

UniST-Pred: A Robust Unified Framework for Spatio-Temporal Traffic Forecasting in Transportation Networks Under Disruptions

Yue Wang¹, Djellel Difallah¹, Areg Karapetyan¹ and Samer Madanat¹

¹New York University Abu Dhabi

{yw3576, dd115, ak7550, smm17}@nyu.edu

Abstract

Spatio-temporal traffic forecasting is a core component of intelligent transportation systems, supporting various downstream tasks such as signal control and network-level traffic management. In real-world deployments, forecasting models must operate under structural and observational uncertainties, conditions that are rarely considered in model design. Recent approaches achieve strong short-term predictive performance by tightly coupling spatial and temporal modeling, often at the cost of increased complexity and limited modularity. In contrast, efficient time-series models capture long-range temporal dependencies without relying on explicit network structure. We propose UniST-Pred, a unified spatio-temporal forecasting framework that first decouples temporal modeling from spatial representation learning, then integrates both through adaptive representation-level fusion. To assess robustness of the proposed approach, we construct a dataset based on an agent-based, microscopic traffic simulator (MATSim) and evaluate UniST-Pred under severe network disconnection scenarios. Additionally, we benchmark UniST-Pred on standard traffic prediction datasets, demonstrating its competitive performance against existing well-established models despite a lightweight design. The results illustrate that UniST-Pred maintains strong predictive performance across both real-world and simulated datasets, while also yielding interpretable spatio-temporal representations under infrastructure disruptions. The source code and generated dataset are publicly available at <https://github.com/yw3576/UniST-Pred>.

1 Introduction

Urban traffic systems operate as large-scale, tightly coupled dynamical networks, where local disturbances can quickly propagate across space and time to produce corridor- or network-level congestion. In many traffic control paradigms [Wei *et al.*, 2019], an implicit but critical requirement is *anticipation*: decisions are made based not only on current observations, but also on expected near-future traffic evolution.

Traffic forecasting is central to operational traffic management, enabling agencies to respond proactively to localized infrastructure impediments such as bridge closures, flooded road segments, or zoned capacity reductions [Li *et al.*, 2022; Manibardo *et al.*, 2021]. By accounting for such disruptions, traffic controllers can prevent cascading congestion effects and maintain stable network operation under non-recurrent events that significantly impact mobility and safety. However, most existing forecasting models assume complete and stable network structure, limiting their practical applicability in real-world deployments.

In this work, we address this gap by rethinking the design of the traffic forecasting module around properties that directly support effective traffic control, including (i) low-latency and lightweight inference suitable for deployment in traffic control centers, (ii) modular integration with existing controllers, and (iii) interpretable spatio-temporal representations that disentangle spatial and temporal features, enabling practitioners to reason about evolving traffic conditions.

Research on traffic forecasting has increasingly focused on deep spatio-temporal models. Existing state-of-the-art approaches can be broadly grouped into two representative directions. The first line focuses on *spatio-temporal graph neural networks and transformers* [Jiang *et al.*, 2023a; Wang *et al.*, 2024], which couple temporal modeling and spatial interaction learning through stacked message passing, recurrent units, or attention mechanisms. Representative methods achieve strong predictive accuracy by modeling long-range and dynamic dependencies on traffic networks. However, these models often rely on intertwined spatio-temporal blocks, leading to increased architectural complexity and computational cost. Such tight coupling also reduces modularity, making it less convenient to adapt or integrate the predictor as a standalone component in larger systems.

The second line emphasizes *efficient temporal modeling* in time-series forecasting [Chen *et al.*, 2023; Nie *et al.*, 2023; Liu *et al.*, 2024]. Recent mixer-style or transformer-based backbones demonstrate that long-range temporal dependencies can be captured effectively without heavy recurrent or attention-based sequence modeling. While these approaches offer favorable scalability, they typically treat each location independently or rely on simple spatial features, and thus do not directly address how to incorporate network structure in a task-adaptive manner.

Most existing methods emphasize either developing expressive spatio-temporal coupling or efficient temporal modeling, but few attempt to address both within a unified and modular design. This motivates a structural reconsideration of how temporal and spatial information should be modeled and combined for traffic forecasting. We propose **UniST-Pred**, a unified spatio-temporal forecasting framework built around a simple but deliberate design principle: *decouple temporal dependency learning from spatial representation extraction, and then fuse them adaptively at the representation level*. By separating temporal modeling, spatial encoding, and fusion into distinct stages, UniST-Pred departs from tightly coupled spatio-temporal architectures, offering a unified yet compositional alternative for traffic forecasting.

Leveraging the proposed framework, we assemble a competitive, lightweight traffic forecasting model suitable for integration into traffic analysis and control pipelines. The temporal component of the model employs the feature-time mixing technique proposed in [Chen *et al.*, 2023] to capture long-range temporal dependencies and cross-feature interactions without recurrent or attention-based sequence modeling. The spatial component follows the idea of Graph Transformer Networks (GTNs) [Yun *et al.*, 2019], learning task-adaptive graph structures by composing multiple candidate relations before applying graph convolution. Finally, temporal and spatial representations are integrated through a squeeze-and-excitation [Hu *et al.*, 2018] residual fusion block, which adaptively reweights the two feature streams while preserving modularity. Overall, the contributions of this work are four-fold:

- ★ **A unified decouple-then-fuse spatio-temporal forecasting framework:** We present a modular modeling paradigm that separates temporal dependency learning from spatial representation extraction and integrates them through adaptive representation-level fusion, providing a competitive and adaptive alternative to tightly intertwined spatio-temporal architectures.
- ★ **A curated simulator-based traffic forecasting dataset with structural network variations:** We provide a traffic flow dataset, referred to as **SimSF-Bay**, derived from MATSim [W Axhausen *et al.*, 2016] simulations of the San Francisco (SF) Bay area’s transportation network (see Sec. A for a brief overview of MATSim). The dataset comprises various disrupted-network scenarios generated by removing selected roads/bridges from the underlying MATSim network, with each scenario corresponding to a specific combination of these roads and/or bridges. These scenarios induce changes in network connectivity and topology, enabling evaluation of model robustness under structural/observational uncertainties that are not captured by existing standard benchmarks.
- ★ **A lightweight and robust traffic forecasting model:** We develop a parameter-efficient traffic forecasting model within the proposed framework and demonstrate its robustness under a range of network disruption scenarios reflecting real-world events such as road closures, incidents, or infrastructure contingencies.
- ★ **Thorough empirical evaluation:** We conduct an ex-

tensive evaluation of the presented model, including quantitative and qualitative performance comparison against well-established traffic forecasting methods on two widely used benchmark datasets, PEMS-Bay [Li *et al.*, 2018] and NYCTaxi [Zheng *et al.*, 2014]. Lastly, we examine model interpretability and provide insights into the learned spatio-temporal representations under simulated network disruptions.

Real-World Value and Collaboration. This work is inherently multidisciplinary, bringing together experts from transportation systems, infrastructure planning, and machine learning communities, including the last author, who is a senior transportation engineer and scientist with first-hand knowledge of the studied San Francisco Bay network. With over 20 years of experience, the last author has contributed to multiple projects and collaborations with transportation agencies both within and outside the U.S., including the Bay Area Metropolitan Transportation Commission (MTC), and is currently leading an active project with the federal infrastructure agency of the UAE. In this work, they helped shape the central motivation of studying traffic forecasting under network disruptions and provided domain expertise for the design and qualitative assessment of the disruption case studies.

2 Related Work

2.1 Spatio-Temporal Traffic Forecasting with Coupled Architectures

Most existing deep learning approaches for traffic forecasting adopt *coupled* spatio-temporal modeling strategies [Jiang *et al.*, 2023b; Yu *et al.*, 2018; Li and Zhu, 2021], where spatial interaction learning and temporal dependency modeling are tightly intertwined. Early models such as STGCN [Yu *et al.*, 2018] combine graph convolution with temporal convolution to capture localized spatial-temporal dependencies, while DCRNN [Li *et al.*, 2018] models traffic diffusion on directed graphs using recurrent encoder-decoder structures. Graph WaveNet [Wu *et al.*, 2019] further improves this line by introducing an adaptive dependency matrix to capture hidden spatial relations beyond predefined road connectivity.

More recent transformer-style architectures further strengthen this coupling by using spatio-temporal attention to dynamically model interactions across nodes and time steps [Zheng *et al.*, 2019; Wang *et al.*, 2024; Jiang *et al.*, 2023a]. ASTGCN [Guo *et al.*, 2019] leverages spatial and temporal attention to identify important nodes and time intervals. STFGNN [Li and Zhu, 2021] builds a spatial-temporal fusion graph to capture correlations across nodes and time steps. AGCRN [Bai *et al.*, 2020] and MTGNN [Wu *et al.*, 2020] learn adaptive graph structures for end-to-end spatial-temporal forecasting. Transformer-based models, such as GMAN [Zheng *et al.*, 2019], STAEformer [Liu *et al.*, 2023], PDFformer [Jiang *et al.*, 2023a], and STGformer [Wang *et al.*, 2024], further improve long-range dependency modeling through spatial-temporal attention, adaptive embeddings, propagation-aware mechanisms, and efficient graph transformer designs. These models demonstrate strong predictive accuracy on large-scale benchmarks, particularly when long-range spatial and temporal dependencies are

prominent. However, their performance often comes at the cost of increased architectural complexity and computational overhead. More importantly, the repeated interleaving of spatial interaction learning and temporal modeling reduces modularity: spatial and temporal components are difficult to isolate, reuse, or simplify, which can hinder flexible deployment when forecasting models are used as standalone modules in larger systems.

2.2 Decoupled Perspectives: Efficient Temporal Modeling and Task-Adaptive Spatial Encoding

Recent studies have explored more modular or decoupled learning strategies for traffic forecasting. ST-SSL [Ji *et al.*, 2023] improves spatial-temporal representations through self-supervised objectives, partially separating representation learning from the final supervised forecasting task. Along this direction, STEP [Shao *et al.*, 2022] and STD-MAE [Gao *et al.*, 2024] further demonstrate the value of pre-training-based representation learning. STEP adopts a pre-training and fine-tuning paradigm to learn transferable spatial-temporal knowledge before adapting it to downstream forecasting tasks. STD-MAE extends this perspective by explicitly decoupling spatial and temporal dependency learning within a masked pre-training framework. Together, these methods illustrate that separating representation acquisition from task-specific prediction can improve transferability and robustness in spatial-temporal forecasting.

In parallel, the time-series forecasting community has shown that long-range temporal dependencies can be modeled efficiently without heavy recurrent or attention-based sequence modeling. Mixer-style architectures [Chen *et al.*, 2023] and efficient transformer variants [Nie *et al.*, 2023; Zhang *et al.*, 2024] demonstrate that temporal dependency learning can be decomposed into simple operations along the time and feature dimensions, achieving competitive performance with significantly reduced complexity. These results suggest that expressive temporal modeling does not inherently require tight coupling with spatial interaction mechanisms.

Separately, another line of research questions the reliance on fixed graph structures for spatial modeling. Graph structure learning and relation composition methods argue that a single predefined adjacency matrix may be suboptimal for downstream tasks. Graph Transformer Networks [Yun *et al.*, 2019] exemplify this perspective by learning task-adaptive graph structures through soft composition of candidate relations. In traffic forecasting, related ideas appear in models that infer dynamic or attention-based spatial dependencies rather than relying solely on physical road connectivity.

3 Preliminaries & Problem Statement

We consider the problem of traffic forecasting on a road transportation network represented as a weighted directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E}, \mathbf{A})$, where $\mathcal{V} = \{v_1, \dots, v_N\}$ denotes the set of road segments or intersections, $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of directed edges describing allowable traffic flow, and $\mathbf{A} \in \mathbb{R}^{N \times N}$ is the weighted adjacency matrix encoding spatial relationships such as distance or connectivity strength.

Each node $v_i \in \mathcal{V}$ is associated with a traffic time series $\mathbf{x}_i = [x_i^1, x_i^2, \dots, x_i^T]$, where x_i^t denotes the observed traffic variable (e.g., speed, flow, or density) at time t . The network state at time t is $\mathbf{X}^t = [x_1^t, \dots, x_N^t]^\top \in \mathbb{R}^N$. In addition, each node may have static features $\mathbf{X}' \in \mathbb{R}^{N \times d}$, such as lane count or capacity.

Given a historical observation window of length H , the task is to predict future traffic states for a horizon H' :

$$f_\theta : (\{\mathbf{X}^{t-H+1}, \dots, \mathbf{X}^t\}, \mathcal{G}, \mathbf{X}') \rightarrow (\hat{\mathbf{X}}^{t+1}, \dots, \hat{\mathbf{X}}^{t+H'})$$

where $f_\theta(\cdot)$ is the predictive model parameterized by θ . The function f_θ must jointly capture temporal dependencies in the historical sequences $\mathbf{X}^{t-H+1:t}$ and spatial correlations encoded in the graph $\mathcal{G}$, as traffic dynamics are influenced by both local temporal trends and interactions among neighboring road segments. In practice, temporal feature extraction (e.g., time mixing) is integrated with graph-based representations to model spatio-temporal dynamics, allowing the model to learn complex dependency patterns that evolve over time and vary across different regions of the network, enabling robust and adaptive traffic prediction.

4 Methodology

We propose a unified spatio-temporal framework for traffic forecasting under dynamic scenarios where network topology may change (e.g., road closures). The model decouples temporal modeling and spatial representation learning, then integrates them via an adaptive fusion module (Fig. 1). For temporal dependencies, we adopt feature/time mixing [Chen *et al.*, 2023] to capture long-range dynamics efficiently, and for spatial dependencies we employ graph-based encoding to leverage road-network structure without assuming a fixed topology. This separation-and-fusion design improves robustness to evolving connectivity and heterogeneous traffic conditions.

Temporal Modeling via Feature and Time Mixing: The temporal block models long-range dependencies and cross-feature interactions using feature/time mixing instead of RNNs or attention. This design captures both short-term variations and long-term patterns efficiently. The time and feature mixing layer is stated as below:

(a) Mixer operators. We use $\mathcal{M}$ to denote structured mixing operators applied to tensor representations. Specifically, $\mathcal{M}_{\text{time}}$ denotes a time-mixing operator that acts along the temporal dimension independently for each node and feature, while $\mathcal{M}_{\text{feat}}$ denotes a feature-mixing operator that models cross-channel interactions at each time step. Both operators preserve the input shape and are implemented as learnable linear-nonlinear transformations.

(b) Stacked temporal operator. The temporal modeling function is defined as a composition of K stacked temporal mixer blocks: $f_{\text{temp}} := \mathcal{T}_{\theta_K}^{(K)} \circ \mathcal{T}_{\theta_{K-1}}^{(K-1)} \circ \dots \circ \mathcal{T}_{\theta_1}^{(1)}$, where each temporal mixer block $\mathcal{T}_{\theta_k}^{(k)}$ is given by

$$\mathcal{T}_{\theta_k}^{(k)} = \mathcal{M}_{\text{feat}}^{(k)} \circ \mathcal{M}_{\text{time}}^{(k)}, \quad k = 1, \dots, K. \quad (1)$$




Figure 1: The architecture of the proposed UniST-Pred model.

Given an input sequence $\mathbf{X}_{t-H+1:t}$, the temporal representation is computed as

$$\mathbf{Y}_1 = f_{\text{temp}}(\mathbf{X}_{t-H+1:t}) = \left(\bigcirc_{k=1}^K \mathcal{M}_{\text{feat}}^{(k)} \circ \mathcal{M}_{\text{time}}^{(k)} \right) (\mathbf{X}_{t-H+1:t}). \quad (2)$$

Here, $\circ$ denotes function composition, and $\bigcirc_{k=1}^K$ denotes ordered composition from block 1 to K , with each block applying time-mixing followed by feature-mixing.

By avoiding sequential inductive biases, the temporal module provides a flexible representation of traffic dynamics that generalizes across traffic regimes and temporal resolutions. The resulting temporal embeddings focus on intrinsic time-series characteristics, independent of spatial structure.

Spatial Representation Learning with Graph-Based Encoding: The spatial block uses the road network as structural context to extract spatial features. Rather than relying on inductive message passing, it learns which road connections are most relevant to the task. Following a GTN-style approach, the model constructs task-adaptive graph structures by softly masking road segments, combining candidate relations, and applying lightweight graph convolutions.

(a) Relation composition (task-adaptive graph generation). For channel $i \in \{1, \dots, C\}$, we generate an adjacency matrix by softly composing two selected adjacency matrices $\mathbf{Q}_{1,i}^{(l)}$ and $\mathbf{Q}_{2,i}^{(l)}$:

$$\mathbf{A}_i^{(l)} = \mathbf{D}_i^{-1} \left(\mathbf{Q}_{1,i}^{(l)} \mathbf{Q}_{2,i}^{(l)} \right), \quad (3)$$

where $\mathbf{D}_i$ is the degree matrix of $\mathbf{Q}_{1,i}^{(l)} \mathbf{Q}_{2,i}^{(l)}$.

(b) Graph propagation (per-channel GCN). Given node features $\mathbf{X}' \in \mathbb{R}^{N \times d}$ (e.g., lane count or capacity), we add self-loops and compute normalized propagation:

$$\tilde{\mathbf{A}}_i^{(l)} = \mathbf{A}_i^{(l)} + \mathbf{I}, \quad \tilde{\mathbf{D}}_i \text{ is the degree matrix of } \tilde{\mathbf{A}}_i^{(l)}.$$

The per-channel representation is

$$\mathbf{H}_i = \sigma \left(\tilde{\mathbf{D}}_i^{-1} \tilde{\mathbf{A}}_i^{(l)} \mathbf{X}' \mathbf{W} \right), \quad (4)$$

where $\sigma(\cdot)$ is an activation function (e.g., ReLU), and $\mathbf{W} \in \mathbb{R}^{d \times d}$ is a trainable weight shared across channels. We concatenate the channel features, $\mathbf{Z} = \parallel_{i=1}^C \mathbf{H}_i$, and apply a fully connected projection to produce the spatial representation:

$$\mathbf{Y}_2 = f_{\text{spat}}(\mathbf{X}') = f_{\theta'}(\mathbf{Z}). \quad (5)$$

Unlike deep or recursive graph convolutions, the spatial block is intentionally shallow, enabling stable spatial feature extraction without attempting to model exact network connectivity. This lightweight design produces effective spatial embeddings that capture essential structural patterns and remain robust to partial topology changes.

Adaptive Fusion via Squeeze-Excitation Residual Integration: To jointly exploit temporal and spatial information, the proposed framework integrates the two representations through a squeeze-excitation [Hu *et al.*, 2018] residual fusion block, which adaptively reweights temporal and spatial features based on their relative importance.

Given the temporal representation $\mathbf{Y}_1 = f_{\text{temp}}(\mathbf{X}_{t-H+1:t})$, and the spatial representation $\mathbf{Y}_2 = f_{\text{spat}}(\mathbf{X}')$, we first stack the two representations to form a joint feature:

$$\mathbf{Y} = \text{Stack}(\mathbf{Y}_1, \mathbf{Y}_2). \quad (6)$$

The fused representation is finally obtained via a squeeze-excitation residual mapping:

$$\hat{\mathbf{X}} = \mathbf{Y} + f_{\theta_{\text{se}}}(f_{\theta_{\text{res}}}(\mathbf{Y})). \quad (7)$$

Let $\mathbf{Z} = f_{\theta_{\text{res}}}(\mathbf{Y})$. The squeeze-excitation function is defined as

$$f_{\theta_{\text{se}}}(\mathbf{Z}) = \mathbf{Z} \odot \text{sigmoid}(\mathbf{W}_2 \sigma(\mathbf{W}_1 \text{avg}(\mathbf{Z}))), \quad (8)$$

where $\odot$ denotes the Hadamard product. Here, $\text{avg}(\cdot)$ denotes a global pooling operator that aggregates $\mathbf{Z}$ along the representation dimension to produce a compact descriptor, $\sigma(\cdot)$ is a nonlinear activation function (e.g., ReLU), and $\text{sigmoid}(\cdot)$ maps the gating coefficients to $(0, 1)$. The matrices $\mathbf{W}_1$ and $\mathbf{W}_2$ are trainable parameters that control the

squeeze and excitation stages, respectively. This formulation adaptively reweights different representation channels while preserving the spatial and temporal structure encoded in $\mathbf{Z}$.

By performing fusion at the representation level rather than tightly coupling spatio-temporal propagation, the framework maintains modularity and interpretability. The residual connection further ensures stable gradient flow and preserves strong temporal representations while incorporating spatial cues, leading to robust and adaptive traffic forecasting across diverse scenarios.

Final prediction: Combining temporal modeling, spatial representation learning, and adaptive fusion, the output prediction is given by

$$\hat{\mathbf{X}}_{t+1:t+H'} = \text{Stack}(f_{\text{temp}}(\mathbf{X}_{t-H+1:t}), f_{\text{spat}}(\mathbf{X}')) + f_{\theta_{\text{se}}}(f_{\theta_{\text{res}}}(\text{Stack}(f_{\text{temp}}(\mathbf{X}_{t-H+1:t}), f_{\text{spat}}(\mathbf{X}')))). \quad (9)$$

The model is trained end-to-end by minimizing the prediction loss:

$$\min_{\theta} \mathcal{L} = \frac{1}{|\mathcal{T}|} \sum_{t \in \mathcal{T}} \ell(\hat{\mathbf{X}}_{t+1:t+H'}, \mathbf{X}_{t+1:t+H'}), \quad (10)$$

where θ collects all learnable parameters in f_{temp} , f_{spat} , $f_{\theta_{\text{res}}}$ and $f_{\theta_{\text{se}}}$.

5 Experiments

5.1 Experimental Setup & Settings

Datasets: We evaluate our method on three different datasets:

- **SimSF-Bay:** SimSF-Bay, introduced in this work, is a traffic flow dataset derived from MATSim simulations of the SF Bay area’s transportation network. The underlying MATSim model and its configurations were adopted from [Papakonstantinou *et al.*, 2020]. The network comprises 7,709 road segments, with traffic flow recorded at 5-minute intervals. The traffic flow from the previous 45 minutes is used as input for predicting the next time step (prediction horizon $H' = 1$).
- **PEMS-Bay** [Li *et al.*, 2018]: PEMS-Bay is a traffic speed dataset collected from the California Transportation Agencies (CalTrans) Performance Measurement System (PeMS) for the Bay Area. It contains data from 325 sensors over a 6-month period, from January 1, 2017 to May 31, 2017. Traffic information is recorded at 5-minute intervals. Traffic data from the previous 7 days around the predicted time are used to forecast the next 12 time steps (prediction horizon $H' = 12$).
- **NYCTaxi** [Zheng *et al.*, 2014]: NYCTaxi is a taxi trajectory dataset for New York City, spanning from January 1, 2015 to March 1, 2015. The data are recorded at 30-minute intervals. Traffic flow from the previous 17.5 hours is used to forecast the flow for the next time step (prediction horizon $H' = 1$).

The PEMS-Bay and NYCTaxi datasets follow the standard setting and assume a fixed transportation network or spatial discretization, with traffic states evolving over time. In

contrast, SimSF-Bay captures scenario-dependent topological variations, including changes in network connectivity and link availability, while maintaining a fixed simulation horizon consistent with standard agent-based transport modeling practice. This setting enables the assessment of model robustness under structural network changes, which is not supported by existing benchmarks. Additional details on the datasets are provided in Sec. B.

Baselines & Metrics: We compare our method against a diverse set of established time series forecasting approaches, covering classical statistical models, temporal deep learning methods, and spatio-temporal graph neural networks. Specifically, the baselines include ARIMA [Kumar and Vanajakshi, 2015], LSTM [Hochreiter and Schmidhuber, 1997], TCN [Lea *et al.*, 2017], TSMixer and its extended variant TSMixer-ext [Chen *et al.*, 2023], as well as several spatio-temporal models, including STGCN [Yu *et al.*, 2018], ASTGCN [Guo *et al.*, 2019], ST-SSL [Ji *et al.*, 2023], DCRNN [Li *et al.*, 2018], and STEP [Shao *et al.*, 2022]. For evaluation, we follow standard practice and report RMSE, MAE, and MAPE, formally defined as follows:

$$\text{RMSE} = \sqrt{\frac{1}{N} \sum_{t=1}^N (y_t - \hat{y}_t)^2}, \quad \text{MAE} = \frac{1}{N} \sum_{t=1}^N |y_t - \hat{y}_t|,$$

$$\text{MAPE} = \frac{1}{N} \sum_{t=1}^N \left| \frac{y_t - \hat{y}_t}{y_t} \right| \times 100\%,$$

where $\{\hat{y}_i\}_{i=1}^N$ and $\{y_i\}_{i=1}^N$ denote the predictions and the ground truth values, respectively. It is noted that for MAPE calculation a mask was applied on the raw ground truth values to filter out zero and near-zero values.

Reproducibility: The source code, dataset, and evaluation scripts associated with this work are publicly available in our project repository¹.

5.2 Performance Comparison

Quantitative Evaluation: Table 1 reports the quantitative comparison results. On SimSF-Bay, UniST-Pred consistently outperforms all baselines, achieving the lowest errors across all evaluation metrics, while TS-Mixer achieves the second best RMSE. On PEMS-Bay, UniST-Pred performs on par with the strongest baseline, STEP, in terms of RMSE, and remains close to the best MAE (1.91 versus 1.79 for STEP). This indicates that the proposed method achieves performance comparable to the leading baseline on this dataset, while its MAPE remains competitive, though lagging behind STEP. On NYC-Taxi, UniST-Pred achieves the lowest RMSE and MAE, with noticeable margins over the closest baseline. Again, TS-Mixer achieves the second best RMSE, while STEP attains better MAPE, with UniST-Pred ranking second. Taken together, these results indicate that UniST-Pred consistently captures absolute error patterns, as reflected by having best results in RMSE and MAE across datasets.

¹<https://github.com/yw3576/UniST-Pred>

Method	SimSF-Bay			PEMS-Bay			NYCTaxi		
	RMSE	MAE	MAPE (%)	RMSE	MAE	MAPE (%)	RMSE	MAE	MAPE (%)
ARIMA	8.5 (-58%)	4.68 (-49%)	45.8 (-25%)	6.50 (-35%)	3.36 (-43%)	8.34 (-46%)	43.22 (-69%)	15.09 (-69%)	45.3 (-8%)
LSTM	6.4 (-44%)	4.23 (-44%)	53.6 (-36%)	4.96 (-15%)	2.37 (-19%)	5.70 (-21%)	22.30 (-40%)	10.34 (-55%)	122 (-66%)
TCN	5.2 (-31%)	3.40 (-30%)	45.3 (-25%)	5.11 (-18%)	2.76 (-31%)	5.97 (-24%)	15.88 (-16%)	5.45 (-14%)	42.3 (-2%)
TS-Mixer	3.9 (-8%)	2.60 (-9%)	37.1 (-8%)	4.91 (-14%)	2.69 (-29%)	6.06 (-25%)	14.82 (-10%)	5.20 (-10%)	45.0 (-8%)
TS-Mixer-ext	4.0 (-10%)	2.58 (-8%)	37.4 (-9%)	4.90 (-14%)	2.69 (-29%)	6.00 (-25%)	14.82 (-10%)	5.20 (-10%)	45.0 (-8%)
STGCN	6.2 (-42%)	4.11 (-42%)	50.6 (-32%)	5.69 (-26%)	2.49 (-23%)	5.79 (-22%)	26.52 (-50%)	13.34 (-65%)	58.5 (-29%)
ASTGCN	4.5 (-20%)	<u>2.51</u> (-6%)	49.8 (-31%)	5.42 (-23%)	2.61 (-27%)	6.00 (-25%)	25.30 (-47%)	9.22 (-49%)	47.6 (-13%)
ST-SSL	8.3 (-57%)	4.39 (-46%)	96.4 (-65%)	5.32 (-21%)	2.25 (-15%)	5.13 (-12%)	36.48 (-63%)	10.02 (-53%)	170 (-76%)
DCRNN	4.0 (-10%)	2.75 (-14%)	37.9 (-10%)	4.74 (-11%)	2.07 (-7%)	4.90 (-7%)	15.84 (-15%)	6.43 (-27%)	60.0 (-31%)
STEP	4.0 (-10%)	2.63 (-10%)	<u>36.9</u> (-7%)	4.20 (0%)	<u>1.91</u> (+7%)	4.18 (+8%)	22.62 (-41%)	10.41 (-55%)	32.6 (+28%)
UniST-Pred	3.6	2.37	34.2	4.20	<u>1.91</u>	<u>4.52</u>	13.39	4.70	<u>41.6</u>

[†] The results for PEMS-Bay evaluate the 12th timestep’s predictions; the results for NYC-Taxi evaluate the outflow predictions.

Table 1: Performance comparison on SimSF-Bay, PEMS-Bay, and NYC-Taxi datasets. The best results are **bold** and the second best results are underlined. The values in the brackets denote the relative performance gap w.r.t. UniST-Pred, computed as $100\% \cdot \frac{M_{\text{UniST-Pred}} - M_{\text{method}}}{M_{\text{method}}}$, where M represents the scores (negative values indicate worse performance).




Figure 2: Per-location MAPE distribution for UniST-Pred, STEP, and TSMixer on three datasets, stratified by high and low traffic flow. (Light colors indicate lower error.)

While average accuracy reflects overall performance, it does not show whether gains are consistent across locations or driven by a subset of easy cases. To examine spatial error patterns, Fig. 2 reports per-location MAPE distributions for TSMixer, STEP, and UniST-Pred across all datasets. Locations are divided by traffic flow percentiles, with **high flow** locations in the top 25% and **low flow** locations in the lower 75%. Each stacked bar shows the proportion of locations with low (0–33%), medium (33–66%), or high (66–100%) MAPE.

On SimSF-Bay, UniST-Pred achieves lower error on more locations than STEP and TSMixer, for both high-flow and low-flow locations. On PEMS-Bay and NYCTaxi, STEP shows a clear advantage on low-flow locations, while UniST-Pred matches or outperforms STEP on high-flow locations. These results suggest that UniST-Pred performs particularly well on high-traffic locations, which are often the most operationally important for traffic management applications.

Qualitative Evaluation: To examine prediction quality, Fig. 3 shows the ground truth and predictions over all roads for three datasets, comparing UniST-Pred with STEP and TSMixer on a representative sample period. For SimSF-Bay, all methods capture the main periodic pattern. In PEMS-Bay, UniST-Pred tracks the broad pattern of the averaged speed and preserves the overall level and long-range trend, whereas STEP and TSMixer exhibit larger deviations near abrupt lo-




Figure 3: Qualitative comparison of ground truth and model predictions.

cal changes. For NYCTaxi, all methods closely follow the ground truth across cycles, with UniST-Pred showing slightly higher fidelity, while the other methods exhibit slightly increased smoothing or variability. These results suggest that UniST-Pred (i) captures dominant temporal dynamics with better phase alignment, (ii) maintains stable predictions without noticeable drift over long horizons, and (iii) provides smooth and consistent estimates of the global signal, which is beneficial when the objective is accurate trend-level forecasting on network-aggregated traffic patterns.

Model Size: Next, we compare the number of model parameters of UniST-Pred against STEP, one of the strongest baseline methods. As shown in Table 2, UniST-Pred is substantially more parameter efficient than STEP across all three datasets. This reduction is particularly noticeable on SimSF-Bay and PEMS-Bay, where our model is approximately $4.0\times$ and $3.7\times$ smaller than STEP, respectively, and becomes especially significant on NYC-Taxi, where UniST-Pred is over $30\times$ smaller. This consistent reduction indicates that the proposed architecture achieves a favorable accuracy to efficiency trade-off, making it more suitable for deployment un-

der computational constraints while maintaining strong predictive performance.

Method	SimSF-Bay	PEMS-Bay	NYCTaxi
STEP	12,952,470	61,138,261	5,670,166
UniST	3,208,208	16,558,467	167,994
Reduction (%)	75.23	72.91	97.04

Table 2: Comparison of model parameter counts between STEP and UniST across datasets.

Multi-horizon Forecasting: We additionally evaluate UniST-Pred on SimSF-Bay under varying prediction horizons, specifically $H' = 3, 6, 9$. The results, listed in Table 3, show that UniST-Pred consistently achieves the lowest error metrics across all considered horizons compared with the benchmark methods. This suggests that the proposed approach generalizes well to multi-horizon forecasting settings while preserving strong predictive performance.

Horizon		TS-Mixer	STEP	UniST-Pred
$H' = 3$	RMSE	4.14	4.38	3.73
	MAE	2.67	2.84	2.40
	MAPE (%)	38.23	41.28	33.29
$H' = 6$	RMSE	4.33	4.52	3.89
	MAE	2.96	2.92	2.49
	MAPE (%)	40.47	41.92	34.27
$H' = 9$	RMSE	4.56	4.75	4.06
	MAE	3.01	3.06	2.60
	MAPE (%)	43.01	43.62	35.37

Table 3: Performance comparison of TS-Mixer, STEP and UniST-Pred on SimSF-Bay across different forecasting horizons.

5.3 Ablation Studies

To analyze the contribution of individual components in the proposed architecture, we conduct an ablation study by systematically removing or replacing key sub-modules. Specifically, we consider variants that remove the spatial block, the temporal block, or the fusion block, as well as variants that replace the spatial block with standard GCN layers or the temporal block with a fully connected (FC) layers.

Method	RMSE ↓	MAE ↓	MAPE (%) ↓
UniST-Pred (full)	3.61	2.37	34.2
<i>Component removal</i>			
w/o spatial	6.06	3.43	65.5
w/o temporal	5.64	4.11	49.8
w/o fusion	4.06	2.57	36.2
<i>Component replacement</i>			
spatial → GCN	4.17	2.65	71.2
temporal → FC	3.96	2.58	38.1

Table 4: Ablation results of UniST-Pred on SimSF-Bay.

Table 4 presents the ablation study results. We observe that removing either the *spatial block* or the *temporal block* consistently degrades performance across all metrics. Similarly, replacing the spatial block with standard GCN layers or the temporal block with fully connected layers increases the error. The *w/o fusion* variant performs competitively among simplified settings, indicating that the backbone alone can learn informative representations. The full model achieves the lowest error across all metrics, confirming that all proposed components contribute to the best overall performance.

5.4 Case Study: Robustness & Explainability

We conclude the empirical evaluation by examining the robustness of UniST-Pred under major infrastructure disruptions. To this end, we use the adopted MATSim model to construct six stress-test scenarios on the SF Bay Area network by removing major bridges or freeway/highway links that play a critical role in the region’s connectivity. Specifically, we consider the removal/closing of (**Scenario 1**) Oakland Bay Bridge; (**Scenario 2**) Carquinez Bridge and San Rafael Bridge; (**Scenario 3**) Golden Gate Bridge and Dumbarton Bridge; (**Scenario 4**) US-101 highway (near Brisbane); (**Scenario 5**) the shared segment of I-80 and I-580 freeways near Albany; and (**Scenario 6**) the links considered in Scenarios 4 and 5.

Robustness Under Disruption: Fig. 4 compares the ground truth traffic flow and the predictions produced by UniST-Pred, both at the level of the network-wide average and for two representative roads, across all six scenarios. At the aggregate level (top row), the predictions closely track the ground truth. For individual roads (middle and bottom rows), the predictions remain trend-consistent, though brief local spikes may be smoothed or missed.

A quantitative comparison against TSMixer and STEP is reported in Table 5, which shows that UniST-Pred attains the lowest RMSE in all six disruption scenarios as well as in the fully connected setting, indicating consistently strong performance across different connectivity conditions.

	TS-Mixer	STEP	UniST-Pred
Scenario 1	3.89	3.93	3.61
Scenario 2	3.87	3.95	3.59
Scenario 3	3.87	3.91	3.60
Scenario 4	3.89	3.91	3.59
Scenario 5	3.87	3.91	3.58
Scenario 6	3.88	3.90	3.58
Fully connected	3.90	3.89	3.62

Table 5: RMSE of UniST-Pred, TSMixer, and STEP across the disruption scenarios and the fully connected network.

Explainability Analysis: Complementary to the above, we study the attribution behavior of UniST-Pred under disruption scenarios by examining which road segments contribute most to its predictions. To this end, we first leverage the trained UniST-Pred model to generate traffic predictions, then apply Integrated Gradients (IG), which quantifies the contribution of each input feature by integrating gradients along a




Figure 4: Comparison of ground-truth traffic flow and UniST-Pred predictions across the six disruption scenarios, shown for the network-wide average (top row) and two representative roads (middle and bottom rows).




Figure 5: Explainability map for (c) Scenario 1: disconnecting Oakland Bay Bridge, (f) Scenario 4: disconnecting US-101, (b) and (e): all connected network. The zoom-in parts in (a) (b) (c) are the Oakland Bay Bridge and its surrounding area. The zoom-in parts in (d) (e) (f) are US-101 and its surrounding area. Darker color indicates higher contribution to the prediction.

straight-line path from a baseline $\mathbf{x}'$ to the input $\mathbf{x}$. More concretely, for a differentiable scalar output $F(\mathbf{x})$, the IG attribution for feature i is calculated as $\text{IG}_i(\mathbf{x}) = (x_i - x'_i) \int_0^1 \frac{\partial F(\mathbf{x}' + \alpha(\mathbf{x} - \mathbf{x}'))}{\partial x_i} d\alpha$.

Fig. 5 visualizes the resulting road-level IG contribution maps, where darker segments correspond to higher contribution to the model’s prediction. The figure contrasts the fully connected network (middle panel) with the disrupted cases in Scenario 1 and Scenario 4 (right panels), revealing a clear redistribution of attribution under disruption. Specifically, in the disruption scenarios, the model’s attribution shifts away from disconnected links toward alternative corridors, suggesting that the prediction mechanism adapts to the altered network topology.

6 Conclusion

This work studies spatio-temporal traffic forecasting as a modular modeling problem and proposes a unified framework that decouples temporal dependency learning from spatial representation extraction with adaptive fusion. By avoiding tightly coupled spatio-temporal propagation, the framework provides a lightweight yet expressive alternative for traffic prediction. Experiments on multiple benchmarks show that the proposed design achieves competitive performance with a compact architecture, while additional analyses offer insights into the contributions of temporal and spatial components. More broadly, this study suggests that a *decouple-then-fuse* strategy can serve as a useful design principle for spatio-temporal learning, with relevance beyond traffic forecasting and practical importance for real world traffic networks.

A MATSim Simulator & SF Bay Transportation Network Model

MATSim (Multi-Agent Transport Simulation) [W Axhausen *et al.*, 2016] is a microscopic, agent-based transportation simulator widely utilized in transportation planning and mobility research. MATSim models travel demand explicitly at the individual level, representing each traveler as an autonomous agent with daily activity plans that evolve through iterative re-planning and network loading. At each iteration, agents execute their plans on a detailed transportation network, generating time-resolved link-level traffic states such as flow, speed, and travel time, which are subsequently used to update agent plans based on experienced costs and utilities defined by user-assigned scoring functions. The simulation is considered converged once agents’ plans and experienced costs stabilize across successive iterations, corresponding to a stochastic user equilibrium. Owing to its realism and flexibility, MATSim is commonly employed by academic researchers, transportation authorities, and planning agencies for scenario analysis, policy evaluation, and infrastructure planning.

Within the scope of current study, we adopt from [Papakonstantinou *et al.*, 2020] the MATSim model calibrated for the San Francisco Bay Area. In this model, the commuters travel plan is based on the vehicular travel information from Pozdnoukhov *et al.* [Pozdnoukhov *et al.*,

2016]. The demand model is based on anonymised cellular network infrastructure data stream and census data from the 2010-2012 California Household Travel Survey data. The demand generation relies on the 1454 Traffic Analysis Zones in the area developed by the Metropolitan Transportation Commission (see https://abag.ca.gov/sites/default/files/pba_2050-regional_growth_forecast_methodology.pdf). The model considers a sample population of 463,938 commuters and the road links in the transportation network are scaled down to 8% of their original capacities to correctly match the population scale.

B Details of Datasets

Property	PEMS-BAY	NYC Taxi	SimSF-Bay
Traffic variable	Speed	Flow	Flow
Spatial structure	Sensor graph	Fixed zones / grid	Dynamic road network
# of nodes	325	200	7709
# of edges	2369	712	8781
# of time steps	5.2k+	2.6m+	168 (per scenario)
Time interval	5 min	30 min	5 min
Time range	01/01/2017 - 05/31/2017	01/01/2015 to 03/01/2015	6 AM to 8 PM (per scenario)
Network topology	Fixed	Fixed	Varies across scenarios

Table 6: Statistics of datasets used in this study.

C Model Parameters & Implementation Details

For the proposed UniST-Pred, the model and implementation details are listed in Table 7. ARIMA was implemented via statsmodels python package [Skipper and Josef, 2010]. LSTM was implemented from Pytorch Neural Networks package [Team, 2023]. Other benchmarks were implemented via the code provided in their project repositories: TCN [Lea *et al.*, 2017], TSMixer and its extended variant TSMixer-ext [Chen *et al.*, 2023], STGCN [Yu *et al.*, 2018], ASTGCN [Guo *et al.*, 2019], ST-SSL [Ji *et al.*, 2023], DCRNN [Li *et al.*, 2018], and STEP [Shao *et al.*, 2022]. For LSTM and TCN, the hidden dimensions were selected via grid search. For TSMixer, we adopted the hyperparameter search procedure provided in its repository. For the graph-based baselines, we used the configurations released with their codebase.

Item	SimSF-Bay	PEMS-Bay	NYCTaxi
Model Parameters			
Input length (H)	9	2016	35
Prediction horizon (H')	1	12	1
Number of channels (C)	4	4	4
Number of GT layers (L_g)	2	2	2
GCN in channel	2	2	1
GCN out channel	100	100	100
Number of TSMixer layers (K)	2	4	2
feature mixing MLP width	100	100	100
Dropout	0.2	0.2	0.2
SE reduction ratio (r)	2	4	2
Implementation Details			
Hardware	NVIDIA A100 80GB PCIe		
Optimizer	ADAM		
Initial learning rate	0.0005	0.0005	0.0005
Weight decay	0.1	0.1	0.1
Batch size	5	64	5
Epochs	20	50	100
Loss function	MAE	Smoothed L1	MAE
Train/val/test split	66%/17%/17%	70%/10%/20%	70%/10%/20%

Table 7: Model parameters and implementation details

Acknowledgments

The authors thank Madeline Sheehan for providing us with the data of the San Francisco Bay Area highway network and the code for processing it. This research was carried out on the High Performance Computing resources at New York University Abu Dhabi.

References

- [Bai et al., 2020] Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. Adaptive graph convolutional recurrent network for traffic forecasting. *Advances in neural information processing systems*, 33:17804–17815, 2020.
- [Chen et al., 2023] Si-An Chen, Chun-Liang Li, Serkan Ö Arik, Nathanael C Yoder, and Tomas Pfister. Tsmixer: An all-mlp architecture for time series forecast-ing. *Trans. Mach. Learn. Res.*, 2023.
- [Gao et al., 2024] Haotian Gao, Renhe Jiang, Zheng Dong, Jinliang Deng, Yuxin Ma, and Xuan Song. Spatial-temporal-decoupled masked pre-training for spatiotemporal forecasting. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pages 3998–4006, 2024.
- [Guo et al., 2019] Shengnan Guo, Youfang Lin, Ning Feng, Chao Song, and Huaiyu Wan. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 922–929, 2019.
- [Hochreiter and Schmidhuber, 1997] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, November 1997.
- [Hu et al., 2018] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [Ji et al., 2023] Jiahao Ji, Jingyuan Wang, Chao Huang, Junjie Wu, Boren Xu, Zhenhe Wu, Junbo Zhang, and Yu Zheng. Spatio-temporal self-supervised learning for traffic flow prediction. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pages 4356–4364, 2023.
- [Jiang et al., 2023a] Jiawei Jiang, Chengkai Han, Wayne Xin Zhao, and Jingyuan Wang. Pdformer: Propagation delay-aware dynamic long-range transformer for traffic flow prediction. *arXiv preprint arXiv:2301.07945*, 2023.
- [Jiang et al., 2023b] Weiwei Jiang, Jiayun Luo, Miao He, and Weixi Gu. Graph neural network for traffic forecasting: The research progress. *ISPRS International Journal of Geo-Information*, 12(3):100, 2023.
- [Kumar and Vanajakshi, 2015] S Vasantha Kumar and Lelitha Vanajakshi. Short-term traffic flow prediction using seasonal arima model with limited input data. *European Transport Research Review*, 7(3):21, 2015.
- [Lea et al., 2017] Colin Lea, Michael D Flynn, Rene Vidal, Austin Reiter, and Gregory D Hager. Temporal convolutional networks for action segmentation and detection. In *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 156–165, 2017.
- [Li and Zhu, 2021] Mengzhang Li and Zhanxing Zhu. Spatial-temporal fusion graph neural networks for traffic flow forecasting. In *AAAI*, 2021.
- [Li et al., 2018] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *International Conference on Learning Representations (ICLR '18)*, 2018.
- [Li et al., 2022] Wenqing Li, Chuhan Yang, and Saif Eddin Jabari. Nonlinear traffic prediction as a matrix completion problem with ensemble learning. *Transportation science*, 56(1):52–78, 2022.
- [Liu et al., 2023] Hangchen Liu, Zheng Dong, Renhe Jiang, Jiewen Deng, Jinliang Deng, Quanjun Chen, and Xuan Song. Spatio-temporal adaptive embedding makes vanilla transformer sota for traffic forecasting. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pages 4125–4129, 2023.
- [Liu et al., 2024] Yong Liu, Tengge Hu, Haoran Zhang, Haixu Wu, Shiyu Wang, Lintao Ma, and Mingsheng Long. itransformer: Inverted transformers are effective for time series forecasting. In *ICLR*. OpenReview.net, 2024.
- [Manibardo et al., 2021] Eric L Manibardo, Ibai Laña, and Javier Del Ser. Deep learning for road traffic forecasting: Does it make a difference? *IEEE Transactions on Intelligent Transportation Systems*, 23(7):6164–6188, 2021.
- [Nie et al., 2023] Yuqi Nie, Nam H. Nguyen, Phanwadee Sinthong, and Jayant Kalagnanam. A time series is worth 64 words: Long-term forecasting with transformers. In *ICLR*. OpenReview.net, 2023.
- [Papakonstantinou et al., 2020] Ilia Papakonstantinou, Alain Tcheukam Siwe, and Samer Michel Madanat. Effects of sea level rise induced land use changes on traffic congestion. *Transportation Research Part D: Transport and Environment*, 87:102515, 2020.
- [Pozdnoukhov et al., 2016] Alexei Pozdnoukhov, Andrew Campbell, Sidney Feygin, Mogeng Yin, and Sudatta Mohanty. San francisco bay area: The smartbay project-connected mobility. In *Multi-Agent Transport Simulation MATSim*. Ubiquity Press, 2016.
- [Shao et al., 2022] Zezhi Shao, Zhao Zhang, Fei Wang, and Yongjun Xu. Pre-training enhanced spatial-temporal graph neural network for multivariate time series forecasting. In *Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining*, pages 1567–1577, 2022.
- [Skipper and Josef, 2010] Seabold Skipper and Perktold Josef. statsmodels: Econometric and statistical modeling with python. *9th Python in Science Conference*, 2010.
- [Team, 2023] PyTorch Team. Pytorch: An open source machine learning library for python. <https://docs.pytorch.org/docs/stable/nn.html>, 2023.

- [W Axhausen *et al.*, 2016] Kay W Axhausen, Andreas Horni, and Kai Nagel. *The multi-agent transport simulation MATSim*. Ubiquity Press, 2016.
- [Wang *et al.*, 2024] Hongjun Wang, Jiyuan Chen, Tong Pan, Zheng Dong, Lingyu Zhang, Renhe Jiang, and Xuan Song. Stgformer: Efficient spatiotemporal graph transformer for traffic forecasting. *arXiv preprint arXiv:2410.00385*, 2024.
- [Wei *et al.*, 2019] Hua Wei, Guanjie Zheng, Vikash Gayah, and Zhenhui Li. A survey on traffic signal control methods. *Transportation Research Part C*, 99:313–334, 2019.
- [Wu *et al.*, 2019] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep spatial-temporal graph modeling. *arXiv preprint arXiv:1906.00121*, 2019.
- [Wu *et al.*, 2020] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, Xiaojun Chang, and Chengqi Zhang. 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*, pages 753–763, 2020.
- [Yu *et al.*, 2018] Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In *IJCAI*, page 3634–3640, 2018.
- [Yun *et al.*, 2019] Seongjun Yun, Minbyul Jeong, Raehyun Kim, Jaewoo Kang, and Hyunwoo J Kim. Graph transformer networks. *Advances in neural information processing systems*, 32, 2019.
- [Zhang *et al.*, 2024] Jian Zhang, Lili Guo, Lei Song, Song Gao, Chuanzhu Hao, and Xuzhi Li. Patchtcn: Patch-based transformer convolutional network for times series analysis. In *Proceedings of the 2024 3rd International Symposium on Computing and Artificial Intelligence*, pages 1–9, 2024.
- [Zheng *et al.*, 2014] Yu Zheng, Licia Capra, Ouri Wolfson, and Hai Yang. Urban computing with taxi trajectories. *ACM Transactions on Intelligent Systems and Technology*, 2014.
- [Zheng *et al.*, 2019] Chuanpan Zheng, Xiaoliang Fan, Cheng Wang, and Jianzhong Qi. Gman: A graph multi-attention network for traffic prediction. *arXiv preprint arXiv:1911.08415*, 2019.