (w/o) only the adaptive hypergraph and linear temporal layers, against the full framework LLMs (STH-SepNet) equipped with different LLM backbones, such as BERT (w/i), GPT (w/i), LLAMA (w/i) and Deepseek (w/i).

The comparison between STH-SepNet-w/o and its LLM-enhanced counterparts reveals substantial performance improvements, particularly in capturing long-range dependencies and multi-scale periodicity. As illustrated in Figure 3, on the BIKE-Outflow dataset,

**Table 2: Comparison with different LLMs of adaptive high-order and low-order spatio-temporal multitime series forecasting results.**

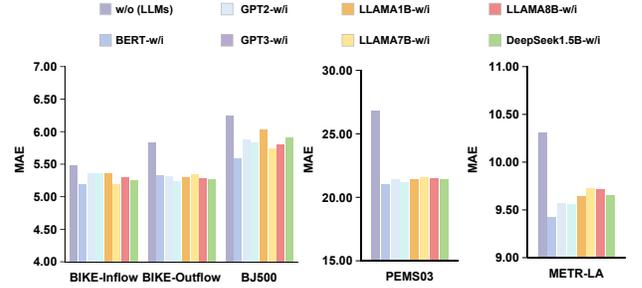
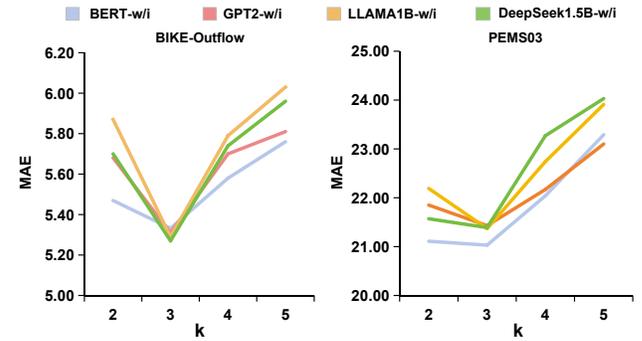
Model		BIKE-Outflow		PEMS03	
		MAE	RMSE	MAE	RMSE
BERT	TIMELLM	6.74	16.13	32.68	50.39
	STH-SepNet-Static	6.34	16.41	29.53	48.94
	STH-SepNet-GNN	<u>5.47</u>	<u>14.36</u>	<u>21.11</u>	<u>34.52</u>
	STH-SepNet	<b>5.33</b>	<b>14.23</b>	<b>21.03</b>	<b>34.17</b>
GPT2	TIMELLM	6.93	16.30	32.63	49.83
	STH-SepNet-Static	6.53	16.61	30.03	49.07
	STH-SepNet-GNN	<u>5.68</u>	<u>14.48</u>	<u>21.85</u>	<u>35.78</u>
	STH-SepNet	<b>5.31</b>	<b>14.24</b>	<b>21.43</b>	<b>35.01</b>
GPT3	TIMELLM	7.01	16.63	34.85	51.89
	STH-SepNet-Static	6.79	16.71	30.64	49.16
	STH-SepNet-GNN	<u>5.66</u>	<u>13.90</u>	<u>21.27</u>	<u>34.78</u>
	STH-SepNet	<b>5.24</b>	<b>14.16</b>	<b>21.13</b>	<b>34.69</b>
LLAMA1B	TIMELLM	7.10	16.74	34.06	51.59
	STH-SepNet-Static	6.27	16.40	29.76	48.63
	STH-SepNet-GNN	<u>5.87</u>	<u>14.80</u>	<u>22.19</u>	<u>35.87</u>
	STH-SepNet	<b>5.29</b>	<b>14.20</b>	<b>21.37</b>	<b>34.92</b>
LLAMA7B	TIMELLM	6.95	16.41	34.17	52.47
	STH-SepNet-Static	6.73	16.87	30.24	49.92
	STH-SepNet-GNN	<u>5.91</u>	<u>14.94</u>	<u>21.64</u>	<u>35.24</u>
	STH-SepNet	<b>5.34</b>	<b>14.31</b>	<b>21.52</b>	<b>35.17</b>
LLAMA8B	TIMELLM	7.02	16.55	35.72	51.97
	STH-SepNet-Static	6.85	16.64	30.47	49.98
	STH-SepNet-GNN	<u>5.70</u>	<u>14.27</u>	<u>21.57</u>	<u>35.23</u>
	STH-SepNet	<b>5.28</b>	<b>14.20</b>	<b>21.51</b>	<b>35.19</b>
DeepSeek1.5B	TIMELLM	6.94	16.25	33.19	50.28
	STH-SepNet-Static	6.79	16.37	30.26	48.81
	STH-SepNet-GNN	<u>5.74</u>	<u>14.55</u>	<u>21.73</u>	<u>35.47</u>
	STH-SepNet	<b>5.27</b>	<b>14.19</b>	<b>21.39</b>	<b>34.96</b>

the LLM-enhanced variant (BERT-w/i) significantly outperforms the non-enhanced version, and the integration of an LLM into STH-SepNet-GNN markedly boosts prediction accuracy. This enhancement can be attributed to the LLM’s capability to learn rich statistical characteristics from batch inputs, including minimum, maximum, median, and trend information. Consequently, the model effectively extracts semantic temporal features and complex periodic trends from traffic spatio-temporal data.

#### 4.7.2 Lightweight LLMs achieve competitive performance while maintaining architectural stability, highlighting the efficiency of the proposed framework. (RQ4)

As shown in Figure 3, on the STH-SepNet models, incorporating a BERT model demonstrates notable accuracy. For example, on the BIKE-Inflow dataset, BERT-w/i not only surpasses larger backbone models such as LLAMA7B-w/i but does so with fewer parameters (BERT:110M, LLAMA7B:6740M), indicating that excessive parameter scaling is not essential for effective temporal modeling. On the BIKE-Inflow dataset, the BERT-w/i not only surpasses larger backbones such as LLAMA7B-w/i but also operates with less than 10% fewer parameters, demonstrating that excessive parameter scaling is not

a prerequisite for effective temporal modeling. Moreover, results on the BJ500 dataset indicate that performance variations across different LLM backbones remain minimal, with fluctuations under 3%. This stability arises from the decoupled nature of the adaptive hypergraph, which independently captures spatial dependencies, thereby reducing reliance on LLM scale and ensuring that lightweight models remain competitive. We also provide a details comparison in Appendix B.3.

**Figure 3: Performance comparison of MAE between STH-SepNet trained on different datasets.****Figure 4: Analysis of effective order on adaptive hypergraph.**

#### 4.7.3 The effective order of the hypergraph can enhance the performance of the model.

Our proposed framework leverages the KNN algorithm to construct hyperedges in hypergraphs, thereby addressing the limitations of spatial dependencies that traditional spatio-temporal convolutions are unable to model. As illustrated in Figure 4, the effective order  $k = 3$  of the adaptive hypergraph significantly enhances model performance. On the BIKE-Outflow and PEMS03 datasets, STH-SepNet models based on BERT, GPT-2, LLAMA1B, and DeepSeek1.5B demonstrate that when  $k = 2$ , the high-order structure of STH-SepNet degenerates into a pairwise relationship. The empirical results indicate that as the order  $k$  increases, the model error initially decreases and then increases. This phenomenon is attributed to the fact that at  $k = 2$ , pairwise relationships fail to capture the underlying dependencies, while higher-order hypergraph structures with increased  $k$  order lead to overfitting of coupled interactions. That is,  $k = 3$  effectively characterizes evolving spatial dependencies.

**4.7.4 Computational Efficiency Analysis.** To analyze the advantages of decoupling in algorithmic efficiency, we test the STH-SepNet model with multiple large language models by comparing their GPU usage and training speed on different datasets using an NVIDIA A6000. Figure 5 shows that the STH-SepNet series outperforms TIMELLM in computational efficiency across most datasets. For instance, STH-SepNet (BERT) showed a GPU memory usage of just 24.6G and a training speed of 392 Epoch/s on the BIKE-Inflow/Outflow dataset, exhibiting superiority over both TIMELLM (BERT) and TIMELLM (GPT2). Furthermore, as the parameter size of LLMs increases, computational efficiency tends to decrease. However, larger model parameters do not enhance the accuracy of spatiotemporal predictions. This indicates that STH-SepNet (BERT) generally outperforms TIMELLM and larger STH-SepNet models in terms of GPU usage, training speed, and overall performance. This advantage stems from STH-SepNet’s decoupled processing of temporal features and its use of average pooling to extract global trend features (Eq. 7), which reduces GPU resource consumption and boosts efficiency.

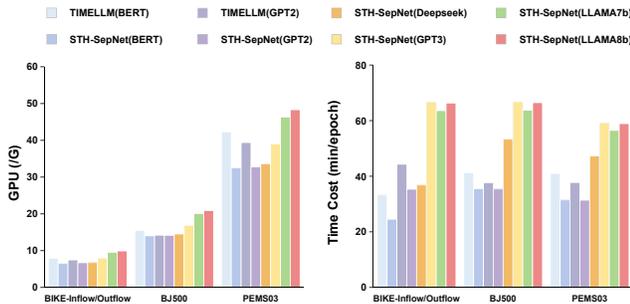


Figure 5: Comparison of GPU and time complexity.

## 5 Conclusion

This paper introduces STH-SepNet, a framework for spatio-temporal prediction that decouples temporal and spatial modeling through two specialized components: lightweight large language models for temporal dynamics and adaptive hypergraphs for spatial dependencies. By employing a spatio-temporal decoupling design, the ability of STH-SepNet to predict spatio-temporal data is significantly enhanced. Experimental results demonstrate improved accuracy across diverse datasets, including traffic networks (e.g., PEMS03, MAE: 21.03 vs. 26.84 for non-LLM variants) and urban mobility systems (e.g., BIKE-Outflow, MAE: 5.33 vs. 6.74 for LLM baselines). The adaptive hypergraph structure dynamically adjusts to spatial distribution shifts, such as policy-driven traffic pattern changes or sudden disruptions, enabling robust predictions in dynamic environments. The improved performance is attributed to the decoupled architecture, which allows temporal and spatial modules to focus on distinct patterns without mutual interference. Adaptive hypergraphs address the limitations of static graph structures by modeling higher-order interactions and real-time spatial drift, while lightweight LLMs efficiently capture temporal trends. This design is shown to generalize across datasets with varying scales

and dynamics, as evidenced by consistent results in both small-scale (e.g., BIKE) and large-scale (e.g., METR-LA) scenarios.

**Limitations and future work.** While STH-SepNet demonstrates strong performance, its current design has limitations. The framework assumes temporal and spatial dependencies can be cleanly decoupled, which may not hold in scenarios where these dimensions are intrinsically intertwined (e.g., rapidly evolving events with coupled spatio-temporal causality). Additionally, the adaptive hypergraph’s reliance on real-time node feature updates could pose challenges in latency-critical applications, where computational overhead for dynamic hyperedge generation might limit responsiveness. In future work, we will tackle these constraints by delving into hybrid architectures that strike a balance between decoupling and controlled interaction mechanisms.

## Acknowledgments

This work is supported by the National Key R&D Program of China under Grant No. 2022ZD0120004, the Zhishan Youth Scholar Program, the National Natural Science Foundation of China under Grant Nos. 62233004, 62273090, and the Jiangsu Provincial Scientific Research Center of Applied Mathematics under Grant No. BK20233002.

## References

- [1] Mohammad Taha Bahadori, Qi Rose Yu, and Yan Liu. 2014. Fast multivariate spatio-temporal analysis via low rank tensor learning. In *Proceedings of the 28th International Conference on Neural Information Processing Systems (Montreal, Canada) (NIPS’14)*. MIT Press, Cambridge, MA, USA, 3491–3499.
- [2] Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. 2020. Adaptive graph convolutional recurrent network for traffic forecasting. In *Proceedings of Advances in Neural Information Processing Systems*, Vol. 33. Curran Associates Inc., Red Hook, NY, USA, 17804–17815.
- [3] Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. 2023. Accurate medium-range global weather forecasting with 3D neural networks. *Nature* 619, 7970 (2023), 533–538.
- [4] Ching Chang, Wen-Chih Peng, and Tien-Fu Chen. 2023. LLM4TS: Two-Stage Fine-Tuning for Time-Series Forecasting with Pre-Trained LLMs. *arXiv preprint arXiv:2308.08469* (2023).
- [5] Jeongwhan Choi, Hwangyong Choi, Jeehyun Hwang, and Noseong Park. 2022. Graph Neural Rough Differential Equations for Traffic Forecasting. In *Proceedings of the AAAI conference on Artificial Intelligence*, Vol. 36. Association for Computing Machinery, New York, NY, USA, 6367–6374.
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 conference of the North American chapter of the association for computational linguistics: human language technologies, volume 1 (long and short papers)*. 4171–4186.
- [7] Zheng Fang, Qingqing Long, Guojie Song, and Kunqing Xie. 2021. Spatial-temporal graph ode networks for traffic flow forecasting. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*. 364–373.
- [8] Haotian Gao, Renhe Jiang, Zheng Dong, Jinliang Deng, Yuxin Ma, and Xuan Song. 2024. Spatial-temporal-decoupled masked pre-training for spatiotemporal forecasting. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*. 3998–4006.
- [9] Adam Goodge, Wee Siong Ng, Bryan Hooi, and See Kiong Ng. 2025. Spatio-Temporal Foundation Models: Vision, Challenges, and Opportunities. *arXiv preprint arXiv:2501.09045* (2025).
- [10] Nate Gruver, Marc Finzi, Shikai Qiu, and Andrew G Wilson. 2024. Large language models are zero-shot time series forecasters. In *Proceedings of Advances in Neural Information Processing Systems*, Vol. 36.
- [11] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025).
- [12] Shengnan Guo, Youfang Lin, Ning Feng, Chao Song, and Huaiyu Wan. 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI conference on Artificial Intelligence*,

- Vol. 33. 922–929.
- [13] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685* (2021).
- [14] Jing Huang and Jie Yang. 2021. UniGNN: a Unified Framework for Graph and Hypergraph Neural Networks. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization.
- [15] Ziyu Jia, Youfang Lin, Jing Wang, Xiaojun Ning, Yuanlai He, Ronghao Zhou, Yuhang Zhou, and H Lehman Li-wei. 2021. Multi-view spatial-temporal graph convolutional networks with domain generalization for sleep stage classification. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 29 (2021), 1977–1986.
- [16] Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-fang Li, Shirui Pan, et al. 2024. Time-LLM: Time Series Forecasting by Reprogramming Large Language Models. In *Proceedings of International Conference on Learning Representations*.
- [17] Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, et al. 2023. Time-llm: Time series forecasting by reprogramming large language models. *arXiv preprint arXiv:2310.01728* (2023).
- [18] Thomas N Kipf and Max Welling. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907* (2016).
- [19] Fuxian Li, Jie Feng, Huan Yan, Guangyin Jin, Fan Yang, Funing Sun, Depeng Jin, and Yong Li. 2023. Dynamic graph convolutional recurrent network for traffic prediction: Benchmark and solution. *ACM Transactions on Knowledge Discovery from Data* 17, 1 (2023), 1–21.
- [20] Jiawei Li, Jingshu Peng, Haoyang Li, and Lei Chen. 2024. UniCL: A Universal Contrastive Learning Framework for Large Time Series Models. *arXiv preprint arXiv:2405.10597* (2024).
- [21] Zhonghang Li, Long Xia, Lei Shi, Yong Xu, Dawei Yin, and Chao Huang. 2024. OpenCity: Open Spatio-Temporal Foundation Models for Traffic Prediction. *CoRR abs/2408.10269* (2024).
- [22] Zhonghang Li, Lianghao Xia, Yong Xu, and Chao Huang. 2024. GPT-ST: generative pre-training of spatio-temporal graph neural networks. In *Proceedings of Advances in Neural Information Processing Systems*, Vol. 36.
- [23] Hangchen Liu, Zheng Dong, Renhe Jiang, Jiawen Deng, Q Chen, and X Song. 2023. STAEformer: Spatio-Temporal Adaptive Embedding Makes Vanilla Transformers SOTA for Traffic Forecasting. In *Proceedings of the ACM International Conference on Information and Knowledge Management*, 21–25.
- [24] Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. 2023. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. *arXiv preprint arXiv:2310.06625* (2023).
- [25] Tung Nguyen, Johannes Brandstetter, Ashish Kapoor, Jayesh K Gupta, and Aditya Grover. 2023. ClimateX: A foundation model for weather and climate. *arXiv preprint arXiv:2301.10343* (2023).
- [26] Tong Nie, Guoyang Qin, Wei Ma, Yuewen Mei, and Jian Sun. 2024. ImputeFormer: Low rankness-induced transformers for generalizable spatiotemporal imputation. In *Proceedings of the Knowledge Discovery and Data Mining*, 2260–2271.
- [27] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2022. A time series is worth 64 words: Long-term forecasting with transformers. *arXiv preprint arXiv:2211.14730* (2022).
- [28] Yuqi Nie, Nam H Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. 2023. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In *Proceedings of International Conference on Learning Representations*.
- [29] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* 1, 8 (2019), 9.
- [30] Zongjiang Shang, Ling Chen, Binqing Wu, and Dongliang Cui. 2024. AdaMSHyper: Adaptive Multi-Scale Hypergraph Transformer for Time Series Forecasting. In *Proceedings of Neural Information Processing Systems*, Vol. 37. 33310–33337.
- [31] Zezhi Shao, Zhao Zhang, Fei Wang, Wei Wei, and Yongjun Xu. 2022. Spatial-temporal identity: A simple yet effective baseline for multivariate time series forecasting. In *Proceedings of the ACM International Conference on Information & Knowledge Management*, 4454–4458.
- [32] Chao Song, Youfang Lin, Shengnan Guo, and Huaiyu Wan. 2020. Spatial-temporal synchronous graph convolutional networks: A new framework for spatial-temporal network data forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 34. 914–921.
- [33] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023).
- [34] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. 2022. Timesnet: Temporal 2d-variation modeling for general time series analysis. *arXiv preprint arXiv:2210.02186* (2022).
- [35] Haixu Wu, Tengge Hu, Yong Liu, Hang Zhou, Jianmin Wang, and Mingsheng Long. 2023. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In *Proceedings of International Conference on Learning Representations*.
- [36] Haixu Wu, Jiehui Xu, Jianmin Wang, and Mingsheng Long. 2021. Autoformer: Decomposition transformers with auto-correlation for long-term series forecasting. In *Proceedings of Advances in Neural Information Processing Systems*, Vol. 34. 22419–22430.
- [37] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. 2020. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 753–763.
- [38] Hao Xue and Flora D Salim. 2023. PromptCast: A new prompt-based learning paradigm for time series forecasting. *IEEE Transactions on Knowledge and Data Engineering* (2023).
- [39] Xiaodong Yan, Tengwei Song, Yifeng Jiao, Jianshan He, Jiaotuan Wang, Ruopeng Li, and Wei Chu. 2023. Spatio-temporal hypergraph learning for next POI recommendation. In *Proceedings of the International ACM SIGIR Conference on Research and Development in Information Retrieval*, 403–412.
- [40] An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. 2024. Qwen2.5-Math Technical Report: Toward Mathematical Expert Model via Self-Improvement. *arXiv preprint arXiv:2409.12122* (2024).
- [41] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2018. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In *Proceedings of International Joint Conference on Artificial Intelligence*. International Joint Conferences on Artificial Intelligence Organization.
- [42] Yuan Yuan, Jingtao Ding, Jie Feng, Depeng Jin, and Yong Li. 2024. Unist: A prompt-empowered universal model for urban spatio-temporal prediction. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 4095–4106.
- [43] Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. 2023. Are transformers effective for time series forecasting?. In *Proceedings of the AAAI conference on Artificial Intelligence*, Vol. 37. 11121–11128.
- [44] Chuanpan Zheng, Xiaoliang Fan, Cheng Wang, and Jianzhong Qi. 2020. Gman: A graph multi-attention network for traffic prediction. In *Proceedings of the AAAI conference on Artificial Intelligence*, Vol. 34. 1234–1241.
- [45] Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. 2021. Informer: Beyond efficient transformer for long sequence time-series forecasting. In *Proceedings of the AAAI conference on Artificial Intelligence*, Vol. 35. 11106–11115.
- [46] Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. 2022. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting. In *Proceedings of International Conference on Machine Learning*. PMLR, 27268–27286.
- [47] Tian Zhou, Peisong Niu, Liang Sun, Rong Jin, et al. 2023. One fits all: Power general time series analysis by pretrained lm. In *Proceedings of Advances in Neural Information Processing Systems*, Vol. 36. 43322–43355.

## A Hypergraph Theory

**DEFINITION A.1. (Hyperedge)** Given a high-order graph  $H = (V, E)$ , a hyperedge  $e \in E$  is a non-empty subset of  $V$ . For each  $e \in E$ ,  $e \neq \emptyset$ , and  $e = (v_{i_1}, v_{i_2}, \dots, v_{i_k}), v_{i_j} \in V$  [14].

**DEFINITION A.2. ( $k$ -uniform Hyperedge)** If a hyperedge  $e \in E$  contains exactly  $k$  vertices, then it is called a  $k$ -uniform hyperedge [14]. Formally, for each  $e \in E$ ,  $|e| = k$ .

**DEFINITION A.3. ( $k$ -hops neighborhoods)** Given a node  $v_i$ , its  $k$ -hops neighborhood  $N_k(v_i)$  comprises all nodes that can be reached from  $v_i$  via at most  $k$  edges from  $v_i$ .

**Proof for Theorem. 1.** The theorem will be proved by showing both directions of the equivalence.

(1) *Sufficiency* ( $\Rightarrow$ ): Assume  $w \in N_{k-1}(v)$ , we need to show that there exists a  $k$ -order hyperedge  $e \in H_v^k$  such that  $w \in e$  and  $e$  contains  $v, w$ , and most  $k-2$  intermediate nodes. By the local connectivity condition, there exists a path from  $v$  to  $w$  using at most  $k-1$  hyperedges. Let this path be represented by the sequence of hyperedges  $e_1, e_2, \dots, e_{k-1}$ . Since each hyperedge can connect

more than two nodes, we can construct a  $k$ -order hyperedge  $e$  that includes  $v$  and  $w$  along with at most  $k - 2$  intermediate nodes. This satisfies the hyperedge coverage condition. If there are multiple such hyperedges, the uniqueness condition ensures that they share the same set of intermediate nodes, thus ensuring consistency.

(2) *Necessity* ( $\Leftarrow$ ): Assume there exists a  $k$ -order hyperedge  $e \in H_v^k$  such that  $w \in e$  and  $e$  contains  $v, w$ , and at most  $k - 2$  intermediate nodes. We need to show that  $w \in N_{k-1}(v)$ . By definition, the hyperedge  $e$  connects  $v$  and  $w$  through at most  $k - 2$  intermediate nodes. This implies that there is a path from  $v$  to  $w$  consisting of at most  $k - 1$  hyperedges (since the hyperedge itself can be considered as part of the path). Thus,  $w$  is within the  $(k - 1)$ -hops neighborhood of  $v$ , satisfying the local connectivity condition. The hyperedge coverage condition is directly satisfied by the existence of  $e$ , and the uniqueness condition ensures that no other hyperedges contradict this structure.

We supplement the pseudocode for the hypergraph generation process in Algorithm A.1 as follows.

---

#### Algorithm A.1 Hyperedge Construction

---

**Require:** Batch data  $[B, L, N, F]$ , high-order parameter  $k$

**Ensure:** Constructed hyperedges and spatial interaction results

- 1: Initialize hyperedges set: hyperedges =  $\emptyset$
  - 2: **for** each node in  $V(|V| = N)$  **do**
  - 3:   Find  $k$  nearest neighbors using KNN
  - 4:   Construct dynamic  $k$ -order hyperedge
  - 5:   Add the hyperedge to hyperedges
  - 6: **end for**
  - 7: Construct adaptive hypergraph based on hyperedges' set
- 

**Table B.1: Comparison of parameter sizes and dimensions across large language models**

Model	Parameters	LLM Dimension
BERT <sup>1</sup>	110M	768
GPT-2 <sup>2</sup>	124M	768
GPT-3 <sup>3</sup>	7580M	4096
LLAMA-1B <sup>4</sup>	1230M	2048
LLAMA-7B <sup>5</sup>	6740M	4096
LLAMA-8B <sup>6</sup>	8000M	4096
DeepSeek-Qwen1.5B <sup>7</sup>	1500M	1536

## B Experiment Settings and Results

Details of datasets description are as follows:

- BIKE-Inflow/Outflow: The dataset captures bicycle demand across 295 traffic nodes in New York, recorded hourly. The dataset spans from 2023-01-01 00:00 to 2024-01-01 23:00.

<sup>1</sup><https://huggingface.co/google-bert/bert-base-uncased>

<sup>2</sup><https://huggingface.co/openai-community/gpt2>

<sup>3</sup><https://huggingface.co/TurkuNLP/gpt3-finnish-large>

<sup>4</sup><https://huggingface.co/meta-llama/Llama-3.2-1B>

<sup>5</sup><https://huggingface.co/huggyllama/llama-7b>

<sup>6</sup><https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct>

<sup>7</sup><https://huggingface.co/deepseek-ai/DeepSeek-R1-Distill-Qwen-1.5B>

- PEMS03 dataset: The dataset contains traffic speed data from 358 stations in the California Highway System, with a 5-minute interval. The time range covers weekdays from 2008-01-01 00:00 to 2008-03-31 23:55:00.
- BJ500: The dataset consists of traffic speed information from 500 stations in the Beijing Highway System, also at 5-minute intervals. The dataset covers weekdays from 2020-07-01 00:00:00 2020-07-31 23:55:00.
- METR-LA: The dataset from the Los Angeles Metropolitan Transportation Authority contains average traffic speed measured by 207 loop detectors on the highways of Los Angeles County ranging from Mar 2012 to Jun 2012.

Details of large language models is displayed in Table B.1. Prefix prompts on the PEMS03 dataset are designed as follows:

[Dataset Description] Pems03 Traffic dataset: The data consists of 358 selected stop data distributed in the California Highway System (CalTrans). The data set time interval is a time interval of 5 minutes. The time range is in 2008-01-01 00:00:00 2008-03-31 23:55:00 weekdays of 358 stations traffic speed information, data shows a strong periodicity. The input consists of adjacency matrix and timing feature matrix. The adjacency matrix is a new matrix with similar traffic flow patterns by analyzing the characteristics of existing spatio-temporal traffic data, and the characteristic matrix is the time series characteristic matrix of each sensor node.

[Task Instruction] Forecast the next  $L$  steps given the previous  $H$  steps.

[Statistical Information] The timestamp information is formatted as [month, day, hour, minutes]. The input time begins from ⟨start time⟩ to ⟨end time⟩, and the prediction time spans from ⟨start prediction time⟩ to ⟨end prediction time⟩. The minimum value is ⟨min value⟩, the maximum value is ⟨max value⟩, and the median value is ⟨median value⟩. The trend of the input is either upward or downward. The top 5 lags are ⟨lag values⟩.

## C Ablation Experiment Results

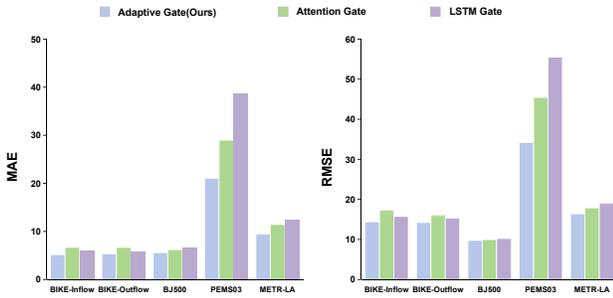
**Ablation Study (RQ3 and RQ4).** We compare two model variants: one without large language models (LLMs) — referred to as STH-SepNet-w/o — which retains only the adaptive hypergraph and linear temporal layers, and another full framework — denoted as STH-SepNet (XXX)-w/i — equipped with various LLM backbones, including BERT, GPT, LLAMA, and Deepseek. As shown in Table C.1 On the PEMS03 dataset, which exhibits dynamic traffic patterns, the exclusion of LLMs results in significantly higher errors (MAE: 26.84, RMSE: 43.44) compared to the full model with a BERT backbone (MAE: 21.03, RMSE: 34.17), corresponding to increases of 21.7% and 21.3%, respectively. Similarly, on the BJ500 dataset, the BERT-equipped variant (MAE: 5.58) outperforms STH-SepNet-w/o (MAE: 6.24), underscoring the LLM’s ability to enhance the performance of traffic data prediction. The results demonstrate the two key findings that LLMs facilitate spatio-temporal prediction and superior performance of our framework.

**Ablation Study for Gating Mechanism.** We evaluate three fusion mechanisms across five datasets. Figure. C.1 displays our proposed adaptive gating mechanism achieves superior prediction accuracy (MAE/RMSE) on all benchmarks. Notably, the adaptive gating exhibits particularly significant performance advantages in

**Table C.1: Performance comparison. Multivariate prediction results for different LLMs with a prediction horizon of 48 time steps and a fixed backtracking window of  $T=48$ . Bolded results indicate the best performance.**

Model	BIKE-Inflow		BIKE-Outflow		PEMS03		BJ500		METR-LA	
	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE	MAE	RMSE
STH-SepNet-w/o (LLMs)	5.47	<b>14.01</b>	5.83	15.17	26.84	43.44	6.24	10.81	10.31	18.26
STH-SepNet (BERT)-w/i	<b>5.18</b>	14.40	5.33	14.23	<b>21.03</b>	<b>34.17</b>	<b>5.58</b>	9.77	<b>9.42</b>	<b>16.41</b>
STH-SepNet (GPT2)-w/i	5.35	14.71	5.31	14.24	21.43	35.01	5.88	9.69	<u>9.57</u>	<u>16.64</u>
STH-SepNet (GPT3)-w/i	5.36	14.78	5.24	14.16	<u>21.13</u>	<u>34.69</u>	5.83	9.92	9.56	16.77
STH-SepNet (LLAMA1B)-w/i	5.36	14.80	5.29	14.20	21.37	34.92	6.03	10.08	9.64	16.84
STH-SepNet (LLAMA7B)-w/i	<b>5.18</b>	14.24	5.34	14.31	21.52	35.27	<u>5.73</u>	<u>9.58</u>	9.72	16.90
STH-SepNet (LLAMA8B)-w/i	5.30	14.66	<u>5.28</u>	<u>14.20</u>	21.51	35.19	<b>5.80</b>	<b>9.65</b>	9.71	16.94
STH-SepNet (DeepSeek-Qwen1.5B)-w/i	<u>5.25</u>	<u>14.48</u>	<u>5.27</u>	<b>14.19</b>	21.39	34.96	5.91	9.94	9.65	16.89

complex traffic scenarios like METR-LA and BJ500 (17.7%–45.6% MAE reduction), compared to cross-attention and LSTM-based gating mechanisms. Further analysis reveals distinct characteristics of alternative approaches: While cross-attention gating achieves suboptimal performance on PEMS03, its global attention computation introduces redundant feature interactions and demonstrates vulnerability to noise under sparse data conditions. The LSTM-based gating shows moderate temporal modeling capability in BIKE-Outflow prediction but fails to effectively capture the dynamic evolution of spatial features. The superiority of adaptive gating stems from its learnable weighting parameters that dynamically balance spatiotemporal feature contributions: 1) abrupt speed changes in METR-LA datasets, 2) tidal flow patterns in BIKE datasets.

**Figure C.1: Ablation studies: comparison of the three fusion mechanisms on various datasets. (LLMs:BERT).**

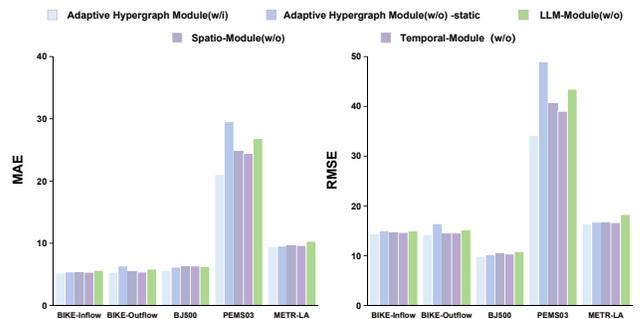
**Effective Order of Adaptive Hypergraph.** Table C.2 illustrates that as the order of the adaptive hypergraph increases, the RMSE error of the model initially decreases and then increases. Specifically, on the Bike-Outflow and PEMS03 datasets, when  $k = 3$ , the STH-SepNet model achieves the smallest MSE values. This finding indicates that an appropriately chosen hypergraph order can significantly boost spatio-temporal prediction performance.

**Ablation Study for Different Modules.** We conduct ablation studies on the STH-SepNet model, which incorporates spatio-temporal modules, static graphs, LLMs, adaptive graphs, and hypergraph modules. Figure C.2 shows that the fully equipped model integrates the adaptive hypergraph module, the spatiotemporal module, and BERT as the language model module, demonstrates

**Table C.2: Analysis of effective order  $k \in \{2, 3, 4, 5\}$  on adaptive hypergraph. Performance metric (RMSE). (LLMs:BERT).**

Dataset		2	3	4	5
BIKE-Outflow	BERT	<b>13.36</b>	14.23	14.89	15.37
	GPT2	14.48	<b>14.24</b>	15.29	15.58
	LLAMA1B	14.80	<b>14.20</b>	15.54	16.19
	DeepSeek1.5B	14.55	<b>14.19</b>	15.46	16.05
PEMS03	BERT	34.52	<b>34.17</b>	35.81	37.84
	GPT2	35.78	<b>35.01</b>	36.22	37.74
	LLAMA1B	35.87	<b>34.92</b>	37.16	39.07
	DeepSeek1.5B	35.47	<b>34.96</b>	38.03	39.27

the most significant improvement in the two key performance metrics of MAE and RMSE. From a spatial perspective, the model with the adaptive hypergraph module (w/i) achieves remarkable performance on most datasets, though slightly inferior to the complete model. In contrast, when the adaptive hypergraph module is removed and only the static graph module is retained (w/o-static), there is a marked increase in MAE and RMSE for tasks such as BIKE-Inflow and BIKE-Outflow. For instance, in the BIKE-Inflow task, the MAE reaches 5.38, and the RMSE soars to 15.01. This indicates that the static graph struggles to adequately depict dynamic and complex spatial associations, thereby reducing the model’s performance and reflecting the irreplaceable nature of the adaptive hypergraph module in learning dynamic spatial structures.

**Figure C.2: Ablation studies: spatio, temporal, static graph, LLMs, adaptive graph and hypergraph module. (LLMs:BERT).**