Frontiers of Information Technology & Electronic Engineering
www.jzus.zju.edu.cn; engineering.cae.cn; www.springerlink.com
ISSN 2095-9184 (print); ISSN 2095-9230 (online)
E-mail: jzus@zju.edu.cn

MltAuxTSPP: a unified benchmark for deep learning-based traffic state prediction with multi-source auxiliary data*

Yusong ZHOU[†1,2], Xiaoyu JIANG[†1,2], Shu SUN[†1,2],
Xinmin ZHANG[†‡1,2], Yuanqiu MO[3], Zhihuan SONG[1,2]

[1]*College of Control Science and Engineering, Zhejiang University, Hangzhou 310027, China*
[2]*State Key Laboratory of Industrial Control Technology, Zhejiang University, Hangzhou 310027, China*
[3]*Mathematics department, Southeast University, Nanjing 211189, China*
[†]E-mail: zhoushichan@zju.edu.cn; jiangxiaoyu@zju.edu.cn; shusun@zju.edu.cn; xinminzhang@zju.edu.cn

Received Mar. 17, 2025; Revision accepted Aug. 17, 2025; Crosschecked

Unedited

Abstract: Deep learning has empowered traffic prediction models to integrate diverse auxiliary data sources, such as weather and temporal information, for enhanced forecasting accuracy. However, existing approaches often suffer from limited generality and scalability, and the field lacks a unified benchmark for fair model comparison. This absence hinders consistent performance evaluation, slows the development of robust and adaptable models, and makes it challenging to quantify the incremental benefits of different auxiliary data sources. To address these issues, we present MltAuxTSPP, a unified benchmark framework for deep learning-based traffic state prediction with multi-source auxiliary data. The framework features a standardized data container and a fusion embedding module, which enable consistent utilization of heterogeneous data and improve scalability. It produces unified hidden representations that can be seamlessly adopted by various downstream models, ensuring fair and reproducible comparisons under identical conditions. Extensive experiments on real-world datasets demonstrate that MltAuxTSPP effectively leverages weather and temporal features to improve long-term forecasting performance and offers a practical and reproducible foundation for advancing research in traffic state prediction.

Key words: Traffic prediction; Benchmark platform; Deep learning; Multi-source auxiliary data
https://doi.org/10.1631/FITEE.2500169 **CLC number:**

1 Introduction

Accurate traffic state prediction is a cornerstone of modern intelligent transportation systems (ITSs) and is essential for applications ranging from traffic management to autonomous driving. Although deep learning models have excelled at capturing complex spatiotemporal dependencies from traffic data (Do et al., 2019; Zhu et al., 2019; Lee et al., 2021), their performance is often constrained by relying solely on traffic-related data, such as fixed position sensor data, trajectory data, network infrastructure data, and trip records data (Nagy and Simon, 2018; Shaygan et al., 2022).

Real-world traffic is significantly influenced by external factors, such as predictable temporal patterns (e.g., weekday rush hours vs. weekends) and adverse weather conditions (Essien et al., 2018; Taghipour and Parsa, 2020). Consequently, recent studies have begun to integrate such auxiliary data to

‡ Corresponding author
* Project supported by the National Key Research and the Development Program of China (No. 2022ZD0120003)
ORCID: Yusong ZHOU, https://orcid.org/0009-0003-4475-623X; Xinmin ZHANG, https://orcid.org/0000-0002-4761-3969
© Zhejiang University Press 2025

enhance prediction accuracy (Taghipour and Parsa, 2020; Essien et al., 2021). However, the integration of multi-source data remains fraught with challenges. Current approaches often handle a limited set of auxiliary variables, which typically lack spatial dimensions. Furthermore, these methods generally lack the generality and scalability required to accommodate diverse and growing data sources. Crucially, the field is hampered by the absence of a unified benchmark platform for systematically evaluating and comparing models under multi-source data conditions.

These limitations collectively pose significant obstacles to advancing the field. The lack of generality hinders the application of existing methods to new datasets or information streams. Poor scalability makes it challenging to handle the increasing volume and variety of data. Crucially, the absence of a unified benchmark leads to fragmented research, impedes objective comparison of models, and ultimately slows ITS innovation.

To overcome these obstacles, we introduce MltAuxTSPP, a novel benchmarking framework for traffic prediction using multi-source auxiliary data. At its core, MltAuxTSPP features a unified data container (UDC) and a versatile fusion embedding module (Fig. 1). This design addresses critical challenges in data handling and scalability, enabling seamless integration with various existing spatiotemporal models. This allows researchers to fairly evaluate and compare different model architectures on a standardized basis. The main contributions of this work are threefold:

We design and propose MltAuxTSPP, a general and scalable framework for enhancing traffic prediction with multi-source data. Its fusion module standardizes the integration of diverse data types for various downstream models.

We construct and release two new multi-source traffic forecasting datasets by augmenting the widely used METR-LA and PEMS-BAY benchmarks with extensive weather and temporal features, facilitating standardized evaluation.

We conduct extensive experiments to systematically evaluate the impact of auxiliary data on state-of-the-art models and provide novel insights into the significant benefits of multi-source data for long-term forecasting.

The rest of the paper is organized as follows. Section 2 reviews relevant research about traffic forecasting, the integration of multi-source data, and benchmark platforms. Section 3 introduces the details of our method, including task formulation, the platform framework, and the pipeline of multi-source data fusion. Section 4 reports the predictive performance of models on real-world traffic datasets under our benchmarking platform, including the design of experiments, prediction results analysis, and auxiliary variable importance. Section 5 is the conclusion of the paper.

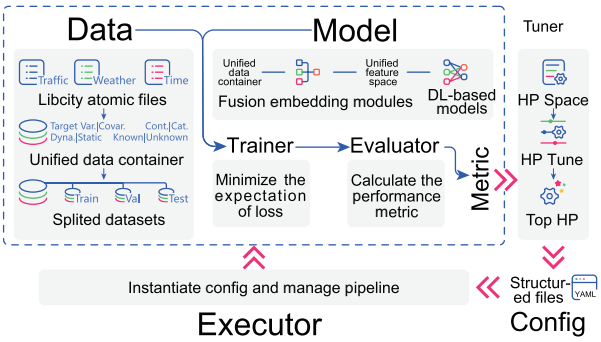

Fig. 1 Benchmarking system in MltAuxTSPP

2 Related work

Traffic forecasting models can be broadly divided into temporal modeling and spatiotemporal modeling. A wide array of temporal models has been applied to traffic prediction, progressing from classical methods (e.g., HA, ARIMA (?Yang et al., 2020)) to sophisticated deep learning architectures like long short-term memory (LSTM) and transformers, which have proven highly effective at capturing complex traffic dynamics (?Ma et al., 2015; Vaswani et al., 2017; Zhou et al., 2021). Spatiotemporal models incorporate spatial dependencies primarily through convolutional neural networks (CNNs) or graph convolutional networks (GCNs), which led to foundational architectures like spatiotemporal GCNs (Yu et al., 2018; Song et al., 2020). Current research, however, focuses on more advanced techniques for simultaneous spatiotemporal modeling. Key innovations include learning dynamic graph structures with adaptive matrices (e.g., Graph WaveNet; (Wu et al., 2019)), modeling continuous dynamics with ordinary differential equations (e.g., STGODE; (Fang et al., 2021)), and developing unified attention mechanisms that concurrently capture spatiotemporal correlations (e.g., T-

Unedited

Graphormer; (Bai and Liu, 2025)).

Although spatiotemporal models excel at learning from traffic data that are commonly sourced from sensor networks like PeMS[1] ((Yu et al., 2018; Guo et al., 2019; Song et al., 2020)), their performance is inherently limited by overlooking external factors. Consequently, although many studies now incorporate auxiliary data such as weather and time features (see more in the Supplementary Materials), their fusion methodologies exhibit critical bottlenecks. Specifically, current methods often rely on simplistic techniques like concatenation, failing to capture complex interdependencies. They also tend to treat auxiliary data as global variables, thereby ignoring crucial spatial heterogeneity (e.g., localized weather effects). Furthermore, these approaches typically lack the scalability and generalizability required to handle diverse or novel auxiliary data streams. These challenges underscore the fact that effective multi-source data utilization requires more than just data availability; advanced, scalable, and spatially aware fusion strategies are paramount.

As shown in Table 1, although several valuable benchmark platforms exist, they are not tailored for multi-source traffic prediction. Platforms like TSPP, TSLib, and BasicTS$^+$ focus on general time-series forecasting (Wu et al., 2023; Baczek et al., 2024; Shao et al., 2025). Conversely, even comprehensive traffic-specific libraries like LibCity (Wang et al., 2021) primarily handle core traffic data and timestamps, offering limited support for integrating diverse auxiliary data sources.

Consequently, a critical gap remains: there is no unified platform designed to systematically evaluate how different models leverage rich auxiliary data (e.g., weather, events) in traffic forecasting. This absence hinders the objective comparison of data fusion strategies and slows progress in developing truly robust models. To address this, we developed MltAuxTSPP as a unified and extensible environment specifically for this purpose.

[1] https://pems.dot.ca.gov/

3 Multi-source auxiliary data fusion framework

3.1 Task formulation

In this work, the task of traffic state prediction with multi-source auxiliary data uses historical traffic states to forecast future traffic states, incorporating auxiliary data from both historical and future perspectives. Traffic state refers to the conditions of traffic flow on a road network, typically characterized by traffic volume, speed, density, and congestion levels. When no auxiliary data are used, the task reduces to single-source spatiotemporal forecasting. If spatial dependencies are ignored or flattened into features, it further simplifies to time-series forecasting.

Given the historical target variable $\boldsymbol{X} \in \mathbb{R}^{T\times n\times f}$, auxiliary variables $\boldsymbol{C} = \{\boldsymbol{C}^{(1)}, \ldots, \boldsymbol{C}^{(m)}\}$ and the road netowrk $\boldsymbol{G} = (\boldsymbol{N}, \boldsymbol{E})$, the task is to predict a future sequence $\boldsymbol{X}^* \in \mathbb{R}^{h\times n\times f}$, where T is the historical input length, and h is the future horizon, n is the road network node count, f is the feature count, and m is the number of auxiliary variables. Each auxiliary variable $\boldsymbol{C}^{(j)} \in \mathbb{R}^{T^{(j)}\times n^{(j)}\times f^{(j)}}$ is characterized by its own temporal span $T^{(j)}$, node count $n^{(j)}$, and feature count $f^{(j)}$. The graph $\boldsymbol{G}$ with node set $\boldsymbol{N}$ and edge set $\boldsymbol{E}$ is encoded as an adjacency matrix $\boldsymbol{A} \in \mathbb{R}^{n\times n}$. Entries A_{ij} indicate either binary connectivity or weighted relationships (e.g., distance, travel time) between nodes N_i and N_j. Notably, $\boldsymbol{A}$ is treated as one of the auxiliary variables in $\boldsymbol{C}$.

The task uses the model $\boldsymbol{M} : \{\boldsymbol{X}, \boldsymbol{C}\} \rightarrow \mathbb{R}^{t\times n\times f}$ to predict the future traffic state $\boldsymbol{X}^*$ and minimize the expected error $\mathbb{E}[\mathcal{L}(\boldsymbol{X}^*, \widehat{\boldsymbol{X}}^*)]$, where $\mathcal{L}$ measures the discrepancy between prediction $\boldsymbol{X}^*$ and ground truth $\widehat{\boldsymbol{X}}^*$.

3.2 Data collection and preprocessing

Our study leverages two standard traffic benchmark datasets, METR-LA and PEMS-BAY (Li et al., 2018), augmented with external data to investigate their impact. The auxiliary data includes three global weather datasets and three discrete temporal features: minute of the hour (MoH), hour of the day (HoD), and day of the week (DoW). To unify the spatiotemporal heterogeneity of these sources, we developed a preprocessing pipeline. First, all

Unedited

Table 1 Comparison of benchmark platforms for time-series and traffic prediction

Framework	Primary Focus	Key Features / Scope	Auxiliary Data Support
TSPP	General Time-Series	Includes 9 models (e.g., Temporal Fusion Transformers (Lim et al., 2021)) and some traffic datasets.	Limited
TSLib	General Time-Series	Provides advanced models like Informer and TimesNet.	Limited
BasicTS$^+$	General Time-Series	Focuses on multi-variable, spatio-temporal, and long-term forecasting.	Limited
LibCity	Traffic Prediction	Comprehensive library with 60 models and 35 datasets for various traffic tasks.	Limited
Ours	**Multi-Source Traffic Prediction**	**Unified environment with data containers and fusion modules for diverse auxiliary inputs.**	**Yes**

Table 2 Summary of datasets

Source	Interval	METR-LA			PEMS-BAY		
		Steps	Nodes	Variables	Steps	Nodes	Variables
Traffic	5 min	34272	207	2	52116	325	2
ERA5-Land			10	10		9	10
ERA5-HEAT			4	2		2	2
GSH			8	1		5	1

weather data are temporally interpolated to a 5-min frequency to match the traffic data. Second, all datasets are spatially converted to point-based representations, with weather data aligned to the nearest traffic sensor. Finally, we developed two enriched datasets in Table 2, and detailed variable descriptions are provided in the Supplementary Materials.

3.3 Framework of the benchmarking system

We designed the benchmarking system in Fig. 1 with the primary goal of integrating multi-source data and fairly comparing the performance of different models. The system adopts a top-down design approach that divides the entire architecture into several modules, each responsible for a specific function to achieve modularity. This design approach not only enhances the maintainability and scalability of the system but also improves its flexibility and comprehensibility.

The system modules are categorized into two main types: core modules and auxiliary modules. The core modules include the data module, model module, trainer module, evaluator module, and executor module, whereas the auxiliary modules consist of the config and tuner modules. The core modules are implemented in Python, and the interfaces and behaviors of the modules are defined through abstract base classes. Through derived classes, these modules can adapt to a variety of scenario requirements. Specifically, the functions of each module are shown in Table 3.

The core and auxiliary modules are interconnected through the executor and Hydra (Yadan, 2019). The executor instantiates the modules based on the configuration files and invokes their interfaces to complete the training and evaluation tasks, then feeds the results back to Optuna for further optimization.

In this work, we address the fusion of multi-source auxiliary data using the pipeline shown in Fig. 2. First, we store the multi-source data in a UDC. Then, the data is transformed into a format suitable for model input through a fusion embedding module. The output of the module is a hidden representation that integrates the multi-source data, ensuring compatibility with the input requirements of subsequent models. A more detailed explanation of this process will be provided in the following sections.

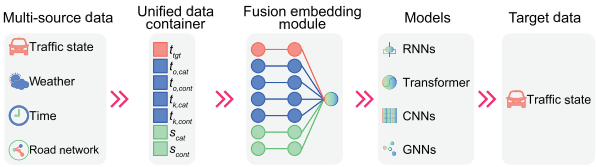

Fig. 2 Multi-source auxiliary data fusion framework for traffic state prediction

3.4 UDC

To manage and unify multi-source data, we use a UDC to store traffic state datasets with auxiliary information, which is an extension of the data aspects in TSPP (Baczek et al., 2024). The dataset can be categorized into the seven components in Table 4 based on the following attributes. Notably, the target variable is dynamic, unknown, and can be either continuous or categorical, whereas covariates may include multiple variables.

The data in our study is classified along four dimensions. By source, we distinguish between the target variable to be predicted and covariates. By temporal nature, data is either static (e.g., network topology) or dynamic. With respect to future availability, inputs can be known at prediction time (e.g., timestamps) or unknown (the target value, weather). Finally, by data type, variables are either continuous

Unedited

Table 3 Description of the framework modules

Module	Description
Data	Responsible for preparing the datasets required by the models, including data integration, filtering, cleaning, and constructing UDCs.
Model	Provides a series of reproducible or custom models to meet the requirements of prediction tasks.
Trainer	Responsible for selecting optimization strategies and customizing relevant decisions to improve model performance.
Evaluator	Assesses the accuracy and effectiveness of the model predictions.
Executor	Controls the instantiation of modules based on the configuration and executes the training, inference, and evaluation processes using standardized interfaces.
Config	Manages all framework-related parameters centrally through YAML files, providing structured and flexible control and adaptability.
Tuner	Uses Optuna (Akiba et al., 2019) to select and apply the optimal configuration parameters after each module execution to optimize model performance.

UDC, unified data container.

Table 4 Components of UDC

Components	Source	Temporal	Observation	Data types
t_{tgt}	Target	Dynamic	Unknown	Continuous Categorical
s_{cont}	Covariates	Static	Known	Continuous
s_{cat}			Known	Categorical
$t_{k,cont}$		Dynamic	Known	Continuous
$t_{k,cat}$			Known	Categorical
$t_{o,cont}$			Unknown	Continuous
$t_{o,cat}$			Unknown	Categorical

UDC, unified data container.

(real-valued) or categorical (discrete).

3.5 Fusion embedding modules

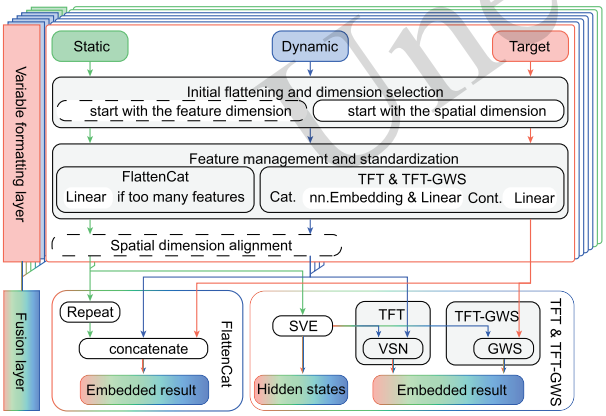

Fig. 3 Fusion embedding modules

To improve the integration of multi-source data and facilitate model reuse, we propose an abstract fusion embedding module designed specifically for handling such data. This module acts as an intermediary between the UDC and the model, with the primary objective of projecting multi-source data into a unified feature space for input into subsequent predictive models. We present three concrete instances of this abstract module: FlattenCat, TFT, and TFT-GWS. FlattenCat standardizes variables through linear transformation, TFT (Lim et al., 2021) incorporates non-linear processing and variable selection, and TFT-GWS is a simplified version of TFT that retains the core concepts while reducing computational overhead.

The abstract module consists of two main components: a variable formatting layer and a fusion layer. The variable formatting layer processes each variable through linear transformations, converting them into standardized variables for subsequent analysis. The fusion layer then integrates these standardized variables along the feature dimension, producing a unified hidden representation that aligns with the model input specifications.

Consider the target variable $\boldsymbol{V}^{\text{tgt}}$ and a set of p variables $\boldsymbol{V} = \{\boldsymbol{V}^1, \ldots, \boldsymbol{V}^p\}$. The target variable $\boldsymbol{V}^{\text{tgt}} \in \mathbb{R}^{L^{\text{tgt}} \times S^{\text{tgt}} \times F^{\text{tgt}}}$ has L^{tgt} time steps, S^{tgt} spatial nodes, and F^{tgt} features. Each variable $\boldsymbol{V}^i \in \mathbb{R}^{L^i \times S^i \times F^i \times \cdots}$ is multidimensional, where L^i is the temporal length, S^i is the number of spatial nodes, and F^i is the number of features. The ellipsis $(\ldots)$ denotes additional higher dimensions. For simplicity, we define the temporal window length of static variables as 1, dynamic variables in the historical dataset as T, and dynamic variables in the prediction window as h. Therefore, L^{tgt} and L^i can take values from $\{1, T, h\}$.

The Fusion Embedding Module processes each auxiliary data source $\boldsymbol{V}^i$ through a three-stage pipeline to produce a standardized tensor $\boldsymbol{\xi}^i$ compatible with downstream models. The stages are as follows:

1. Initial flattening and dimension selection: First, each high-dimensional input tensor $\boldsymbol{V}^i \in \mathbb{R}^{S^i \times F^i \times \cdots}$ is vectorized. This can be configured in two ways: (a) preserving the spatial dimension to yield a tensor $\boldsymbol{V'}^i \in \mathbb{R}^{S^i \times F'^i}$, or (b) collapsing the spatial dimension to yield a flat feature vector $\boldsymbol{V'}^i \in \mathbb{R}^{F'^i}$.

2. Feature dimension alignment: Next, the resulting feature dimension F'^i is projected to a target model dimension F_m. We provide two methods for this projection from $\boldsymbol{V'}^i_f$ to $\widetilde{V}^i_f$. (1) FlattenCat

Unedited

conditionally applies a linear projection to reduce dimensionality as defined in Eq. (1). (2) TFT, inspired by Temporal Fusion Transformers processes categorical and continuous features separately (using embeddings or linear layers) before projecting them to the target dimension F_m, as detailed in Eqs (2) and (3).

$$\widetilde{\boldsymbol{V}}_f^i = \begin{cases} FC(\boldsymbol{V'}_f^i) \in \mathbb{R}^{F_m} & if F'^i \geq F_m, \\ \boldsymbol{V'}_f^i \in \mathbb{R}^{F'^i} & if F'^i < F_m. \end{cases} \tag{1}$$

$$\widetilde{\boldsymbol{V}}_f^i \in \mathbb{R}^{F_m} = FC(Flatten(Embed(\boldsymbol{V'}_{f,cat}^i))), \text{ where } Embed \text{ is nn.Embedding} \tag{2}$$

$$\widetilde{\boldsymbol{V}}_f^i \in \mathbb{R}^{F_m} = FC(\boldsymbol{V'}_{f,cont}^i). \tag{3}$$

3. Spatial dimension alignment (optional): Finally, if the spatial dimension S^i was preserved in the first stage, it is aligned to the target model's required spatial dimension S^{tgt} via a linear layer ($FC : \mathbb{R}^{S^i} \to \mathbb{R}^{S^{tgt}}$). This step is essential for downstream spatio-temporal models (e.g., T-GCN) that require spatially consistent inputs, but it is bypassed for purely temporal models (e.g., LSTM).

In the fusion layer, the three modules diverge in their methodologies. Consider the feature dimension of the formatted variable $\boldsymbol{\xi}_f^i \in \mathbb{R}^{F_f^i}$. FlattenCat extends the static values along the temporal axis and combines the formatted outputs from all variables to form a unified hidden representation $\boldsymbol{H_f} = [\boldsymbol{\xi}_f^1, \ldots, \boldsymbol{\xi}_f^p] \in \mathbb{R}^{F_H}$, where F_H is the sum of the individual feature dimensions, calculated as $F_H = \sum_{i=1}^p F_f^i = \sum_{i=1}^p \min(F^i, F_m)$.

TFT processes formatted static variables using a static variable encoder (SVE), as shown in Fig. 5c. The SVE incorporates a variable selection network (VSN) and four gated residual networks (GRN) to generate four static encodings: $\boldsymbol{c_s}$, $\boldsymbol{c_e}$, $\boldsymbol{c_h}$, and $\boldsymbol{c_c}$, as shown in Figs 5a and 5d. Subsequently, TFT employs an additional VSN to process dynamic variables, including the target, and combines these processed dynamic variables with the static encoding $\boldsymbol{c_s}$.

The VSN implementation within TFT, crucial for identifying salient variables, follows three main steps, whose logical relationship is key to its function:

Individual feature transformation: each formatted variable $\boldsymbol{\xi}_f^i$ is individually processed by a GRN, which applies a gating mechanism to filter out irrelevant variables for the prediction task. That is, each GRN generates a transformed variable $\widetilde{\boldsymbol{\xi}}_f^i \in \mathbb{R}^{F_m}$, as shown in Eq. (4).

$$\widetilde{\boldsymbol{\xi}}_f^i \in \mathbb{R}^{F_m} = GRN(\boldsymbol{\xi}_f^i, \boldsymbol{c_s}). \tag{4}$$

Variable weight generation: to determine the relevance of each variable, the VSN computes dynamic attention weights. As described in Eq. (5), all original formatted input variables $\{\boldsymbol{\xi}_f^1, \ldots, \boldsymbol{\xi}_f^p\}$ are first stacked and flattened into a single vector $\boldsymbol{\Xi}_f \in \mathbb{R}^{p*F_m}$. This vector is then passed through another GRN with a context vector $\boldsymbol{c_s}$ if providing static variables, and the output is fed into a Softmax function. This yields a set of dynamic weights $\widetilde{\boldsymbol{W}} \in \mathbb{R}^p = \widetilde{\boldsymbol{w}}^1, \ldots, \widetilde{\boldsymbol{w}}^p$, where each $\widetilde{\boldsymbol{w}}^i$ corresponds to the importance of the i-th variable. These weights are data-dependent, meaning they are recalculated for each input instance.

$$\begin{aligned} \boldsymbol{\Xi}_f \in \mathbb{R}^{p*F_m} &= Flatten([\boldsymbol{\xi}_f^1, \ldots, \boldsymbol{\xi}_f^p]) \\ \widetilde{\boldsymbol{W}} \in \mathbb{R}^p &= Softmax(GRN(\widetilde{\boldsymbol{\Xi}}_f)) \\ &= \{\widetilde{\boldsymbol{w}}^1, \ldots, \widetilde{\boldsymbol{w}}^p\}. \end{aligned} \tag{5}$$

Weighted summation: finally, the VSN produces its output $\boldsymbol{H_f}$ by performing a weighted summation of the transformed features $\widetilde{\boldsymbol{\xi}}_f^i$ (from Eq. (4)), using the generated variable weights $\widetilde{\boldsymbol{w}}^i$ (from Eq. (5)). This process, illustrated in Eq. (6), effectively selects and combines the most relevant features.

$$\boldsymbol{H_f} = \sum_{i=1}^p \widetilde{\boldsymbol{w}}^i \widetilde{\boldsymbol{\xi}}_f^i. \tag{6}$$

In practice, we observed that TFT's VSN, particularly the GRNs within it, contributes significantly to memory usage, especially with a large number of variables. To address this inefficiency, we introduced the gated weighted summation (GWS) module, as depicted in Fig. 5b. The GWS module modifies the VSN architecture for improved efficiency. The key distinctions between the GWS module and TFT's VSN are as follows: Modification to individual feature transformation (cf. Eq. (4)): Replaces the computationally intensive GRN with a simple gated unit, GELU activation function. The transformation becomes $\widetilde{\boldsymbol{\xi}}_f^i = \text{GELU}(\boldsymbol{\xi}_f^i)$. This significantly reduces parameters and computational load compared to a full GRN.

Unedited

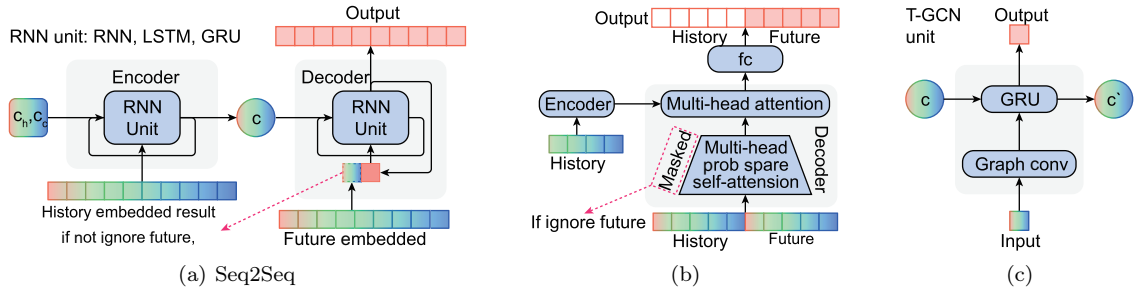

(a) Seq2Seq (b) (c)

Fig. 4 Fusion embedding modules and models: (a) Seq2Seq; (b) Informer; (c) T-GCN unit. GCN, graph convolutional network

Modification to variable weight generation (cf. Eq. (5)): Simplifies weight generation by making weights data-independent. It uses a learnable (trainable) parameter vector $\boldsymbol{W} \in \mathbb{R}^p$. The variable selection weights $\widetilde{\boldsymbol{W}}$ are then obtained by applying Softmax directly to this vector: $\widetilde{\boldsymbol{W}} = \mathrm{Softmax}(\boldsymbol{W})$. These weights are learned during the training process and remain fixed thereafter for inference, rather than being re-computed for each input.

No change in weighted summation (cf. Eq. (6)).

3.6 Models with fusion embedding results

We investigate the integration of embedding modules in canonical time series and spatiotemporal prediction models, focusing on three key models. We select Seq2Seq (Sutskever et al., 2014) as a representative of RNNs and Informer (Zhou et al., 2021) as an example of Transformer-based approaches. Given the graph-structured nature of our data, we choose T-GCN (Zhao et al., 2020) as the main spatiotemporal model and adapt its architecture like Seq2Seq for multi-step forecasting, as shown in Fig. 4c. We also choose GraphWaveNet (Wu et al., 2019).

In our model design, we observe that the original Seq2Seq and Informer models intentionally disregard or occlude variables within the future window. This design choice stems from their primary use in forecasting a single unknown variable, for which the true value in the future window is unavailable. However, in traffic-state prediction tasks that involve multiple auxiliary variables, covariates from various data sources within the future window are accessible and can potentially improve predictive accuracy. To address this, we refine the model to better utilize information from the future window as shown in Fig. 4. Specifically, we use both the hidden state and the output from the embedding module as the input to the decoder of Seq2Seq instead of only the hidden state. For Informer, we remove the mask in the decoder's first layer, allowing it to receive information from the future window. Additionally, when extracting samples from the dataset for the future window, we assign a value of 0 to unknown variables to simulate the uncertainty encountered in real-world scenarios.

Unedited

4 Experiments

4.1 Experimental setup

We conduct experiments on the augmented PEMS-BAY and METR-LA datasets detailed in Subsection 3.2. All input variables are normalized using Z-score, and each dataset is divided into training, validation, and test sets using a 70%:20%:10% ratio. Model performance is evaluated using mean absolute error (MAE) and mean squared error (MSE).

All experiments are conducted on a server equipped with an NVIDIA RTX 4090 GPU (24GB) and an Intel Xeon Gold 6430 CPU, running Ubuntu 22.04. Our framework is built with Python 3.12.3 and PyTorch 2.3.0 (CUDA 12.1). To enhance training efficiency, we utilize automatic mixed precision (AMP) via the `torch.amp` package

Our experimental process consists of two main stages: parameter tuning and final evaluation. To ensure a fair comparison, we first tune hyperparameters for baseline models (Seq2Seq, Informer, and T-GCN). In this stage, models are trained for up to three epochs using only traffic variables as input, consistent with their original designs. We use the Adam optimizer with a learning rate of 0.001 and

Table 5 Experimental conditions for model evaluation

Condition	Settings
Number of Nodes (N_s)	1-1, 4-1, 16-x, all($n-i$ denotes n nodes with index intervals i, and $x \in \{1,2,4,8\}$)
Time Windows (L_w)	$L_w \in \{12,24,48,96\}$ (input and output windows are identical)
Auxiliary Variables (Co.)	Traffic-only ($covariates = No$); All auxiliary variables ($covariates = All$)
Embedding Modules (E.M.)	FlattenCat, TFT, TFT-GWS
Model Perception of Future Window	With or without future window information
Model After Fusion Embedding (M.)	Seq2Seq, Informer, T-GCN, GWNet

GWNet, Graph WaveNet; GCN, graph convolutional network; GWS, gated weighted summation.

a batch size of 512, and employ early stopping to prevent overfitting. For the main evaluation, we assess model performance under the various conditions summarized in Table 5. To ensure robustness, each experiment is performed 10 times with different random seeds. Unless otherwise specified, all models are trained for up to 15 epochs using the Adam optimizer (learning rate 0.001, batch size 512) with early stopping. Minor adjustments to batch size and AMP usage were made on a case-by-case basis to prevent out-of-memory errors.

4.2 Performance and resource utilization of fusion embedding modules

In Subsection 3.5, we introduced three fusion embedding modules: FlattenCat, TFT, and TFT-GWS. To evaluate their performance and resource utilization, we measure the time and memory costs of each module during both inference and training, as detailed in Table 6. Our analysis focuses on varying sequence lengths, spatial dimensions, and the number of external features to provide a comprehensive understanding of their performance characteristics. We do not reiterate the computational overhead of subsequent models (e.g., Informer, T-GCN), because this has been extensively covered in previous research (Shao et al., 2025). When running the experiments, each module receives multiple time series variables as input, spanning historical and future windows of length L_w. Specifically, each variable has a shape of $(16, 2L_w, N_s, 1)$, encompassing one target component and N_{cv} variables for each covariate component as detailed in Table 4. When retaining spatial dimension (R.S.D.) is enabled, the module outputs a variable of shape $(16, 2L_w, N_s, 512)$; otherwise, the output shape is $(16, 2L_w, 512)$.

4.3 Basic performance of forecasting

The baseline model performance on the PEMS-BAY and METR-LR datasets using only traffic data is presented in Table 7 (more in Supplementary Materials). The analysis highlights four primary findings: (1) The prediction error consistently increases with the length of the forecast horizon. (2) The model performance is highly sensitive to the number of nodes used, with different architectures responding differently to network scale. (3) The choice of embedding module is critical because nonlinear embeddings provide a significant advantage for long-term forecasting over simpler methods. (4) As the number of nodes increases, the complex spatial dependencies can act as noise, and the resulting modeling

Unedited

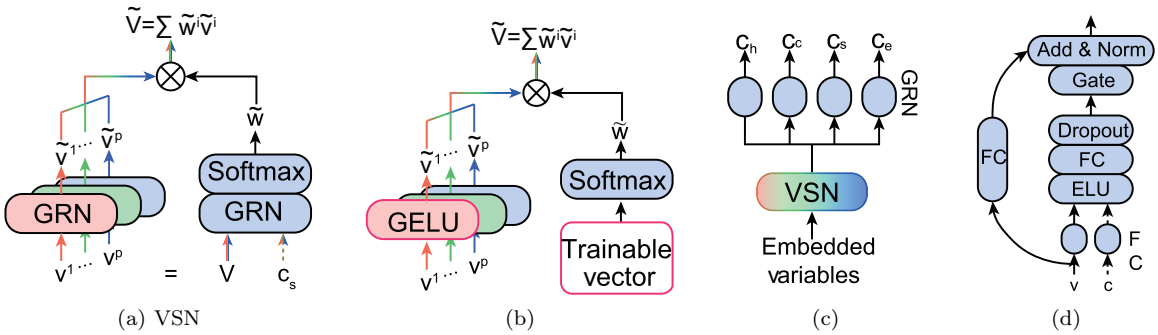

Fig. 5 The structure of (a) VSN; (b) GWS; (c) SVE; (d) GRN. GRN, gated residual network; GWS, gated weighted summation; SVE, static variable encoder; VSN, variable selection network

Table 6 Performance and resource utilization of fusion embedding modules

E.M.[d]	N_s[e]	N_{cv}[a]	Parameters						Inference																Training											
									Time (ms)						Memory (MB)						Time (ms)						Memory (MB)					
			0		1		4		0		1		4		0		1		4		0		1		4		0		1		4	
		L_w[c]	6	48	6	48	6	48	6	48	6	48	6	48	6	48	6	48	6	48	6	48	6	48	6	48	6	48	6	48	6	48
		R.S.D.[b]																														
FlattenCat	8	False	0	0	0	0	0	0	0.1	0.1	0.4	0.4	0.7	0.7	0.0	0.2	0.2	1.3	0.6	4.6	-	-	-	-	-	-	-	-	-	-	-	-
		True	0	0	432	432	1.7K	1.7K	0.1	0.1	0.6	0.6	1.6	1.5	0.0	0.2	1.1	1.9	1.4	4.6	-	-	2.5	2.7	5.5	5.7	-	-	1.1	2.7	1.5	10.2
	16	False	0	0	0	0	0	0	0.1	0.1	0.4	0.4	0.7	0.7	0.1	0.5	0.3	2.6	1.1	9.9	-	-	-	-	-	-	-	-	-	-	-	-
		True	0	0	1.6K	1.6K	6.5K	6.5K	0.1	0.1	0.7	0.6	1.5	1.4	0.1	0.5	1.2	2.9	1.8	9.9	-	-	2.6	2.4	6.0	5.3	-	-	1.3	5.4	2.5	20.9
TFT	8	False	2.1M	2.1M	17.3M	17.3M	50.7M	50.7M	1.2	1.2	4.4	5.3	13.3	10.4	2.7	16.6	8.9	62.4	26.6	208.8	5.7	6.4	21.3	20.0	53.7	48.3	9.2	32.8	55.9	121.7	183.4	409.1
		True	2.1M	2.1M	16.9M	16.9M	49.3M	49.3M	1.2	2.0	4.8	11.6	14.6	40.9	16.5	132.1	63.3	482.1	209.5	1637.0	6.1	6.4	20.9	35.4	60.2	146.0	32.3	324.4	151.8	1462.6	510.5	5049.4
	16	False	2.1M	2.1M	17.7M	17.7M	52.3M	52.3M	1.0	1.0	5.4	5.4	13.4	11.9	2.7	16.7	9.0	62.9	26.8	210.4	6.6	5.5	20.0	23.1	54.8	51.4	9.2	32.9	57.5	126.1	189.8	428.6
		True	2.1M	2.1M	16.9M	16.9M	49.3M	49.3M	1.1	3.7	6.1	21.8	14.3	81.3	33.0	264.3	126.9	964.2	418.2	3274.0	5.7	11.6	24.2	67.4	56.2	283.3	76.6	654.9	351.4	2976.4	1179.2	10273.7
TFT-GWS	8	False	4.6K	4.6K	6.7M	6.7M	14.3M	14.3M	0.3	0.4	2.2	2.8	4.8	4.7	1.8	12.1	9.2	67.9	29.9	232.7	1.9	2.1	6.2	7.8	15.3	17.8	3.4	27.1	11.3	87.3	34.3	293.8
		True	1.0K	1.0K	6.4M	6.4M	12.9M	12.9M	0.4	0.6	3.3	5.6	5.9	22.1	12.0	96.1	69.3	530.0	233.4	1828.6	1.8	2.1	7.8	17.8	20.4	97.5	27.0	216.1	117.4	840.8	396.2	2961.1
	16	False	8.7K	8.7K	7.1M	7.1M	16.0M	16.0M	0.4	0.4	2.8	2.8	4.7	5.6	1.8	12.2	9.3	68.4	30.1	234.3	1.9	2.1	7.2	7.5	16.9	15.8	3.4	27.2	10.5	91.2	31.3	313.3
		True	1.0K	1.0K	6.4M	6.4M	12.9M	12.9M	0.4	1.4	4.0	11.3	7.4	46.9	24.0	192.2	138.9	1059.9	466.1	3657.2	2.2	4.4	9.9	35.6	28.8	197.3	54.0	432.2	235.2	1680.3	788.5	5921.0

[a] N_{cv}: Number of variables per covariate component in Table 4 (0 for no covariates, 1 for 6 covariates, 4 for 24 covariates).
[b] **R.S.D.**: retaining spatial dimension (False if not retained, True if retained).
[cde] These abbreviations are defined in Table 5.
GWS, gated weighted summation.

Table 7 Forecasting performance without covariates and future information

Dataset	Model	E.M.[b]	FlattenCat								TFT								TFT-GWS							
		L_w[c]	12		24		48		96		12		24		48		96		12		24		48		96	
		N_s[d]	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
PEMS-BAY	Seq2Seq	1-1	0.2967	0.2702	0.3750	0.3148	0.5221	0.3949	0.6479	0.4398	0.3065	0.2673	0.3663	0.3105	0.5422	0.4024	0.6146	0.4245	0.2910	0.2629	0.3645	0.3078	**0.5069**	0.3850	**0.5767**	**0.3964**
		16-1	0.2864	0.2681	0.4540	0.3360	0.5574	0.3812	0.5541	0.3880	0.3067	0.2766	0.4266	0.3234	0.5638	0.3879	0.5541	0.3851	0.3096	0.2772	0.4507	0.3295	0.5516	0.3833	**0.5266**	**0.3687**
		16-4	0.1875	0.2050	0.2624	0.2421	0.3250	0.2810	0.3782	0.3083	0.1964	0.2065	0.2645	0.2430	0.3378	0.2882	0.3873	0.3132	0.1974	0.2094	0.2602	0.2418	**0.3217**	**0.2778**	**0.3698**	**0.3055**
		all	**0.2841**	**0.2762**	**0.3178**	**0.2905**	**0.3810**	**0.3281**	0.3867	0.3271	0.2890	0.2780	0.3194	0.2917	0.3817	0.3296	**0.3803**	**0.3224**	0.2916	0.2767	0.3216	0.2912	0.3834	0.3302	0.3857	0.3247
	Informer	1-1	0.3140	0.2770	0.4252	0.3352	0.5931	0.4249	0.6496	0.4424	0.3267	0.2879	0.4361	0.3421	0.5885	0.4160	0.6416	0.4393	0.3120	0.2775	0.4245	0.3369	0.5708	0.4161	0.6294	0.4319
		16-1	0.3215	0.2936	0.4859	0.3569	0.5775	0.3990	0.5835	0.4030	0.3509	0.3066	0.4943	0.3543	0.5993	0.4080	0.5940	0.4074	0.4076	0.3147	0.5071	0.3612	0.5663	0.3951	0.5714	0.3950
		16-4	0.2136	0.2234	0.2885	0.2605	0.3921	0.3184	0.4073	0.3359	0.2205	0.2275	0.2880	0.2587	0.4117	0.3247	0.4179	0.3412	0.2242	0.2321	0.2902	0.2618	0.3488	0.2950	0.4020	0.3283
		all	0.3061	0.2908	0.3481	0.3118	0.3901	0.3342	0.3847	0.3300	0.3106	0.2914	0.3386	0.3043	0.4220	0.3491	0.3841	0.3318	0.3175	0.2964	0.3464	0.3120	0.3912	0.3357	0.3814	0.3291
	T-GCN	1-1	0.3035	0.2777	0.3730	0.2959	0.5411	0.3951	0.7115	0.4678	0.2892	0.2667	**0.3531**	**0.2927**	0.5205	0.3748	0.6151	0.3995	**0.2839**	**0.2586**	0.3554	0.2946	0.5209	**0.3731**	0.6579	0.4447
		16-1	0.2880	0.2677	0.4683	0.3677	0.7011	0.4772	0.9316	0.5682	0.2614	0.2534	0.4213	0.3456	0.6394	0.4411	0.8651	0.5424	0.2636	0.2512	0.4308	0.3482	0.6521	0.4644	0.8797	0.5495
		16-4	0.2454	0.2265	0.4009	0.3071	0.6237	0.4003	0.8116	0.4885	0.2242	0.2153	0.3771	0.2894	0.5850	0.3710	0.7847	0.4747	0.2313	0.2159	0.3867	0.2943	0.5994	0.3922	0.7958	0.4853
	GWNet	1-1	0.3002	0.2633	0.4449	0.3372	0.5529	0.4001	0.6642	0.4382	0.3278	0.2889	0.4142	0.3178	0.5237	0.3938	0.6281	0.4163	0.3274	0.3075	0.4355	0.3341	0.5874	0.4215	0.6572	0.4244
		16-1	**0.2299**	**0.2283**	**0.3449**	**0.2892**	**0.4979**	**0.3468**	0.5751	0.4011	0.2386	0.2331	0.3579	0.2897	0.5262	0.3614	0.5566	0.3867	0.2493	0.2564	0.3869	0.3133	0.5337	0.3673	0.5710	0.3948
		16-4	**0.1794**	**0.1897**	**0.2383**	0.2242	0.3592	0.2845	0.4367	0.3344	0.1933	0.1943	0.2422	**0.2222**	0.3740	0.2920	0.4395	0.3277	0.1998	0.2134	0.2604	0.2391	0.3676	0.2919	0.4366	0.3292

[a] Best results within the same window and node configuration are in **bold**.
[bcd] These abbreviations are defined in Table 5.
GCN, graph convolutional network; GWS, gated weighted summation; MAE, mean absolute error; MSE, mean squared error.

Unedited

challenges can overwhelm the informational benefits, ultimately harming predictive accuracy.

4.4 Improvement with covariates

We evaluated the impact of covariates by comparing prediction performance with and without them under identical experimental configurations (Subsection 4.3). As shown in Table 9, covariates significantly improve accuracy for long-term forecasts and in low-complexity, single-node settings, highlighting their value for extended prediction horizons.

The impact of covariates varies with prediction horizon and node complexity. For short horizons (e.g., 12 and 24 steps), covariates offer minimal or even detrimental effects. This is likely due to the sampling frequency mismatch between high-frequency traffic data and low-frequency covariates like weather data. Conversely, their benefits are more pronounced for longer horizons (e.g., 48 and 96 steps), where their inclusion consistently enhances model performance. The positive impact of covariates is most significant in single-node tasks. In multi-node scenarios, however, models can still effectively leverage covariate information by capturing interactions across different nodes, leading to performance gains.

The FlattenCat method is less stable and effective than TFT and TFT-GWS, especially in short-term and full-node scenarios. FlattenCat exhibits inconsistent performance; although effective in some long-term, low-node-count tasks, it shows marked instability in short-term and full-node scenarios. Specifically, it causes significant performance degradation in short-term predictions (12 and 24 steps) and leads to training anomalies in full-node configurations. In contrast, TFT and TFT-GWS consistently outperform FlattenCat and remain stable across most scenarios. Their advantage is particularly notable in single-node configurations. Although their performance gains are less pronounced in full-node tasks, they avoid the anomalous behavior seen with FlattenCat. It is worth noting that TFT-GWS occasionally underperforms TFT, particularly in short-term predictions, which we attribute to potential inefficiencies in its gating mechanism.

4.5 Impact of covariates in the future window

We conduct a comparative analysis of prediction tasks that consider future window information with those that do not. The analysis in Table 8 shows that incorporating future window information, particularly with known covariates, significantly improves

Table 8 Relative forecast performance gain with future information

Dataset	CoV.[b]	E.M.[c]	FlattenCat								TFT								TFT-GWS							
		L_w[d]	12		24		48		96		12		24		48		96		12		24		48		96	
		Model	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
PEMS-BAY	All	Seq2Seq	**2.22**	**1.21**	**2.79**	**1.49**	**3.41**	**1.81**	**9.16**	**6.74**	**9.19**	**5.52**	**12.85**	**8.18**	**14.27**	**11.23**	**10.29**	**9.50**	**4.87**	**3.09**	**10.46**	**6.61**	**12.93**	**10.17**	**13.41**	**11.30**
		Informer	-0.34	-0.17	**3.57**	**1.75**	**4.50**	**2.21**	-11.24	-4.55	**7.95**	**3.84**	**4.85**	**3.74**	**6.42**	**5.43**	**12.00**	**10.14**	**0.46**	**0.21**	**1.50**	**1.82**	**2.53**	**3.17**	**10.40**	**8.56**
		T-GCN	**4.78**	**3.51**	**5.26**	**3.64**	**8.93**	**6.46**	**19.95**	**12.49**	**10.72**	**6.73**	**15.04**	**10.73**	**20.22**	**15.94**	**21.89**	**16.90**	**8.76**	**5.79**	**15.02**	**10.58**	**21.64**	**16.38**	**31.37**	**22.21**
		GWNet	**5.26**	**0.07**	**3.19**	-0.70	**0.26**	-3.09	-9.11	-3.62	-3.44	-1.91	**4.91**	-0.36	**7.41**	**1.71**	**15.75**	**10.76**	**7.14**	**4.26**	**1.56**	-1.10	**5.67**	**1.02**	**11.21**	**6.31**
	No	Seq2Seq	-0.61	-0.18	-0.10	-0.09	**0.78**	**0.18**	**1.32**	**1.22**	-1.28	-1.87	-0.98	-1.00	**3.35**	**2.20**	**1.58**	**2.01**	-6.10	-3.16	-2.23	-1.58	**0.44**	**0.30**	-2.64	-1.46
		Informer	**1.19**	**0.65**	**0.17**	-0.24	-0.54	-0.62	-0.64	-0.53	**3.38**	**1.43**	-2.02	-1.76	-5.73	-4.23	-5.53	-4.37	**4.80**	**1.99**	**0.52**	**0.03**	-2.16	-1.29	-1.00	-1.32
		T-GCN	**0.02**	**0.07**	-0.71	-0.79	-0.07	**0.12**	**0.89**	**0.14**	-0.69	-0.39	**1.43**	**1.01**	**0.82**	-0.57	**4.60**	**1.07**	**0.02**	**0.86**	**0.03**	**0.66**	-0.53	-0.76	**4.99**	**3.03**
		GWNet	-4.25	-5.32	-1.61	-1.86	**0.22**	-0.07	**0.53**	**0.71**	**2.11**	-0.93	**1.40**	-3.88	-2.56	-1.57	-0.64	-0.74	-0.84	**2.83**	**3.00**	**0.44**	**1.54**	**0.04**	-2.64	-3.20

[a] The percentage improvement is calculated as $v = -(m_{w/} - m_{w/o})/m_{w/o} * 100.0\%$, where $m_{w/}$ and $m_{w/o}$ are metrics with and without future information, respectively. Positive improvements are in **bold**.
[bcd] These abbreviations are defined in Table 5.
GCN, graph convolutional network; GWS, gated weighted summation; MAE, mean absolute error; MSE, mean squared error.

prediction performance.

TFT and TFT-GWS benefit from future covariates, whereas FlattenCat struggles, often showing adverse effects due to weak covariate-target correlations. When using the TFT and TFT-GWS embedding modules, known covariates in the future window generally improve prediction accuracy, enabling the model to more effectively capture the evolving patterns of the target data. However, for the FlattenCat module, the inclusion of known covariates in the future window exhibits minimal impact on prediction performance, and in some instances, it even adversely affects performance, especially in the Informer model. This can be attributed to covariates with low correlation to the target data, which tend to interfere more with the model than provide valuable informatio Future window information offers minimal benefits without covariates and may introduce biases that hinder prediction performance. Moreover, we observe that, in the absence of covariates, the benefits of future window information are marginal, and in some cases, it even results in a decline in prediction performance. This can be attributed to biases inherent in the embedding module, which interfere with the ability of the model to learn the target data.

4.6 Variable importance

To understand variable importance, we first note that simple statistical measures, such as Pearson correlation and maximum information coefficient (MIC) (Reshef et al., 2011; Albanese et al., 2013), reveal weak relationships between traffic speed and most auxiliary variables (Fig. 6), except for the MIC between traffic speed and HoD or DoW. However, analyzing the learned variable selection weights from our fusion modules reveals a more nuanced and interpretable picture, as shown in Fig. 7.

The model correctly assigns the highest impor-

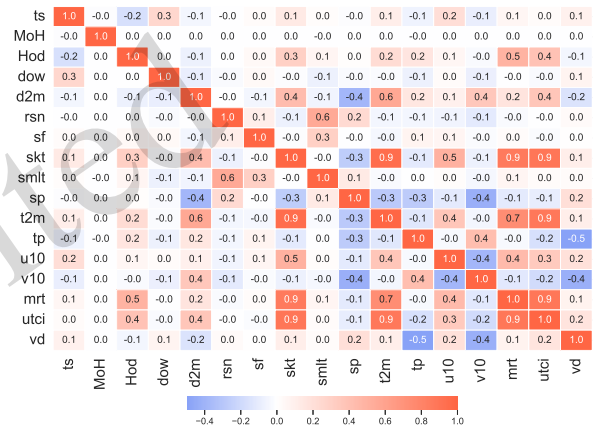

(a)

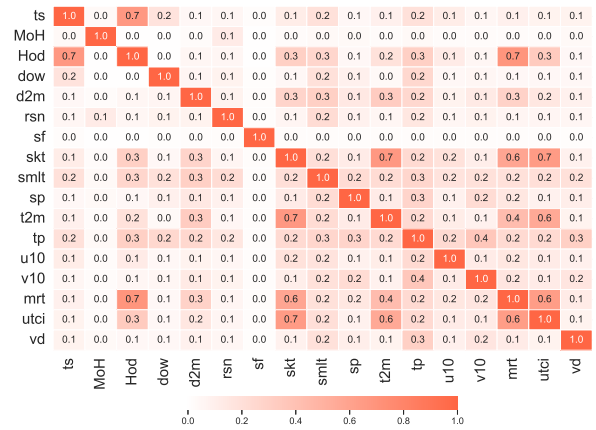

(b)

Fig. 6 Correlation between traffic state and covariates: (a) Pearson; (b) MIC. MIC, maximum information coefficient

Unedited

Table 9 Relative forecast performance gain with covariates

Dataset	Model	E.M.[b]	FlattenCat								TFT								TFT-GWS							
		L_w[c]	12		24		48		96		12		24		48		96		12		24		48		96	
		N_a[d]	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE	MSE	MAE
PEMS-BAY	Seq2Seq	1-1	**2.28**	-5.56	**0.08**	-3.05	**4.96**	**0.71**	**15.99**	-0.55	**13.21**	**5.88**	**5.79**	**3.34**	**31.39**	**17.95**	**43.00**	**22.74**	**5.59**	**0.46**	**5.41**	-1.08	**24.07**	**9.69**	**36.15**	**13.33**
		16-1	-46.26	-16.85	-7.69	-3.56	**2.56**	**0.72**	**5.88**	**4.98**	**0.19**	-0.00	-3.23	**0.09**	**13.05**	**10.14**	**13.26**	**11.88**	-41.74	-14.29	-6.95	-2.22	**9.05**	**8.04**	**7.49**	**5.98**
		all	-275.81	-110.98	-237.09	-100.99	-183.22	-78.76	-176.47	-78.77	**0.42**	**0.66**	**3.15**	**1.95**	**10.96**	**6.96**	**10.13**	**5.29**	-1.80	-1.88	**0.61**	-0.63	**9.75**	**6.63**	**10.18**	**3.89**
	Informer	1-1	-3.83	-10.12	-0.50	-4.46	**14.35**	**6.69**	**22.63**	**3.90**	**10.85**	**4.36**	**10.34**	**2.25**	**32.53**	**16.46**	**37.18**	**18.11**	**5.68**	-1.34	**17.03**	**4.96**	**35.30**	**20.55**	**35.88**	**15.16**
		16-1	-38.03	-11.90	-2.84	-0.91	**7.38**	**5.04**	**7.43**	**6.44**	-10.36	-1.20	**6.61**	**4.74**	**19.64**	**15.00**	**16.31**	**12.79**	-10.08	-3.52	**7.18**	**7.18**	**16.55**	**14.01**	**14.10**	**11.91**
		all	-40.14	-23.43	-50.60	-29.83	-139.86	-61.33	-46.46	-24.30	**4.56**	**2.50**	**3.90**	**1.64**	**19.23**	**11.53**	**8.26**	**4.43**	**3.35**	**2.00**	**3.34**	**2.17**	**9.08**	**5.95**	**4.01**	**1.61**
	T-GCN	1-1	**7.69**	**1.15**	**7.95**	-4.44	**26.67**	**13.73**	**39.14**	**22.71**	**18.83**	**9.13**	**10.08**	**1.13**	**30.54**	**14.73**	**35.22**	**14.92**	-1.76	-9.90	**3.40**	-8.58	**23.22**	**7.08**	**39.22**	**23.18**
		16-1	**5.79**	-3.43	**8.48**	**0.57**	**17.87**	**10.78**	**23.99**	**16.52**	**11.60**	**4.64**	**21.70**	**13.58**	**30.42**	**19.21**	**40.49**	**30.26**	**0.63**	-8.72	**5.89**	-0.07	**14.80**	**9.95**	**19.05**	**12.80**
	GWNet	1-1	**5.46**	-3.11	**13.89**	**2.16**	**12.18**	**8.38**	**15.27**	**5.07**	**16.16**	**4.43**	-14.48	-16.44	-6.30	-3.58	-27.52	-19.43	**0.20**	-2.42	**11.75**	**4.75**	**7.58**	**10.61**	-40.02	-18.22
		16-1	-0.39	-2.02	**1.09**	**1.73**	**12.60**	**6.67**	**15.59**	**11.83**	**2.12**	**0.07**	**6.36**	**0.58**	**12.75**	**5.66**	**2.13**	**1.67**	-19.56	-7.75	-3.10	-1.76	**14.42**	**5.92**	**3.77**	**3.17**

[a] The percentage improvement is calculated as $v = -(m_{w/} - m_{w/o})/m_{w/o} * 100.0\%$, where $m_{w/}$ and $m_{w/o}$ are metrics with and without future covariates, respectively. Positive improvements are in **bold**.
[bcd] These abbreviations are defined in Table 5.
GCN, graph convolutional network; GWS, gated weighted summation; MAE, mean absolute error; MSE, mean squared error.

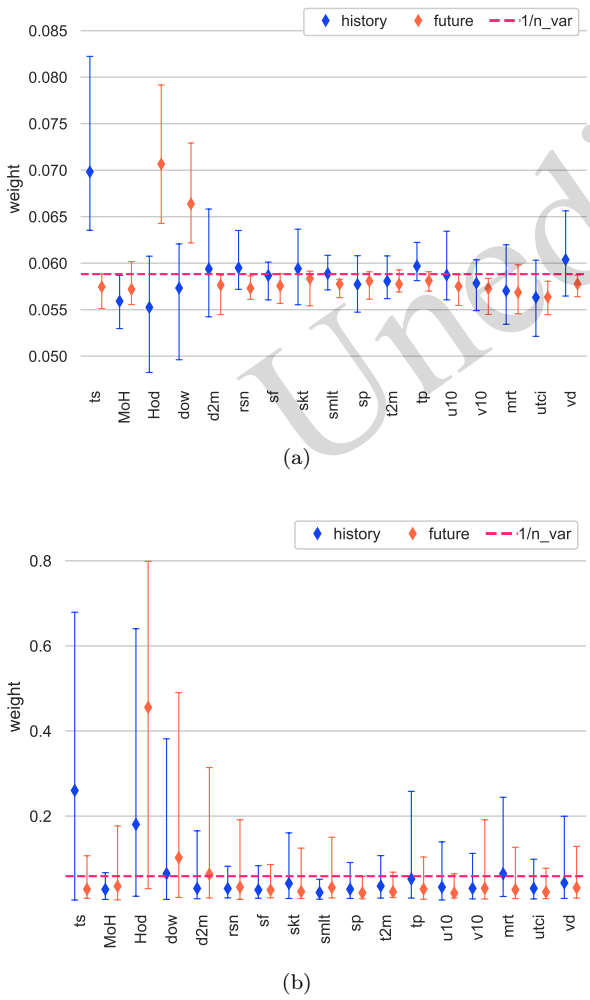

Fig. 7 Variable selection weights in Informer: (a) with TFT-GWS; (b) with TFT. GWS, gated weighted summation

Unedited

tance to the historical target variable itself. More revealingly, it identifies key predictive features that simple correlations miss. For instance, it gives high weight to future HoD and DoW, and to historical precipitation (tp) and visibility (vd). The importance of these variables is empirically supported: traffic patterns are strongly tied to DoW (Fig. 8), and adverse weather is known to impact driving speed (FHWA Road Weather Management, 2024). This demonstrates that the model learns meaningful, non-linear relationships beyond simple statistical correlation. Furthermore, these weights allow for a comparative analysis of the embedding modules. The weights in the TFT-GWS module exhibit a smaller dynamic range than those in the TFT module. This suggests that TFT is more selective in focusing on critical variables, whereas TFT-GWS gives more even consideration to all inputs, which could make it more susceptible to noise from irrelevant features and negatively impact performance.

5 Conclusions

In this study, we introduced MltAuxTSPP, an innovative benchmarking framework designed to enhance traffic state prediction by integrating multi-source auxiliary data. Through a UDC and adaptable fusion embedding modules, MltAuxTSPP effectively addresses key challenges in data utilization and scalability. Our experimental results demonstrate that incorporating auxiliary data, particularly weather and temporal features, significantly improves long-term forecast accuracy. This work validates the efficacy of fusion embedding modules (e.g., FlattenCat, TFT, and TFT-GWS) in transforming heterogeneous data into a unified, predictive feature space. The comparative analysis of these methods offers valuable insights for model selection. Ultimately, MltAuxTSPP provides a robust and exten-

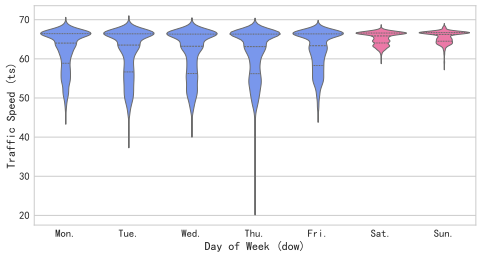

Fig. 8 Average traffic speed distribution by DoW on PEMS-BAY. DoW, day of the week

sible framework that not only advances prediction accuracy but also serves as a critical tool for developing more resilient and efficient strategies for ITS.

Although this study shows promising results, it has several limitations. First, computational overhead increases significantly with large-scale, multi-source data. Second, the model's sensitivity to the quality of auxiliary data poses a challenge. Although our experiments demonstrated improved performance, longer prediction horizons are needed to fully assess the influence of external factors such as weather and special events. Furthermore, the current dataset lacks examples of extreme weather conditions, and the model architecture is limited to classical deep learning approaches.

Future work will address these limitations by focusing on more efficient data processing techniques and adaptive modeling strategies. We also plan to evaluate our methods on larger datasets, such as LargeST (Liu et al., 2023), with extended time horizons and an increased number of traffic nodes, building upon related research in traffic patterns (Shao et al., 2025; Zheng et al., 2025). Additionally, we recognize the importance of collecting data on extreme weather events to ensure the algorithm's robustness in real-world scenarios.

Contributors

Yusong ZHOU conceptualized the research, performed all data collection and analysis, conducted all experiments, and wrote the original draft of the paper. Xiaoyu JIANG and Shu SUN provided valuable suggestions during the conceptualization phase. Xinmin ZHANG, Yuanqiu MO and Zhihuan SONG performed critical review, edited of the paper, and helped with its finalization.

Conflict of interest

All the authors declare that they have no conflict of interest.

Data availability

The data that supporting the findings of this study are available through the open-source LibCity project. Datasets are already uploaded in network disks BaiduDisk with code 1231 or Google Drive.

References

Akiba T, Sano S, Yanase T, et al., 2019. Optuna: A next-Generation Hyperparameter Optimization Framework. Proc 25th ACM SIGKDD Int Conf on Knowledge Discovery & Data Mining. https://doi.org/10.1145/3292500.3330701

Albanese D, Filosi M, Visintainer R, et al., 2013. Minerva and minepy: a C engine for the MINE suite and its R, Python and MATLAB wrappers. *Bioinformatics*, 29(3):407–408. https://doi.org/10.1093/bioinformatics/bts707

Baczek J, Zhylko D, Titericz G, et al., 2024. TSPP: A Unified Benchmarking Tool for Time-series Forecasting. https://doi.org/10.48550/arXiv.2312.17100

Bai HY, Liu X, 2025. T-Graphormer: Using Transformers for Spatiotemporal Forecasting. arXiv. https://doi.org/10.48550/arXiv.2501.13274

Chen K, Chen F, Lai B, et al., 2020. Dynamic Spatio-Temporal Graph-Based CNNs for Traffic Flow Prediction. *IEEE Access*, 8:185136–185145. https://doi.org/10.1109/ACCESS.2020.3027375

Chen YY, Lv Y, Li Z, et al., 2016. Long Short-Term Memory Model for Traffic Congestion Prediction with Online Open Data. 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC), p.132–137. https://doi.org/10.1109/ITSC.2016.7795543

Chen M, Yu X, Liu Y, 2018. PCNN: Deep Convolutional Networks for Short-Term Traffic Congestion Prediction. *IEEE Transactions on Intelligent Transportation Systems*, 19(11):3550–3559. https://doi.org/10.1109/TITS.2018.2835523

Cho K, Merrienboer B, Gulcehre C, et al., 2014. Learning Phrase Representations Using RNN Encoder-Decoder for Statistical Machine Translation. https://doi.org/10.48550/arXiv.1406.1078

Do LNN, Taherifar N, Vu HL, 2019. Survey of Neural Network-Based Models for Short-Term Traffic State Prediction. *WIREs Data Min Knowl Disc*, 9(1):e1285. https://doi.org/10.1002/widm.1285

Essien A, Petrounias I, Sampaio P, et al., 2021. A Deep-Learning Model for Urban Traffic Flow Prediction with Traffic Events Mined from Twitter. *World Wide Web*, 24(4):1345–1368. https://doi.org/10.1007/s11280-020-00800-3

Essien A, Petrounias I, Sampaio P, et al., 2018. The Impact of Rainfall and Temperature on Peak and Off-Peak Urban Traffic. Database and Expert Systems Applications, p.399–407. https://doi.org/10.1007/978-3-319-98812-2_36

Essien A, Petrounias I, Sampaio P, et al., 2019. Improving Urban Traffic Speed Prediction Using Data Source Fusion and Deep Learning. 2019 IEEE International

Unedited

Conference on Big Data and Smart Computing (BigComp), p.1–8.
https://doi.org/10.1109/BIGCOMP.2019.8679231

Fang Z, Long QQ, Song GJ, et al., 2021. Spatial-temporal graph ode networks for traffic flow forecasting. Proc 27th ACM SIGKDD Conf on Knowledge Discovery & Data Mining, p.364–373.
https://doi.org/10.1145/3447548.3467430

FHWA Road Weather Management, 2024. How Do Weather Events Affect Roads?.
https://ops.fhwa.dot.gov/weather/q1_roadimpact.htm [Accessed on May. 20, 2025].

Guo SN, Lin YF, Feng N, et al., 2019. Attention Based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting. Proc 39th AAAI Conf on Artificial Intelligence, p.922–929.
https://doi.org/10.1609/aaai.v33i01.3301922

Guo K, Hu Y, Sun Y, et al., 2021. Hierarchical Graph Convolution Network for Traffic Forecasting. Proceedings of the AAAI Conference on Artificial Intelligence, 35(1):151–159.
https://doi.org/10.1609/aaai.v35i1.16088

Han J, Zhang W, Liu H, et al., 2024. BigST: Linear Complexity Spatio-Temporal Graph Neural Network for Traffic Forecasting on Large-Scale Road Networks. *Proc. VLDB Endow.*, 17(5):1081–1090.
https://doi.org/10.14778/3641204.3641217

He Y, Jia Y, Jia Y, et al., 2024. An Integrated Intra-View and Inter-View Framework for Multiple Traffic Variable Data Simultaneous Recovery. *IEEE Transactions on Intelligent Transportation Systems*, 25(11):17200–17217.
https://doi.org/10.1109/TITS.2024.3414506

Hochreiter S, Schmidhuber J, 1997. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780.
https://doi.org/10.1162/neco.1997.9.8.1735

Huang Z, Xia J, Li F, et al., 2019. A Peak Traffic Congestion Prediction Method Based on Bus Driving Time. *Entropy*, 21(7):709.
https://doi.org/10.3390/e21070709

Kilpelainen M, Summala H, 2007. Effects of Weather and Weather Forecasts on Driver Behaviour. *Transportation Research Part F: Traffic Psychology and Behaviour*, 10(4):288–299.
https://doi.org/10.1016/j.trf.2006.11.002

Lam WHK, Shao H, Sumalee A, 2008. Modeling Impacts of Adverse Weather Conditions on a Road Network with Uncertainties in Demand and Supply. *Transportation Research Part B: Methodological*, 42(10):890–910.
https://doi.org/10.1016/j.trb.2008.02.004

Lee K, Eo M, Jung E, et al., 2021. Short-Term Traffic Prediction With Deep Neural Networks: A Survey. *IEEE Access*, 9:54739–54756.
https://doi.org/10.1109/ACCESS.2021.3071174

Lee C, Kim Y, Jin S, et al., 2020. A Visual Analytics System for Exploring, Monitoring, and Forecasting Road Traffic Congestion. *IEEE Transactions on Visualization and Computer Graphics*, 26(11):3133–3146.
https://doi.org/10.1109/TVCG.2019.2922597

Li Y, Yu R, Shahabi C, et al., 2018. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. 6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings.
https://openreview.net/forum?id=SJiHXGWAZ

Li Q, Xu P, Yang X, et al., 2024. Spatial-Temporal Traffic Prediction With an Interactive Spatial-Enhanced Graph Convolutional Network Model. *IEEE Transactions on Intelligent Transportation Systems*, 25(12):20767–20778. https://doi.org/10.1109/TITS.2024.3467172

Liao B, Zhang J, Wu C, et al., 2018. Deep Sequence Learning with Auxiliary Information for Traffic Prediction. Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, p.537–546.
https://doi.org/10.1145/3219819.3219895

Lim B, Arık SÖ, Loeff N, et al., 2021. Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. *Int J Forecast*, 37(4):1748-1764.
https://doi.org/10.1016/j.ijforecast.2021.03.012

Liu L, Zhang R, Peng J, et al., 2018. Attentive Crowd Flow Machines. Proceedings of the 26th ACM International Conference on Multimedia, p.1553–1561.
https://doi.org/10.1145/3240508.3240681

Liu Z, Zhang Y, Huang R, et al., 2021. EXPERIENCE: Algorithms and Case Study for Explaining Repairs with Uniform Profiles over IoT Data. *J. Data and Information Quality*, 13(3).
https://doi.org/10.1145/3436239

Liu T, Jiang A, Zhou J, et al., 2023. GraphSAGE-Based Dynamic Spatial–Temporal Graph Convolutional Network for Traffic Prediction. *IEEE Transactions on Intelligent Transportation Systems*, 24(10):11210–11224.
https://doi.org/10.1109/TITS.2023.3279929

Liu J, Li T, Xie P, et al., 2020. Urban Big Data Fusion Based on Deep Learning: An Overview. *Information Fusion*, 53:123–133.
https://doi.org/10.1016/j.inffus.2019.06.016

Liu X, Xia YT, Liang YX, et al., 2023. LargeST: A Benchmark Dataset for Large-Scale Traffic Forecasting. Proc 37th Int Conf on Neural Information Processing Systems, p.75354-75371.
https://doi.org/10.5555/3666122.3669415

Ma XL, Yu HY, Wang YP, et al., 2015. Large-Scale Transportation Network Congestion Evolution Prediction Using Deep Learning Theory. *PLOS ONE*, 10(3):e0119044.
https://doi.org/10.1371/journal.pone.0119044

Ma X, Dai Z, He Z, et al., 2017. Learning Traffic as Images: A Deep Convolutional Neural Network for Large-Scale Transportation Network Speed Prediction. *Sensors*, 17(4):818.
https://doi.org/10.3390/s17040818

Nagy AM, Simon V, 2018. Survey on Traffic Prediction in Smart Cities. *Pervasive Mob Comput*, 50:148–163.
https://doi.org/10.1016/j.pmcj.2018.07.004

Reshef DN, Reshef YA, Finucane HK, et al., 2011. Detecting Novel Associations in Large Data Sets. *Science*, 334(6062):1518–1524.
https://doi.org/10.1126/science.1205438

Scarselli F, Gori M, Tsoi AC, et al., 2009. The Graph Neural Network Model. *IEEE Transactions on Neural Networks*, 20(1):61–80.
https://doi.org/10.1109/TNN.2008.2005605

Unedited

Shao ZZ, Wang F, Xu YJ, et al., 2025. Exploring Progress in Multivariate Time Series Forecasting: Comprehensive Benchmarking and Heterogeneity Analysis. *IEEE Trans Knowl Data Eng*, 37(1):291–305. https://doi.org/10.1109/TKDE.2024.3484454

Shaygan M, Meese C, Li WX, et al., 2022. Traffic Prediction Using Artificial Intelligence: Review of Recent Advances and Emerging Opportunities. *Transport Research Part C: Emerg Technol*, 145:103921. https://doi.org/10.1016/j.trc.2022.103921

Song C, Lin YF, Guo SN, et al., 2020. Spatial-Temporal Synchronous Graph Convolutional Networks: A New Framework for Spatial-Temporal Network Data Forecasting. Proc 39th AAAI Conf on Artificial Intelligence, p.914–921. https://doi.org/10.1609/aaai.v34i01.5438

Sun F, Dubey A, White J, 2017. DxNAT — Deep Neural Networks for Explaining Non-Recurring Traffic Congestion. 2017 IEEE International Conference on Big Data (Big Data), p.2141–2150. https://doi.org/10.1109/BigData.2017.8258162

Sun S, Chen J, Sun J, 2019. Traffic Congestion Prediction Based on GPS Trajectory Data. *International Journal of Distributed Sensor Networks*, 15(5):1550147719847440. https://doi.org/10.1177/1550147719847440

Sutskever I, Vinyals O, Le QV, 2014. Sequence to Sequence Learning with Neural Networks. Proc 27th Int Conf on Neural Information Processing Systems, p.3104–3112. https://doi.org/10.5555/2969033.2969173

Taghipour H, Parsa AB, 2020. A Dynamic Approach to Predict Travel Time in Real Time Using Data Driven Techniques and Comprehensive Data Sources. *Transport Eng*, 2:100025. https://doi.org/10.1016/j.treng.2020.100025

Tang J, Chen X, Hu Z, et al., 2019. Traffic Flow Prediction Based on Combination of Support Vector Machine and Data Denoising Schemes. *Physica A: Statistical Mechanics and its Applications*, 534:120642. https://doi.org/10.1016/j.physa.2019.03.007

Vaswani A, Shazeer N, Parmar N, et al., 2017. Attention is all you need. Proc 31st Int Conf on Neural Information Processing Systems, p.6000–6010. https://doi.org/10.5555/3295222.3295349

Wang JY, Jiang JW, Jiang WJ, et al., 2021. LibCity: An Open Library for Traffic Prediction. Proc 29th Int Conf on Advances in Geographic Information Systems, p.145–148. https://doi.org/10.1145/3474717.3483923

Wu Z, Pan S, Long G, et al., 2020. Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks. Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, p.753–763. https://doi.org/10.1145/3394486.3403118

Wu ZH, Pan SR, Long GD, et al., 2019. Graph WaveNet for Deep Spatial-Temporal Graph Modeling. Proc 28th Int Joint Conf on Artificial Intelligence, p.1907–1913. https://doi.org/10.24963/IJCAI.2019/264

Wu HX, Hu TG, Liu Y, et al., 2023. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. Int Conf on Learning Representations. https://doi.org/10.48550/arXiv.2210.02186

Yadan O, 2019. Hydra - A Framework for Elegantly Configuring Complex Applications. https://github.com/facebookresearch/hydra

Yang XX, Zou YJ, Tang JJ, et al., 2020. Evaluation of Short-Term Freeway Speed Prediction Based on Periodic Analysis Using Statistical Models and Machine Learning Models. *J Adv Transport*, 2020(1):9628957. https://doi.org/10.1155/2020/9628957

Yao H, Tang X, Wei H, et al., 2019. Revisiting Spatial-Temporal Similarity: A Deep Learning Framework for Traffic Prediction. Proceedings of the AAAI Conference on Artificial Intelligence, 33(01):5668–5675. https://doi.org/10.1609/aaai.v33i01.33015668

Yin X, Wu G, Wei J, et al., 2022. Deep Learning on Traffic Prediction: Methods, Analysis, and Future Directions. *IEEE Transactions on Intelligent Transportation Systems*, 23(6):4927–4943. https://doi.org/10.1109/TITS.2021.3054840

Ying C, Cai T, Luo S, et al., 2021. Do Transformers Really Perform Badly for Graph Representation?. Advances in Neural Information Processing Systems, 34:28877–28888.

Yuan H, Li G, 2021. A Survey of Traffic Prediction: From Spatio-Temporal Data to Intelligent Transportation. *Data Science and Engineering*, 6(1):63–85. https://doi.org/10.1007/s41019-020-00151-z

Yu B, Yin HT, Zhu ZX, 2018. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. Proc 27th Int Joint Conf on Artificial Intelligence, p.3634–3640. https://doi.org/10.24963/ijcai.2018/505

Zhang S, Yao Y, Hu J, et al., 2019. Deep Autoencoder Neural Networks for Short-Term Traffic Congestion Prediction of Transportation Networks. *Sensors*, 19(10):2229. https://doi.org/10.3390/s19102229

Zhang J, Chen F, Cui Z, et al., 2021. Deep Learning Architecture for Short-Term Passenger Flow Forecasting in Urban Rail Transit. *IEEE Transactions on Intelligent Transportation Systems*, 22(11):7004–7014. https://doi.org/10.1109/TITS.2020.3000761

Zhang J, Zheng Y, Qi D, 2017. Deep Spatio-Temporal Residual Networks for Citywide Crowd Flows Prediction. Proceedings of the AAAI Conference on Artificial Intelligence, 31(1). https://doi.org/10.1609/aaai.v31i1.10735

Zhang L, Liu Q, Yang W, et al., 2013. An Improved K-Nearest Neighbor Model for Short-Term Traffic Flow Prediction. *Procedia - Social and Behavioral Sciences*, 96:653–662. https://doi.org/10.1016/j.sbspro.2013.08.076

Zhang J, Chen F, Guo Y, et al., 2020. Multi-Graph Convolutional Network for Short-Term Passenger Flow Forecasting in Urban Rail Transit. *IET Intelligent Transport Systems*, 14(10):1210–1217. https://doi.org/10.1049/iet-its.2019.0873

Zhao L, Song YJ, Zhang C, et al., 2020. T-GCN: A Temporal Graph ConvolutionalNetwork for Traffic Prediction. *IEEE Trans Intell Transport Syst*, 21(9):3848–3858. https://doi.org/10.1109/TITS.2019.2935152

Unedited

Zheng X, Bagloee SA, Sarvi M, 2024. TRECK: Long-Term Traffic Forecasting With Contrastive Representation Learning. *IEEE Transactions on Intelligent Transportation Systems*, 25(11):16964–16977. https://doi.org/10.1109/TITS.2024.3421328

Zheng X, Bagloee SA, Sarvi M, 2025. A comparative study of deep learning and ensemble learning to extend the horizon of traffic forecasting. https://doi.org/10.48550/arXiv.2504.21358

Zhou HY, Zhang SH, Peng JQ, et al., 2021. Informer: Beyond Efficient Transformer for Long Sequence Time-Series Forecasting. Proc 39th AAAI Conf on Artificial Intelligence, p.11106–11115. https://doi.org/10.1609/aaai.v35i12.17325

Zhu L, Yu FR, Wang YG, et al., 2019. Big Data Analytics in Intelligent Transportation Systems: A Survey. *IEEE Trans Intell Trans Syst*, 20(1):383–398. https://doi.org/10.1109/TITS.2018.2815678

Zou G, Lai Z, Wang T, et al., 2024. MT-Stnet: A Novel Multi-Task Spatiotemporal Network for Highway Traffic Flow Prediction. *IEEE Transactions on Intelligent Transportation Systems*, 25(7):8221–8236. https://doi.org/10.1109/TITS.2024.3411638

Unedited