

HyperD: Hybrid Periodicity Decoupling Framework for Traffic Forecasting

Minlan Shao¹, Zijian Zhang^{2,*}, Yili Wang¹, Yiwei Dai¹, Xu Shen¹, Xin Wang^{1,*}

¹ School of Artificial Intelligence, Jilin University, Changchun, China

² College of Computer Science and Technology, Jilin University, Changchun, China

{shaoml24, daiyw23, shenxu23}@mails.jlu.edu.cn, {zhangzijian, wangyili, xinwang}@jlu.edu.cn

Abstract

Accurate traffic forecasting plays a vital role in intelligent transportation systems, enabling applications such as congestion control, route planning, and urban mobility optimization. However, traffic forecasting remains challenging due to two key factors: (1) complex spatial dependencies arising from dynamic interactions between road segments and traffic sensors across the network, and (2) the coexistence of multi-scale periodic patterns (e.g., daily and weekly periodic patterns driven by human routines) with irregular fluctuations caused by unpredictable events (e.g., accidents or weather disruptions). To tackle these challenges, we propose **HyperD** (Hybrid Periodicity Decoupling), a novel framework that decouples traffic data into **periodic** and **residual components**. The periodic component is handled by the **Hybrid Periodic Representation Module**, which extracts fine-grained daily and weekly patterns using learnable periodic embeddings and spatial-temporal attention. The residual component, which captures non-periodic, high-frequency fluctuations, is modeled by the **Frequency-Aware Residual Representation Module**, leveraging complex-valued MLP in frequency domain. To enforce semantic separation between the two components, we further introduce a **Dual-View Alignment Loss**, which aligns low-frequency information with the periodic branch and high-frequency information with the residual branch. Extensive experiments on four real-world traffic datasets demonstrate that HyperD achieves state-of-the-art prediction accuracy, while offering superior robustness under disturbances and improved computational efficiency compared to existing methods.

Code — <https://github.com/ll121202/HyperD>

Introduction

The rapid development of urbanization and sensing technologies drives abundant spatial-temporal data across domains such as traffic systems (Liu et al. 2024b; Zhou et al. 2020; Ma et al. 2023), power grids (Li et al. 2025), and environmental monitoring (Liu et al. 2025c). Among these, traffic forecasting is a core task in intelligent transportation systems, as accurate flow prediction is essential for understanding and managing urban mobility dynamics (Xia et al. 2025;




Figure 1: Illustration of traffic flow in the PEMS04 dataset. (a) Traffic flow for sensors 54, 73, and 299 over one week. (b) Traffic flow of sensor 54 over three consecutive weeks. Both images demonstrate the strong presence of periodic patterns in traffic data.

Zhou et al. 2025). It requires traffic forecasting models to well capture rich spatial dependencies and diverse temporal patterns inherent in real-world road networks.

Driven by consistent human routines and recurring activity patterns, periodicity plays a fundamental role in the temporal pattern of traffic data (Zhang, Zheng, and Qi 2017; Yu et al. 2019). As illustrated in Figure 1, traffic flow curves often repeat in consistent patterns across days and weeks, while also showing notable variation across different sensor locations. This indicates that **periodicity is not only prominent but also spatially diverse**, making it a critical predictive signal. As a result, many existing methods attempt to leverage these patterns through trend-seasonality decomposition (Yuan and Qiao 2024; Wang et al. 2024a; Cao et al. 2025), which separates the data into a trend component and a seasonal component. However, these methods suffer from two key limitations: ① **Implicit Modeling of Periodicity** — Periodic information is split across trend and seasonal components, overlooking its unified structure across time scales and spatial locations. This fragmentation hinders direct modeling of coherent periodic patterns. ② **Inaccurate Decomposition** — Traditional decomposition methods are non-learnable and inefficient, often resulting in misalignment between trend and seasonal components, and poor performance in capturing complex temporal patterns (Yu et al. 2024; Kim et al. 2025).

To move beyond implicit periodic modeling, recent methods like CycleNet (Lin et al. 2024) make initial strides toward explicit periodic modeling. However, these approaches

*Corresponding authors: Zijian Zhang and Xin Wang.
Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

can hardly tackle traffic forecasting because they fail to model interactions between spatial locations and temporal dynamics. Additionally, they consider only a single periodic scale, leaving rich multi-scale periodic patterns underutilized. Addressing these limitations requires a flexible framework that can capture both multi-scale periodic patterns and fine-grained spatial-temporal interactions. Furthermore, periodic patterns in traffic data are often overlapping and spatially heterogeneous, while irregular fluctuations challenge conventional modeling approaches.

To cope with limitations above, we propose HyperD, a **Hybrid Periodic Decoupling** framework that decouples traffic data into **periodic and residual components**, each modeled by tailored mechanisms. (1) For implicit modeling of periodicity, the **Hybrid Periodic Representation Module** uses multiple learnable embeddings and a spatial-temporal attention mechanism to model multi-scale periodic patterns. The learnable embeddings are leveraged according to the timestamps of traffic states and maintain multi-scale periodic patterns explicitly, addressing the challenge of overlapping periodic patterns that existing methods often fail to resolve. (2) For inaccurate decomposition, HyperD uses the **Frequency-Aware Residual Representation Module** to model residual dynamics, and further introduce the **Dual-View Alignment Loss** to ensure separation between the two branches. Specifically, this loss aligns low-frequency information with the periodic branch and high-frequency information with the residual branch, encouraging each component to focus on its respective frequency band for more precise decoupling. This promotes divergence between periodic and residual components, enabling HyperD to capture both **multi-scale periodic patterns** and **irregular fluctuations**. Our main contributions are as follows:

- HyperD is proposed as a novel framework that decouples traffic data into periodic and residual components, addressing multi-scale periodic patterns and irregular fluctuations for more accurate forecasting.
- A Hybrid Periodic Representation Module captures multi-scale periodic patterns using learnable embeddings and spatial-temporal attention mechanism, while a Frequency-Aware Residual Representation Module models non-periodic fluctuations through a spatial-temporal frequency encoder.
- The Dual-View Alignment Loss ensures effective decoupling between periodic and residual components, preventing semantic redundancy. Extensive experiments demonstrate that HyperD outperforms existing methods in accuracy, robustness, and computational efficiency.

Related Work

Spatial-Temporal Forecasting

Spatial-temporal graph neural networks (STGNNs) (Zhang et al. 2023a; Miao et al. 2024; Liu et al. 2024a; Wang et al. 2024c) have become a dominant approach in traffic forecasting, as they jointly model spatial correlations and temporal dynamics. The integration of LLMs with spatial-temporal models has also shown strong potential for multi-variate time-series forecasting (Liu et al. 2025a; Zhang et al.

2023b; Liu et al. 2025b; Shen et al. 2025). For spatial modeling, recent approaches typically employ graph neural networks (Wang et al. 2025; Yang et al. 2025; Fu et al. 2025; Shen et al. 2024; He et al. 2025; Wang et al. 2024b; Shao et al. 2025), to encode spatial correlations among sensors. Representative examples include STGCN (Yu, Yin, and Zhu 2018) and DCRNN (Li et al. 2018), which build on fixed road network structures, while GWNNet (Wu et al. 2019) and AGCRN (Bai et al. 2020) further introduce adaptive graphs to capture time-varying spatial dependencies.

For temporal modeling, various architectures have been explored. RNN-based models (Jiang et al. 2023b) capture short-term sequential patterns, while TCNs (Wu et al. 2019; Fang et al. 2021) employ dilated convolutions to model longer temporal dependencies. More recently, Transformer-based models (Jiang et al. 2023a; Gao et al. 2024) use attention mechanisms to capture complex temporal dynamics.

Spatial-Temporal Decoupling Methods

To handle the complexity of traffic dynamics, some recent methods attempt to decouple spatial-temporal data into interpretable components. D²STGNN (Shao et al. 2022) decouples traffic signals into diffusion and inherent components to model different propagation patterns. STWave (Fang et al. 2023) disentangles representations of traffic time series, separating complex traffic dynamics into stable trends and fluctuating events. STDN (Cao et al. 2025) decomposes traffic flow into the trend-cyclical and seasonal components in view of spatial-temporal embeddings.

Methodology

Problem Definition. We represent the traffic network as a directed graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}, \mathbf{A}\}$, where \mathcal{V} is a set of N nodes, corresponding to the traffic sensors deployed on the road network. \mathcal{E} denotes the set of edges indicating the connectivity between sensors. $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the adjacency matrix that describes whether a connection exists between nodes.

Given a historical traffic time series $\mathbf{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{T_1}\} \in \mathbb{R}^{T_1 \times N}$, where $\mathbf{x}_t \in \mathbb{R}^N$ represents the observation at time t across N nodes, our goal is to predict the future traffic states over the next T_2 time steps. We denote the predicted time series as $\hat{\mathbf{Y}} = \{\hat{\mathbf{y}}_{T_1+1}, \hat{\mathbf{y}}_{T_1+2}, \dots, \hat{\mathbf{y}}_{T_1+T_2}\} \in \mathbb{R}^{T_2 \times N}$, which is obtained by:

$$\mathbf{X}, \mathcal{G} \xrightarrow{f_\theta} \hat{\mathbf{Y}}, \quad (1)$$

where f_θ is a mapping function parameterized by θ .

Hybrid Periodic Representation Module

The overall framework of HyperD is shown in Figure 2. To effectively capture multi-scale periodic patterns in traffic data, we explicitly model periodicity using learnable embeddings. This allows HyperD to leverage prior knowledge of daily and weekly temporal structures, which are among the most dominant patterns in traffic data. This approach can also be easily extended to other temporal patterns, such as hourly or monthly. The Hybrid Periodic Representation Module consists of three main components: Learnable

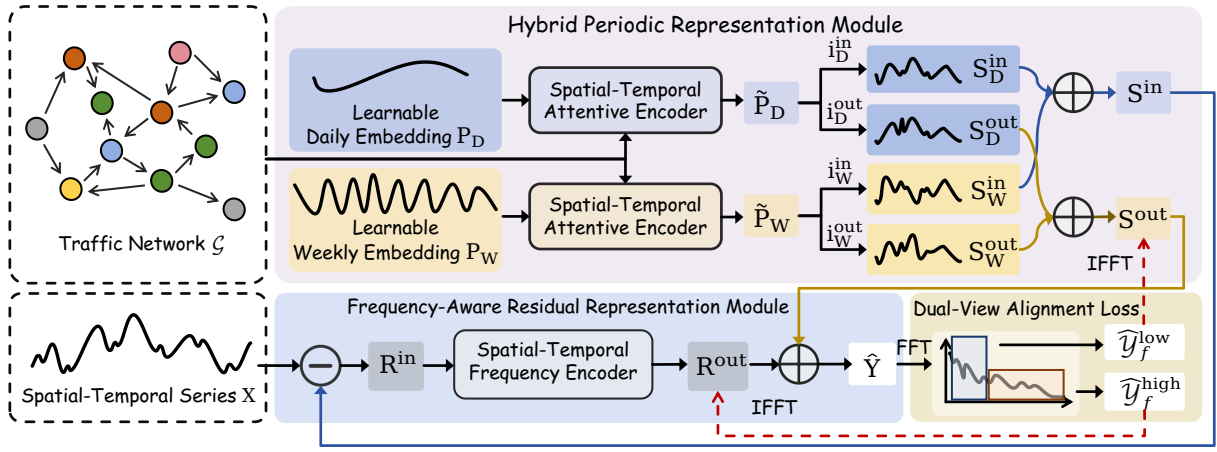


Figure 2: Overview of HyperD, which comprises three main components: (1) The Hybrid Periodic Representation Module encodes daily and weekly embeddings using the Spatial-Temporal Attentive Encoder and generates hybrid periodic patterns. (2) The Frequency-Aware Residual Representation Module encodes the residual using a Spatial-Temporal Frequency Encoder and combines it with the periodic component to yield the final prediction. (3) The Dual-View Alignment Loss separates the prediction into low- and high-frequency parts, which are then aligned with the periodic and residual branches, respectively.

Daily Embeddings (LDE) and Learnable Weekly Embeddings (LWE), Spatial-Temporal Attentive Encoder (STAE), and Hybrid Periodic Pattern.

Learnable Daily and Weekly Embedding The goal of LDE and LWE is to explicitly capture the dominant daily and weekly periodic patterns in traffic data, which are essential for accurate traffic forecasting. To achieve this, we introduce two learnable embeddings: $\mathbf{P}_D \in \mathbb{R}^{L_D \times N}$ for daily patterns and $\mathbf{P}_W \in \mathbb{R}^{L_W \times N}$ for weekly patterns, where L_D and L_W are the respective period lengths. These embeddings allow the model to effectively represent periodic patterns at multiple temporal scales.

The period lengths are determined by data sampling frequency. For instance, with a 5-minute sampling rate, the daily period length is $L_D = 288$ and the weekly period length is $L_W = 2016$. To accelerate the convergence and incorporate temporal prior knowledge, we initialize these embeddings using a statistical prior initialization strategy. We normalize the training data and compute the mean values for each node at every time step within the day and week. These mean values are used to initialize the embeddings, providing a reliable starting point for learning periodic patterns.

As a core component of the Hybrid Periodic Representation Module, these learnable embeddings captures the primary periodic patterns in the traffic data. To refine these embeddings and capture long-range dependencies, we introduce the STAE, which combines spatial correlations and temporal dynamics through attention mechanisms.

Spatial-Temporal Attentive Encoder To model interactions among nodes and capture long-range dependencies, we refine the learned periodic embeddings using the STAE. It combines GCNs (Kipf and Welling 2017) with spatial and temporal self-attention mechanisms to integrate both spatial correlations and temporal dynamics. This encoder enhances the embeddings, ensuring that both local and global dependencies

are effectively captured.

Given a periodic embedding $\mathbf{P} \in \mathbb{R}^{L \times N}$, where L denotes the period length (e.g., L_D for daily, L_W for weekly), we first apply GCN to incorporate spatial correlations:

$$\mathbf{H} = \sigma(\hat{\mathbf{A}}\mathbf{P}\mathbf{W}), \quad (2)$$

where $\hat{\mathbf{A}} \in \mathbb{R}^{N \times N}$ is the normalized Laplacian matrix, $\mathbf{W} \in \mathbb{R}^{N \times N}$ is the trainable weight matrix, and σ denotes the ReLU activation function.

To further capture long-range dependencies within the periodic patterns, we apply the self-attention mechanism (Vaswani et al. 2017) along both temporal and spatial dimensions:

$$\begin{aligned} \mathbf{Q} &= \mathbf{H}\mathbf{W}^q & \mathbf{K} &= \mathbf{H}\mathbf{W}^k & \mathbf{V} &= \mathbf{H}\mathbf{W}^v \\ \text{Attention}(\mathbf{H}) &= \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^\top}{\sqrt{d}}\right)\mathbf{V}, \end{aligned} \quad (3)$$

where \mathbf{H} is the input feature matrix, and $\mathbf{W}^q, \mathbf{W}^k, \mathbf{W}^v$ are the trainable projection matrices.

We perform temporal self-attention by treating each time step as a query, with trainable projections $\mathbf{W}_t^q, \mathbf{W}_t^k, \mathbf{W}_t^v \in \mathbb{R}^{N \times N}$. We perform spatial self-attention by transposing \mathbf{H} and treating each node as a query, with trainable projections $\mathbf{W}_s^q, \mathbf{W}_s^k, \mathbf{W}_s^v \in \mathbb{R}^{L \times L}$. Let $\mathbf{H}_t = \text{Attention}_t(\mathbf{H})$ and $\mathbf{H}_s = \text{Attention}_s(\mathbf{H}^\top)^\top$ denote the outputs of temporal and spatial self-attention, respectively. The two outputs are concatenated and fused through a linear projection:

$$\tilde{\mathbf{P}} = \mathbf{W}_o[\mathbf{H}_t; \mathbf{H}_s], \quad (4)$$

where $\tilde{\mathbf{P}} \in \mathbb{R}^{L \times N}$, $\mathbf{W}_o \in \mathbb{R}^{2N \times N}$ is the trainable weight matrix. The same operations are independently applied to the daily and weekly embeddings, producing $\tilde{\mathbf{P}}_D \in \mathbb{R}^{L_D \times N}$ and $\tilde{\mathbf{P}}_W \in \mathbb{R}^{L_W \times N}$, respectively. As the operations are architecturally identical, we present a unified formulation above for clarity. The refined embeddings, $\tilde{\mathbf{P}}_D$ and $\tilde{\mathbf{P}}_W$, are integrated to form the Hybrid Periodic Pattern.

Hybrid Periodic Pattern To integrate daily and weekly periodic patterns, we construct a hybrid periodic pattern by aggregating the corresponding segments from the daily and weekly embeddings. For each time step, we use two temporal metadata: the **time of day** and the **day of week**, to compute the indices in the daily and weekly embeddings. The **daily index** is given by $\mathbf{i}_D = \text{time of day}$. The **weekly index** is computed as $\mathbf{i}_W = \text{time of day} + \text{day of week} \times L_D$. Given the input time series $\mathbf{X} \in \mathbb{R}^{T_1 \times N}$, we retrieve the corresponding segments from the daily and weekly embeddings, $\tilde{\mathbf{P}}_D$ and $\tilde{\mathbf{P}}_W$, using the indices \mathbf{i}_D^{in} and \mathbf{i}_W^{in} . These segments are denoted as:

$$\mathbf{S}_D^{\text{in}}(t) = \tilde{\mathbf{P}}_D[\mathbf{i}_D^{\text{in}}(t), :] \quad \mathbf{S}_W^{\text{in}}(t) = \tilde{\mathbf{P}}_W[\mathbf{i}_W^{\text{in}}(t), :], \quad (5)$$

where $t \in \{1, \dots, T_1\}$ represents the time steps in the historical time series. These two segments are then aggregated to form the hybrid periodic pattern for the historical time series: $\mathbf{S}^{\text{in}} = \mathbf{S}_D^{\text{in}} + \mathbf{S}_W^{\text{in}}$.

Similarly, for the predicted time series $\hat{\mathbf{Y}}$, we retrieve the corresponding daily and weekly segments for each time steps $t \in \{T_1 + 1, \dots, T_1 + T_2\}$:

$$\mathbf{S}_D^{\text{out}}(t) = \tilde{\mathbf{P}}_D[\mathbf{i}_D^{\text{out}}(t), :] \quad \mathbf{S}_W^{\text{out}}(t) = \tilde{\mathbf{P}}_W[\mathbf{i}_W^{\text{out}}(t), :] \quad (6)$$

They are aggregated to form the hybrid periodic pattern for the predicted time series: $\mathbf{S}^{\text{out}} = \mathbf{S}_D^{\text{out}} + \mathbf{S}_W^{\text{out}}$.

The Hybrid Periodic Pattern captures the primary periodic dynamics in the data, setting a solid foundation for forecasting. However, traffic data also contains irregular fluctuations that cannot be captured by periodic patterns alone. To address this, we introduce the **Frequency-Aware Residual Representation Module**.

Frequency-Aware Residual Representation Module

While daily and weekly embeddings capture stable periodic patterns, traffic data also contains irregular fluctuations and disruptions that deviate from these periodic behaviors. To model these variations, we compute the residual component as the difference between the original traffic data \mathbf{X} and the hybrid periodic pattern \mathbf{S}^{in} :

$$\mathbf{R}^{\text{in}} = \mathbf{X} - \mathbf{S}^{\text{in}}. \quad (7)$$

These residual fluctuations often exhibit high-frequency variations across both space and time. To better capture these dynamics, the Frequency-Aware Residual Representation Module applies a **Spatial-Temporal Frequency Encoder** (STFE), which transforms the residual component into the frequency domain. This encoder models frequency-specific spatial and temporal behaviors, allowing the network to capture the high-frequency residual variations and complement the periodic patterns for more accurate forecasting.

Spatial-Temporal Frequency Encoder To capture spatial-temporal interactions in the residual component, we transform it into the frequency domain along both the spatial and temporal dimensions. In each dimension, we apply a complex-valued MLP to model frequency-specific behaviors. The transformed residual is then converted back to the time domain.

We first project the residual component $\mathbf{R}^{\text{in}} \in \mathbb{R}^{T_1 \times N}$ into a high-dimensional space, resulting in $\tilde{\mathbf{R}}^{\text{in}} \in \mathbb{R}^{T_1 \times N \times D}$, where D is the embedding dimension. To enable frequency-domain modeling, we apply Fast Fourier Transform (FFT) along both dimensions:

$$\mathcal{R}_f^s = \text{FFT}_s(\tilde{\mathbf{R}}^{\text{in}}, \text{dim} = \text{spatial}), \quad (8)$$

where $\mathcal{R}_f^s \in \mathbb{C}^{T_1 \times N' \times D}$ is the spatial frequency representation, and $N' = \lfloor \frac{N}{2} \rfloor + 1$ corresponds to the truncated spectrum due to symmetry in FFT.

Motivated by recent advances in frequency modeling (Yi et al. 2023), we introduce a complex-valued MLP (C-MLP) to refine the frequency representation. Operating in the frequency domain with C-MLP is equivalent to convolution in the time domain, enabling more efficient and global modeling while reducing computational cost by focusing on the essential frequency components:

$$\begin{aligned} \text{Re}(\mathcal{R}_1) &= \sigma(\text{Re}(\mathcal{R}_f^s) \cdot \mathbf{W}_1^r - \text{Im}(\mathcal{R}_f^s) \cdot \mathbf{W}_1^i + \mathbf{b}_1^r) \\ \text{Im}(\mathcal{R}_1) &= \sigma(\text{Im}(\mathcal{R}_f^s) \cdot \mathbf{W}_1^r + \text{Re}(\mathcal{R}_f^s) \cdot \mathbf{W}_1^i + \mathbf{b}_1^i) \\ \text{Re}(\mathcal{R}_2) &= \sigma(\text{Re}(\mathcal{R}_1) \cdot \mathbf{W}_2^r - \text{Im}(\mathcal{R}_1) \cdot \mathbf{W}_2^i + \mathbf{b}_2^r) \\ \text{Im}(\mathcal{R}_2) &= \sigma(\text{Im}(\mathcal{R}_1) \cdot \mathbf{W}_2^r + \text{Re}(\mathcal{R}_1) \cdot \mathbf{W}_2^i + \mathbf{b}_2^i), \end{aligned} \quad (9)$$

where $\mathcal{R}_2 \in \mathbb{C}^{T_1 \times N' \times D}$ is the output after applying C-MLP to the spatial frequency representation, and $\text{Re}(\cdot)$ and $\text{Im}(\cdot)$ denote the real and imaginary parts of a complex-valued tensor, respectively. The weights $\mathbf{W}_1^r, \mathbf{W}_1^i \in \mathbb{R}^{D \times D'}$, $\mathbf{W}_2^r, \mathbf{W}_2^i \in \mathbb{R}^{D' \times D}$ and $b_1^r, b_1^i \in \mathbb{R}^{D'}$, $b_2^r, b_2^i \in \mathbb{R}^D$ are the trainable parameters, and σ denotes the ReLU activation function. D' is the hidden layer dimension.

After the frequency domain refinement, we apply the Inverse Fast Fourier Transform (IFFT) to convert the frequency representation back into the time domain:

$$\mathbf{R}^s = \text{IFFT}_s(\mathcal{R}_2, \text{dim} = \text{spatial}), \quad (10)$$

where $\mathbf{R}^s \in \mathbb{R}^{T_1 \times N \times D}$ is the refined residual representation in the spatial domain. Similarly, we repeat the same process (FFT, C-MLP, IFFT) along the temporal dimension to capture the temporal frequency behaviors:

$$\begin{aligned} \mathcal{R}_f^t &= \text{FFT}_t(\mathbf{R}^s, \text{dim} = \text{temporal}) \\ \mathcal{R}_2^t &= \text{C-MLP}(\mathcal{R}_f^t) \\ \mathbf{R}^t &= \text{IFFT}_t(\mathcal{R}_2^t, \text{dim} = \text{temporal}). \end{aligned} \quad (11)$$

Finally, we add the residual connection and apply a linear projection to forecast T_2 future time steps:

$$\mathbf{R}^{\text{out}} = \text{Proj}(\mathbf{R}^t + \tilde{\mathbf{R}}^{\text{in}}). \quad (12)$$

The final prediction is then computed by combining the outputs of the hybrid periodic component \mathbf{S}^{out} and the residual component \mathbf{R}^{out} :

$$\hat{\mathbf{Y}} = \mathbf{S}^{\text{out}} + \mathbf{R}^{\text{out}}. \quad (13)$$

After obtaining the final prediction $\hat{\mathbf{Y}}$ by combining the hybrid periodic component \mathbf{S}^{out} and the residual component \mathbf{R}^{out} , we need to ensure the two components remain distinct to maximize their complementary information. To achieve this, we introduce the Dual-View Alignment Loss (DVA).

Dual-View Alignment Loss

Despite being processed separately, the hybrid periodic and residual components may still overlap in their representations, reducing the decoupling’s effectiveness. To address this, the DVA explicitly enforces separation between the low-frequency periodic component and the high-frequency residual component.

We apply FFT to the predicted output $\hat{\mathbf{Y}}$ to obtain its frequency representation $\hat{\mathcal{Y}}_f = \text{FFT}(\hat{\mathbf{Y}})$, and divide the frequency spectrum into low- and high-frequency parts using a predefined threshold F_{low} :

$$\hat{\mathcal{Y}}_f^{\text{low}} = \hat{\mathcal{Y}}_f[0 : F_{\text{low}}] \quad \hat{\mathcal{Y}}_f^{\text{high}} = \hat{\mathcal{Y}}_f[F_{\text{low}} :], \quad (14)$$

where $\hat{\mathcal{Y}}_f^{\text{low}}$ represents the low-frequency part, which predominantly captures the hybrid periodic patterns (i.e., the periodic component), and $\hat{\mathcal{Y}}_f^{\text{high}}$ represents the high-frequency part, which captures sharp fluctuations and non-periodic variations (i.e., the residual component).

Next, we apply the IFFT to each part, recovering the corresponding time domain representations. The low frequency part is compared with the hybrid periodic pattern \mathbf{S}^{out} , while the high frequency part is compared with the residual \mathbf{R}^{out} , both using the Mean Squared Errors (MSE):

$$\mathcal{L}_{\text{dva}} = \text{MSE}(\text{IFFT}(\hat{\mathcal{Y}}_f^{\text{low}}), \mathbf{S}^{\text{out}}) + \text{MSE}(\text{IFFT}(\hat{\mathcal{Y}}_f^{\text{high}}), \mathbf{R}^{\text{out}}). \quad (15)$$

This loss enforces frequency alignment by encouraging the low-frequency part of the prediction to match the periodic component and the high-frequency part to match the residual component, ensuring that the two branches capture distinct and complementary information.

The prediction loss $\mathcal{L}_{\text{pred}}$ is computed as:

$$\mathcal{L}_{\text{pred}} = \sum_{t=T_1+1}^{T_1+T_2} \sum_{n=1}^N \left| \hat{\mathbf{Y}}_{t,n} - \mathbf{Y}_{t,n}^{\text{gt}} \right|. \quad (16)$$

where $\hat{\mathbf{Y}}_{t,n}$ denotes the predicted value at time step t for node n , and $\mathbf{Y}_{t,n}^{\text{gt}}$ is the corresponding ground-truth.

Finally, the total training loss \mathcal{L} combines the prediction loss with the alignment loss:

$$\mathcal{L} = \mathcal{L}_{\text{pred}} + \alpha * \mathcal{L}_{\text{dva}}, \quad (17)$$

where α is a weighting coefficient that balances the contribution of the alignment loss.

Experiments

In this section, we conduct comprehensive experiments to evaluate HyperD from multiple perspectives. We aim to answer the following key research questions:

Q1: How does HyperD perform compared to SOTA prediction methods across diverse real-world traffic datasets?

Q2: What is the impact of each module on the overall performance of HyperD?

Q3: How does HyperD perform in terms of reliability and computational performance under real-world disturbances?

Datasets

To assess the performance of HyperD, we conduct experiments on four commonly used real-world traffic flow datasets: PEMS03/04/07/08 (Song et al. 2020).

Baselines

To comprehensively evaluate our HyperD, we compare with two lines of state-of-the-art methods, including **(a) spatial-temporal prediction methods:** STGCN (Yu, Yin, and Zhu 2018), DCRNN (Li et al. 2018), GWNet (Wu et al. 2019), ASTGCN (Guo et al. 2019), MTGNN (Wu et al. 2020), STGODE (Fang et al. 2021), ST-WA (Cirstea et al. 2022), DGCRN (Li et al. 2023) and STPGNN (Kong, Guo, and Liu 2024), and **(b) spatial-temporal decoupling methods:** D²STGNN (Shao et al. 2022), STWave (Fang et al. 2023), CycleNet (Lin et al. 2024), and STDN (Cao et al. 2025). We compare two CycleNet variants: CycleNet-W and CycleNet-D, both built on an MLP backbone.

Experimental Setup

To ensure a fair comparison, we adopt the experimental settings commonly used in previous studies (Shao et al. 2024). The datasets are divided into training, validation, and test sets with a ratio of 6:2:2, respectively. We use the past 12 time steps (previous hour) to forecast the next 12 time steps (upcoming hour).

Main Results

To answer Q1, we evaluate the overall prediction performance of HyperD against existing methods, as shown in Table 1. HyperD outperforms all compared models across the datasets, demonstrating the effectiveness of our approach. Key findings are as follows:

① **Decoupled methods generally outperform conventional spatial-temporal prediction methods.** While models like D²STGNN and STWave perform suboptimally on three out of four datasets, their strong performance can be attributed to their decoupled architectures, which separate trend and seasonal components for more focused modeling. Among non-decoupled methods, DGCRN performs well, leveraging graph convolutions to capture spatial dependencies and recurrent layers for temporal dependencies.

② **HyperD outperforms all decoupled methods, showcasing the advantage of explicitly modeling periodic patterns.** HyperD reduces average MAE by 22.63% and 23.27% compared to CycleNet-D and CycleNet-W, respectively. This highlights that HyperD’s hybrid periodic pattern, aided by spatial-temporal modeling, provides a more comprehensive representation of multi-scale periodic patterns compared to the single-scale design in CycleNet.

Ablation Study

To answer Q2, we conduct ablation studies in the PEMS04 and PEMS08 datasets to assess the contribution of each module in HyperD.

③ **Each module in HyperD significantly contributes to its overall performance.** Removing any module leads to a noticeable decrease in accuracy, highlighting the importance

Method	Dataset	PEMS03			PEMS04			PEMS07			PEMS08		
	Metric	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE
Spatial-Temporal Prediction Methods	STGCN	17.49	30.12	17.15%	22.70	35.55	14.59%	25.38	38.78	11.08%	18.02	27.83	11.40%
	DCRNN	18.18	30.31	18.91%	24.70	38.12	17.12%	25.30	38.58	11.66%	17.86	27.83	11.45%
	GWNet	19.85	32.94	19.31%	25.45	39.70	17.29%	26.85	42.78	12.12%	19.13	31.05	12.68%
	ASTGCN	17.69	29.66	19.40%	22.93	35.22	16.56%	28.05	42.57	13.92%	18.61	28.16	13.08%
	MTGNN	17.23	25.89	17.35%	19.98	31.92	14.13%	23.92	35.86	12.43%	15.03	23.89	10.23%
	STGOE	16.50	27.84	16.69%	20.84	32.82	13.77%	22.99	37.54	10.14%	16.81	25.97	10.62%
	ST-WA	15.17	26.63	15.83%	19.06	31.02	12.52%	20.74	34.05	8.77%	15.41	24.62	9.94%
	DGCRN	<u>14.74</u>	25.97	<u>15.42%</u>	18.80	30.65	12.82%	20.48	33.25	9.06%	14.60	24.16	9.33%
STPGNN	14.87	<u>25.89</u>	15.54%	18.86	30.13	13.14%	21.77	35.28	9.37%	14.69	23.85	9.55%	
Spatial-Temporal Decoupling Methods	D ² STGNN	15.10	26.57	15.23%	<u>18.42</u>	<u>29.97</u>	12.81%	<u>19.68</u>	<u>33.24</u>	8.43%	14.35	24.18	9.33%
	STWave	14.84	26.20	15.86%	18.57	30.24	<u>12.57%</u>	19.72	33.72	<u>8.19%</u>	<u>13.84</u>	<u>23.60</u>	<u>9.19%</u>
	CycleNet-D	17.54	27.48	20.02%	23.34	36.11	16.94%	25.55	39.82	11.82%	18.89	29.20	13.42%
	CycleNet-W	18.34	28.57	20.25%	24.29	38.13	18.00%	24.99	40.25	11.01%	18.27	29.83	11.89%
	STDN	16.05	27.51	17.71%	18.67	30.92	13.16%	22.94	36.06	10.32%	14.79	24.60	10.26%
	HyperD	14.69	23.66	15.76%	18.20	29.94	12.44%	19.37	33.22	8.05%	13.59	23.24	8.91%

Table 1: Performance comparison of all models on four real-world datasets. The best results are highlighted in **bold**, while the second-best results are underlined. All results are the average value of 5 repetitions.

Dataset	PEMS04			PEMS08		
	MAE	RMSE	MAPE	MAE	RMSE	MAPE
w/o LDE	18.30	30.02	12.59%	13.64	23.37	9.01%
w/o LWE	19.80	31.37	13.92%	15.83	25.05	10.43%
w/o STAE	21.07	33.46	15.82%	15.97	25.96	11.12%
w/o STFE	25.24	41.07	17.57%	18.47	31.27	12.12%
Re-MLP	18.57	30.52	12.75%	13.73	23.42	9.07%
w/o DVA	18.47	30.16	12.61%	13.70	23.44	9.08%
HyperD	18.20	29.94	12.44%	13.59	23.24	8.91%

Table 2: Ablation study results in the PEMS04 and PEMS08 datasets.

of each component. As shown in Table 2, we evaluate several variants: **w/o LDE**, **w/o LWE**, **w/o STAE**, **w/o STFE**, **Re-MLP** (replacing STFE with a simple MLP), and **w/o DVA**.

Notably, removing the LWE leads to greater degradation than LDE. For example, in the PEMS04 dataset, the MAE rises from 18.20 to 19.80 without LWE, whereas it increases only to 18.30 without LDE. This is because weekly patterns capture longer periodic structures, including both weekday and weekend dynamics, thereby providing richer information. Additionally, removing either STFE or STAE significantly degrades performance. For instance, in the PEMS04 dataset, removing the STAE increases the MAE from 18.20 to 21.07, while removing the STFE leads to an even larger increase to 25.24. This suggests that the spatial-temporal modeling of the residual component plays a more critical role than that of the periodic patterns, which captures sudden and irregular fluctuations that are not explained by regular periodic patterns.

④ **Our statistical prior initialization outperforms all**

Dataset	PEMS04			PEMS08		
	MAE	RMSE	MAPE	MAE	RMSE	MAPE
Zero	18.49	30.22	12.79%	14.21	23.90	9.64%
Uniform	18.48	30.28	12.67%	13.80	23.45	9.19%
Normal	18.76	30.48	13.17%	13.89	23.59	9.26%
Xavier	18.43	30.20	12.72%	13.82	23.47	9.22%
He	18.45	30.19	12.70%	13.74	23.33	9.19%
HyperD	18.20	29.94	12.44%	13.59	23.24	8.91%

Table 3: Performance comparison of different initialization strategies in the PEMS04 and PEMS08 datasets.

other initialization strategies, effectively learning meaningful periodic embeddings. As shown in Table 3, HyperD consistently achieves the best performance across both the PEMS04 and PEMS08 datasets. This superior performance, reflected in the lower MAE, RMSE, and MAPE values, demonstrates that statistical prior initialization significantly enhances model convergence and the ability to learn accurate periodic patterns. Compared to other initialization strategies, HyperD ensures faster and more stable training, leading to better overall results. Detailed descriptions of these initialization strategies are provided in Appendix C.4.

Reliability and Computational Performance

To answer Q3, we evaluate the reliability and computational performance of HyperD under real-world disturbances and measure its efficiency in terms of computational time and memory consumption.

⑤ **HyperD outperforms all other models, maintaining the lowest performance drop under various traffic perturbations, highlighting its superior robustness.** We simulate three types of traffic perturbations to evaluate the




Figure 3: Robustness testing in the PEMS03 datasets.

Model	Max Mem.(GB)	Epoch Time(s)	Infer Time(CPU)
D ² STGNN	17.90	57.29	8.00
STWave	10.28	49.65	7.11
STDN	15.65	59.20	7.10
HyperD	1.64	7.31	1.49
<i>Reduction</i>	6.27×↓	6.79×↓	4.77×↓

Table 4: Efficiency comparison of HyperD and other decoupled models in the PEMS04 dataset with a batch size of 64.

robustness of HyperD: (1) **Sudden surge** (values $\times 1.5$), (2) **Sudden interruption** (values set to zero), (3) **Segment shuffling** (values over four time steps reordered). Results for PEMS03 are shown in Figure 3, and additional results and experimental settings are provided in Appendix C.5. While all models suffer performance degradation under these disturbances, HyperD consistently outperforms the other decoupled models, achieving the lowest MAE and RMSE across all types of perturbations. Notably, HyperD exhibits the smallest relative performance drop, indicating its robustness in maintaining performance even when faced with different traffic disruptions. In contrast, D²STGNN and STDN show significant performance declines, particularly under the sudden interruption perturbation. Although STWave performs better than the others, it still falls short of HyperD in handling disruptions, confirming the advantage of our approach in handling real-world traffic disturbances.

⑥ **HyperD demonstrates clear advantages in memory usage, training speed, and inference time, confirming its efficiency for real-time applications.** We assess the computational performance of HyperD by comparing it with other decoupled models on the PEMS04 dataset, focusing on three key metrics: **Maximum Memory** (peak GPU usage), **Epoch Time** (average training time per epoch), **Infer Time** (average prediction time per instance). As presented in Table 4, HyperD demonstrates clear advantages in computational efficiency. It requires significantly less GPU memory than other models, trains faster per epoch, and achieves lower inference latency on the CPU. These improvements reflect the lightweight design of HyperD and the effectiveness of its architectural simplifications, making it well-suited for real-world and resource-constrained applications.

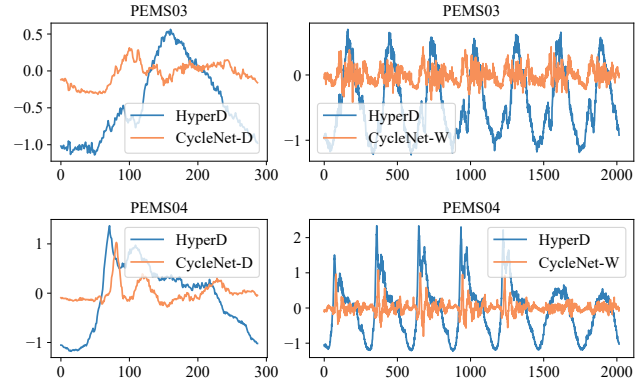


Figure 4: Visualization of the daily and weekly embeddings learned by HyperD, as well as the daily embedding from CycleNet-D and the weekly embedding from CycleNet-W, in the PEMS03 and PEMS04 datasets.

Visualization of Periodic Embeddings

Although not central to our main research questions, we also investigate whether HyperD can learn more expressive periodic representations compared to prior methods. ⑦ **HyperD captures more expressive and detailed periodic patterns, setting a new standard in modeling temporal dynamics and outperforming existing methods.** As shown in Figure 4, HyperD successfully learns detailed daily and weekly embeddings, whereas CycleNet-D and CycleNet-W yield compressed ones. The embeddings learned by HyperD are more stretched and expanded along the temporal axis, capturing clearer and more detailed periodic patterns. This superior expressiveness, enabling clearer periodic patterns, stems from HyperD’s spatial-temporal modeling that facilitates long-range interactions.

Conclusion

In this paper, we introduce HyperD, a novel Hybrid Periodicity Decoupling Framework for traffic forecasting. HyperD decouples traffic data into periodic and residual components, each processed by the Hybrid Periodic Representation Module and the Frequency-Aware Residual Representation Module, respectively. We further introduce a Dual-View Alignment Loss to promote effective and thorough decoupling between the two components. Extensive experiments on four real-world datasets demonstrate that HyperD surpasses state-of-the-art methods in forecasting accuracy, robustness, and computational efficiency, highlighting the importance of explicitly modeling periodic patterns.

Acknowledgments

This work was supported by a grant from the National Natural Science Foundation of China under grants (No. 62372211, 62272191), the Science and Technology Development Program of Jilin Province (No. 20250102216JC), and the China Postdoctoral Science Foundation under Grant Number (2025M771587).

References

- Bai, L.; Yao, L.; Li, C.; Wang, X.; and Wang, C. 2020. Adaptive graph convolutional recurrent network for traffic forecasting. *Advances in neural information processing systems*, 33: 17804–17815.
- Cao, L.; Wang, B.; Jiang, G.; Yu, Y.; and Dong, J. 2025. Spatiotemporal-aware Trend-Seasonality Decomposition Network for Traffic Flow Forecasting. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 11463–11471.
- Cirstea, R.-G.; Yang, B.; Guo, C.; Kieu, T.; and Pan, S. 2022. Towards spatio-temporal aware traffic time series forecasting. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, 2900–2913. IEEE.
- Fang, Y.; Qin, Y.; Luo, H.; Zhao, F.; Xu, B.; Zeng, L.; and Wang, C. 2023. When spatio-temporal meet wavelets: Disentangled traffic forecasting via efficient spectral graph attention networks. In *2023 IEEE 39th international conference on data engineering (ICDE)*, 517–529. IEEE.
- Fang, Z.; Long, Q.; Song, G.; and Xie, K. 2021. Spatial-temporal graph ode networks for traffic flow forecasting. In *Proceedings of the 27th ACM SIGKDD conference on knowledge discovery & data mining*, 364–373.
- Fu, L.; Deng, B.; Huang, S.; Liao, T.; Zhang, C.; and Chen, C. 2025. Learn from Global Rather Than Local: Consistent Context-Aware Representation Learning for Multi-View Graph Clustering. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2025, Montreal, Canada, August 16-22, 2025*, 5145–5153.
- Gao, H.; Jiang, R.; Dong, Z.; Deng, J.; Ma, Y.; and Song, X. 2024. Spatial-temporal-decoupled masked pre-training for spatiotemporal forecasting. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 3998–4006.
- Guo, S.; Lin, Y.; Feng, N.; Song, C.; and Wan, H. 2019. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, 922–929.
- He, X.; Wang, Y.; Fan, W.; Shen, X.; Juan, X.; Miao, R.; and Wang, X. 2025. Mamba-based graph convolutional networks: Tackling over-smoothing with selective state space. *arXiv preprint arXiv:2501.15461*.
- Jiang, J.; Han, C.; Zhao, W. X.; and Wang, J. 2023a. Pdformer: Propagation delay-aware dynamic long-range transformer for traffic flow prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 4365–4373.
- Jiang, R.; Wang, Z.; Yong, J.; Jeph, P.; Chen, Q.; Kobayashi, Y.; Song, X.; Fukushima, S.; and Suzumura, T. 2023b. Spatio-temporal meta-graph learning for traffic forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 8078–8086.
- Kim, J.; Kim, H.; Kim, H.; Lee, D.; and Yoon, S. 2025. A comprehensive survey of deep learning for time series forecasting: architectural diversity and open challenges. *Artificial Intelligence Review*, 58(7): 1–95.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *International Conference on Learning Representations (ICLR)*.
- Kong, W.; Guo, Z.; and Liu, Y. 2024. Spatio-temporal pivotal graph neural networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 8627–8635.
- Li, F.; Feng, J.; Yan, H.; Jin, G.; Yang, F.; Sun, F.; Jin, D.; and Li, Y. 2023. Dynamic graph convolutional recurrent network for traffic prediction: Benchmark and solution. *ACM Transactions on Knowledge Discovery from Data*, 17(1): 1–21.
- Li, S.; Li, H.; Li, X.; Xu, Y.; Lin, Z.; and Jiang, H. 2025. Causal Intervention Is What Large Language Models Need for Spatio-Temporal Forecasting. *IEEE Transactions on Cybernetics*.
- Li, Y.; Yu, R.; Shahabi, C.; and Liu, Y. 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *International Conference on Learning Representations*.
- Lin, S.; Lin, W.; Hu, X.; Wu, W.; Mo, R.; and Zhong, H. 2024. CycleNet: Enhancing Time Series Forecasting through Modeling Periodic Patterns. In *Thirty-eighth Conference on Neural Information Processing Systems*.
- Liu, C.; Xiao, Z.; Long, C.; Wang, D.; Li, T.; and Jiang, H. 2024a. MVCAR: Multi-view collaborative graph network for private car carbon emission prediction. *IEEE Transactions on Intelligent Transportation Systems*.
- Liu, C.; Xu, Q.; Miao, H.; Yang, S.; Zhang, L.; Long, C.; Li, Z.; and Zhao, R. 2025a. TimeCMA: Towards LLM-Empowered Multivariate Time Series Forecasting via Cross-Modality Alignment. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 18780–18788.
- Liu, C.; Yang, S.; Xu, Q.; Li, Z.; Long, C.; Li, Z.; and Zhao, R. 2024b. Spatial-temporal large language model for traffic prediction. In *2024 25th IEEE International Conference on Mobile Data Management (MDM)*, 31–40. IEEE.
- Liu, C.; Zhou, S.; Xu, Q.; Miao, H.; Long, C.; Li, Z.; and Zhao, R. 2025b. Towards Cross-Modality Modeling for Time Series Analytics: A Survey in the LLM Era. In *IJCAI*.
- Liu, G.; Zhang, Y.; Zhang, P.; and Gu, J. 2025c. Spatiotemporal graph contrastive learning for wind power forecasting. *IEEE Transactions on Sustainable Energy*.
- Ma, Q.; Zhang, Z.; Zhao, X.; Li, H.; Zhao, H.; Wang, Y.; Liu, Z.; and Wang, W. 2023. Rethinking sensors modeling: Hierarchical information enhanced traffic forecasting. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 1756–1765.
- Miao, H.; Liu, Z.; Zhao, Y.; Guo, C.; Yang, B.; Zheng, K.; and Jensen, C. S. 2024. Less is more: Efficient time series dataset condensation via two-fold modal matching. *PVLDB*, 18(2): 226–238.
- Shao, M.; Wang, Y.; Shen, X.; and Wang, X. 2025. Enhanced Molecular Property Prediction with SMILES and Graph Aligned Contrastive Learning. In *Pacific-Asia Conference on Knowledge Discovery and Data Mining*, 117–129. Springer.

- Shao, Z.; Wang, F.; Xu, Y.; Wei, W.; Yu, C.; Zhang, Z.; Yao, D.; Sun, T.; Jin, G.; Cao, X.; et al. 2024. Exploring progress in multivariate time series forecasting: Comprehensive benchmarking and heterogeneity analysis. *IEEE Transactions on Knowledge and Data Engineering*, 37(1): 291–305.
- Shao, Z.; Zhang, Z.; Wei, W.; Wang, F.; Xu, Y.; Cao, X.; and Jensen, C. S. 2022. Decoupled Dynamic Spatial-Temporal Graph Neural Network for Traffic Forecasting. *Proceedings of the VLDB Endowment*, 15(11): 2733–2746.
- Shen, X.; Lio, P.; Yang, L.; Yuan, R.; Zhang, Y.; and Peng, C. 2024. Graph rewiring and preprocessing for graph neural networks based on effective resistance. *IEEE Transactions on Knowledge and Data Engineering*, 36(11): 6330–6343.
- Shen, X.; Liu, Y.; Dai, Y.; Wang, Y.; Miao, R.; Tan, Y.; Pan, S.; and Wang, X. 2025. Understanding the Information Propagation Effects of Communication Topologies in LLM-based Multi-Agent Systems. *arXiv preprint arXiv:2505.23352*.
- Song, C.; Lin, Y.; Guo, S.; and Wan, H. 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*, volume 34, 914–921.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Wang, B.; Wang, P.; Zhang, Y.; Wang, X.; Zhou, Z.; Bai, L.; and Wang, Y. 2024a. Towards dynamic spatial-temporal graph learning: A decoupled perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9089–9097.
- Wang, L.; He, D.; Zhang, H.; Liu, Y.; Wang, W.; Pan, S.; Jin, D.; and Chua, T.-S. 2024b. Goodat: Towards test-time graph out-of-distribution detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15537–15545.
- Wang, Y.; Liu, Y.; Liu, N.; Miao, R.; Wang, Y.; and Wang, X. 2025. AdaGCL+: An Adaptive Subgraph Contrastive Learning Towards Tackling Topological Bias. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Wang, Y.; Liu, Y.; Shen, X.; Li, C.; Ding, K.; Miao, R.; Wang, Y.; Pan, S.; and Wang, X. 2024c. Unifying unsupervised graph-level anomaly detection and out-of-distribution detection: A benchmark. *arXiv preprint arXiv:2406.15523*.
- Wu, Z.; Pan, S.; Long, G.; Jiang, J.; Chang, X.; and Zhang, C. 2020. Connecting the dots: Multivariate time series forecasting with graph neural networks. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 753–763.
- Wu, Z.; Pan, S.; Long, G.; Jiang, J.; and Zhang, C. 2019. Graph wavenet for deep spatial-temporal graph modeling. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, 1907–1913.
- Xia, J.; Yang, Y.; Shen, J.; Wang, S.; and Cao, J. 2025. FairTP: A Prolonged Fairness Framework for Traffic Prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 26391–26399.
- Yang, L.; Cai, Y.; Ning, H.; Zhuo, J.; Jin, D.; Ma, Z.; Guo, Y.; Wang, C.; and Wang, Z. 2025. Universal Graph Self-Contrastive Learning. In *IJCAI*, 3534–3542.
- Yi, K.; Zhang, Q.; Fan, W.; Wang, S.; Wang, P.; He, H.; An, N.; Lian, D.; Cao, L.; and Niu, Z. 2023. Frequency-domain mlps are more effective learners in time series forecasting. *Advances in Neural Information Processing Systems*, 36: 76656–76679.
- Yu, B.; Yin, H.; and Zhu, Z. 2018. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence*, 3634–3640. International Joint Conferences on Artificial Intelligence Organization.
- Yu, G.; Zou, J.; Hu, X.; Aviles-Rivero, A. I.; Qin, J.; and Wang, S. 2024. Revitalizing Multivariate Time Series Forecasting: Learnable Decomposition with Inter-Series Dependencies and Intra-Series Variations Modeling. In *International Conference on Machine Learning*, 57818–57841. PMLR.
- Yu, Y.; Tang, X.; Yao, H.; Yi, X.; and Li, Z. 2019. Citywide traffic volume inference with surveillance camera records. *IEEE Transactions on Big Data*, 7(6): 900–912.
- Yuan, X.; and Qiao, Y. 2024. Diffusion-TS: Interpretable Diffusion for General Time Series Generation. In *The Twelfth International Conference on Learning Representations*.
- Zhang, J.; Zheng, Y.; and Qi, D. 2017. Deep spatio-temporal residual networks for citywide crowd flows prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Zhang, Z.; Huang, Z.; Hu, Z.; Zhao, X.; Wang, W.; Liu, Z.; Zhang, J.; Qin, S. J.; and Zhao, H. 2023a. Mlpst: Mlp is all you need for spatio-temporal prediction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 3381–3390.
- Zhang, Z.; Zhao, X.; Liu, Q.; Zhang, C.; Ma, Q.; Wang, W.; Zhao, H.; Wang, Y.; and Liu, Z. 2023b. Promptst: Prompt-enhanced spatio-temporal multi-attribute prediction. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, 3195–3205.
- Zhou, Z.; Huang, Q.; Wang, B.; Hou, J.; Yang, K.; Liang, Y.; Zheng, Y.; and Wang, Y. 2025. Coms2t: A complementary spatiotemporal learning system for data-adaptive model evolution. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Zhou, Z.; Wang, Y.; Xie, X.; Chen, L.; and Zhu, C. 2020. Foresee urban sparse traffic accidents: A spatiotemporal multi-granularity perspective. *IEEE Transactions on Knowledge and Data Engineering*, 34(8): 3786–3799.