

# LIGHTWEIGHT AND INTERPRETABLE TRANSFORMER VIA MIXED GRAPH ALGORITHM UNROLLING FOR TRAFFIC FORECAST

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Unlike conventional “black-box” transformers with classical self-attention mechanism, we build a lightweight and interpretable transformer-like neural net by unrolling a mixed-graph-based optimization algorithm to forecast traffic with spatial and temporal dimensions. We construct two graphs: an undirected graph  $\mathcal{G}^u$  capturing spatial correlations across geography, and a directed graph  $\mathcal{G}^d$  capturing sequential relationships over time. We predict future samples of signal  $\mathbf{x}$ , assuming it is “smooth” with respect to both  $\mathcal{G}^u$  and  $\mathcal{G}^d$ , where we design new  $\ell_2$  and  $\ell_1$ -norm variational terms to quantify and promote signal smoothness (low-frequency reconstruction) on a directed graph. We design an iterative algorithm based on alternating direction method of multipliers (ADMM), and unroll it into a feed-forward network for data-driven parameter learning. We insert graph learning modules for  $\mathcal{G}^u$  and  $\mathcal{G}^d$  that play the role of self-attention. Experiments show that our unrolled networks achieve competitive traffic forecast performance as state-of-the-art prediction schemes, while reducing parameter counts drastically.

## 1 INTRODUCTION

Transformer, based on the classical self-attention mechanism (Vaswani et al., 2017), is now a dominant deep learning (DL) architecture that has achieved state-of-the-art (SOTA) performance across multiple fields (Woźniak et al., 2020; Zamir et al., 2022). However, like other “off-the-shelf” DL models, transformer has a massive number of parameters, and its internal structure is not easily interpretable. Parameter reduction is a major industrial concern for practical computation and/or memory-constrained environments<sup>1</sup>. *Algorithm unrolling* (Monga et al., 2021) offers an alternative hybrid model-based / data-driven paradigm; by “unrolling” iterations of a model-based algorithm minimizing a well-defined optimization objective into neural layers stacked together to form a feed-forward network for data-driven parameter learning, the resulting network can be *both* mathematically interpretable<sup>2</sup> and competitive in performance. Notably, Yu et al. (2023) designs an algorithm minimizing a sparse rate reduction (SRR) objective that unrolls into a transformer-like neural net for image classification. Thuc et al. (2024) designs graph-based algorithms minimizing graph smoothness priors that unroll into transformer-like neural nets for image interpolation while reducing parameters. However, only a positive *undirected* graph model was used in Thuc et al. (2024) to capture simple pairwise correlations between neighboring pixels in a static image.

*Spatial-temporal data* (e.g., traffic and weather) contains more complex node-to-node relationships: geographically near stations exhibit *spatial correlations*, while past observations influences future data, resulting in *sequential relationships*. In this paper, we study the design of transformers via graph algorithm unrolling for spatial-temporal data, using traffic forecast as a concrete example application.

We roughly categorize traffic forecast methods into model-based and DL-based methods. Among model-based methods, one approach to traffic forecast is to filter the signal along the spatial and temporal dimensions independently (Ramakrishna et al., 2020)—e.g., employ graph spectral filters in the *graph signal processing* (GSP) literature (Ortega et al., 2018; Cheung et al., 2018) along the

<sup>1</sup>See a brief review of lightweight learning models for different applications in Appendix A.1.

<sup>2</sup>Common in algorithm unrolling (Monga et al., 2021), “interpretability” here means that each neural layer corresponds to an iteration of an optimization algorithm minimizing a mathematically-defined objective.

irregular spatial dimension, and traditional Fourier filters along the regular time dimension. Another approach is to model the spatial correlations across time as a sequence of slowly time-varying graphs, and then filter each spatial signal at time  $t$  separately, where the changes in consecutive graph adjacency matrices in time are constrained by the Frobenius norm (Kalofolias et al., 2017),  $\ell_1$ -norm (Yamada et al., 2019), or low-rankness (Bagheri et al., 2024). However, neither approach exploits spatial and temporal relationships simultaneously for optimal performance.

Recent DL-based efforts address traffic forecast by building elaborate transformer architectures to capture spatial and temporal relations (Feng & Tassioulas, 2022; Gao et al., 2022; Jin et al., 2024)<sup>3</sup>. Like transformers used in other fields (Zamir et al., 2022), these are also black boxes with large parameter counts. *Graph attention networks* (GAT) (Velickovic et al., 2018) and *graph transformers* (Dwivedi & Bresson, 2021) are adaptations of the self-attention mechanism in transformers for graph-structured data (e.g., computing output embeddings using only input embeddings in local graph neighborhoods). However, they are still uninterpretable and maintain large parameter sizes.

In contrast, we extend Thuc et al. (2024) to build lightweight “white-box” transformers that are fully interpretable via *mixed-graph* algorithm unrolling for traffic forecast. Specifically, we learn *two* graphs from data: i) *undirected* graph  $\mathcal{G}^u$  to capture spatial correlations across geography, and ii) a *directed* graph  $\mathcal{G}^d$  to capture sequential relationships over time. We show that our graph learning modules play the role of self-attention (Bahdanau et al., 2014), and thus our unrolled graph-based neural nets are transformers. Given the two learned graphs, we define a prediction objective for future samples in signal  $\mathbf{x}$ , assuming  $\mathbf{x}$  is “smooth” with respect to (w.r.t.) both  $\mathcal{G}^u$  and  $\mathcal{G}^d$  in variational terms: *graph Laplacian regularizer* (GLR) (Pang & Cheung, 2017) for undirected  $\mathcal{G}^u$ , and newly designed *directed graph Laplacian regularizer* (DGLR) and *directed graph total variation* (DGTV) for directed  $\mathcal{G}^d$ . We design a corresponding linear-time optimization algorithm based on *alternating direction method of multipliers* (ADMM) (Boyd et al., 2011) that unrolls into neural layers for parameter learning. **For the first time, our DGLR and DGTV variational terms quantify and promote signal smoothness on a directed graph, with intuitive low-pass filter interpretations**<sup>4</sup>. Experiments show that our networks achieve competitive traffic forecast performance as SOTA prediction schemes, while employing drastically fewer parameters (**7.2% of transformer-based PDFormer** (Jiang et al., 2023)).

## 2 PRELIMINARIES

**Undirected Graph Definitions:** Denote by  $\mathcal{G}^u(\mathcal{V}, \mathcal{E}^u, \mathbf{W}^u)$  an *undirected* graph with node set  $\mathcal{V} = \{1, \dots, N\}$  and undirected edge set  $\mathcal{E}^u$ , where  $(i, j) \in \mathcal{E}^u$  implies that an undirected edge exists connecting nodes  $i$  and  $j$  with weight  $w_{i,j}^u = W_{i,j}^u$ , and  $(i, j) \notin \mathcal{E}^u$  implies  $W_{i,j}^u = 0$ .  $\mathbf{W}^u \in \mathbb{R}^{N \times N}$  is the symmetric *adjacency matrix* for  $\mathcal{G}^u$ . Denote by  $\mathbf{D}^u \in \mathbb{R}^{N \times N}$  a diagonal *degree matrix*, where  $D_{i,i}^u = \sum_j W_{j,i}^u$ . The symmetric *graph Laplacian matrix* is  $\mathbf{L}^u \triangleq \mathbf{D}^u - \mathbf{W}^u$  (Ortega et al., 2018).  $\mathbf{L}^u$  is *positive semi-definite* (PSD) if all edge weights are non-negative  $w_{i,j}^u \geq 0, \forall i, j$  (Cheung et al., 2018). The symmetric *normalized graph Laplacian matrix* is  $\mathbf{L}_n^u \triangleq \mathbf{I} - (\mathbf{D}^u)^{-1/2} \mathbf{W}^u (\mathbf{D}^u)^{-1/2}$ , while the asymmetric *random-walk graph Laplacian matrix* is  $\mathbf{L}_r^u \triangleq \mathbf{I} - (\mathbf{D}^u)^{-1} \mathbf{W}^u$ .

One can define a *graph spectrum* by eigen-decomposing the real and symmetric graph Laplacian  $\mathbf{L}^u$  (or normalized graph Laplacian  $\mathbf{L}_n^u$ ), where the  $k$ -th eigen-pair  $(\lambda_k, \mathbf{v}_k)$  is interpreted as the  $k$ -th graph frequency and Fourier mode, respectively (Ortega et al., 2018). Given  $\{\mathbf{v}_k\}$  are orthonormal vectors, one can decompose a signal  $\mathbf{x}$  into its graph frequency components as  $\alpha = \mathbf{V}^\top \mathbf{x}$ , where  $\mathbf{L}^u = \mathbf{V} \text{diag}(\{\lambda_k\}) \mathbf{V}^\top$ , and  $\mathbf{V}^\top$  is the *graph Fourier transform* (GFT).

The  $\ell_2$ -norm *graph Laplacian regularizer* (GLR) (Pang & Cheung, 2017),  $\mathbf{x}^\top \mathbf{L}^u \mathbf{x} = \sum_{(i,j) \in \mathcal{E}^u} w_{i,j}^u (x_i - x_j)^2$ , is used to regularize an ill-posed signal restoration problem, like denoising, to bias low-frequency signal reconstruction (i.e., signals consistent with similarity graph  $\mathcal{G}^u$ )

<sup>3</sup>A detailed review of DL models on traffic forecast is included in Appendix A.2.

<sup>4</sup>See a detailed review of related works on directed graph Laplacians and discussion of our novelty in directed graph frequency analysis in Appendix A.3.

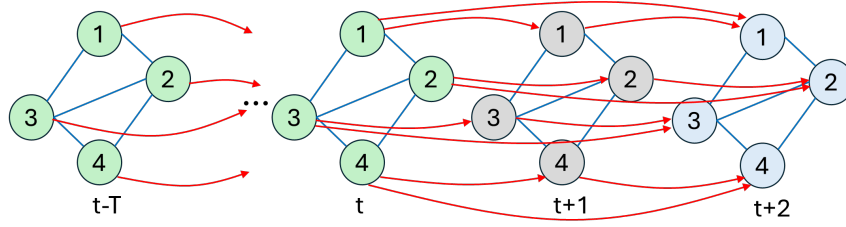


Figure 1: Example of a mixed graph with undirected edges (blue) connecting nodes of the same time instants, and directed edges (red) connecting nodes at  $t$  to nodes in window  $\{t+1, t+2\}$ .

given observation  $\mathbf{y}$ :

$$\begin{aligned} \mathbf{x}^* &= \arg \min_{\mathbf{x}} \|\mathbf{y} - \mathbf{x}\|_2^2 + \mu \mathbf{x}^\top \mathbf{L}^u \mathbf{x} \\ &= (\mathbf{I} + \mu \mathbf{L}^u)^{-1} \mathbf{y} = \mathbf{V} \text{diag} \left( \frac{1}{1 + \mu \lambda_1}, \dots, \frac{1}{1 + \mu \lambda_N} \right) \mathbf{V}^\top \mathbf{y}. \end{aligned} \quad (1)$$

$\mu$  is a bias-variance tradeoff parameter, and the low-pass filter response is  $f(\lambda) = (1 + \mu \lambda)^{-1}$ .

**Directed Graph Definitions:** We define analogous notations for a *directed* graph, denoted by  $\mathcal{G}^d(\mathcal{V}, \mathcal{E}^d, \mathbf{W}^d)$ .  $[i, j] \in \mathcal{E}^d$  implies a directed edge exists from node  $i$  to  $j$  with weight  $w_{j,i}^d = W_{j,i}^d$ .  $\mathbf{W}^d \in \mathbb{R}^{N \times N}$  is the asymmetric *adjacency matrix*. Denote by  $\mathbf{D}^d \in \mathbb{R}^{N \times N}$  a diagonal *in-degree matrix*, where  $D_{i,i}^d = \sum_j W_{i,j}^d$ . The *directed graph Laplacian* matrix is defined as  $\mathbf{L}^d \triangleq \mathbf{D}^d - \mathbf{W}^d$ . The *directed random-walk graph Laplacian* matrix is  $\mathbf{L}_r^d \triangleq \mathbf{I} - (\mathbf{D}^d)^{-1} \mathbf{W}^d$ .

Similar frequency notion cannot be simply defined for directed graphs via eigen-decomposition of directed graph Laplacian  $\mathbf{L}^d$ , because  $\mathbf{L}^d$  is asymmetric, and thus the Spectral Theorem (Hawkins, 1975) does not apply. We circumvent this problem by designing new variational terms in Section 3.2.

### 3 OPTIMIZATION FORMULATION & ALGORITHM

#### 3.1 MIXED GRAPH FOR SPATIAL/TEMPORAL DATA

We describe the construction of a mixed graph for spatial/temporal data. A measurement station  $i$  observes samples  $x_i^{t-T}, \dots, x_i^t$  at past and present instants  $t-T, \dots, t$  and samples  $x_i^{t+1}, \dots, x_i^{t+S}$  at future instants  $t+1, \dots, t+S$ . Denote by  $\mathbf{x} = [\mathbf{x}^{t-T}; \dots; \mathbf{x}^{t+S}] \in \mathbb{R}^{N(T+S+1)}$  the target graph signal—a concatenation of samples from all  $N$  nodes across all  $T+S+1$  instants. A *product graph*  $\mathcal{G}$  of  $N \times (T+S+1)$  nodes represents samples at  $T+S+1$  instants.  $\mathcal{G}$  is a mixture of: i) undirected graph  $\mathcal{G}^u$  connecting spatially node pairs of the same instants, and ii) directed graph  $\mathcal{G}^d$  connecting temporally each node  $i$  at instant  $\tau$  to the same node at instants  $\tau+1, \dots, \tau+W$ , where  $W$  denotes the pre-defined *time window*. See Fig. 1 for an illustration of a mixed graph with undirected spatial edges (blue) and directed temporal edges (red) for  $W=2$ .

#### 3.2 VARIATIONAL TERMS FOR DIRECTED GRAPHS

Our constructed  $\mathcal{G}^d$  is a *directed acyclic graph* (DAG); a directed edge always stems from a node at instant  $\tau$  to a node at a future instant  $\tau+s$ , and thus no cycles exist. We first define a symmetric variational term for  $\mathcal{G}^d$  called *directed graph Laplacian regularizer* (DGLR) to quantify variation of a signal  $\mathbf{x}$  on  $\mathcal{G}^d$ . Denote by  $\mathcal{S} \subset \mathcal{N}$  the set of *source nodes* with zero in-degrees, and  $\bar{\mathcal{S}} \triangleq \mathcal{N} \setminus \mathcal{S}$  its complement. First, we add a self-loop of weight 1 to each node in  $\mathcal{S}$ ; this ensures that  $(\mathbf{D}^d)^{-1}$  is well-defined. Next, we define the row-stochastic *random-walk adjacency matrix*  $\mathbf{W}_r^d \triangleq (\mathbf{D}^d)^{-1} \mathbf{W}^d$ .

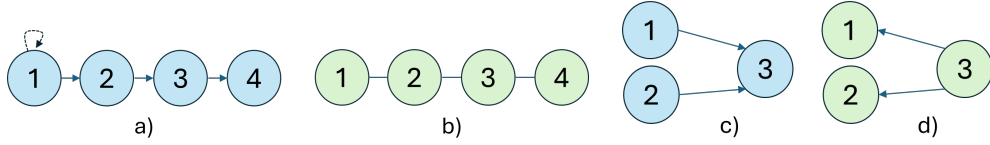


Figure 2: Example of a 4-node DAG  $\mathcal{G}^d$ , with an added self-loop at node 1, specified by  $\mathbf{L}_r^d$  (a), and corresponding undirected graph  $\mathcal{G}^u$ , specified by  $\mathcal{L}_r^d = (\mathbf{L}_r^d)^\top \mathbf{L}_r^d = \mathbf{L}^u$  (b). Example of a 3-node DAG (c), and a 3-node DAG with opposite directional edges (d).

### 3.2.1 $\ell_2$ -NORM VARIATIONAL TERM

Interpreting  $\mathbf{W}_r^d$  as a *graph shift operator* (GSO) (Chen et al., 2015) on  $\mathcal{G}^d$ , we first define an  $\ell_2$ -norm variational term<sup>5</sup> for  $\mathcal{G}^d$  as the squared difference between  $\mathbf{x}$  and its graph-shifted version  $\mathbf{W}_r^d \mathbf{x}$ :

$$\|\mathbf{x} - \mathbf{W}_r^d \mathbf{x}\|_2^2 = \|(\mathbf{I} - \mathbf{W}_r^d) \mathbf{x}\|_2^2 = \mathbf{x}^\top \underbrace{(\mathbf{L}_r^d)^\top \mathbf{L}_r^d}_{\mathcal{L}_r^d} \mathbf{x} \quad (2)$$

where *symmetrized directed graph Laplacian* matrix  $\mathcal{L}_r^d \triangleq (\mathbf{L}_r^d)^\top \mathbf{L}_r^d$  is symmetric and PSD. We call  $\mathbf{x}^\top \mathcal{L}_r^d \mathbf{x}$  the DGLR, which computes the sum of squared differences between each child  $j \in \bar{\mathcal{S}}$  and its parents  $i$  for  $[i, j] \in \mathcal{E}^d$ . Note that DGLR computes to zero for the constant vector  $\mathbf{1}$ :

$$\mathbf{1}^\top \mathcal{L}_r^d \mathbf{1} = \mathbf{1}^\top (\mathbf{L}_r^d)^\top \mathbf{L}_r^d \mathbf{1} = \mathbf{1}^\top (\mathbf{L}_r^d)^\top (\mathbf{I} - \mathbf{W}_r^d) \mathbf{1} \stackrel{(a)}{=} \mathbf{0}$$

where (a) follows from  $\mathbf{W}_r^d$  being row-stochastic. This makes sense, as constant  $\mathbf{1}$  is the smoothest signal and has no variation across any graph kernels.

We interpret eigen-pairs  $\{(\xi_k, \mathbf{u}_k)\}$  of  $\mathcal{L}_r^d$  as graph frequencies and graph Fourier modes for directed graph  $\mathcal{G}^d$ , respectively, similar to eigen-pairs  $\{(\lambda_k, \mathbf{v}_k)\}$  of graph Laplacian  $\mathbf{L}^u$  for an undirected graph. Importantly, we prove that our frequency definition using DGLR defaults to pure sinusoids as frequencies in the unweighted directed line graph case. Thus, our algorithm minimizing DGLR includes GSP schemes like Ramakrishna et al. (2020) that analyze signals along the time dimension using classical Fourier filters as special cases.

**Theorem 3.1.** Consider a directed line graph  $\mathcal{G}^d$  of  $N$  nodes with directed edge weights equal to 1, where the first (source) node is augmented with a self-loop of weight 1. The symmetrized directed graph Laplacian  $\mathcal{L}_r^d = (\mathbf{L}_r^d)^\top \mathbf{L}_r^d$ , where  $\mathbf{L}_r^d$  is the random-walk graph Laplacian for  $\mathcal{G}^d$ , is the same as graph Laplacian  $\mathbf{L}^u$  for an undirected line graph  $\mathcal{G}^u$  of  $N$  nodes with edge weights equal to 1.

See Appendix B for a formal proof.

**Remark:** Theorem 3.1 only gives a special case when the DGLR of a graph is equal to the GLR of its undirected base graph. In fact, using our symmetrized  $\mathcal{L}_r^d$  does not mean directionality has been lost. As an example, for the 3-node DAG in Fig. 2(c), the variational term is  $\|\mathbf{L}_r^d \mathbf{x}\|_2^2 = (x_3 - \frac{1}{2}(x_1 + x_2))^2$  and computes to 0 for  $\mathbf{x} = [2 \ 0 \ 1]^\top$ . On the other hand, for the DAG in Fig. 2(d) with opposite directions, the variational term is  $\|\mathbf{L}_r^d \mathbf{x}\|_2^2 = (x_1 - x_3)^2 + (x_2 - x_3)^2$  and computes to 0 only if  $\mathbf{x}$  is a constant vector. With our directed graph construction with *time windows* as described in Section 3.1, since all nodes after the second time step has more than one predecessors, the DGLR of the undirected graph does not equal the GLR of its base graph.

### 3.2.2 $\ell_1$ -NORM VARIATIONAL TERM

Using  $\mathbf{W}_r^d$  again as a GSO, we define an  $\ell_1$ -norm variational term called *directed graph total variation* (DGTV) as

$$\|\mathbf{x} - \mathbf{W}_r^d \mathbf{x}\|_1 = \|\mathbf{L}_r^d \mathbf{x}\|_1 = \sum_{j \in \bar{\mathcal{S}}} |x_j - \sum_i w_{j,i} x_i|. \quad (3)$$

Unlike DGLR in eq. (2), DGTV is not symmetric and has no obvious frequency interpretation. Nonetheless, we provide a *two-channel filterbank* interpretation in the sequel.

<sup>5</sup>A similar directed graph variational term was defined in Li et al. (2023) towards an objective for directed graph sampling. We focus instead on directed graph signal restoration.

### 3.3 OPTIMIZATION FORMULATION

Denote by  $\mathbf{y} \in \mathbb{R}^M$  the observation,  $\mathbf{x} \in \mathbb{R}^{N(T+S+1)}$  the target signal, and  $\mathbf{H} \in \{0, 1\}^{M \times N(T+S+1)}$  the sampling matrix that selects  $M$  observed samples from  $N(T+S+1)$  entries in  $\mathbf{x}$ . Given defined undirected and directed edges, we write an objective for  $\mathbf{x}$  containing a squared-error fidelity term and graph smoothness terms for the two graphs  $\mathcal{G}^u$  and  $\mathcal{G}^d$ —GLR, DGLR, and DGTV:

$$\min_{\mathbf{x}} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \mu_u \mathbf{x}^\top \mathbf{L}^u \mathbf{x} + \mu_{d,2} \mathbf{x}^\top \mathcal{L}_r^d \mathbf{x} + \mu_{d,1} \|\mathbf{L}_r^d \mathbf{x}\|_1 \quad (4)$$

where  $\mu_u, \mu_{d,2}, \mu_{d,1} \in \mathbb{R}_+$  are weight parameters for the three regularization terms. Combination of  $\ell_2$ - and  $\ell_1$ -norm penalties—DGLR and DGTV for directed graph  $\mathcal{G}^d$  in our case<sup>6</sup>—is called *elastic net regularization* in statistics (Zou & Hastie, 2005) with demonstrable improved robustness. Further, employing both regularization terms means that, after algorithm unrolling, we can adapt an appropriate mixture by learning weights  $\mu_{d,2}$  and  $\mu_{d,1}$  per neural layer.

Convex minimization in eq. (4) is composed of smooth  $\ell_2$ -norm terms and one non-smooth  $\ell_1$ -norm term. We pursue a divide-and-conquer approach and solve eq. (4) via an ADMM framework (Boyd et al., 2011) by minimizing  $\ell_2$ - and  $\ell_1$ -norm terms in eq. (4) alternately.

### 3.4 ADMM OPTIMIZATION ALGORITHM

We first introduce *auxiliary variable*  $\phi$  and rewrite the optimization in eq. (4) as

$$\min_{\mathbf{x}, \phi} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \mu_u \mathbf{x}^\top \mathbf{L}^u \mathbf{x} + \mu_{d,2} \mathbf{x}^\top \mathcal{L}_r^d \mathbf{x} + \mu_{d,1} \|\phi\|_1, \quad \text{s.t. } \phi = \mathbf{L}_r^d \mathbf{x}. \quad (5)$$

Using the augmented Lagrangian method (Boyd & Vandenberghe, 2004), we rewrite eq. (5) in an unconstrained form:

$$\min_{\mathbf{x}, \phi} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \mu_u \mathbf{x}^\top \mathbf{L}^u \mathbf{x} + \mu_{d,2} \mathbf{x}^\top \mathcal{L}_r^d \mathbf{x} + \mu_{d,1} \|\phi\|_1 + \gamma^\top (\phi - \mathbf{L}_r^d \mathbf{x}) + \frac{\rho}{2} \|\phi - \mathbf{L}_r^d \mathbf{x}\|_2^2 \quad (6)$$

where  $\gamma \in \mathbb{R}^{N(T+S+1)}$  is a Lagrange multiplier vector, and  $\rho \in \mathbb{R}_+$  is a non-negative ADMM parameter. We solve eq. (6) iteratively by minimizing  $\mathbf{x}$  and  $\phi$  in alternating steps till convergence.

#### 3.4.1 MINIMIZING $\mathbf{x}^{\tau+1}$

Fixing  $\phi^\tau$  and  $\gamma^\tau$  at iteration  $\tau$ , optimization in eq. (6) for  $\mathbf{x}^{\tau+1}$  becomes an unconstrained convex quadratic objective. The solution is a system of linear equations:

$$\left( \mathbf{H}^\top \mathbf{H} + \mu_u \mathbf{L}^u + \left( \mu_{d,2} + \frac{\rho}{2} \right) \mathcal{L}_r^d \right) \mathbf{x}^{\tau+1} = (\mathbf{L}_r^d)^\top \left( \frac{\rho}{2} \phi^\tau + \frac{\gamma^\tau}{2} \right) + \mathbf{H}^\top \mathbf{y}. \quad (7)$$

Because coefficient matrix  $\mathbf{H}^\top \mathbf{H} + \mu_u \mathbf{L}^u + \left( \mu_{d,2} + \frac{\rho}{2} \right) \mathcal{L}_r^d$  is sparse, symmetric and PD, eq. (7) can be solved in linear time via *conjugate gradient* (CG) (Shewchuk, 1994) without matrix inversion.

**Splitting  $\ell_2$ -norm Terms:** Instead of solving eq. (7) directly for  $\mathbf{x}^{\tau+1}$  at iteration  $\tau$ , we introduce again auxiliary variables  $\mathbf{z}_u$  and  $\mathbf{z}_d$  in eq. (5) for the  $\ell_2$ -norm GLR and DGLR terms respectively, and rewrite the optimization as

$$\min_{\mathbf{x}, \mathbf{z}_u, \mathbf{z}_d} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \mu_u \mathbf{z}_u^\top \mathbf{L}^u \mathbf{z}_u + \mu_{d,2} \mathbf{z}_d^\top \mathcal{L}_r^d \mathbf{z}_d + (\gamma^\tau)^\top (\phi^\tau - \mathbf{L}_r^d \mathbf{x}) + \frac{\rho}{2} \|\phi^\tau - \mathbf{L}_r^d \mathbf{x}\|_2^2 \quad (8)$$

s.t.  $\mathbf{x} = \mathbf{z}_u = \mathbf{z}_d$ .

Using again the augmented Lagrangian method, we rewrite the unconstrained version as

$$\min_{\mathbf{x}, \mathbf{z}_u, \mathbf{z}_d} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \mu_u \mathbf{z}_u^\top \mathbf{L}^u \mathbf{z}_u + \mu_{d,2} \mathbf{z}_d^\top \mathcal{L}_r^d \mathbf{z}_d + (\gamma^\tau)^\top (\phi^\tau - \mathbf{L}_r^d \mathbf{x}) + \frac{\rho}{2} \|\phi^\tau - \mathbf{L}_r^d \mathbf{x}\|_2^2 \quad (9)$$

$$+ (\gamma_u^\tau)^\top (\mathbf{x} - \mathbf{z}_u) + \frac{\rho_u}{2} \|\mathbf{x} - \mathbf{z}_u\|_2^2 + (\gamma_d^\tau)^\top (\mathbf{x} - \mathbf{z}_d) + \frac{\rho_d}{2} \|\mathbf{x} - \mathbf{z}_d\|_2^2$$

<sup>6</sup>We can employ both  $\ell_2$ - and  $\ell_1$ -norm regularization terms for the undirected graph  $\mathcal{G}^u$  as well. For simplicity, we focus on designing new regularization terms for directed graphs in this paper.

where  $\gamma_u, \gamma_d \in \mathbb{R}^{N(T+S+1)}$  are multipliers, and  $\rho_u, \rho_d \in \mathbb{R}_+$  are ADMM parameters. Optimizing  $\mathbf{x}$ ,  $\mathbf{z}_u$  and  $\mathbf{z}_d$  in eq. (9) in turn at iteration  $\tau$  lead to three linear systems:

$$\left( \mathbf{H}^\top \mathbf{H} + \frac{\rho}{2} \mathcal{L}_r^d + \frac{\rho_u + \rho_d}{2} \mathbf{I} \right) \mathbf{x}^{\tau+1} = (\mathbf{L}_r^d)^\top \left( \frac{\gamma^\tau}{2} + \frac{\rho}{2} \phi^\tau \right) - \frac{\gamma_u^\tau}{2} + \frac{\rho_u}{2} \mathbf{z}_u^\tau - \frac{\gamma_d^\tau}{2} + \frac{\rho_d}{2} \mathbf{z}_d^\tau + \mathbf{H}^\top \mathbf{y} \quad (10)$$

$$\left( \mu_u \mathbf{L}^u + \frac{\rho_u}{2} \mathbf{I} \right) \mathbf{z}_u^{\tau+1} = \frac{\gamma_u^\tau}{2} + \frac{\rho_u}{2} \mathbf{x}^{\tau+1} \quad (11)$$

$$\left( \mu_{d,2} \mathcal{L}_r^d + \frac{\rho_d}{2} \mathbf{I} \right) \mathbf{z}_d^{\tau+1} = \frac{\gamma_d^\tau}{2} + \frac{\rho_d}{2} \mathbf{x}^{\tau+1}. \quad (12)$$

Thus, instead of solving eq. (7) for  $\mathbf{x}^{\tau+1}$  directly at iteration  $\tau$ ,  $\mathbf{x}^{\tau+1}$ ,  $\mathbf{z}_u^{\tau+1}$  and  $\mathbf{z}_d^{\tau+1}$  are optimized in turn via eq. (10) through eq. (12) using CG. This splitting induces more network parameters for data-driven learning after algorithm unrolling (see Section 4.1) towards better performance, and affords us spectral filter interpretations of the derived linear systems.

**Interpretation:** To interpret solution  $\mathbf{z}_u^{\tau+1}$  spectrally, we rewrite eq. (11) as

$$\mathbf{z}_u^{\tau+1} = \left( \frac{2\mu_u}{\rho_u} \mathbf{L}^u + \mathbf{I} \right)^{-1} \left( \frac{\gamma_u}{\rho_u} + \mathbf{x}^{\tau+1} \right) = \mathbf{V} \text{diag} \left( \left\{ \frac{1}{1 + \frac{2\mu_u}{\rho_u} \lambda_k} \right\} \right) \mathbf{V}^\top \left( \frac{\gamma_u}{\rho_u} + \mathbf{x}^{\tau+1} \right) \quad (13)$$

where  $\mathbf{L}^u = \mathbf{V} \text{diag}(\{\lambda_k\}) \mathbf{V}^\top$  is the eigen-decomposition. Thus, using frequencies defined by eigen-pairs  $\{(\lambda_k, \mathbf{v}_k)\}$  of  $\mathbf{L}^u$ ,  $\mathbf{z}_u^{\tau+1}$  is the *low-pass* filter output of  $\mathbf{x}^{\tau+1}$  offset by  $\frac{\gamma_u}{\rho_u}$ , with frequency response  $f(\lambda) = (1 + \frac{2\mu_u}{\rho_u} \lambda)^{-1}$ .

As done for  $\mathbf{z}_u^{\tau+1}$ , we can rewrite  $\mathbf{z}_d^{\tau+1}$  in eq. (12) as

$$\mathbf{z}_d^{\tau+1} = \mathbf{U} \text{diag} \left( \left\{ \frac{1}{1 + \frac{2\mu_{d,2}}{\rho_d} \xi_k} \right\} \right) \mathbf{U}^\top \left( \frac{\gamma_d}{\rho_d} + \mathbf{x}^{\tau+1} \right) \quad (14)$$

where  $\mathcal{L}_r^d = \mathbf{U} \text{diag}(\{\xi_k\}) \mathbf{U}^\top$  is the eigen-decomposition. Analogously,  $\mathbf{z}_d^{\tau+1}$  is the *low-pass* filter output of  $\mathbf{x}^{\tau+1}$  offset by  $\frac{\gamma_d}{\rho_d}$ , with frequency response  $f(\xi) = (1 + \frac{2\mu_{d,2}}{\rho_d} \xi)^{-1}$ .

To interpret solution  $\mathbf{x}^{\tau+1}$  in eq. (10) as a low-pass filter output, see Appendix C.

### 3.4.2 MINIMIZING $\phi^{\tau+1}$

Fixing  $\mathbf{x}^{\tau+1}$  at iteration  $\tau$ , optimizing  $\phi^{\tau+1}$  in eq. (6) means

$$\phi^{\tau+1} = \arg \min_{\phi} g(\phi) = \mu_{d,1} \|\phi\|_1 + (\gamma^\tau)^\top (\phi - \mathbf{L}_r^d \mathbf{x}^{\tau+1}) + \frac{\rho}{2} \|\phi - \mathbf{L}_r^d \mathbf{x}^{\tau+1}\|_2^2. \quad (15)$$

We can compute eq. (15) element-wise as follows (see Appendix D for a derivation):

$$\begin{aligned} \delta &= (\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1} - \rho^{-1} \gamma_i^\tau \\ \phi_i^{\tau+1} &= \text{sign}(\delta) \cdot \max(|\delta| - \rho^{-1} \mu_{d,1}, 0). \end{aligned} \quad (16)$$

**Interpretation:** Soft-thresholding in eq. (16) to compute  $\phi_i^{\tau+1}$  attenuates the  $i$ -th entry of  $\mathbf{L}_r^d \mathbf{x}^{\tau+1}$ . Given that graph Laplacians—second difference matrices on graphs—are high-pass filters (Ortega et al., 2018), we interpret  $\mathbf{L}_r^d$  as the *high-pass channel* of a two-channel filterbank (Vetterli & Kovacevic, 2013), where the corresponding *low-pass channel* is  $\mathbf{I} - \mathbf{L}_r^d = \mathbf{W}_r^d$ . Using eq. (16) for processing means that the low-pass channel  $\mathbf{W}_r^d$  is unused and thus preserved. For example, the constant (low-frequency) signal  $\mathbf{1}$  passes through the low-pass channel undisturbed, *i.e.*,  $\mathbf{1} = \mathbf{W}_r^d \mathbf{1}$ , while it is filtered out completely by the high-pass channel, *i.e.*,  $\mathbf{0} = \mathbf{L}_r^d \mathbf{1}$ . Thus, attenuation of the high-pass channel means eq. (16) is a low-pass filter.

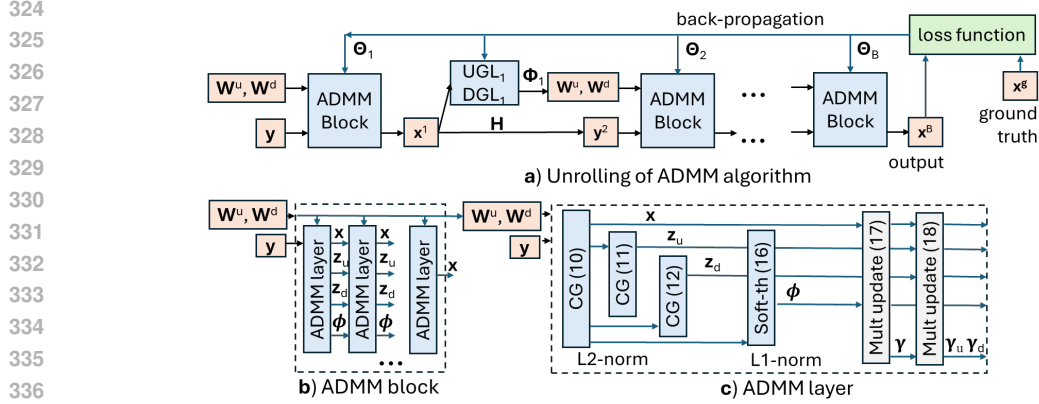


Figure 3: Unrolling of proposed iterative ADMM algorithm into blocks and neural layers.

### 3.4.3 UPDATING LAGRANGE MULTIPLIERS

For objective eq. (6), we follow standard ADMM procedure to update Lagrange multiplier  $\gamma^{\tau+1}$  after each computation of  $\mathbf{x}^{\tau+1}$  and  $\phi^{\tau+1}$ :

$$\gamma^{\tau+1} = \gamma^{\tau} + \rho(\phi^{\tau+1} - \mathbf{L}_r^d \mathbf{x}^{\tau+1}). \quad (17)$$

Splitting the  $\ell_2$ -norm terms means that additional multipliers  $\gamma_u^{\tau+1}$  and  $\gamma_d^{\tau+1}$  in objective eq. (9) must also be updated after each computation of  $\mathbf{x}^{\tau+1}$ ,  $\mathbf{z}_u^{\tau+1}$ ,  $\mathbf{z}_d^{\tau+1}$ , and  $\phi^{\tau+1}$ :

$$\gamma_u^{\tau+1} = \gamma_u^{\tau} + \rho_u(\mathbf{x}^{\tau+1} - \mathbf{z}_u^{\tau+1}), \quad \gamma_d^{\tau+1} = \gamma_d^{\tau} + \rho_d(\mathbf{x}^{\tau+1} - \mathbf{z}_d^{\tau+1}). \quad (18)$$

The ADMM loop is repeated till  $\mathbf{x}^{\tau+1}$  and  $\phi^{\tau+1}$  converge.

## 4 UNROLLED NEURAL NETWORK

We unroll our iterative ADMM algorithm to an interpretable neural net. The crux to our network is two periodically inserted graph learning modules for undirected and directed edge weight computation that are akin to the classical self-attention mechanism in transformers (Bahdanau et al., 2014).

### 4.1 ADMM ALGORITHM UNROLLING

We implement each iteration of our algorithm in Section 3.4 as a neural layer. Specifically, each ADMM iteration is implemented as four sub-layers in sequence: solving linear systems eq. (10) for  $\mathbf{x}^{\tau+1}$ , eq. (11) for  $\mathbf{z}_u^{\tau+1}$ , and eq. (12) for  $\mathbf{z}_d^{\tau+1}$  via CG, and computing  $\phi^{\tau+1}$  via eq. (16). CG is itself an iterative descent algorithm, where parameters  $\alpha$  and  $\beta$  corresponding to step size and momentum can be tuned end-to-end (see Appendix E.1 for detailed implementation). Weight parameters  $\mu_u, \mu_{d,2}, \mu_{d,1}$  for prior terms and ADMM parameters  $\rho, \rho_u, \rho_d$  are also tuned per layer. Multipliers  $\gamma^{\tau+1}, \gamma_u^{\tau+1}, \gamma_d^{\tau+1}$  are then updated via eq. (17) and eq. (18) to complete one layer. Parameters learned during back-propagation for each ADMM block  $b$  are denoted by  $\Theta_b$ . See Fig. 3 for an illustration, and Appendix E.2 for the detailed algorithms.

### 4.2 SELF-ATTENTION OPERATOR IN TRANSFORMER

We first review the basic self-attention mechanism: a scaled dot product following by a softmax operation (Bahdanau et al., 2014). Denote by  $\mathbf{x}_i \in \mathbb{R}^E$  an *embedding* for token  $i$ . The *affinity* between tokens  $i$  and  $j$ ,  $e(i, j)$ , is the dot product of linear-transformed  $\mathbf{K}\mathbf{x}_i$  and  $\mathbf{Q}\mathbf{x}_j$ , where  $\mathbf{Q}, \mathbf{K} \in \mathbb{R}^{E \times E}$  are the *query* and *key* matrices, respectively. Using softmax, non-negative *attention weight*  $a_{i,j}$  is computed from  $e(i, j)$ 's as

$$a_{i,j} = \frac{\exp(e(i, j))}{\sum_{l=1}^N \exp(e(i, l))}, \quad e(i, j) = (\mathbf{Q}\mathbf{x}_j)^\top (\mathbf{K}\mathbf{x}_i). \quad (19)$$

Given attention weights  $a_{i,j}$ , output embedding  $\mathbf{y}_i$  for token  $i$  is computed as

$$\mathbf{y}_i = \sum_{l=1}^N a_{i,l} \mathbf{x}_l \mathbf{V} \quad (20)$$

where  $\mathbf{V} \in \mathbb{R}^{E \times E}$  is a *value* matrix. By “self-attention”, we mean that output embeddings are computed from weighted input embeddings. A transformer is thus a sequence of embedding-to-embedding mappings via self-attention operations defined by learned  $\mathbf{Q}$ ,  $\mathbf{K}$  and  $\mathbf{V}$  matrices.

### 4.3 GRAPH LEARNING MODULES

**Undirected Graph Learning:** We learn an undirected graph  $\mathcal{G}^u$  in a module  $\text{UGL}_b$  at block  $b$ . For each node  $i$ , we compute a low-dimensional feature vector  $\mathbf{f}_i^u = F^u(\mathbf{e}_i) \in \mathbb{R}^K$ , where  $\mathbf{e}_i \in \mathbb{R}^E$  is a vector of signal values, position embeddings in time and Laplacian embeddings of the physical graph as done in Feng & Tassiulas (2022), and  $F^u(\cdot)$  is a chosen learned nonlinear *feature function*  $F^u: \mathbb{R}^E \mapsto \mathbb{R}^K$ , e.g., a shallow GNN (see Appendix F.3 for details). As done in Thuc et al. (2024), the feature distance between nodes  $i$  and  $j$  is computed as the *Mahalanobis distance*  $d^u(i, j)$ :

$$d^u(i, j) = (\mathbf{f}_i^u - \mathbf{f}_j^u)^\top \mathbf{M} (\mathbf{f}_i^u - \mathbf{f}_j^u) \quad (21)$$

where  $\mathbf{M} \succeq 0$  is a symmetric PSD *metric matrix* of dimension  $\mathbb{R}^{K \times K}$ . Distance in eq. (21) is non-negative and *symmetric*, i.e.,  $d^u(i, j) = d^u(j, i)$ . Weight  $w_{i,j}^u$  of undirected edge  $(i, j) \in \mathcal{E}^u$  is computed as

$$w_{i,j}^u = \frac{\exp(-d^u(i, j))}{\sqrt{\sum_{l \in \mathcal{N}_i} \exp(-d^u(i, l))} \sqrt{\sum_{k \in \mathcal{N}_j} \exp(-d^u(k, j))}} \quad (22)$$

where the denominator summing over 1-hop neighborhood  $\mathcal{N}_i$  is inserted for normalization.

**Remark:** Interpreting distance  $d^u(i, j)$  as  $-e(i, j)$ , edge weights  $w_{i,j}^u$ ’s in eq. (22) are essentially attention weights  $a_{i,j}$ ’s in eq. (19), and thus the undirected graph learning module constitutes a self-attention mechanism, albeit requiring fewer parameters, since parameters for function  $F^u(\cdot)$  and metric matrix  $\mathbf{M}$  can be much smaller than large and dense query and key matrices,  $\mathbf{Q}$  and  $\mathbf{K}$ . Further, no value matrix  $\mathbf{V}$  is required, since output signal  $\mathbf{x}^{\tau+1}$  is computed through a set of low-pass filters—eq. (10), eq. (11), eq. (12) and eq. (16)—derived from our optimization of objective eq. (4).

**Directed Graph Learning:** Similarly, we learn a directed graph  $\mathcal{G}^d$  in a module  $\text{DGL}_b$  at block  $b$ . For each node  $i$ , we compute a feature vector  $\mathbf{f}_i^d = F^d(\mathbf{e}_i) \in \mathbb{R}^K$  using feature function  $F^d(\cdot)$ . Like eq. (21), the Mahalanobis distance  $d^d(i, j)$  between nodes  $i$  and  $j$  is computed from  $\mathbf{f}_i^d$  and  $\mathbf{f}_j^d$  with a learned PSD metric  $\mathbf{P}$ . Weight  $w_{j,i}^d$  of directed edge  $[i, j] \in \mathcal{E}^d$  is then computed using exponentials and normalization as

$$w_{j,i}^d = \frac{\exp(-d^d(j, i))}{\sum_{(i,k) \in \mathcal{E}^d} \exp(-d^d(k, i))}. \quad (23)$$

Parameters learned via back-propagation by the two graph learning modules at block  $b$ —parameters of  $F^u(\cdot)$ ,  $F^d(\cdot)$ ,  $\mathbf{M}$  and  $\mathbf{P}$ —are denoted by  $\Phi_b$ .

## 5 EXPERIMENTS

### 5.1 EXPERIMENTAL SETUP

We evaluate our model’s performance on commonly used traffic speed dataset METR-LA (Li et al., 2017) and traffic flow dataset PEMS03 (Guo et al., 2022), each of which contains city-size traffic data with a sampling interval of 5 minutes in a local region, together with the real connections and distances (or travel cost) between sensors. We reduce the size of the original dataset by sub-sampling 1/3 of all data uniformly<sup>7</sup>, and split it into training, validation, and test sets by 6:2:2. We predict the

<sup>7</sup>With a drastically smaller parameter count, our model focuses on the efficiency in utilizing the observations and the effectiveness with small training sets. Uniform subsampling reduces the size of the training set while making full use of the entire observation. See a detailed analysis in Appendix G.4.

Table 1: Comparison of categorized baselines and model parameters for 60-minute forecasting on the PEMS03 dataset against ours.

Category	Selected Models	Params #
Model-based	VAR (Reinsel, 2003)	-
GNNs	STGCN <sup>†</sup> (Yu et al., 2018), STSGCN (Song et al., 2020)	321K, 3,496K
GATs	GMAN (Zheng et al., 2020), ST-Wave (Fang et al., 2023)	210K, 883K
Transformers	PDFormer (Jiang et al., 2023), STAEformer (Liu et al., 2023)	531K, 1,404K
Adaptive-graph models	Graph WaveNet (GWN) (Wu et al., 2019), AGCRN (Bai et al., 2020)	277K, 749K
MLP-based models	STID (Shao et al., 2022), SimpleTM (Chen et al., 2025)	123K, 540K
Mixed-graph unrolling	Ours	<b>38K</b>

<sup>†</sup> STGCN predicts only one time step at a time. We train the model to predict only the next step, and predict the entire sequence recursively.

Table 2: Comparison of RMSE/MAE/MAPE(%) metrics of our lightweight transformer to baseline models on 30/60-minute forecasting in PeMS03 and METR-LA datasets. We use **boldface** for the smallest error, color the 2nd and 3rd small error in **blue**, and underline the next 2 smallest errors.

Dataset & Horizon	PEMS03 (358 nodes, 547 edges)			METR-LA (207 nodes, 1,315 edges)		
	30 minutes	60 minutes	120 minutes	30 minutes	60 minutes	120 minutes
VAR	28.07/16.53/17.49	30.54/18.31/19.61	36.65/22.40/24.64	10.72/5.55/11.29	12.59/6.99/13.46	<b>14.83</b> /8.90/16.54
STGCN	28.06/18.24/17.76	37.31/24.29/23.00	54.34/34.83/30.52	11.26/5.08/10.93	13.91/6.35/13.67	16.92/8.11/17.69
STSGCN	26.41/16.75/16.86	30.62/19.35/19.15	38.04/23.86/23.62	10.25/4.05/9.18	12.65/5.18/11.48	15.74/ <b>6.84</b> /15.33
GMAN	25.79/16.23/20.04	27.57/17.48/24.33	<b>30.69</b> /19.20/26.69	11.97/5.34/10.66	14.49/6.93/13.12	15.53/7.59/14.70
ST-Wave	<u>25.57</u> / <b>15.11</b> / <b>15.04</b>	28.65/16.81/19.24	<b>29.88</b> / <b>17.11</b> / <b>17.71</b>	10.81/4.11/9.16	13.24/5.33/11.32	23.18/11.22/ <b>12.71</b>
PDFormer	<b>23.71</b> / <b>15.05</b> /18.16	<b>27.16</b> / <u>17.26</u> /21.21	35.77/22.25/25.01	<u>10.21</u> / <b>3.89</b> / <u>8.50</u>	<b>12.17</b> / <b>4.81</b> / <b>10.92</b>	17.27/9.55/19.10
STAEformer	30.22/18.85/26.62	38.36/23.68/29.21	48.72/31.96/44.71	<u>10.16</u> / <b>3.73</b> / <b>8.25</b>	12.58/ <b>4.79</b> / <b>10.08</b>	<b>14.63</b> / <b>5.82</b> / <b>11.87</b>
GWN	<u>25.76</u> / <u>15.22</u> / <b>16.83</b>	28.15/18.11/ <b>17.56</b>	34.60/21.10/22.45	11.42/5.18/ <b>8.21</b>	<u>12.46</u> / <u>5.17</u> /11.97	23.13/11.32/ <u>13.53</u>
AGCRN	27.40/ <u>15.19</u> / <b>14.35</b>	29.90/ <b>16.76</b> / <b>15.32</b>	32.79/ <b>18.72</b> / <b>17.03</b>	<b>10.11</b> / <b>3.82</b> /8.61	<u>12.56</u> / <b>5.00</b> / <u>11.25</u>	<b>14.77</b> / <b>6.33</b> / <u>14.05</u>
STID	26.50/17.27/18.59	31.88/20.82/24.89	42.98/28.03/42.41	<u>10.21</u> /4.05/9.47	12.61/5.21/12.22	15.48/6.98/16.78
SimpleTM	<b>23.75</b> / <b>15.13</b> / <b>15.59</b>	<b>25.56</b> / <b>15.97</b> / <b>15.49</b>	<b>30.58</b> / <b>18.73</b> / <b>18.18</b>	<b>8.76</b> /4.12/ <b>7.63</b>	<b>11.96</b> /5.49/ <b>9.65</b>	<u>15.44</u> / <u>7.64</u> / <b>12.27</b>
<b>Ours</b>	<b>25.05</b> /15.85/16.49	<b>26.96</b> / <b>16.58</b> /18.07	34.06/20.23/23.86	<b>10.06</b> / <u>4.05</u> /9.23	<b>12.17</b> /5.27/11.78	<u>15.19</u> / <u>7.14</u> /16.44

traffic speed/flow in the following 30/60/120 minutes (6/12/24 steps) from the previously observed 60 minutes (12 steps) for all datasets. We restrict our model to learn a *sparse* graph based on the real road connections and a limited time window to model only the *local* spatial/temporal influence. For undirected graphs  $\mathcal{G}^u$ , we connect each node to its  $k$  nearest neighbors (NNs). We construct 5 ADMM blocks, each contains 25 ADMM layers, and insert a pair of graph learning modules before each ADMM block. We set feature dimensions  $K = 6$ , and stack vertically 4 graph learning modules to learn in parallel as an implementation of the *multi-head* attention mechanism. We minimize the Huber loss with  $\delta = 1$  for the *entire* reconstructed sequence and the groundtruths. See Appendix F for more details.

## 5.2 EXPERIMENTAL RESULTS

We train each run of our model on an entire NVIDIA GeForce RTX 3090. We evaluate our unrolling model against the baselines shown in Table 1. Each model is trained for 70 epochs on the reduced dataset, and the metrics are the Root Mean Squared Error (RMSE), Mean Absolute Error (MAE), and Mean Absolute Percentage Error (MAPE) in the predicted 6/12/24 steps.

Table 2 shows the performance of our models and baselines in traffic forecasting, and the last column of Table 1 compares the number of parameters. We observe that our models achieve comparable performance to most baseline models, **achieving top-3 performance in at least one metric for 30- and 60-minute forecasting**, while employing drastically fewer parameters (**7.2% of transformer-based PDFormer**). Among all the models, ST-Wave and SimpleTM generally achieve the best performance. However, ST-Wave relies on an application-specific usage of wavelet transforms, and SimpleTM employs linear models with no inductive bias, which is parameter-heavy (Battaglia et al., 2018) and **lacks interpretability**; both models require even more parameters than PDFormer. The largest GCN model, STSGCN, fails to achieve a satisfying performance for PEMS03, where the real connections are highly sparse. STAEformer, the largest transformer, also overfits heavily in the reduced PEMS03 datasets. **We also observe that** our unrolled model

486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539

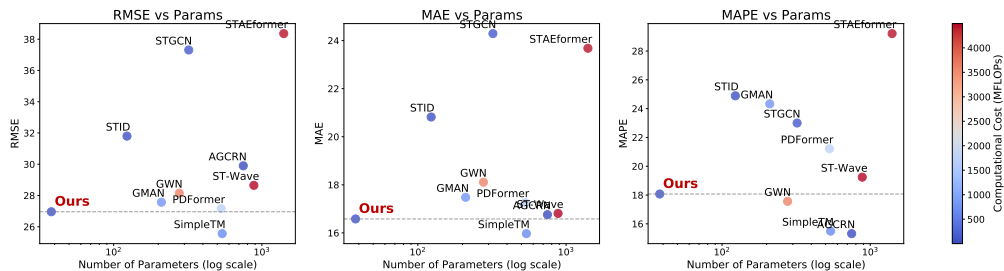


Figure 4: The trade-off between the number of parameters and performance for 60-minute forecast on PEMS03 dataset. The color of each scatter point represents the inference computational cost (in MFLOPs) in Table 6 in Appendix G.3.

also generally outperforms the multi-attention model GMAN. Though both our model and GMAN leverage local graph neighborhoods when computing attention, we dramatically reduce parameters by replacing classical self-attention mechanisms in transformers with unrolled ADMM iterations with no observable degradation in performance. *In summary, none of the DL-based baseline schemes unambiguously achieves the best performance across the two datasets, while all require significantly more parameters than our models. In contrast, our model achieves top-3 forecasting performance in two of three prediction windows for both datasets (top-5 in the remaining prediction window), in at least one metric.*

We plot the metrics with respect to the parameter size and computational cost for 60-minute forecast on the PEMS03 datasets in Figure 4. The conventional baseline models with larger parameter counts generally exhibit lower prediction errors, showing an overall downward trend in the plot despite the outlier STAEFormer which overfits heavily. In contrast, our model lies distinctly apart from the main trend in the lower-left corner in each plot, delivering the accuracy of models over 10× larger while using dramatically fewer parameters.

Further experiments are presented in Appendix G, and a comprehensive ablation study is provided in Appendix H. We further validate the *pattern-level* interpretability of the learned undirected graph through eigenvector centrality and traffic pattern analysis in Appendix J.

## 6 CONCLUSION

To capture complex node-to-node relations in spatial-temporal data, we unroll a mixed-graph optimization algorithm into a lightweight transformer-like neural net, where an undirected graph models spatial correlations and a directed graph models temporal relationships. We show that the two graph learning modules play the role of self-attention, meaning that our unrolled network is a transformer. We design new  $\ell_2$  and  $\ell_1$ -norm regularizers to quantify and promote signal smoothness on directed graphs. We interpret the derived processing operations from a designed ADMM algorithm as low-pass filters. Experiments demonstrate competitive prediction performance at drastically reduced parameter counts. One limitation of our approach is that our computed Mahalanobis distances  $d^u(i, j)$  and  $d^d(i, j)$  for respective undirected and directed graphs are non-negative (due to learned PSD metric matrices  $\mathbf{M}$  and  $\mathbf{P}$ ), whereas affinity  $e(i, j)$  in traditional self-attention can be negative. We will study extensions to signed distances and more complex graph modeling as future work.

## REFERENCES

- Saghar Bagheri, Gene Cheung, Tim Eadie, and Antonio Ortega. Joint signal interpolation / time-varying graph estimation via smoothness and low-rank priors. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 9646–9650, 2024. doi: 10.1109/ICASSP48485.2024.10447459.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *CoRR*, abs/1409.0473, 2014. URL <https://api.semanticscholar.org/CorpusID:11212020>.

- 540 Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. Adaptive graph convolutional recurrent  
541 network for traffic forecasting. In *Proceedings of the 34th International Conference on Neural*  
542 *Information Processing Systems, NIPS '20*, Red Hook, NY, USA, 2020. Curran Associates Inc.  
543 ISBN 9781713829546.
- 544
- 545 Y. Bai, G. Cheung, X. Liu, and W. Gao. Graph-based blind image deblurring from a single photograph.  
546 *IEEE Transactions on Image Processing*, 28, no.3:1404–1418, 2019.
- 547
- 548 Peter Battaglia, Jessica Blake Chandler Hamrick, Victor Bapst, Alvaro Sanchez, Vinicius Zambaldi,  
549 Mateusz Malinowski, Andrea Tacchetti, David Raposo, Adam Santoro, Ryan Faulkner, Caglar  
550 Gulcehre, Francis Song, Andy Ballard, Justin Gilmer, George E. Dahl, Ashish Vaswani, Kelsey  
551 Allen, Charles Nash, Victoria Jayne Langston, Chris Dyer, Nicolas Heess, Daan Wierstra, Pushmeet  
552 Kohli, Matt Botvinick, Oriol Vinyals, Yujia Li, and Razvan Pascanu. Relational inductive biases,  
553 deep learning, and graph networks. *arXiv*, 2018. URL <https://arxiv.org/pdf/1806.01261.pdf>.
- 554
- 555 Phillip Bonacich. Power and centrality: A family of measures. *American Journal of Sociology*, 92(5):  
556 1170–1182, 1987. ISSN 00029602, 15375390. URL <http://www.jstor.org/stable/2780000>.
- 557
- 558 S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge, 2004.
- 559
- 560 S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learn-  
561 ing via the alternating direction method of multipliers. In *Foundations and Trends in Optimization*,  
562 volume 3, no.1, pp. 1–122, 2011.
- 563
- 564 Alexandru Brateanu, Raul Balmez, Adrian Avram, Ciprian Orhei, and Cosmin Ancuti. Lyt-net:  
565 Lightweight yuv transformer-based network for low-light image enhancement. *IEEE Signal*  
566 *Processing Letters*, 32:2065–2069, 2025. doi: 10.1109/LSP.2025.3563125.
- 567
- 568 Changlu Chen, Yanbin Liu, Ling Chen, and Chengqi Zhang. Bidirectional spatial-temporal adaptive  
569 transformer for urban traffic flow forecasting. *IEEE Transactions on Neural Networks and Learning*  
570 *Systems*, 34(10):6913–6925, 2022.
- 571
- 572 Hui Chen, Viet Luong, Lopamudra Mukherjee, and Vikas Singh. SimpleTM: A simple baseline  
573 for multivariate time series forecasting. In *The Thirteenth International Conference on Learning*  
574 *Representations*, 2025. URL <https://openreview.net/forum?id=oANkBaVci5>.
- 575
- 576 S. Chen, A. Sandryhaila, J. Moura, and J. Kovacevic. Signal recovery on graphs: Variation minimiza-  
577 tion. In *IEEE Transactions on Signal Processing*, volume 63, no.17, pp. 4609–4624, September  
2015.
- 578
- 579 Yang Chen, Cheng Cheng, and Qiyu Sun. Graph fourier transform based on singular value decom-  
580 position of the directed laplacian. *Sampling Theory, Signal Processing, and Data Analysis*, 21  
581 (2):1–19, 2023. doi: 10.1007/s43670-023-00062-w. URL <https://doi.org/10.1007/s43670-023-00062-w>.
- 582
- 583 G. Cheung, E. Magli, Y. Tanaka, and M. Ng. Graph spectral image processing. In *Proceedings of the*  
584 *IEEE*, volume 106, no.5, pp. 907–930, May 2018.
- 585
- 586 Vijay Prakash Dwivedi and Xavier Bresson. A generalization of transformer networks to graphs,  
587 2021. URL <https://arxiv.org/abs/2012.09699>.
- 588
- 589 Aref Einizade, Fragkiskos Malliaros, and Jhony H Giraldo. Continuous product graph neural networks.  
590 *Advances in Neural Information Processing Systems*, 37:90226–90252, 2024.
- 591
- 592 Yuchen Fang, Yanjun Qin, Haiyong Luo, Fang Zhao, Bingbing Xu, Liang Zeng, and Chenxing Wang.  
593 When spatio-temporal meet wavelets: Disentangled traffic forecasting via efficient spectral graph  
attention networks. In *2023 IEEE 39th international conference on data engineering (ICDE)*, pp.  
517–529. IEEE, 2023.

- 594 Aosong Feng and Leandros Tassioulas. Adaptive graph spatial-temporal transformer network for  
595 traffic forecasting. In *Proceedings of the 31st ACM International Conference on Information &  
596 Knowledge Management, CIKM '22*, pp. 3933–3937, New York, NY, USA, 2022. Association  
597 for Computing Machinery. ISBN 9781450392365. doi: 10.1145/3511808.3557540. URL  
598 <https://doi.org/10.1145/3511808.3557540>.
- 599 Elias Frantar, Saleh Ashkboos, Torsten Hoefler, and Dan Alistarh. Gptq: Accurate post-training  
600 quantization for generative pre-trained transformers. In *International Conference on Learning  
601 Representations (ICLR)*, 2023.
- 602 Haotian Gao, Renhe Jiang, Zheng Dong, Jinliang Deng, Yuxin Ma, and Xuan Song. Spatial-temporal-  
603 decoupled masked pre-training for spatiotemporal forecasting. In *Proceedings of the Thirty-  
604 Third International Joint Conference on Artificial Intelligence*. International Joint Conferences on  
605 Artificial Intelligence Organization, 2024.
- 607 Zhihan Gao, Xingjian Shi, Hao Wang, Yi Zhu, Yuyang (Bernie) Wang, Mu Li, and Dit-Yan  
608 Yeung. Earthformer: Exploring space-time transformers for earth system forecasting. In  
609 S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh (eds.), *Advances  
610 in Neural Information Processing Systems*, volume 35, pp. 25390–25403. Curran Asso-  
611 ciates, Inc., 2022. URL [https://proceedings.neurips.cc/paper\\_files/paper/  
612 2022/file/a2affd71d15e8fedffe18d0219f4837a-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2022/file/a2affd71d15e8fedffe18d0219f4837a-Paper-Conference.pdf).
- 613 Shengnan Guo, Youfang Lin, Huaiyu Wan, Xiucheng Li, and Gao Cong. Learning dynamics  
614 and heterogeneity of spatial-temporal graph data for traffic forecasting. *IEEE Transactions on  
615 Knowledge and Data Engineering*, 34(11):5415–5428, 2022. doi: 10.1109/TKDE.2021.3056502.
- 617 Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs.  
618 *Advances in neural information processing systems*, 30, 2017.
- 619 Thomas Hawkins. Cauchy and the spectral theory of matrices. *Historia Mathematica*, 2(1):1–29,  
620 1975. ISSN 0315-0860. doi: [https://doi.org/10.1016/0315-0860\(75\)90032-4](https://doi.org/10.1016/0315-0860(75)90032-4). URL <https://www.sciencedirect.com/science/article/pii/0315086075900324>.
- 623 Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 2 edition,  
624 2012.
- 625 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,  
626 and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *International  
627 Conference on Learning Representations (ICLR)*, 2022.
- 628 Guangyu Huo, Yong Zhang, Boyue Wang, Junbin Gao, Yongli Hu, and Baocai Yin. Hierarchical  
629 spatio-temporal graph convolutional networks and transformer network for traffic flow forecasting.  
630 *IEEE Transactions on Intelligent Transportation Systems*, 24(4):3855–3867, 2023.
- 632 Jiawei Jiang, Chengkai Han, Wayne Xin Zhao, and Jingyuan Wang. Pdformer: Propagation delay-  
633 aware dynamic long-range transformer for traffic flow prediction. In *Proceedings of the AAAI  
634 conference on artificial intelligence*, volume 37, pp. 4365–4373, 2023.
- 635 Qizhao Jin, Xingang Zhang, Xinyu Xiao, Ying Wang, Shiming Xiang, and Chunhong Pan. Preformer:  
636 Simple and efficient design for precipitation nowcasting with transformers. *IEEE Geoscience and  
637 Remote Sensing Letters*, 21:1–5, 2024. doi: 10.1109/LGRS.2023.3325628.
- 639 Vassilis Kalofolias, Andreas Loukas, Dorina Thanou, and Pascal Frossard. Learning time varying  
640 graphs. In *2017 IEEE International Conference on Acoustics, Speech and Signal Processing  
641 (ICASSP)*, pp. 2826–2830, 2017. doi: 10.1109/ICASSP.2017.7952672.
- 642 Semin Kwak, Laura Shimabukuro, and Antonio Ortega. Frequency analysis and filter design  
643 for directed graphs with polar decomposition. In *ICASSP 2024 - 2024 IEEE International  
644 Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 9661–9665, 2024. doi:  
645 10.1109/ICASSP48485.2024.10446657.
- 647 Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion convolutional recurrent neural network:  
Data-driven traffic forecasting. *arXiv preprint arXiv:1707.01926*, 2017.

- 648 Yuejiang Li, H. Vicky Zhao, and Gene Cheung. Eigen-decomposition-free directed graph sampling  
649 via Gershgorin disc alignment. In *ICASSP 2023 - 2023 IEEE International Conference on*  
650 *Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2023. doi: 10.1109/ICASSP49357.  
651 2023.10096903.
- 652 Hangchen Liu, Zheng Dong, Renhe Jiang, Jiewen Deng, Jinliang Deng, Qunjun Chen, and Xuan  
653 Song. Spatio-temporal adaptive embedding makes vanilla transformer sota for traffic forecast-  
654 ing. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge*  
655 *Management*, pp. 4125–4129, 2023.
- 656 Hou-I Liu, Marco Galindo, Hongxia Xie, Lai-Kuan Wong, Hong-Han Shuai, Yung-Hui Li, and  
657 Wen-Huang Cheng. Lightweight deep learning for resource-constrained environments: A survey.  
658 *ACM Comput. Surv.*, 56(10), June 2024. ISSN 0360-0300. doi: 10.1145/3657282. URL <https://doi.org/10.1145/3657282>.
- 659 Antonio G. Marques, Santiago Segarra, and Gonzalo Mateos. Signal processing on directed graphs:  
660 The role of edge directionality when processing and learning from network data. *IEEE Signal*  
661 *Processing Magazine*, 37(6):99–116, 2020. doi: 10.1109/MSP.2020.3014597.
- 662 Rahul Mishra and Hari Prabhat Gupta. Designing and training of lightweight neural networks on edge  
663 devices using early halting in knowledge distillation. *IEEE Transactions on Mobile Computing*, 23  
664 (5):4665–4677, 2024. doi: 10.1109/TMC.2023.3297026.
- 665 Vishal Monga, Yuelong Li, and Yonina C. Eldar. Algorithm unrolling: Interpretable, efficient deep  
666 learning for signal and image processing. *IEEE Signal Processing Magazine*, 38(2):18–44, 2021.  
667 doi: 10.1109/MSP.2020.3016905.
- 668 Koen Oostermeijer, Qing Wang, and Jun Du. Lightweight causal transformer with local self-attention  
669 for real-time speech enhancement. In *Proceedings of Interspeech 2021*, pp. 1614–1618, Brno,  
670 Czechia, August 2021. ISCA. URL [https://www.isca-archive.org/interspeech\\_](https://www.isca-archive.org/interspeech_2021/oostermeijer21_interspeech.pdf)  
671 [2021/oostermeijer21\\_interspeech.pdf](https://www.isca-archive.org/interspeech_2021/oostermeijer21_interspeech.pdf).
- 672 A. Ortega, P. Frossard, J. Kovacevic, J. M. F. Moura, and P. Vandergheynst. Graph signal processing:  
673 Overview, challenges, and applications. In *Proceedings of the IEEE*, volume 106, no.5, pp.  
674 808–828, May 2018.
- 675 Ishak Pacal. Maxcervixt: A novel lightweight vision transformer-based approach for precise cer-  
676 vical cancer detection. *Knowledge-Based Systems*, 289:111482, 2024. ISSN 0950-7051. doi:  
677 <https://doi.org/10.1016/j.knosys.2024.111482>. URL [https://www.sciencedirect.com/](https://www.sciencedirect.com/science/article/pii/S0950705124001175)  
678 [science/article/pii/S0950705124001175](https://www.sciencedirect.com/science/article/pii/S0950705124001175).
- 679 J. Pang and G. Cheung. Graph Laplacian regularization for inverse imaging: Analysis in the  
680 continuous domain. In *IEEE Transactions on Image Processing*, volume 26, no.4, pp. 1770–1785,  
681 April 2017.
- 682 Prajit Ramachandran, Barret Zoph, and Quoc V Le. Swish: a self-gated activation function. *arXiv*  
683 *preprint arXiv:1710.05941*, 7(1):5, 2017.
- 684 Raksha Ramakrishna, Hoi-To Wai, and Anna Scaglione. A user guide to low-pass graph signal  
685 processing and its applications: Tools and applications. *IEEE Signal Processing Magazine*, 37(6):  
686 74–85, 2020. doi: 10.1109/MSP.2020.3014590.
- 687 Gregory C Reinsel. *Elements of multivariate time series analysis*. Springer Science & Business  
688 Media, 2003.
- 689 Luana Ruiz, Fernando Gama, and Alejandro Ribeiro. Gated graph recurrent neural networks. *IEEE*  
690 *Transactions on Signal Processing*, 68:6303–6318, 2020.
- 691 Aliaksei Sandryhaila and José M. F. Moura. Discrete signal processing on graphs: Frequency  
692 analysis. *IEEE Transactions on Signal Processing*, 62(12):3042–3054, 2014. doi: 10.1109/TSP.  
693 2014.2321121.
- 694 Bastian Seifert and Markus Püschel. Digraph signal processing with generalized boundary conditions.  
695 *IEEE Transactions on Signal Processing*, 69:1422–1437, 2021. doi: 10.1109/TSP.2021.3051267.

- 702 Rasoul Shafipour, Ali Khodabakhsh, Gonzalo Mateos, and Evdokia Nikolova. A directed graph  
703 fourier transform with spread frequency components. *IEEE Transactions on Signal Processing*, 67  
704 (4):946–960, 2019. doi: 10.1109/TSP.2018.2886151.
- 705  
706 Zezhi Shao, Zhao Zhang, Fei Wang, Wei Wei, and Yongjun Xu. Spatial-temporal identity: A simple  
707 yet effective baseline for multivariate time series forecasting. In *Proceedings of the 31st ACM*  
708 *International Conference on Information & Knowledge Management*, pp. 4454–4458, 2022.
- 709 J. R Shewchuk. An introduction to the conjugate gradient method without the agonizing pain.  
710 Technical report, USA, 1994.
- 711  
712 Chao Song, Youfang Lin, Shengnan Guo, and Huaiyu Wan. Spatial-temporal synchronous graph  
713 convolutional networks: A new framework for spatial-temporal network data forecasting. In  
714 *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 914–921, 2020.
- 715  
716 Le Sun, Xinyu Wang, Yuhui Zheng, Zebin Wu, and Liyong Fu. Multiscale 3-d–2-d mixed cnn and  
717 lightweight attention-free transformer for hyperspectral and lidar classification. *IEEE Transactions*  
*on Geoscience and Remote Sensing*, 62:1–16, 2024. doi: 10.1109/TGRS.2024.3367374.
- 718  
719 Xuxiang Ta, Zihan Liu, Xiao Hu, Le Yu, Leilei Sun, and Bowen Du. Adaptive spatio-temporal graph  
720 neural network for traffic forecasting. *Knowledge-based systems*, 242:108199, 2022.
- 721  
722 Tam Thuc, Parham Eftekhari, Seyed Alireza Hosseini, Gene Cheung, and Philip A. Chou. Interpretable  
723 lightweight transformer via unrolling of learned graph smoothness priors. In A. Globerson,  
724 L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural*  
*Information Processing Systems*, volume 37, pp. 6393–6416. Curran Associates, Inc., 2024.
- 725  
726 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz  
727 Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing*  
*systems*, 30, 2017.
- 728  
729 Ashish Vaswani, Samy Bengio, Eugene Brevdo, Francois Chollet, Aidan N. Gomez, Stephan Gouws,  
730 Llion Jones, Łukasz Kaiser, Nal Kalchbrenner, Niki Parmar, Ryan Sepassi, Noam Shazeer, and  
731 Jakob Uszkoreit. Tensor2tensor for neural machine translation. *CoRR*, abs/1803.07416, 2018.  
732 URL <http://arxiv.org/abs/1803.07416>.
- 733  
734 Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua  
Bengio. Graph attention networks, 2018. URL <https://arxiv.org/abs/1710.10903>.
- 735  
736 Martin Vetterli and Jelena Kovacevic. Wavelets and subband coding. In *Prentice Hall Signal Process-*  
*ing Series*, 2013. URL <https://api.semanticscholar.org/CorpusID:10506924>.
- 737  
738 Michał Woźniak, Bartosz Zawadzki, Michał Słowik, and Tomasz Kajdanowicz. Overview of the  
739 transformer-based models for nlp tasks. In *2020 15th Conference on Computer Science and*  
*Information Systems (FedCSIS)*, pp. 407–414. IEEE, 2020. ISBN 978-83-955416-7-4. doi:  
740 10.15439/2020F20. URL <https://ieeexplore.ieee.org/document/9222960>.
- 741  
742 Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. Graph wavenet for deep  
743 spatial-temporal graph modeling. In *Proceedings of the 28th International Joint Conference on*  
*Artificial Intelligence, IJCAI’19*, pp. 1907–1913. AAAI Press, 2019. ISBN 9780999241141.
- 744  
745 Mingxing Xu, Wenrui Dai, Chunmiao Liu, Xing Gao, Weiyao Lin, Guo-Jun Qi, and Hongkai  
746 Xiong. Spatial-temporal transformer networks for traffic flow forecasting. *arXiv preprint*  
747 *arXiv:2001.02908*, 2020.
- 748  
749 Koki Yamada, Yuichi Tanaka, and Antonio Ortega. Time-varying graph learning based on sparseness  
750 of temporal variation. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech*  
*and Signal Processing (ICASSP)*, pp. 5411–5415, 2019. doi: 10.1109/ICASSP.2019.8682762.
- 751  
752 Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-temporal graph convolutional networks: A deep  
753 learning framework for traffic forecasting. In *Proceedings of the Twenty-Seventh International Joint*  
*Conference on Artificial Intelligence, IJCAI-2018*, pp. 3634–3640. International Joint Conferences  
754 on Artificial Intelligence Organization, July 2018. doi: 10.24963/ijcai.2018/505. URL <http://dx.doi.org/10.24963/ijcai.2018/505>.
- 755

756 Yaodong Yu, Sam Buchanan, Druv Pai, Tianzhe Chu, Ziyang Wu, Shengbang Tong, Ben-  
757 jamin Haeffele, and Yi Ma. White-box transformers via sparse rate reduction. In A. Oh,  
758 T. Neumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neu-  
759 ral Information Processing Systems*, volume 36, pp. 9422–9457. Curran Associates, Inc.,  
760 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/  
761 file/1e118ba9ee76c20df728b42a35fb4704-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/1e118ba9ee76c20df728b42a35fb4704-Paper-Conference.pdf).

762 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and  
763 Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In  
764 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
765 5728–5739, 2022.

766 Ailing Zeng, Muxi Chen, Lei Zhang, and Qiang Xu. Are transformers effective for time series  
767 forecasting? 2023.

769 Haixin Zhao and Nilesh Madhu. Study of lightweight transformer architectures for single-  
770 channel speech enhancement. In *Proceedings of the 33rd European Signal Processing Con-  
771 ference (EUSIPCO)*, Isola delle Femmine - Palermo, Italy, September 2025. EURASIP. URL  
772 <https://arxiv.org/abs/2505.21057>. Accepted for publication.

774 Chuanpan Zheng, Xiaoliang Fan, Cheng Wang, and Jianzhong Qi. Gman: A graph multi-attention  
775 network for traffic prediction. In *Proceedings of the AAAI conference on artificial intelligence*,  
776 volume 34, pp. 1234–1241, 2020.

777 Hui Zou and Trevor Hastie. Regularization and Variable Selection Via the Elastic Net. *Journal of  
778 the Royal Statistical Society Series B: Statistical Methodology*, 67(2):301–320, 03 2005. ISSN  
779 1369-7412. doi: 10.1111/j.1467-9868.2005.00503.x. URL [https://doi.org/10.1111/j.  
780 1467-9868.2005.00503.x](https://doi.org/10.1111/j.1467-9868.2005.00503.x).

## 783 A RELATED WORKS

### 785 A.1 LIGHTWEIGHT LEARNING MODELS

786 Given the importance of parameter reduction for increasingly large deep learning models, there exist  
787 various studies towards this common goal across application domains. For large language models  
788 (LLMs), smaller transformer models have been achieved via low-rank assumption (Hu et al., 2022)  
789 or parameter quantization (Frantar et al., 2023). For speech enhancement, Oostermeijer et al. (2021)  
790 proposes local causal self-attention, while Zhao & Madhu (2025) proposes parameter sharing across  
791 spectral and time dimensions. For low-light image enhancement, Brateanu et al. (2025) processes  
792 the Y-channel via convolution, pooling, and multi-head self-attention, separately from the U- and  
793 V-channels. For hyperspectral image classification, Sun et al. (2024) proposes a novel multiscale  
794 3D-2D mixed CNN design. For cervical cancer detection, Pacal (2024) substitutes larger MBConv  
795 blocks in the MaxViT architecture with smaller ConvNeXtv2 blocks, and MLP blocks with GRN-  
796 based MLPs. Similar architecture miniaturization techniques have been proposed for edge devices  
797 (Mishra & Gupta, 2024) and resource-constrained environments (Liu et al., 2024).

798 For traffic prediction, there are practical use cases where parameter reduction is important. For  
799 example, drivers in a city predict local traffic conditions (in an area-specific neighborhood graph)  
800 on their memory-constrained mobile devices every five minutes to compute optimal routes to their  
801 destinations, based on updated congestion information shared in their local neighborhoods, instead of  
802 sending individual requests to the cloud and overwhelming the central server.

803 Thuc et al. (2024) achieves a lightweight model for image interpolation in a more principled manner:  
804 first design a linear-time optimization algorithm minimizing a chosen *undirected* graph smoothness  
805 prior (such as *graph Laplacian regularizer* (GLR) (Pang & Cheung, 2017) or *graph total variation*  
806 (GTV) (Bai et al., 2019)), then unroll its iterations into neural layers to compose a “white-box”  
807 transformer-like neural net for data-driven parameter learning. This algorithm unrolling approach  
808 results in a lightweight *and* mathematically interpretable network. Our current work differs from Thuc  
809 et al. (2024) in that we model the more complex node-to-node relationships in spatial-temporal data  
using a *mixed* graph: i) an *undirected* graph  $\mathcal{G}^u$  to capture spatial correlations among geographically

near stations, and ii) a *directed* graph  $\mathcal{G}^d$  to capture sequential relationships along the temporal dimension. In particular, the notion of “frequencies” for directed graphs is not well understood in the *graph signal processing* (GSP) community. We are the first to define this notion and promote low-pass signal reconstruction via our designed  $\ell_2$ -norm *directed graph Laplacian regularizer* (DGLR) variational term in eq. (2) and  $\ell_1$ -norm *directed graph total variation* (DGTV) variational term in eq. (3), both derived from the *graph shift operator* (GSO) (Chen et al., 2015). To the best of our knowledge, we are the first define and promote smooth signals on a directed graph in a data-driven transformer setting.

## A.2 DEEP MODELS FOR TRAFFIC FORECASTING

Traffic forecasting problems contain spatial and temporal correlations between stations and time steps; thus, Graph Neural Network (GNN) is naturally employed in early research. Early research, such as STGCN (Yu et al., 2018), STSGCN (Song et al., 2020) and GGRNN (Ruiz et al., 2020), employs Graph Convolutional Networks (GCNs) to capture spatial relations and uses 1D convolutions and RNNs to process temporal relationships. ASTGNN (Ta et al., 2022), Graph WaveNet (Wu et al., 2019) and AGCRN (Bai et al., 2019) explored adaptive graph settings to capture potential non-local signal relations over the graph, trading in interpretability for improved performance. Graph Attention Networks (GATs) learn graph weights from embedded signals to enable flexible and data-targeted graph learning, typical examples include GMAN (Zheng et al., 2020) and ST-Wave (Fang et al., 2023). All GNNs above focus only on the spatial correlations, while the product graph (PG) is an intuitive model that describes local spatial-temporal relations at the same time. Einizade et al. (2024) presents a continuous product graph neural network for spatial-temporal forecasting, which achieves superior performance to the previous most-popular graph methods.

Transformers extend the powerful self-attention mechanism to all spatial-temporal data points, which enhances model performance with an extra-large number of parameters. Popular transformer-based models for traffic forecasting employ two transformers to process spatial and temporal dimensions separately (Xu et al., 2020; Chen et al., 2022; Huo et al., 2023), while state-of-the-art transformers introduce adaptive local attentions (Liu et al., 2023), time-delay (Jiang et al., 2023), or pre-training (Gao et al., 2024).

In recent years, simple linear models have been proposed to challenge the effectiveness and efficiency of heavy transformers. DLinear (Zeng et al., 2023) simply separates the input signals as trends and residuals, and employs 2 simple linear layers shared across nodes for prediction, but beats some of the transformer baselines in *long-term* forecasting. Other representative simple linear models include STID (Shao et al., 2022) and SimpleTM (Chen et al., 2025).

As shown in Table 1, we select two most representative and well-performing models from each category as baselines for comparison.

## A.3 STUDIES ON DIRECTED GRAPH LAPLACIANS

In the DGSP literature, frequencies of a general graph (directed or undirected) were first defined in Sandryhaila & Moura (2014) via Jordan decomposition of the adjacency matrix, which is known to be numerically unstable. In Shafipour et al. (2019), a set of orthonormal vectors (deemed directed graph frequencies) are computed via a minimization of spectral dispersion in a constrained quadratic programming formulation in eq. (8), which is computationally expensive. In Marques et al. (2020), a graph shift operator (GSO)  $\mathbf{S}$  is assumed to be diagonalizable into  $\mathbf{S} = \mathbf{V}\text{diag}\{\{\lambda_k\}\}\mathbf{V}^{-1}$ , where the generally non-orthogonal eigenvectors in  $\mathbf{V}$  are deemed frequency modes and ordered using a total variation measure in eq. (2), given eigenvalues  $\lambda_k$ ’s can be complex-valued. However, asymmetric adjacency and graph Laplacian matrices are not always diagonalizable via the Spectral Theorem. Seifert & Püschel (2021) proposes to augment a directed graph with extra edges to destroy nontrivial Jordan blocks, so that the resulting adjacency matrix becomes diagonalizable. Adding edges, however, fundamentally changes the graph representation of the original pairwise relationships. Chen et al. (2023) proposes to use singular value decomposition (SVD) on the directed Laplacian to derive frequencies for directed graphs, which is computationally expensive. Kwak et al. (2024) proposes to decompose the asymmetric Laplacian matrix into a unitary part (capturing rotation or

cycles) and a positive semi-definite part (capturing diffusion-like behavior) also using SVD and polar decomposition, thus suffering the same computation burden. None of these related works proposed variational terms to define directed graph frequencies, which in turn can be simply added to an objective to promote signal smoothness during optimization without first computing frequency components, as we have done.

## B PROOF OF THEOREM 3.1

We prove Theorem 3.1 by construction as follows. A directed line graph  $\mathcal{G}^d$  of  $N$  nodes with inter-node edge weights 1 (and self-loop of weight 1 at node 1) has adjacency matrix  $\mathbf{W}_r^d$  and random-walk Laplacian matrix  $\mathbf{L}_r^d = \mathbf{I}_N - \mathbf{W}_r^d$  of the following forms:

$$\mathbf{W}_r^d = \begin{bmatrix} 1 & 0 & \dots & 0 \\ 1 & 0 & \dots & 0 \\ 0 & 1 & 0 & \dots & 0 \\ \vdots & \ddots & \ddots & & 0 \\ 0 & \dots & 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{L}_r^d = \begin{bmatrix} 0 & 0 & \dots & 0 \\ -1 & 1 & \dots & 0 \\ 0 & -1 & 1 & \dots & 0 \\ \vdots & \ddots & \ddots & & 0 \\ 0 & \dots & 0 & -1 & 1 \end{bmatrix}. \quad (24)$$

Note that rows 2 to  $N$  of  $\mathbf{L}_r^d$  actually compose the *incidence* matrix of an undirected line graph  $\mathcal{G}^u$  with edge weights 1. Thus, the symmetrized directed graph Laplacian  $\mathcal{L}_r^d = (\mathbf{L}_r^d)^\top \mathbf{L}_r^d$  is also the combinatorial graph Laplacian  $\mathbf{L}^u$ :

$$\mathcal{L}_r^d = \mathbf{L}^u = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 \\ -1 & 2 & -1 & 0 & \dots & 0 \\ 0 & -1 & 2 & -1 & 0 \dots & 0 \\ \vdots & & \ddots & \ddots & \ddots & \\ 0 & \dots & 0 & -1 & 2 & -1 \\ 0 & \dots & 0 & 0 & -1 & 1 \end{bmatrix}. \quad (25)$$

As an illustration, consider a 4-node directed line graph  $\mathcal{G}^d$  with edge weights 1; see Fig. 2(a) for an illustration. We see that the symmetrized directed graph Laplacian  $\mathcal{L}_r^d$  defaults to the graph Laplacian  $\mathbf{L}^u$  for a 4-node *undirected* line graph  $\mathcal{G}^u$  (Fig. 2(b)):

$$\mathbf{W}_r^d = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad \mathbf{L}_r^d = \begin{bmatrix} 0 & 0 & 0 & 0 \\ -1 & 1 & 0 & 0 \\ 0 & -1 & 1 & 0 \\ 0 & 0 & -1 & 1 \end{bmatrix}, \quad \mathcal{L}_r^d = \begin{bmatrix} 1 & -1 & 0 & 0 \\ -1 & 2 & -1 & 0 \\ 0 & -1 & 2 & -1 \\ 0 & 0 & -1 & 1 \end{bmatrix}. \quad (26)$$

## C INTERPRETING LINEAR SYSTEM EQ. (10)

For notation simplicity, we first define

$$\mathbf{r}^\tau \triangleq (\mathbf{L}_r^d)^\top \left( \frac{\boldsymbol{\gamma}^\tau}{2} + \frac{\rho}{2} \boldsymbol{\phi}^\tau \right) - \frac{\boldsymbol{\gamma}_u^\tau}{2} + \frac{\rho_u}{2} \mathbf{z}_u^\tau - \frac{\boldsymbol{\gamma}_d^\tau}{2} + \frac{\rho_d}{2} \mathbf{z}_d^\tau. \quad (27)$$

Solution  $\mathbf{x}^{\tau+1}$  to linear system eq. (10) can now be written as

$$\mathbf{x}^{\tau+1} = \left( \mathbf{H}^\top \mathbf{H} + \frac{\rho}{2} \mathcal{L}_r^d + \frac{\rho_u + \rho_d}{2} \mathbf{I} \right)^{-1} (\mathbf{r}^\tau + \mathbf{H}^\top \mathbf{y}). \quad (28)$$

Define  $\tilde{\mathcal{L}}_r^d \triangleq \mathbf{H}^\top \mathbf{H} + \frac{\rho}{2} \mathcal{L}_r^d$ . Note that  $\mathbf{H}^\top \mathbf{H}$  is a diagonal matrix with zeros and ones (corresponding to chosen sample indices) along its diagonal, and thus  $\tilde{\mathcal{L}}_r^d$  is a scaled variant of  $\mathcal{L}_r^d$  with the addition of selected self-loops of weight 1. Given that  $\mathcal{L}_r^d = (\mathbf{L}_r^d)^\top \mathbf{L}_r^d$  is real and symmetric,  $\tilde{\mathcal{L}}_r^d$  is also real and symmetric. By eigen-decomposing  $\tilde{\mathcal{L}}_r^d = \tilde{\mathbf{U}} \text{diag}(\{\tilde{\xi}_k\}) \tilde{\mathbf{U}}^\top$ , we can rewrite  $\mathbf{x}^{\tau+1}$  as

$$\mathbf{x}^{\tau+1} = \tilde{\mathbf{U}} \text{diag} \left( \left\{ \frac{1}{\frac{\rho_u + \rho_d}{2} + \tilde{\xi}_k} \right\} \right) \tilde{\mathbf{U}}^\top (\mathbf{r}^\tau + \mathbf{H}^\top \mathbf{y}) \quad (29)$$

which demonstrates that  $\mathbf{x}^{\tau+1}$  is a low-pass filter output of up-sampled  $\mathbf{H}^\top \mathbf{y}$  (with bias  $\mathbf{r}^\tau$ ).

## D DERIVING SOLUTION TO OPTIMIZATION EQ. (15)

The optimization for  $\phi$  in objective eq. (15) can be performed entry-by-entry. Specifically, for the  $i$ -th entry  $\phi_i$ :

$$\min_{\phi_i} g(\phi_i) = \mu_{d,1}|\phi_i| + \gamma_i^\tau (\phi_i - (\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1}) + \frac{\rho}{2} (\phi_i - (\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1})^2 \quad (30)$$

where  $(\mathbf{L}_r^d)_i$  denotes the  $i$ -row of  $\mathbf{L}_r^d$ .  $g(\phi_i)$  is convex and differentiable when  $\phi_i \neq 0$ . Specifically, when  $\phi_i > 0$ , setting the derivative of  $g(\phi_i)$  w.r.t.  $\phi_i$  to zero gives

$$\begin{aligned} \mu_{d,1} + \gamma_i^\tau + \rho (\phi_i^* - (\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1}) &= 0 \\ \phi_i^* &= (\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1} - \rho^{-1} \gamma_i^\tau - \rho^{-1} \mu_{d,1} \end{aligned} \quad (31)$$

which holds only if  $(\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1} - \rho^{-1} \gamma_i^\tau > \rho^{-1} \mu_{d,1}$ . On the other hand, when  $\phi_i < 0$ ,

$$\begin{aligned} -\mu_{d,1} + \gamma_i^\tau + \rho (\phi_i^* - (\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1}) &= 0 \\ \phi_i^* &= (\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1} - \rho^{-1} \gamma_i^\tau + \rho^{-1} \mu_{d,1} \end{aligned} \quad (32)$$

which holds only if  $(\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1} - \rho^{-1} \gamma_i^\tau < -\rho^{-1} \mu_{d,1}$ . Summarizing the two cases, we get

$$\begin{aligned} \delta &= (\mathbf{L}_r^d)_i \mathbf{x}^{\tau+1} - \rho^{-1} \gamma_i^\tau \\ \phi_i^* &= \text{sign}(\delta) \cdot \max(|\delta| - \rho^{-1} \mu_{d,1}, 0) \end{aligned} \quad (33)$$

**Alternative solution with subgradients:** Denote  $g_1(\phi) = \mu_{d,1} \|\phi\|_1$ ,  $g_2(\phi) = (\gamma^\tau)^\top (\phi - \mathbf{L}_r^d \mathbf{x}^{\tau+1}) + \frac{\rho}{2} \|\phi - \mathbf{L}_r^d \mathbf{x}^{\tau+1}\|_2^2$ , the subgradient of  $g(\phi)$  is

$$\partial g(\phi) = \{\mathbf{g} + \nabla g_2(\phi) | \mathbf{g} \in \partial g_1(\phi)\} = \{\mu_{d,1} \sigma + \gamma^\tau + \rho(\phi - \mathbf{L}_r^d \mathbf{x}^{\tau+1}) | \sigma \in \partial(\|\phi\|_1)\}, \quad (34)$$

where

$$\partial_i(\|\phi\|_1) = \begin{cases} \text{sign}(\phi_i), & \phi_i \neq 0, \\ [-1, 1], & \phi_i = 0. \end{cases} \quad (35)$$

The minimal  $\phi^*$  satisfies  $\mathbf{0} \in \partial g(\phi^*)$ , thus

$$\begin{aligned} \mu_{d,1} + \gamma_i^\tau + \rho(\phi_i^* - (\mathbf{L}_r^d)_i \mathbf{x}) &= 0, & \text{if } \phi_i^* > 0, \\ -\mu_{d,1} + \gamma_i^\tau + \rho(\phi_i^* - (\mathbf{L}_r^d)_i \mathbf{x}) &= 0, & \text{if } \phi_i^* < 0, \\ \exists \sigma \in [-1, 1], \text{ s.t. } \mu_{d,1} \sigma + \gamma_i^\tau + \rho(\phi_i^* - (\mathbf{L}_r^d)_i \mathbf{x}) &= 0, & \text{if } \phi_i^* = 0, \end{aligned} \quad (36)$$

which gives

$$\phi_i^* = \begin{cases} (\mathbf{L}_r^d)_i \mathbf{x} - \frac{\gamma_i^\tau}{\rho} - \frac{\mu_{d,1}}{\rho}, & (\mathbf{L}_r^d)_i \mathbf{x} - \frac{\gamma_i^\tau}{\rho} > \frac{\mu_{d,1}}{\rho} \\ 0, & -\frac{\mu_{d,1}}{\rho} \leq (\mathbf{L}_r^d)_i \mathbf{x} - \frac{\gamma_i^\tau}{\rho} \leq \frac{\mu_{d,1}}{\rho} \\ (\mathbf{L}_r^d)_i \mathbf{x} - \frac{\gamma_i^\tau}{\rho} + \frac{\mu_{d,1}}{\rho}, & (\mathbf{L}_r^d)_i \mathbf{x} - \frac{\gamma_i^\tau}{\rho} < -\frac{\mu_{d,1}}{\rho} \end{cases}. \quad (37)$$

Thus the solution of eq. (15) is

$$\phi^{\tau+1} = \phi^* = \text{soft}_{\frac{\mu_{d,1}}{\rho}} \left( \mathbf{L}_r^d \mathbf{x} - \frac{\gamma_i^\tau}{\rho} \right) = \text{sign} \left( \mathbf{L}_r^d \mathbf{x} - \frac{\gamma_i^\tau}{\rho} \right) \max \left( \left| \mathbf{L}_r^d \mathbf{x} - \frac{\gamma_i^\tau}{\rho} \right| - \frac{\mu_{d,1}}{\rho}, 0 \right). \quad (38)$$

## E ALGORITHMS FOR THE UNROLLED NETWORK

### E.1 PARAMETERIZING CONJUGATED GRADIENT METHOD IN SOLVING LINEAR EQUATIONS

A linear system  $\mathbf{A}\mathbf{x} = \mathbf{b}$  where  $\mathbf{A}$  is a PD matrix can be solved iteratively by the conjugated gradient (CG) method, as is described in Algorithm 1. We parameterize the coefficients of the ‘‘gradients’’  $\alpha_k$  and ‘‘momentum’’  $\beta_k$  in each CG iteration to transform the iterative algorithm to stacked neural net layers. Note that  $\alpha, \beta \geq 0$  in every iteration in the original algorithm, and should be kept non-negative in the unrolled version.

**Algorithm 1** Conjugate Gradient (CG) Method to Solve  $\mathbf{Ax} = \mathbf{b}$  (Unrolled version  $\text{uCG}(\mathbf{A}, \mathbf{b}, \mathbf{x}_0)$ )

**Require:** PSD matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$ , vector  $\mathbf{b} \in \mathbb{R}^n$ , initial guess  $\mathbf{x}_0$   
**Ensure:** Approximate solution  $\mathbf{x}$  such that  $\|\mathbf{Ax} - \mathbf{b}\| \leq \epsilon$ , where  $\epsilon$  is the tolerance of convergence

- 1:  $\mathbf{r}_0 \leftarrow \mathbf{b} - \mathbf{Ax}_0$   $\triangleright$  *Initial residual*
- 2:  $\mathbf{p}_0 \leftarrow \mathbf{r}_0$
- 3:  $k \leftarrow 0$
- 4: **while**  $\|\mathbf{Ax}_k - \mathbf{b}\| > \epsilon$  **do**
- 5:  $\alpha_k \leftarrow \frac{\mathbf{r}_k^\top \mathbf{r}_k}{\mathbf{p}_k^\top \mathbf{A} \mathbf{p}_k}$   $\triangleright$  *Learnable in unrolling*
- 6:  $\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k + \alpha_k \mathbf{p}_k$
- 7:  $\mathbf{r}_{k+1} \leftarrow \mathbf{r}_k - \alpha_k \mathbf{A} \mathbf{p}_k$   $\triangleright$  *Dual residual*
- 8:  $\beta_k \leftarrow \frac{\mathbf{r}_{k+1}^\top \mathbf{r}_{k+1}}{\mathbf{r}_k^\top \mathbf{r}_k}$   $\triangleright$  *Learnable in unrolling*
- 9:  $\mathbf{p}_{k+1} \leftarrow \mathbf{r}_{k+1} + \beta_k \mathbf{p}_k$
- 10:  $k \leftarrow k + 1$
- 11: **end while**
- 12: **return**  $\mathbf{x}_{k+1}$

**E.2 DETAILED ALGORITHMS FOR THE PROPOSED MODEL**

Algorithms 2 detailedly explained the operations in ADMM blocks corresponding to Figure 3(b) and (c). An overall algorithm of the unrolled network corresponding to Figure 3(a) is shown in Algorithm 3. All matrix multiplication is computed *locally* to ensure low computational and memory cost with

$$(\mathbf{L}^u \mathbf{x})_j = D_{j,j}^u x_j - \sum_{i \in \mathcal{N}_j} W_{i,j} x_i, \quad (\mathbf{L}_r^u \mathbf{x})_j = x_j - \sum_{i \in \mathcal{N}_i} \frac{W_{i,j}}{\sqrt{D_{i,i}^u D_{j,j}^u}} x_i, \quad (39)$$

$$(\mathbf{L}_r^d \mathbf{x})_j = x_j - \sum_{(i,j) \in \mathcal{E}^d} \frac{W_{j,i}^d}{D_{j,j}^d} x_i, \quad ((\mathbf{L}_r^d)^\top \mathbf{x})_j = x_j - \sum_{(j,i) \in \mathcal{E}^d} \frac{W_{i,j}^d}{D_{i,i}^d} x_i, \quad \mathcal{L}_r^d \mathbf{x} = (\mathbf{L}_r^d)^\top \mathbf{L}_r^d \mathbf{x}. \quad (40)$$

**Algorithm 2** Forward Propagation of the Unrolled ADMM\_BLOCK

**Require:** Initial signal output of the last ADMM block  $\mathbf{x} \in \mathbb{R}^{T+S+1}$ , undirected graph Laplacian matrix  $\mathbf{L}^u$ , directed graph Laplacian matrix  $\mathbf{L}_r^d$ .  
**Ensure:** Network output of the unrolled ADMM block with  $M$  layers.

$\triangleright$  **Learnable parameters:** layer-wise  $\{\mu_u^\tau, \mu_{d,1}^\tau, \mu_{d,2}^\tau, \rho^\tau, \rho_u^\tau, \rho_d^\tau\}_{\tau=1}^M$ , parameter sets  $\{(\alpha_k, \beta_k)\}$  in each  $\text{uCG}$  module.

- 1: Select real observations with mask:  $\mathbf{y} \leftarrow \mathbf{H}\mathbf{x}$ ,  $\mathbf{y} \in \mathbb{R}^{T+1}$ .
- 2: **for**  $\tau$  in  $0 : M$  **do**  $\triangleright$  *M ADMM layers*
- 3: Initialize  $\phi \leftarrow \mathbf{L}_r^d \mathbf{x}$
- 4:  $\mathbf{x} \leftarrow \text{uCG}(\mathbf{H}^\top \mathbf{H} + \frac{\rho^\tau}{2} \mathcal{L}_r^d + \frac{\rho_u^\tau + \rho_d^\tau}{2} \mathbf{I}, (\mathbf{L}_r^d)^\top (\frac{\gamma}{2} + \frac{\rho}{2} \phi) - \frac{\gamma_u}{2} + \frac{\rho_u^\tau}{2} \mathbf{z}_u - \frac{\gamma_d}{2} + \frac{\rho_d^\tau}{2} \mathbf{z}_d + \mathbf{H}^\top \mathbf{y}, \mathbf{x})$
- 5:  $\mathbf{z}_u \leftarrow \text{uCG}(\mu_u^\tau \mathbf{L}^u + \frac{\rho_u^\tau}{2} \mathbf{I}, \frac{\gamma_u}{2} + \frac{\rho_u^\tau}{2} \mathbf{x}, \mathbf{z}_u)$
- 6:  $\mathbf{z}_d \leftarrow \text{uCG}(\mu_{d,2}^\tau \mathcal{L}_r^d + \frac{\rho_d^\tau}{2} \mathbf{I}, \frac{\gamma_d}{2} + \frac{\rho_d^\tau}{2} \mathbf{x}, \mathbf{z}_d)$   $\triangleright$  *Minimize  $\mathbf{x}$  with eqs. (10) to (12)*
- 7:  $\phi \leftarrow \text{soft}_{\mu_{d,1}^\tau / \rho^\tau}(\mathbf{L}_r^d \mathbf{x} - \gamma / \rho^\tau)$   $\triangleright$  *Minimize  $\phi$  with eq. (16)*
- 8:  $\gamma \leftarrow \gamma + \rho^\tau (\phi^{\tau+1} - \mathbf{L}_r^d \mathbf{x}^{\tau+1})$
- 9:  $\gamma_u \leftarrow \gamma_u + \rho_u^\tau (\mathbf{x} - \mathbf{z}_u)$
- 10:  $\gamma_d \leftarrow \gamma_d + \rho_d^\tau (\mathbf{x} - \mathbf{z}_d)$   $\triangleright$  *Update multipliers with eqs. (17) and (18)*
- 11: **end for**
- 12: **return**  $\mathbf{x}$

**Algorithm 3 Overall Algorithm of the Proposed Unrolled Network**


---

1026  
1027  
1028 **Require:** An observation from the training set  $\mathbf{y} \in \mathbb{R}^{T+1}$ .  
1029 **Ensure:** The reconstructed signal  $\mathbf{x} \in \mathbb{R}^{T+S+1}$  containing both the observation and prediction.  
1030      $\triangleright$  *Notations: Observing stations count  $N$ , time-stamp series for the entire reconstructed signal*  
1031      $\mathbf{t} \in \mathbb{R}^{T+S+1}$ .  
1032     1:  $\mathbf{e}_{\text{st}} \leftarrow \text{STEmb}(N, \mathbf{t})$       $\triangleright$  *Shared spatial-temporal embeddings (Appendix F.3.1)*  
1033     2:  $\mathbf{x}_{\text{pred}}^0 \leftarrow \text{Extrapolation}(\mathbf{y}), \mathbf{x}_{\text{pred}}^0 \in \mathbb{R}^S$       $\triangleright$  *Initial prediction (Appendix F.4)*  
1034     3:  $\mathbf{x} \leftarrow [\mathbf{y}; \mathbf{x}_{\text{pred}}^0], \mathbf{x} \in \mathbb{R}^{T+S+1}$   
1035     4: **for**  $i$  in  $0 : B$  **do**      $\triangleright$  *B ADMM blocks and graph learning modules*  
1036         5:  $\mathbf{e}_i \leftarrow [\mathbf{x}_i; \mathbf{e}_{\text{st},i}], \mathbf{f}_i \leftarrow F^u(\mathbf{e}_i)$  for  $i = 1, \dots, N$       $\triangleright$  *Feature extraction (Appendix F.3.2)*  
1037         6:  $\mathbf{L}^{u(h)} \leftarrow \text{UGL}_h(\mathbf{f}_1, \dots, \mathbf{f}_N), \mathbf{L}_r^{d(h)} \leftarrow \text{DGL}_h(\mathbf{f}_1, \dots, \mathbf{f}_N), h = 1, 2, \dots, H$   
1038              $\triangleright$  *Multi-head graph learning (Section 4.3)*  
1039         7:  $\mathbf{x}_{\text{new}}^{(h)} \leftarrow \text{ADMM\_Block}_h(\mathbf{x}, \mathbf{L}^{u(h)}, \mathbf{L}_r^{d(h)}), h = 1, 2, \dots, H$   
1040         8:  $\mathbf{x}_{\text{new}} = \sum_{h=1}^H a_h \mathbf{x}_{\text{new}}^{(h)}$       $\triangleright$  *Learnable merge of multi-head results*  
1041         9:  $\mathbf{x} \leftarrow p_i \mathbf{x}_{\text{new}} + (1 - p_i) \mathbf{x}$       $\triangleright$  *Residual connections with learnable parameter  $p_i$*   
1042     10: **end for**  
1043     11: **return**  $\mathbf{x}$

---

**F MODEL SETUP****F.1 ADMM BLOCKS SETUP**

1050 We initialize  $\mu_u, \mu_{d,1}, \mu_{d,2}$  with 3, and the  $\rho, \rho_u, \rho_d$  with  $\sqrt{N/(T+S+1)}$ , where  $N$  is the number  
1051 of sensors and  $T+S+1$  is the full length of the recovered signal (both the observed part and the part  
1052 to be predicted). Those parameters are trained every ADMM layer. We initialize the CGD parameters  
1053 ( $\alpha, \beta$ )s as 0.08, and tune them for every CGD iteration. We clamp each element of the CGD step  
1054 size  $\alpha$  to  $[0, 0.8]$  for stability. Each element of the momentum term coefficient  $\beta$  is clamped to be  
1055 non-negative.

**F.2 GRAPH LEARNING MODULES SETUP**

1060 We define multiple PD matrices  $\mathbf{P}$ s and  $\mathbf{M}$ s in our unrolling model due to the different impact  
1061 strengths at different times. For the undirected graph, we define  $\mathbf{M}^t$  as the Mahanobolis distance  
1062 matrix for the undirected graph slice at time  $t$  to vary the impacting strength between spatial  
1063 neighbors in different time steps. For the directed graph, we assume that the impact strength varies  
1064 with different *intervals* but remains uniform for different monitoring stations. Thus, we define  $\mathbf{P}^w$   
1065 as the Mahanobolis distance matrix to represent the temporal influence of interval  $w$  for each node.  
1066 Therefore, we define  $(T+S+1)$  learnable PSD matrices  $\mathbf{M}^0, \dots, \mathbf{M}^{T+S}$  for undirected graph  
1067 learning, and  $W$  PSD learnable matrices  $\mathbf{P}^1, \dots, \mathbf{P}^W$  for directed graph learning in total.

1068 We learn each PSD matrix  $\mathbf{M}$  in each UGL by learning a square matrix  $\mathbf{M}_0 \in \mathbb{R}^{K \times K}$  by  $\mathbf{M} =$   
1069  $\mathbf{M}_0^\top \mathbf{M}_0$ , and initialize  $\mathbf{M}_0$  with a matrix whose diagonal is filled with 1.5. Similarly, we learn  
1070 each PSD matrix  $\mathbf{P}^w$  in DGL by learning  $\mathbf{P}_0^w$  where  $\mathbf{P}^w = (\mathbf{P}_0^w)^\top \mathbf{P}_0^w$ . We initialize  $\mathbf{P}_0^w$  with  
1071  $(1 + 0.2w/W)\mathbf{I}$  to set the edge weights representing influences of longer intervals slightly smaller.

1072 We learn 4 mixed graphs in parallel with 4 sets of matrices  $\mathbf{M}^0, \dots, \mathbf{M}^{T+S}$  and  $\mathbf{P}^1, \dots, \mathbf{P}^W$  to  
1073 implement the *multi-head* attention mechanism in conventional transformers. The output signals of  
1074 all 4 channels are integrated to one by a linear layer at the end of each ADMM block.

**F.3 FEATURE EXTRACTORS SETUP**

1075  
1076  
1077 We detailedly introduce our design of the feature extractor described in Section 4.3. Figure 5b shows  
1078 the overall structure of our feature extractors.  
1079

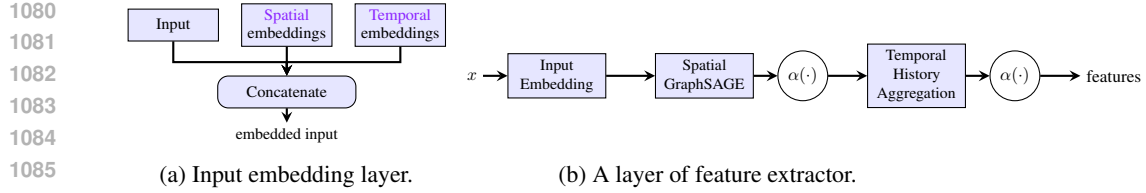


Figure 5: Framework of our feature extractor. (a) We first propose a non-parametric input embedding layer to combine signal with spatial/temporal embeddings. (b) Then aggregate embedded inputs of spatial neighbors and history information in a time window to generate a layer of feature extractor.

### F.3.1 INPUT EMBEDDINGS

We generate *low-parametric* spatial and temporal embeddings for each node in our mixed spatio-temporal graph, see Figure 5a. For spatial embeddings, we embed each monitoring station with a 5-dimensional learnable vector, and broadcast the embeddings to every time step in the product graph. For temporal embeddings, we follow Feng & Tassiulas (2022) to use real-time-stamp, time-in-day and day-in-week embeddings. To achieve better fitting capacity while keeping the model size small, we set only the time-in-day and day-in-week embeddings *learnable*, and use *fixed* embeddings for real time stamps. We use a fixed 10-dimensional sine-cosine position embeddings (Vaswani et al., 2018) for each time stamp:

$$\begin{aligned} e[t, 2i] &= \sin(t/10000^i) \\ e[t, 2i + 1] &= \cos(t/10000^i), i = 0, 1, 2, 3, 4, \end{aligned} \quad (41)$$

and set dimensions of learnable time-in-day and day-in-week embeddings as 6 and 4. All embeddings are *shared* across all graph learning modules.

We directly concatenate these spatial and temporal embedding to our signals  $\mathbf{x}_j^t$  for node  $j$  at time  $t$  as  $\mathbf{e}_j^t = [\mathbf{x}_j^t; \mathbf{e}_j^S; \mathbf{e}_t^T]$  for feature extractors, where  $\mathbf{e}_j^S$  denote the spatial Laplacian embeddings of node  $j$  of the physical graph, and  $\mathbf{e}_t^T$  denotes the temporal position embeddings of time step  $t$ .

### F.3.2 FEATURE EXTRACTION FROM EMBEDDINGS

We use a variant of GraphSAGE (Hamilton et al., 2017) to aggregate spatial features and introduce Temporal History Aggregation (THA) to aggregate temporal features. The GraphSAGE module aggregates the inputs of each node’s  $k$ -NNs with a simple and uniformed linear layer to generate  $K$ -dimension spatial features for each node of each time step. The THA module aggregates the embedded inputs of past  $W$  time steps for each node with a uniformed linear layer, where  $W$  is the size of the time window. For both GraphSAGE and THA modules, we pad zeros to inputs as placeholders, and use Swish (Ramachandran et al., 2017) as activation functions with  $\beta = 0.8$ .

## F.4 CONSTRUCTING GRAPHS $\mathbf{W}^{u,0}, \mathbf{W}^{d,0}$ FOR THE FIRST ADMM BLOCK

Our ADMM block recovers the full signal  $\mathbf{x}$  from observed signal  $\mathbf{y}$  with mixed graph  $\mathcal{G}^u, \mathcal{G}^d$  constructed from embedded *full* signal  $\mathbf{e} = [\mathbf{x}; \mathbf{e}^S; \mathbf{e}^T]$ . Thus, for the first block in ADMM, we need an *initial extrapolation* for the signals to predict.

We propose a simple module to make an initial guess of the signals: we modify the feature extractor described in Appendix F.3.2 to generate intermediate features from the embedded input of the observed part. Then we combine the feature and temporal dimensions, and use a simple linear layer with Swish activation to generate the initially extrapolated signal.

**Note:** Both the feature extractor and the initial prediction module are carefully designed to involve as few parameters as possible to make the entire model lightweight. Furthermore, low-parameter and simple feature engineering ensures that the effectiveness of our model is attributed to the unrolled blocks and graph learning modules, rather than to the trivial feature engineering parts.

## F.5 TRAINING SETUP

We preprocess the dataset by *standardizing* the entire time series for each node in space, but minimize a loss between the real groundtruth signal and the *rescaled* whole output signal. We use an Adam optimizer with learning rate  $5 \times 10^{-4}$ , together with a scheduler of reduction on validation loss with a reduction rate of 0.2 and patience of 5. We set different batch sizes for different datasets to leverage the most of the GPU memory (in this case, we use an entire Nvidia GeForce RTX 3090). Detailed selection of parameters such as number of nearest spatial neighbors  $k$  and time window size  $W$  is shown in Table 3.

Table 3: Training setup of graph construction and batch size for all traffic datasets.

Dataset	PeMS03	PeMS08	METR-LA
$k$	4	6	6
$W$	6	6	6
Batch size	12	16	16

## G FURTHER EXPERIMENTS

### G.1 ADDITIONAL EXPERIMENTAL RESULTS

Table 4 shows the performance comparison between our model and baselines in 30-/60-minute forecast on the PEMS08 (Guo et al., 2022) dataset which contains 170 nodes and 295 edges. Results show that our model also ranks top-3 in at least one metric of all experiments, and even gives the smallest MAPE among all models

Table 4: Experimental results on additional dataset PEMS08 for 30- and 60-minute forecast. The indication of bolded, highlighted and underlined entries is the same as Table 2.

Horizons Model	30 minutes (6 steps)			60 minutes (12 steps)		
	RMSE	MAE	MAPE (%)	RMSE	MAE	MAPE (%)
VAR	26.81	17.74	12.69	<u>28.59</u>	<u>18.99</u>	13.52
STGCN	36.58	24.39	16.83	50.61	33.01	20.57
STSGCN	26.73	17.43	11.48	30.22	19.74	<u>12.90</u>
GMAN	<b>24.84</b>	<b>15.73</b>	<u>10.87</u>	<b>26.22</b>	<b>16.83</b>	13.64
ST-Wave	26.85	17.19	11.08	<b>27.28</b>	<b>16.85</b>	<b>11.77</b>
PDFormer	<b>22.92</b>	<b>13.91</b>	<b>9.17</b>	<b>27.36</b>	<u>17.61</u>	<u>12.90</u>
STAEFormer	<u>26.54</u>	<u>16.62</u>	<u>10.88</u>	32.68	20.78	14.86
GWN	30.95	21.65	13.56	40.24	29.12	17.79
AGCRN	34.31	19.91	<u>10.88</u>	38.09	22.78	<u>12.75</u>
SITD	27.40	18.11	12.10	33.12	22.66	20.08
SimpleTM	<b>24.44</b>	<b>15.07</b>	<b>9.09</b>	32.48	20.26	<b>12.34</b>
<b>Ours</b>	<u>25.92</u>	<u>16.22</u>	<b>10.17</b>	<u>27.78</u>	<b>17.22</b>	<b>11.11</b>

### G.2 EXPERIMENTS ON THE ENTIRE DATASETS

To further strengthen our proposed method, we ran the 60-minute forecast experiment on the PEMS03 dataset without sampling with 70 epochs of training, and the results are shown in Table 5. Our unrolled

Table 5: Performance comparison of our model to the baselines on a 60-minute forecast for the entire PEMS03 dataset. The indication of bolded, highlighted and underlined entries is the same as Table 2.

Dataset	PEMS03			METR-LA		
Model	RMSE	MAE	MAPE (%)	RMSE	MAE	MAPE (%)
VAR	30.53	18.30	19.70	12.57	7.05	13.48
STGCN	36.77	21.81	19.37	<u>12.27</u>	5.88	11.79
STSGCN	28.17	17.46	16.86	12.73	<u>5.11</u>	<u>10.66</u>
PDFormer	<b>25.05</b>	<b>14.78</b>	<b>15.53</b>	<u>12.32</u>	<b>4.77</b>	<b>10.20</b>
GMAN	28.44	17.04	22.82	13.75	6.44	12.72
GWN	28.68	17.08	<u>16.49</u>	13.82	6.41	11.92
AGCRN	<u>27.65</u>	<u>15.77</u>	<b>15.02</b>	<b>12.18</b>	<b>4.74</b>	<u>10.50</u>
STID	<u>26.80</u>	<b>15.14</b>	<u>16.45</u>	<b>12.00</b>	<b>4.63</b>	<b>10.14</b>
SimpleTM	<b>24.51</b>	<b>15.32</b>	<b>15.55</b>	12.86	5.67	<b>9.89</b>
<b>Ours</b>	<b>26.71</b>	<u>16.36</u>	17.18	<b>12.22</b>	<u>5.35</u>	11.87

network still ranks top-3 in at least one of the metrics for each task, indicating the consistency of our performance with subsampling.

### G.3 INFERENCE-TIME COMPUTATIONAL AND MEMORY COST

We compute the inference-time computational and memory cost on 60-minute forecasting for the PEMS03 dataset. We calculate the total floating-point operations (FLOPs) during inference as the computational cost with a batch size of 1, and evaluate the peak allocated memory by CUDA for each batch of size 32 as the memory cost. Results in Table 6 show that our unrolled network has the least computation cost among all the baselines except AGCRN and STID, which are generally consistently outperformed by our model. Our model’s inference-time memory cost is only larger than SimpleTM and the two baselines mentioned above. In conclusion, our lightweight transformer reduces the computational cost to only 4.9% of transformer-based PDFormer and 11.2% of the GAT-based GMAN with reduced or at least comparable memory cost.

Table 6: Comparison of computational cost (GFLOPs) for each data point and inference memory cost (GB) for a batch of size 32 in 60-minute forecasting for the PEMS03 dataset.

Model*	<b>Ours</b>	STGCN	GMAN	ST-Wave	PDFormer
Computation (GFLOPs)	0.087	$0.013 \times 12^\dagger$	0.777	4.496	1.771
Memory (GB)	1.154	$1.065 \times 12$	3.382	1.526	1.910
Model	STAEFormer	GWN	AGCRN	STID	SimpleTM
Computation (GFLOPs)	4.428	$0.300 \times 12$	0.0002	0.036	0.670
Memory (GB)	1.846	$0.411 \times 12$	0.523	0.034	0.430

\* The computational and memory cost for STSGCN is not included due to the limitation of applying tools such as `thop` in the public code.

† The  $\times 12$  notation in the table means that the model predicts the outputs recursively.

### G.4 TRAINING PERFORMANCE W.R.T. TRAINING DATA COST

Due to the parameter efficiency of our model, our unrolling model is also data-efficient under the same observation length. We trained our models and baselines with reduced training datasets by selecting sample windows from the sequence with larger strides of 5 and 10 in the PEMS03 dataset for 60-minute forecasting. The results below show that baseline models with large parameter counts suffer from steep performance degradation when the training data becomes scarce. In contrast, our model maintains stable performance under limited data conditions, demonstrating both lower data requirements and better generalizability. This makes it particularly suitable for deployment in data-scarce or resource-constrained environments.

Table 7: Performance under different sampling strides (metrics: MAE / RMSE / MAPE(%)) on 60-minute forecasting for PEMS03 dataset.

Model	Stride = 3 (original)	Stride = 5	Stride = 10
Ours	26.96 / 16.58 / 18.07	28.98 / 17.73 / 18.63	29.63 / 18.58 / 19.76
VAR	30.54 / 18.31 / 19.61	30.68 / 18.37 / 19.74	31.25 / 18.71 / 21.22
STSGCN	30.62 / 19.35 / 19.15	32.37 / 20.67 / 20.66	33.22 / 21.30 / 21.58
PDFormer	27.16 / 17.26 / 21.21	25.70 / 15.37 / 15.99	77.29 / 44.36 / 126.63
AGCRN	29.90 / 16.76 / 15.32	39.25 / 20.60 / 16.95	62.58 / 109.27 / 30.72

## H ABLATION STUDIES

### H.1 ABLATION STUDIES FOR $\ell_2$ - AND $\ell_1$ - NORMS IN DIRECT GRAPH MODELING

First, we discuss the effectiveness of the new directed graph priors, DGLR and DGTV, as well as the benefit of directed-graph modeling of temporal relations. We separately test the effects of DGLR and DGTV by removing the prior and unrolling the ADMM algorithm. We denote these models as “w/o DGTV” and “w/o DGLR”. To test the benefit of the directed-graph modeling, we change the temporal graph  $\mathcal{G}^d$  into a fully undirected graph  $\mathcal{G}^n$  while maintaining the connections, and replace the directed graph priors with a GLR prior  $\mathbf{x}^\top \mathbf{L}^n \mathbf{x}$  with the same coefficient  $\mu_n = \mu_{d,2}$ . We denote the undirected-time unrolling model as “UT”. Detailed derivations of the ADMM algorithm for all the above ablation studies are in Appendix I.

Table 8: Performance comparison in RMSE / MAE / MAPE(%) of the ablation studies of directed-graph modeling and directed graph priors in 60-minute forecasting with PEMS03 and METR-LA.

Dataset	Mixed-Graph	UT	w/o DGTV	w/o DGLR
PEMS03	26.96 / 16.58 / 18.07	29.00 / 18.21 / 19.99	27.62 / 17.54 / 19.04	30.53 / 18.69 / 19.06
METR-LA	12.17 / 5.27 / 11.78	12.53 / 5.37 / 11.97	12.59 / 5.39 / 12.26	12.45 / 5.34 / 11.98

As is summarized in Table 8, removing the DGLR or DGTV term leads to noticeable drops in performance, which underscores their importance as directed graph priors. Moreover, our model outperforms the fully-undirected-graph-based model, demonstrating that sequential relationships—such as time—are better modeled as directed graphs.

### H.2 SENSITIVITY OF $k$ , $W$ IN GRAPH CONSTRUCTION

We change the number of spatial nearest neighbors  $k = 4, 6, 8$  and time window size  $W = 4, 6, 8$  and present the results for PEMS03 and METR-LA dataset in Table 9. The selected hyperparameters are marked with “\*”. Decreasing or increasing  $k, W$  leads to performance degradation. Our selection of  $k, W$  are optimal among all test hyperparameters, considering the tradeoff with the cost of time and memory.

Table 9: Performance on 60-minute forecasting with PEMS03 (left) and METR-LA (right) under different  $k, W$  settings.

PEMS03, 60-minute forecast					METR-LA, 60-minute forecast				
$k$	$W$	RMSE	MAE	MAPE(%)	$k$	$W$	RMSE	MAE	MAPE(%)
4*	6*	26.96	16.58	18.07	6*	6*	12.17	5.27	11.78
4	4	27.23	16.80	17.73	4	6	12.28	5.24	11.86
4	8	27.30	17.34	19.34	8	6	12.45	5.26	11.91
6	6	29.57	18.00	19.43	6	4	12.38	5.23	11.84
8	6	28.40	17.99	18.48	6	8	12.52	5.42	12.29

### H.3 SENSITIVITY OF NUMBER OF BLOCKS AND LAYERS IN THE UNROLLING MODEL

We vary the number of blocks and ADMM layers for PEMS03 dataset and present the results in Table 10. The results show that reducing number of blocks or layers will increase the error in prediction. Our selected settings  $n_{\text{blocks}} = 5, n_{\text{layers}} = 25$  is the optimal setting so far in consideration of the trade-off in computing memory and time cost.

Table 10: Performance on 60-minute forecast with PEMS03 dataset under different number of blocks and layers.

$n_{\text{blocks}}$	$n_{\text{layers}}$	PEMS03, 60-minute forecast			METR-LA, 60-minute forecast		
		RMSE	MAE	MAPE(%)	RMSE	MAE	MAPE(%)
5*	25*	<b>26.96</b>	<b>16.58</b>	18.07	<b>12.17</b>	5.27	<b>11.78</b>
5	20	27.52	17.36	18.97	12.51	5.30	11.99
5	30	27.25	17.27	18.77	12.51	5.29	11.96
4	25	27.84	17.23	<b>17.96</b>	12.43	5.32	11.93
6	25	26.97	16.92	17.97	12.46	<b>5.25</b>	11.80

### H.4 ABLATION STUDY FOR THE SPLITTING- $\ell_2$ -TERM STRATEGY

We also ran ablations of the splitting-term strategy by optimizing  $\mathbf{x}^{\tau+1}$  using only eq. (7). CG in solving eq. (7) converges much slower than in solving eq. (10)-eq. (12), and the unrolled model can become numerically unstable in training, indicating that splitting the  $\ell_2$  terms can stabilize the training process.

## I OPTIMIZATION FORMULATION AND ADMM ALGORITHM IN ABLATION STUDIES

### I.1 OPTIMIZATION & ALGORITHM WITHOUT DGTV TERM

By removing the DGTV term  $\|\mathbf{L}_r^d \mathbf{x}\|_1$  in eq. (4), and introducing the auxiliary variables  $\mathbf{z}_u, \mathbf{z}_d$  as eq. (9), we rewrite the unconstrained version of the optimization by augmented Lagrangian method as

$$\begin{aligned} \min_{\mathbf{x}, \mathbf{z}_u, \mathbf{z}_d} & \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \mu_u \mathbf{z}_u^\top \mathbf{L}^u \mathbf{z}_u + \mu_{d,2} \mathbf{z}_d^\top \mathcal{L}_r^d \mathbf{z}_d + (\gamma_u^\tau)^\top (\mathbf{x} - \mathbf{z}_u) + \frac{\rho_u}{2} \|\mathbf{x} - \mathbf{z}_u\|_2^2 \\ & + (\gamma_d^\tau)^\top (\mathbf{x} - \mathbf{z}_d) + \frac{\rho_d}{2} \|\mathbf{x} - \mathbf{z}_d\|_2^2. \end{aligned} \quad (42)$$

We optimize  $\mathbf{x}, \mathbf{z}_u$  and  $\mathbf{z}_d$  in turn at iteration  $\tau$  as follows:

$$\left( \mathbf{H}^\top \mathbf{H} + \frac{\rho_u + \rho_d}{2} \mathbf{I} \right) \mathbf{x}^{\tau+1} = \mathbf{H}^\top \mathbf{y} - \frac{\gamma_u^\tau + \gamma_d^\tau}{2} + \frac{\rho_u}{2} \mathbf{z}_u^\tau + \frac{\rho_d}{2} \mathbf{z}_d^\tau, \quad (43)$$

$$\left( \mu_u \mathbf{L}^u + \frac{\rho_u}{2} \mathbf{I} \right) \mathbf{z}_u^{\tau+1} = \frac{\gamma_u^\tau}{2} + \frac{\rho_u}{2} \mathbf{x}^{\tau+1}, \quad (44)$$

$$\left( \mu_{d,2} \mathcal{L}_r^d + \frac{\rho_d}{2} \mathbf{I} \right) \mathbf{z}_d^{\tau+1} = \frac{\gamma_d^\tau}{2} + \frac{\rho_d}{2} \mathbf{x}^{\tau+1}. \quad (45)$$

The updating of  $\gamma_u^\tau, \gamma_d^\tau$  follows eq. (18).

### I.2 OPTIMIZATION & ALGORITHM WITHOUT DGLR TERM

By removing the DGLR term  $\mathbf{x} \mathcal{L}_r^d \mathbf{x}$  in eq. (4) and introduce only the auxiliary variables  $\phi$  in eq. (5) and  $\mathbf{z}_u$  in eq. (9), we rewrite the unconstrained version of the optimization of  $\mathbf{x}^{\tau+1}$  by augmented Lagrangian method as

$$\min_{\mathbf{x}, \mathbf{z}_u} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \mu_u \mathbf{z}_u^\top \mathbf{L}^u \mathbf{z}_u + (\gamma^\tau)^\top (\phi^\tau - \mathbf{L}_r^d \mathbf{x}) + \frac{\rho}{2} \|\phi^\tau - \mathbf{L}_r^d \mathbf{x}\|_2^2 + (\gamma_u^\tau)^\top (\mathbf{x} - \mathbf{z}_u) + \frac{\rho_u}{2} \|\mathbf{x} - \mathbf{z}_u\|_2^2 \quad (46)$$

We optimize  $\mathbf{x}$  and  $\mathbf{z}_u$  in turn at iteration  $\tau$  as follows:

$$\left(\mathbf{H}^\top \mathbf{H} + \frac{\rho}{2} \mathcal{L}_r^d + \frac{\rho_u}{2} \mathbf{I}\right) \mathbf{x}^{\tau+1} = \mathbf{H}^\top \mathbf{y} + (\mathbf{L}_r^d)^\top \left(\frac{\gamma^\tau}{2} + \frac{\rho}{2} \phi^\tau\right) - \frac{\gamma_u^\tau}{2} + \frac{\rho_u}{2} \mathbf{z}_u^\tau \quad (47)$$

$$\left(\mu_u \mathbf{L}^u + \frac{\rho_u}{2} \mathbf{I}\right) \mathbf{z}_u^{\tau+1} = \frac{\gamma_u^\tau}{2} + \frac{\rho_u}{2} \mathbf{x}^{\tau+1} \quad (48)$$

The optimization of  $\phi^{\tau+1}$  is the same as eq. (15), and the solution is the same as eq. (16). The updating of  $\gamma, \gamma_u$  follows eq. (17) and eq. (18).

### I.3 ALGORITHM UNROLLING WITH AN UNDIRECTED TEMPORAL GRAPH

In this section, we are changing the directed temporal graph  $\mathcal{G}^d$  into an *undirected* graph while maintaining the connections. Denote the undirected temporal graph as  $\mathcal{G}^n$ . We introduce the GLR operator  $\mathbf{x}^\top \mathbf{L}^n \mathbf{x}$  for  $\mathcal{G}^n$  on signal  $\mathbf{x}$  where  $\mathbf{L}^n$  is the normalized Laplacian matrix of  $\mathcal{G}^n$ . Thus the optimization problem is as follows:

$$\min_{\mathbf{x}} \|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2 + \mu_u \mathbf{x}^\top \mathbf{L}^u \mathbf{x} + \mu_n \mathbf{x}^\top \mathbf{L}^n \mathbf{x} \quad (49)$$

where  $\mu_u, \mu_n \in \mathbb{R}_+$  are weight parameters for the two GLR priors. The optimization is similar to that in Appendix I.1. By a similar approach which introduces auxiliary variables  $\mathbf{z}_u = \mathbf{x}, \mathbf{z}_n = \mathbf{x}$ , we optimize  $\mathbf{x}, \mathbf{z}_u, \mathbf{z}_n$  as follows:

$$\left(\mathbf{H}^\top \mathbf{H} + \frac{\rho_u + \rho_n}{2} \mathbf{I}\right) \mathbf{x}^{\tau+1} = \mathbf{H}^\top \mathbf{y} - \frac{\gamma_u^\tau + \gamma_n^\tau}{2} + \frac{\rho_u}{2} \mathbf{z}_u^\tau + \frac{\rho_n}{2} \mathbf{z}_n^\tau, \quad (50)$$

$$\left(\mu_u \mathbf{L}^u + \frac{\rho_u}{2} \mathbf{I}\right) \mathbf{z}_u^{\tau+1} = \frac{\gamma_u^\tau}{2} + \frac{\rho_u}{2} \mathbf{x}^{\tau+1}, \quad (51)$$

$$\left(\mu_n \mathbf{L}^n + \frac{\rho_n}{2} \mathbf{I}\right) \mathbf{z}_n^{\tau+1} = \frac{\gamma_n^\tau}{2} + \frac{\rho_n}{2} \mathbf{x}^{\tau+1}. \quad (52)$$

The updating of  $\gamma_u^\tau, \gamma_n^\tau$  is also similar to eq. (18) as follows:

$$\gamma_u^{\tau+1} = \gamma_u^\tau + \rho_u (\mathbf{x}^{\tau+1} - \mathbf{z}_u^{\tau+1}), \quad \gamma_n^{\tau+1} = \gamma_n^\tau + \rho_n (\mathbf{x}^{\tau+1} - \mathbf{z}_n^{\tau+1}). \quad (53)$$

The undirected temporal graph  $\mathcal{G}^n$  is learned by an UGL described in Section 4.3, and is initialized similar to the spatial UGL described in Appendix F.2. The initialization of  $\mu_n, \rho_n$  and the CGD parameters in solving eq. (52) follows the setup in Appendix F.1.

## J INTERPRETING THE LEARNED GRAPH VIA EIGENVECTOR CENTRALITY

We analyze the relationship between the last undirected graph learned by the model and the observed traffic flows. For the best trained model on a 60-minute forecast on PEMS03, we calculate the learned weighted adjacency matrices  $\mathbf{W}^u$  of the undirected graphs on the first 288 time steps (24 hours) in the test set. We calculate the normalized positive eigenvector corresponding to the largest eigenvalue of  $\mathbf{W}^u$ , known as the *Perron vector*, which is strictly positive by the Perron-Frobenius Theorem (Horn & Johnson, 2012). The Perron vector  $\mathbf{v} \in \mathbb{R}^N$  reflects the *eigenvector centrality* (Bonacich, 1987) of the individual nodes: the ‘‘importance’’ of individual nodes, where a node is deemed important if it is connected to other important nodes, *i.e.*,  $\lambda v_i = \sum_j W_{i,j}^u v_j$  for  $v_j > 0$ .

Figure 6 shows the evolution of the top-10 Perron vector entries and the corresponding traffic flow. First, we observe that evident disturbances in the Perron vector entries only exist when there is a rapid lift in the traffic flows from near zero to hundreds, practically corresponding to the beginning of a new morning. Furthermore, the eigenvector centrality returns to the initial state once the traffic flow becomes stable, and there are no significant changes in the ranking of the eigenvector centrality, indicating that **the node set with the largest Perron vector entries is generally consistent throughout the observation**, which marks the critical congestion hubs.

To further examine the role of the discovered important nodes in short-range interactions, we then select the four most important nodes #328, #27, #350 and #171 according to the elements in the Perron vectors, and plot the traffic flow of each node and the physical neighbors in selected

1404  
1405  
1406  
1407  
1408  
1409  
1410  
1411  
1412  
1413  
1414  
1415  
1416  
1417  
1418  
1419  
1420  
1421  
1422  
1423  
1424  
1425  
1426  
1427  
1428  
1429  
1430  
1431  
1432  
1433  
1434  
1435  
1436  
1437  
1438  
1439  
1440  
1441  
1442  
1443  
1444  
1445  
1446  
1447  
1448  
1449  
1450  
1451  
1452  
1453  
1454  
1455  
1456  
1457

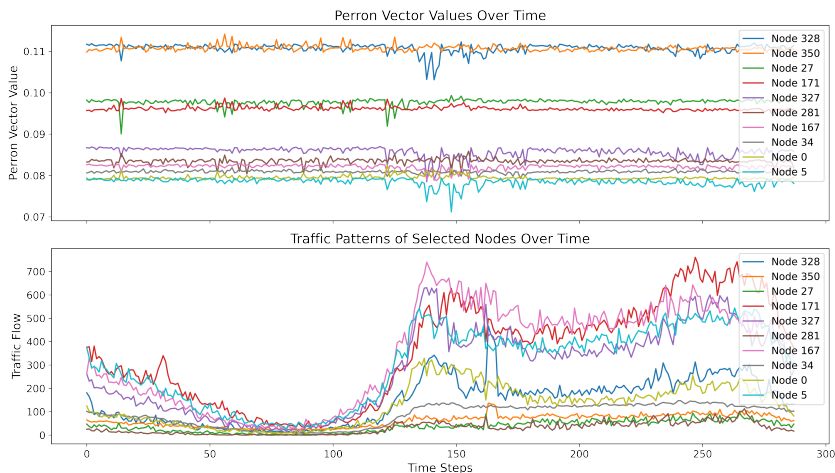
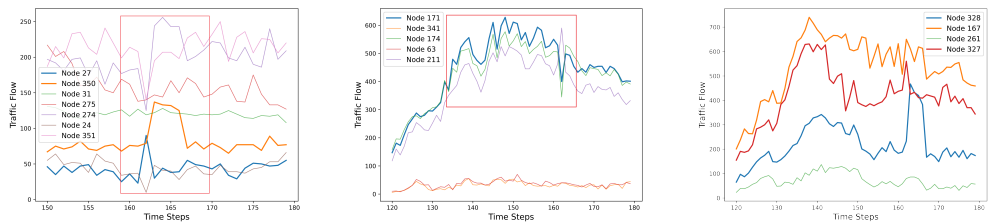


Figure 6: The evolution of eigenvector centrality in the learned undirected graph in the last block and traffic flow in the first 24 hours in the test split of the PEMS03 dataset.



(a) Neighborhood of nodes #27 and #350. (b) Neighborhood of node #171. (c) Neighborhood of node #328.

Figure 7: The traffic patterns of the nodes in the neighborhood of the selected important nodes by Perron vectors. The traffic flows of selected nodes are bolded.

windows of 30 and 60 steps in the daytime in Figure 7, where the selected important nodes are plotted with bolder lines (note that the neighborhood of nodes #27 and #350 are exactly the same set). In the red box in Figure 7a, the spike in node #27 results in the negative spike in the neighboring nodes #274, #24 and #351. Further, the combined effect of this spike and the bump in node #350 results in the overshoot after the downward spike in node #274. The signal trends in the red box in Figure 7b clearly reflect the time delay of the traffic flow in the neighboring nodes except the “quiet sequences” to the selected node #171. Both figures 7a and 7b demonstrate clear spatiotemporal dependency patterns captured by the learned undirected graph, where **high-centrality nodes play a dominant role in propagating local disturbances and temporal delays to their physical neighbors**. Additionally, as shown in Figure 7c, for the traffic patterns in the neighborhood of node #328, which has the largest Perron vector entry, two of its three neighboring nodes exhibit the 5th (#171) and 7th (#167) largest eigenvector centralities. This observation is consistent with the rationale that the most important node tends to be connected to other highly important nodes. Consequently, the traffic flow trends in the neighborhood of node #350 tend to be less affected by the central node, as its neighbors exhibit relatively high eigenvector centralities themselves.

**In conclusion, our model successfully discovers physically interpretable graph structures from data, and the learned node importance is highly consistent with actual traffic dynamics.**

Practically, our model can identify the key intersections that critically influence traffic dynamics and further guide traffic engineering in alleviating future traffic congestion.

## K ERROR BAR FOR OUR MODEL IN TABLE 2

Table 11 shows the  $2\sigma$  error bars of our model’s performance in the 60-minute forecast in Table 2 in Section 5, each of which is computed from 3 runs on different random seeds.

Table 11:  $2\sigma$  error bar of the performance metrics of our model in 60-minute traffic forecast in all three datasets.

Dataset	PEMS03	METR-LA	PEMS08
$2\sigma_{\text{RMSE}}$	0.05	0.04	0.04
$2\sigma_{\text{MAE}}$	0.02	0.02	0.02
$2\sigma_{\text{MAPE}(\%)}$	0.03	0.04	0.04

## L REPRODUCIBILITY

For reproducibility, our code is available in <https://anonymous.4open.science/r/Unrolling-GSP-STForecast-6AEE>.

## M LLM USAGE DECLARATION

We affirm that no LLMs were used in research ideation, code implementation, or writing of this paper; their use was limited to correcting grammatical errors.