

Localised Adaptive Spatial-Temporal Graph Neural Network

Wenyang Duan
School of Mathematics and Computer
Science
Nanchang University
Nanchang, China
wenyangduan@ncu.edu.cn

Xiaoxi He*
Faculty of Science and Technology
University of Macau
Macau, China
hexiaoxi@um.edu.mo

Zimu Zhou
School of Data Science
City University of Hong Kong
Hong Kong, China
zimuzhou@cityu.edu.hk

Lothar Thiele
D-ITET
ETH Zurich
Zurich, Switzerland
thiele@ethz.ch

Hong Rao
School of Software
Nanchang University
Nanchang, China
raohong@ncu.edu.cn

ABSTRACT

Spatial-temporal graph models are prevailing for abstracting and modelling spatial and temporal dependencies. In this work, we ask the following question: *whether and to what extent can we localise spatial-temporal graph models?* We limit our scope to adaptive spatial-temporal graph neural networks (ASTGNNs), the state-of-the-art model architecture. Our approach to localisation involves sparsifying the spatial graph adjacency matrices. To this end, we propose Adaptive Graph Sparsification (AGS), a graph sparsification algorithm which successfully enables the localisation of ASTGNNs to an extreme extent (fully localisation). We apply AGS to two distinct ASTGNN architectures and nine spatial-temporal datasets. Intriguingly, we observe that spatial graphs in ASTGNNs can be sparsified by over 99.5% without any decline in test accuracy. Furthermore, even when ASTGNNs are fully localised, becoming graphless and purely temporal, we record no drop in accuracy for the majority of tested datasets, with only minor accuracy deterioration observed in the remaining datasets. However, when the partially or fully localised ASTGNNs are reinitialised and retrained on the same data, there is a considerable and consistent drop in accuracy. Based on these observations, we reckon that (i) in the tested data, the information provided by the spatial dependencies is primarily included in the information provided by the temporal dependencies and, thus, can be essentially ignored for inference; and (ii) although the spatial dependencies provide redundant information, it is vital for the effective training of ASTGNNs and thus cannot be ignored during training. Furthermore, the localisation of ASTGNNs holds the potential to reduce the heavy computation overhead required on large-scale spatial-temporal data and further enable the distributed deployment of ASTGNNs.

*corresponding author

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

KDD '23, August 6–10, 2023, Long Beach, CA, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0103-0/23/08...\$15.00
<https://doi.org/10.1145/3580305.3599418>

CCS CONCEPTS

• **Computing methodologies** → **Neural networks.**

KEYWORDS

graph sparsification, spatial-temporal graph neural network, spatial-temporal data

ACM Reference Format:

Wenyang Duan, Xiaoxi He, Zimu Zhou, Lothar Thiele, and Hong Rao. 2023. Localised Adaptive Spatial-Temporal Graph Neural Network. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '23)*, August 6–10, 2023, Long Beach, CA, USA. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3580305.3599418>

1 INTRODUCTION

An increasing number of modern intelligent applications [3, 6, 24] rely on *spatial-temporal data*, i.e., data collected across both space and time. Spatial-temporal data often contain *spatial and temporal dependencies*, i.e., the current measurement at a particular location has causal dependencies on the historical status at the same and other locations. Learning these spatial and temporal dependencies is usually the essence of spatial-temporal data mining and plays a vital role in spatial-temporal inference [1].

Spatial-temporal data can be effectively represented by *spatial-temporal graph models*, which can depict complex relationships and interdependencies between objects in non-Euclidean domains [2, 11, 12, 19, 21, 24, 24, 26, 28]. Of our particular interest are *adaptive spatial-temporal graph neural networks* (ASTGNNs), a popular class of spatial-temporal graph models which have demonstrated outstanding performance in applications involving spatial-temporal data such as traffic forecasting, blockchain price prediction, and biosurveillance forecasting [3, 6–8, 24, 25].

In this work, we ask the following question: **whether and to what extent can we localise spatial-temporal graph models?** We limit our investigations to ASTGNNs, as they represent the state-of-the-art spatial-temporal graph model architecture. ASTGNNs usually model spatial dependencies with adaptive graph convolutional layers. The spatial dependencies are captured by learning the graph adjacency matrices. Therefore, the localisation of an ASTGNN is achieved via sparsifying the adjacency matrices, which can also be understood as pruning edges in the spatial graph. Note that

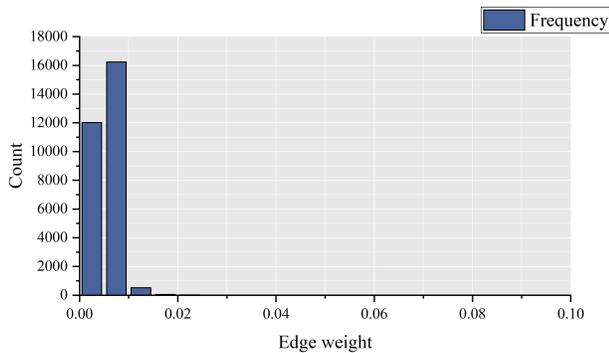


Figure 1: Histogram of edge weights of the spatial graph from an ASTGNN trained on the PeMSD8 dataset.

the localisation of an ASTGNN refers to the sparsification of only the adjacency matrices capturing the spatial dependencies. It is crucial not to be confused with sparsifying other weight matrices, such as those used in the temporal modules, often seen in GNN sparsification with pre-defined graph architectures [5, 16, 27, 29]. An initial clue which indicates that the localisation of ASTGNNs could be feasible, probably even to an extreme extent, is shown in Figure. 1). Here we show the distribution of the elements in the spatial graph’s adjacency matrix from an ASTGNN trained on the PeMSD8 dataset [4]. It is obvious that most edges in the spatial graph have weights close to zero.

We are interested in the localisation of spatial-temporal graph models for the following reasons:

- *A deeper understanding of the spatial and temporal dependencies in the data.* Although it is commonly accepted that both spatial and temporal dependencies are vital for inference, it is unclear whether and to what extent the information provided by these dependencies overlaps. If the localisation induces a marginal accuracy drop, then the information provided by the spatial dependencies is already largely included in the information contained in the temporal dependencies and, thus, unnecessary for inference.
- *Resource-efficient ASTGNN designs.* ASTGNNs are notoriously computation heavy as the size of the spatial graph grows quadratically with the number of vertices, thus limiting their usage on large-scale data and applications. The localisation of ASTGNNs may significantly reduce the resource requirement of these spatial-temporal graph models and enable new spatial-temporal applications.
- *Distributed deployment of spatial-temporal graph models.* In many cases, the data to construct the spatial-temporal graph models are collected via distributed sensing systems *e.g.*, sensor networks. However, making predictions of each vertex using these models requires the history of other vertices, thereby involving data exchange between sensor nodes. Localisation of spatial-temporal graph models may enable individual sensor nodes to make predictions autonomously without communicating with each other, which saves bandwidth and protects privacy in a distributed system.

We explore the localisation of ASTGNNs via Adaptive Graph Sparsification (AGS), a novel algorithm dedicated to the sparsification of adjacency matrices in ASTGNNs. The core of AGS is a differentiable approximation of the L_0 -regularization of a mask matrix, which allows the back-propagation to go past the regularizer and thus enables a progressive sparsification while training. We apply AGS to two representative ASTGNN architectures and nine different spatial-temporal datasets. The experiment results are surprising. (i) The spatial adjacency matrices can be sparsified to over 99.5% without deterioration in test accuracy on all datasets. (ii) Even fully localised ASTGNNs, which effectively degenerate to purely temporal models, can still provide decent test accuracy (no deterioration on most tested datasets while only minor accuracy drops on the rest). (iii) When we reinitialise the weights of the localised ASTGNNs and retrain them on the spatial-temporal datasets, we cannot reinstate the same inference accuracy. Figure. 2 summarises our experiments and observations.

Our empirical study implies two hypotheses. (i) In the tested spatial-temporal datasets, the information provided by the spatial dependencies is primarily included in the information provided by the temporal dependencies. Therefore, **the spatial dependencies can be safely ignored for inference** without a noteworthy loss of accuracy. (ii) Although the information contained in the spatial and temporal dependencies overlaps, such overlapping provides the vital redundancy necessary for properly training a spatial-temporal graph model. Thus, **the spatial dependencies cannot be ignored during training.**

Our main contributions are summarised as follows:

- To the best of our knowledge, this is the first study on the localisation of spatial-temporal graph models. We surprisingly observed that spatial dependencies could be largely ignored during inference without losing accuracy. Extensive experiments on common spatial-temporal datasets and representative ASTGNN architectures demonstrated that only a few edges (less than 0.5% on all tested datasets) are required to maintain the inference accuracy. More surprisingly, when the spatial dependencies are completely ignored, *i.e.*, the ASTGNNs are fully localised, they can still maintain a decent inference accuracy (no deterioration on most tested datasets, minor drops on the rest).
- With further investigations, we suggest the hypothesis that, although spatial dependencies can be primarily ignored during inference, they can drastically improve training effectiveness. This is supported by the observation that, if we reinitialise all parameters in the sparsified ASTGNNs and retrain them with the same data, the retrained networks yield considerably and consistently worse accuracy.
- To enable the localisation of ASTGNNs, we propose Adaptive Graph Sparsification (AGS), a novel graph sparsification algorithm dedicated to ASTGNNs. The core of AGS is a differentiable approximation of the L_0 -regularization of a mask matrix, which allows the back-propagation to go past the regularizer and thus enables a progressive sparsification while training.

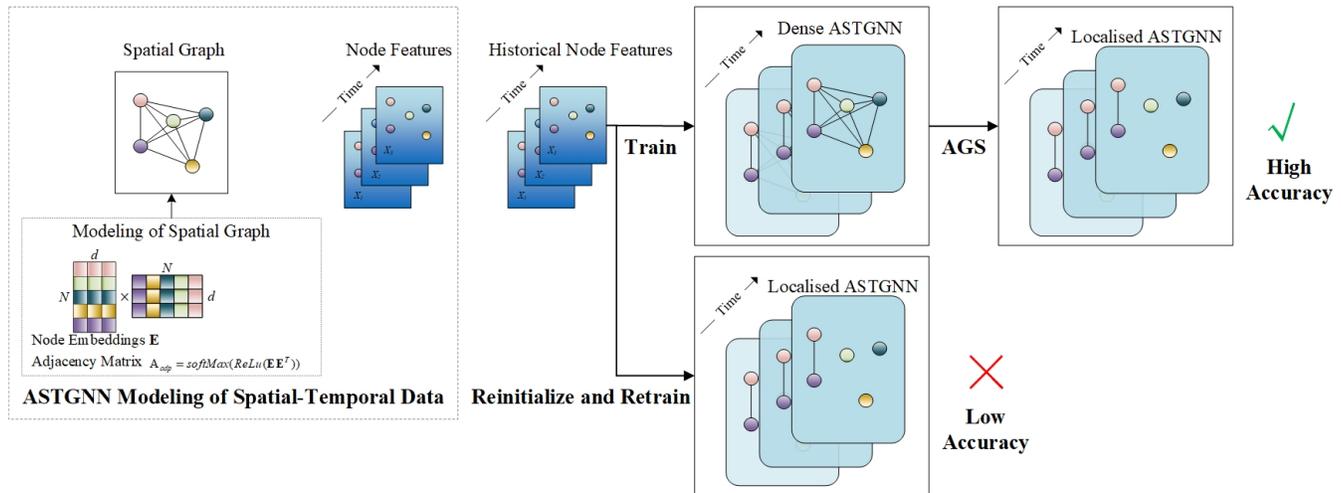


Figure 2: An overview of our experiments and observations. We train ASTGNNs on spatial-temporal datasets, achieving baseline accuracies. Then we localise the ASTGNNs with the proposed algorithm AGS, achieving accuracies comparable to the dense graph baselines. Finally, we reinitialise the localised ASTGNNs and retrain them on the same datasets, resulting in considerably and consistently deteriorated accuracies.

2 RELATED WORK

Our work is relevant to the following threads of research.

2.1 Spatial-Temporal Graph Neural Networks

Spatial-temporal graph neural networks (STGNNs) play an essential role in spatial-temporal data analysis for their ability to learn hidden patterns of spatial irregular signals varying across time [24]. These models often combine graph convolutional networks and recurrent neural networks. For example, Graph Convolutional Recurrent Network (GCRN) [21] combines an LSTM with ChebNet. Diffusion Convolutional Recurrent Neural Network [19] incorporates a proposed diffusion graph convolutional layer into GRU in an encoder-decoder manner to make multi-step predictions. Alternatively, CNN-based models can represent the temporal relations in spatial-temporal data in a non-recursive manner. For instance, CGCN [28] combines 1D convolutional layers with GCN layers. STGCN [26] composes a spatial-temporal model for skeleton-based action recognition using a 1D convolutional layer and a Partition Graph Convolution (PGC) layer. More recent proposals such as ASTGCN [11], STG2Seq [2], and LSGCN [12] further employ attention mechanisms to model dynamic spatial dependencies and temporal dependencies. In addition, some researchers consider the out-of-distribution generalisation of STGNN, and propose a domain generalisation framework based on hypernetworks to solve this problem [10]. However, these models adopt a predefined graph structure, which may not reflect the complete spatial dependency.

To capture the dynamics in graph structures of spatial-temporal data, an emerging trend is to utilize adaptive spatial-temporal graph neural networks (ASTGNNs). Graph WaveNet [25] proposes an AGCN layer to learn a normalized adaptive adjacency matrix without a pre-defined graph. ASTGAT introduces a network generator model that generates an adaptive discrete graph with the Gumbel-Softmax technique [15]. The network generator can adaptively infer

the hidden correlations from data. AGCRN [3] designs a Node Adaptive Parameter Learning enhanced AGCN (NAPL-AGCN) to learn node-specific patterns. Due to its start-of-the-art performance, NAPL-AGCN has been integrated into various recent models such as Z-GCNets [7], STG-NCDE [8], and TAMP-S2GCNets [6].

Despite the superior performance of ASTGNNs, they incur tremendous computation overhead, mainly because (i) learning an adaptive adjacency matrix involves calculating the edge weight between each pair of nodes, and (ii) the aggregation phase is computationally intensive. We aim at efficient ASTGNN inference, particularly for large graphs.

2.2 Graph Sparsification for GNNs

With graphs rapidly growing, the training and inference cost of GNNs has become increasingly expensive. The prohibitive cost has motivated growing interest in graph sparsification. The purpose of graph sparsification is to extract a small sub-graph from the original large one. SGCN [16] is the first to investigate graph sparsification for GNNs, *i.e.*, pruning input graph edges, and learned an extra DNN surrogate. NeuralSparse [29] prunes task-irrelevant edges from downstream supervision signals to learn robust graph representation. More recent works such as UGS [5] and GBET [27] explore graph sparsification from the perspective of the winning tickets hypothesis.

The aforementioned works only explore graph sparsification for vanilla GNNs and non-temporal data with pre-defined graphs. Our work differs by focusing on *spatial-temporal* GNNs with *adaptive* graph architectures.

3 PRELIMINARIES

This section provides a quick review of the representative architectures of ASTGNNs.

3.1 Spatial-Temporal Data as Graph Structure

Following the conventions in spatial-temporal graph neural network research [2, 3, 6–8, 11, 12, 19, 21, 24–26, 28], we represent the spatial-temporal data as a sequence of discrete frames \mathcal{X} with $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$, where $\mathcal{X} = \{X_1, X_2, \dots, X_T\}$. The graph \mathcal{G} is also known as the spatial network, which consists of a set of nodes \mathcal{V} and a set of edges \mathcal{E} . Let $|\mathcal{V}| = N$. Then the edges are presented with an adjacency matrix $A_t \in \mathbb{R}^{N \times N}$, where $X_t \in \mathbb{R}^{N \times C}$ is the node feature matrix with dimension C at timestep t , for $t = 1, \dots, T$.

Given graphs \mathcal{G} and \mathcal{T} historical observations $\mathcal{X}^{\mathcal{T}} = \{X_{t-\mathcal{T}}, \dots, X_{t-1}\} \in \mathbb{R}^{\mathcal{T} \times N \times C}$, we aim to learn a function \mathcal{F} which maps the historical observations into the future observations in the next \mathcal{H} timesteps:

$$\{X_t, \dots, X_{t+\mathcal{H}}\} = \mathcal{F}(X_{t-\mathcal{T}}, \dots, X_{t-1}; \theta, \mathcal{G}) \quad (1)$$

where θ denotes all the learnable parameters.

3.2 Modeling the Spatial Network \mathcal{G}

Since we focus on the role of spatial dependencies, we explain the state-of-the-art modelling of the spatial network in spatial-temporal graph neural networks (STGNNs).

The basic method to model the spatial network \mathcal{G} at timestep t with its node feature matrix X_t in representative STGNNs [2, 12, 25] is Graph Convolution Networks (GCNs). A one-layer GCN can be defined as:

$$Z_t = \sigma \left(\tilde{D}^{-\frac{1}{2}} \tilde{A}_t \tilde{D}^{-\frac{1}{2}} X_t W \right) \quad (2)$$

where $\tilde{A} = A + I_N$ is the adjacency matrix of the graph with added self-connections. I_N is the identity matrix. \tilde{D} is the degree matrix. $W \in \mathbb{R}^{C \times F}$ is a trainable parameter matrix. $\sigma(\cdot)$ is the activation function. $Z_t \in \mathbb{R}^{N \times F}$ is the output. All information regarding the input X_t at timestep t is aggregated in Z_t .

A key improvement to model the spatial network is to adopt Adaptive Graph Convolution Networks (AGCNs) to capture the dynamics in the graph \mathcal{G} , which leads to the adaptive spatial-temporal graph neural networks (ASTGNNs) [3, 6–8, 25]. In the following, we briefly explain Adaptive Graph Convolutional Recurrent Network (AGCRN) [3] and the extension with transformers, denoted as AGFormer, two representative ASTGNN models.

- **AGCRN.** It enhances the GCN layer by combining the normalized self-adaptive adjacency matrix with a Node Adaptive Parameter Learning (NAPL), which is known as NAPL-AGCN.

$$\begin{aligned} \mathbf{A}_{adp} &= \text{SoftMax} \left(\text{ReLU} \left(\mathbf{E} \mathbf{E}^T \right) \right) \\ Z_t &= \sigma \left(\mathbf{A}_{adp} X_t \mathbf{E} W_{\mathcal{G}} \right) \end{aligned} \quad (3)$$

where $W_{\mathcal{G}} \in \mathbb{R}^{d \times C \times F}$ and $\mathbf{E} W_{\mathcal{G}} \in \mathbb{R}^{N \times C \times F}$. $\mathbf{A}_{adp} \in \mathbb{R}^{N \times N}$ is the normalized self-adaptive adjacency matrix [25]. d is the embedding dimension for $d \ll N$. Each row of \mathbf{E} presents the embedding of the node. During training, \mathbf{E} is updated to learn the spatial dependencies among all nodes. Instead of directly learning W in (2) shared by all nodes, NAPL-AGCN uses $\mathbf{E} W_{\mathcal{G}}$ to learn node-specific parameters. From the view of one node (e.g., node i), $\mathbf{E}_i W_{\mathcal{G}}$ are the corresponding node-specific parameters according to its node embedding \mathbf{E}_i . Finally, to capture both spatial and temporal dependencies, AGCRN

integrates NAPL-AGCN and Gated Recurrent Units (GRU) by replacing the MLP layers in GRU with NAPL-AGCN.

- **AGFormer.** It extends AGCRN by modelling the temporal dependencies with transformers [17, 22]. A Transformer is a stack of transformer blocks. One block contains a multi-head self-attention mechanism and a fully connected feed-forward network. We replace the MLP layers in the multi-head self-attention mechanism with NAPL-AGCN to construct a transformer-based ASTGNN model, which we call AGFormer.

Modelling the spatial network with NAPL-AGCN achieves state-of-the-art performances on multiple benchmarks, and it has been widely adopted in diverse ASTGNN variants [3, 6–8]. However, NAPL-AGCN is much more inefficient than GCNs, since \mathbf{A}_{adp} is a matrix without zeros, while the adjacency matrix of a pre-defined graph in GCNs is far more sparse than \mathbf{A}_{adp} . This motivates us to explore the localisation of spatial-temporal graph models taking ASTGNNs as examples.

4 ADAPTIVE GRAPH SPARSIFICATION

This section presents Adaptive Graph Sparsification (AGS), a new algorithm dedicated to the sparsification of adjacency matrices in ASTGNNs.

Formulation. The NAPL-AGCN-based ASTGNNs with normalized self-adaptive adjacency matrix \mathbf{A}_{adp} can be trained using the following objective:

$$\mathcal{L}(\theta, \mathbf{A}_{adp}) = \frac{\sum_{\tau \in \mathcal{T}} \sum_{v \in \mathcal{V}} \left\| \mathbf{y}^{(\tau, v)} - \hat{\mathbf{y}}^{(\tau, v)} \right\|_1}{|\mathcal{V}| \times |\mathcal{T}|} \quad (4)$$

where \mathcal{T} is a training set, τ is a train sample, and $\mathbf{y}^{(\tau, v)}$ is the ground-truth of node v in τ .

Given a pre-trained model $\mathcal{F}(\cdot; \theta, \mathbf{A}_{adp})$, we introduce a mask $\mathbf{M}_{\mathbf{A}}$ to prune the adjacency matrix \mathbf{A}_{adp} . The shape of $\mathbf{M}_{\mathbf{A}}$ is identical to \mathbf{A}_{adp} . Specifically, given $\mathcal{F}(\cdot; \theta, \mathbf{A}_{adp})$, we obtain $\mathbf{M}_{\mathbf{A}}$ by optimizing the following objective:

$$\begin{aligned} \mathcal{L}_{AGS} &= \mathcal{L}(\theta, \mathbf{A}_{adp} \odot \mathbf{M}_{\mathbf{A}}) + \lambda \|\mathbf{M}_{\mathbf{A}}\|_0, \\ \|\mathbf{M}_{\mathbf{A}}\|_0 &= \sum_{i=1}^N \sum_{j=1}^N m_{(i,j)}, m_{(i,j)} \in \{0, 1\} \end{aligned} \quad (5)$$

where \odot is the element-wise product, $m_{(i,j)}$ corresponds to binary ‘‘gate’’ that indicates whether an edge is pruned, and λ is a weighting factor for L_0 -regularization of $\mathbf{M}_{\mathbf{A}}$.

Sparsification Algorithm. An intuitive way to get $\mathbf{M}_{\mathbf{A}}$ is to initialize a trainable weight matrix $\mathbf{U} \in \mathbb{R}^{N \times N}$ and map the entry $u_{(i,j)} \in \mathbf{U}$ into binary ‘‘gates’’ using a Bernoulli distribution: $m_{(i,j)} = \mathcal{B}(u_{(i,j)})$, where \mathcal{B} is a Bernoulli distribution. However, directly introducing \mathbf{U} to model $\mathcal{F}(\cdot; \theta, \mathbf{A}_{adp})$ has two problems.

- It may be unscalable for large-scale graphs.
- The L_0 sparsity penalty is non-differentiable.

For the scalability issue, we adopt a node adaptive weight learning technique to reduce the computation cost by simply generating \mathbf{U} with the node embedding \mathbf{E} :

$$\mathbf{U} = \mathbf{E} \mathbf{W}_{\mathbf{E}} \quad (6)$$

Algorithm 1: Adaptive Graph Sparsification (AGS)

Input: \mathcal{X} : input data, $\mathcal{F}(\cdot; \theta, \mathbf{A}_{adp})$: Spatial-temporal GNN with initialization self-adaptive adjacency matrix \mathbf{A}_{adp} , N_1 : number of pre-training iterations, N_2 : number of sparsification iterations, s_g : pre-defined sparsity level for graph.

Output: $\mathcal{F}(\cdot; \theta, \mathbf{A}_{adp} \odot \mathbf{M}_A)$

- 1: **while** iteration $i < N_1$ **do**
- 2: Forward to compute the loss in Eq.(4).
- 3: Back-propagate to update θ and \mathbf{A}_{adp} .
- 4: **end while**
- 5: Obtain pre-trained $\mathcal{F}(\cdot; \theta, \mathbf{A}_{adp})$.
- 6: Sort entries in \mathbf{A}_{adp} by magnitude in an ascending order then obtain list $K = \{k_i\}_{i=1}^{N \times N}$.
- 7: **while** iteration $i < N_2$ and $1 - \frac{\|\mathbf{M}_A\|_0}{\|\mathbf{A}_{adp}\|_0} < s_g$ **do**
- 8: set $\mathbf{M}_A^{(i,j)} = 1$ if $\mathbf{A}_{adp}^{(i,j)} \notin K_{:N^2 s_g}$.
- 9: Forward to compute the loss in Eq.(5).
- 10: Back-propagate to update θ and $\mathbf{A}_{adp} \odot \mathbf{M}_A$.
- 11: **end while**

where $W_E \in \mathbb{R}^{d \times N}$ is a trainable matrix.

For the non-differentiable issue, we introduce the hard concrete distribution instead of the Bernoulli distribution [9], which is a contiguous relaxation of a discrete distribution and can approximate binary values.

Accordingly, the computation of binary ‘‘gates’’ $m(i, j)$ can be formulated as:

$$z \sim \mathcal{U}(0, 1), s_{(i,j)} = \text{Sig} \left(\log z - \log(1 - z) + \log(u_{(i,j)}) \right) / \beta \quad (7)$$

$$\bar{s}_{(i,j)} = s_{(i,j)} (\zeta - \gamma) + \gamma, m_{(i,j)} = \min \left(1, \max \left(0, \bar{s}_{(i,j)} \right) \right)$$

where z is a uniform distribution, Sig is a sigmoid function, β is a temperature value and $(\zeta - \gamma)$ is the interval with $\zeta < 0$ and $\gamma > 1$. We set $\zeta = -0.1$ and $\gamma = 1.1$ in practice. Then, \mathbf{M}_A is applied to prune the lowest-magnitude entries in \mathbf{A}_{adp} , w.r.t. pre-defined ratios p_g .

Alg. 1 outlines the procedure of AGS. The pruning begins after the network is pre-trained (Line 5). First, we sort edges in the adjacency matrix by magnitude in ascending order (Line 6). Then we perform pruning iteratively (Line 7). The top p_g of the edges are removed, and the rest are retained. We identify the remaining edge by setting its corresponding entry in \mathbf{M}_A to 1 in each iteration (Line 8). The edges to be pruned are removed using Eq.(5) (Line 9).

Discussions. We make two notes on our AGS algorithm.

- AGS differs from prior magnitude-based pruning methods designed for graphs [5, 27] in that (i) it sparsifies *adaptive graphs* in spatial-temporal GNNs rather than vanilla GNNs and non-temporal data with pre-defined graphs; and (ii) it does not require iterative retraining on the pruned graphs to restore model accuracy.
- Pruning the adjacency matrix with AGS notably reduces the complexity of ASTGNNs for inference. The inference time

Table 1: Summary of datasets used in experiments.

Datasets	#Nodes	Range
PeMSD3	358	09/01/2018 - 30/11/2018
PeMSD4	307	01/01/2018 - 28/02/2018
PeMSD7	883	01/07/2017 - 31/08/2017
PeMSD8	170	01/07/2016 - 31/08/2016
Bytom	100	27/07/2017 - 07/05/2018
Decentral	100	14/10/2017 - 07/05/2018
Golem	100	18/02/2017 - 07/05/2018
CA	55	01/02/2020 - 31/12/2020
TX	251	01/02/2020 - 31/12/2020

complexity of unpruned NAPL-AGCN layers is $O(N^2 d + LN\mathcal{T}F^2 + L\mathcal{T}\|\mathbf{A}_{adp}\|_0 F + NdF)$. After sparsification, the inference time complexity of NAPL-AGCN layers is $O(N^2 d + LN\mathcal{T}F^2 + L\mathcal{T}\|\mathbf{A}_{adp} \odot \mathbf{M}_A\|_0 F + NdF)$, where $N^2 d$ is the time complexity of computing adaptive adjacency matrix, d is the embedding dimension, N is the number of nodes, $\|\mathbf{A}_{adp} \odot \mathbf{M}_A\|_0$ is the number of remaining edges, F is the representation dimension of node features, \mathcal{T} is the length of input, L is the number of layers.

5 EXPERIMENTS

To answer the question of whether and to what extent we can localise a spatial-temporal graph model, we conducted extensive experiments explained in this section.

5.1 Neural Network Architecture

We evaluate the performance of AGS over two representative NAPL-AGCN-based ASTGNN architectures: **AGCRN** [3] and its extension **AGFormer**. AGCRN is a state-of-the-art ASTGNN architecture combining AGCN and RNN layers. AGCN layers are used to capture the spatial dependencies, while RNN layers are there to model the temporal dependencies. AGFormer, on the other hand, can be regarded as an alternative version of the AGCRN, in which the RNN layers are substituted by Transformer layers. We intentionally chose these two ASTGNN architectures sharing the same spatial module but using different temporal modules, to show that both the effectiveness of AGS and our observations on the learned spatial dependencies are orthogonal to the temporal modules involved.

5.2 Datasets and Configurations

The localisation of ASTGNNs is evaluated on nine real-world spatial-temporal datasets from three application domains: transportation, blockchain and biosurveillance. Table 1 summarizes the specifications of the datasets used in our experiments. The detailed datasets and configurations are provided in Appendix A.1. The details on tuning hyperparameters are provided in Appendix A.2.

5.3 Main Experimental Results

Our major experiments are illustrated in Figure. 2. We first train AGCRNs and AGFormers on the nine spatial-temporal datasets,

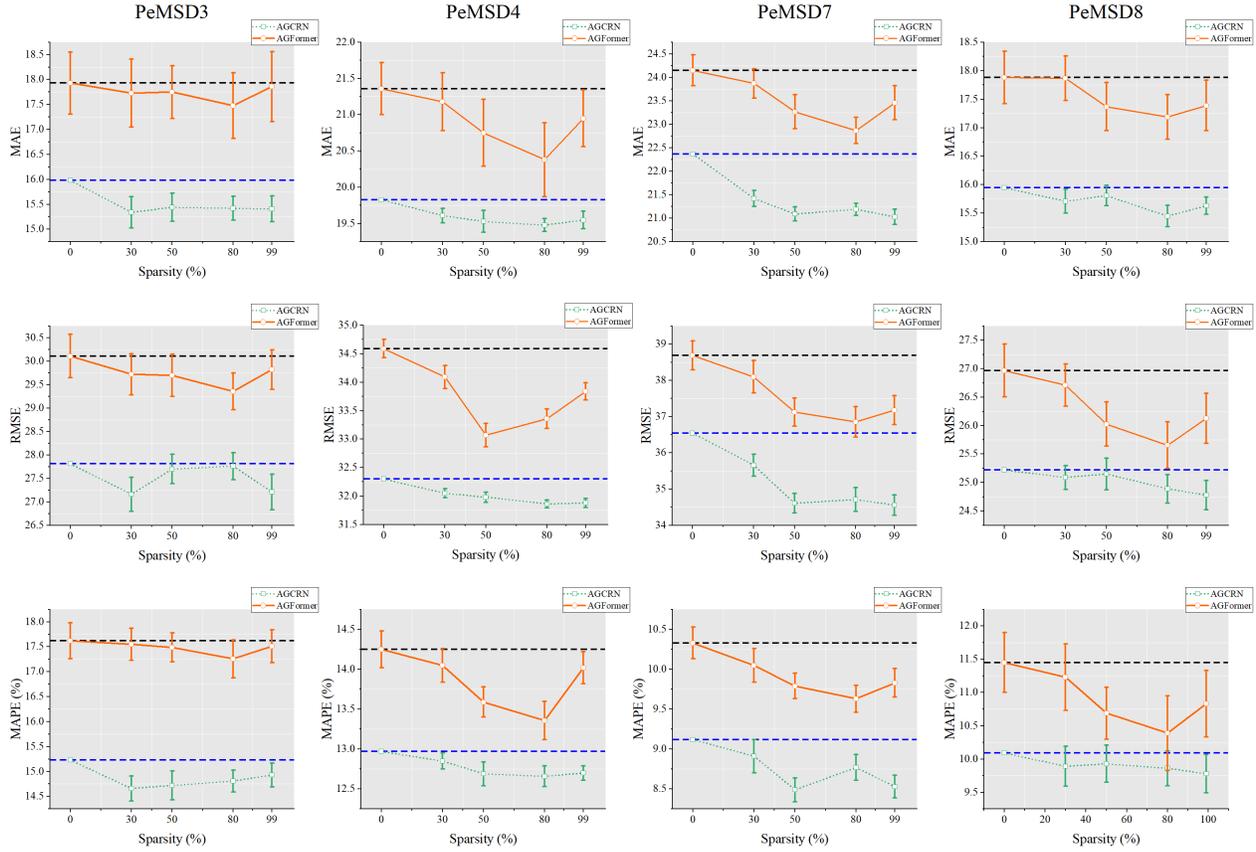


Figure 3: Test accuracies of original and localised (up to 99%) AGCRNs and AGFormers, tested on transportation datasets (PeMSD3, PeMSD4, PeMSD7 and PeMSD8). Horizontal dash lines represent the baselines of non-localised AGCRNs and AGFormers.

achieving baseline accuracies. Then we conduct localisation of these well-trained AGCRNs and AGFormers using AGS. Finally, we reinitialise all weights in the localised AGCRNs and retrain them on the original datasets with the same training settings.

The experiment results are organised as follows:

- The test accuracies of the non-localised AGCRNs and AGFormers and those with localisation degrees until 99% are collected in Figure. 3 (on transportation datasets), Figure. 4 (on biosurveillance datasets) and Figure. 5 (blockchain datasets). These figures contain the test accuracies at graph sparsity of 0%, 30%, 50%, 80% and 99%.
- The test accuracies of AGCRNs with localisation degrees between 99.1% and 100% are illustrated in Figure. 6, shown in dotted green curves.
- The test accuracies of localised AGCRNs that are reinitialised and retrained are also collected in Figure. 6, shown in solid purple curves.

Error bars in these figures show the standard deviation of five runs. We made the following observations from these results:

- **The localisation of AGCRNs and AGFormers is possible.** Applying AGS on AGCRNs and AGFormers and localising them to a localisation degree of 99% incurs no performance degradation across all datasets. On the contrary, in many experiments, the test accuracy keeps improving until 99%-localisation. Further localisation of AGCRNs up to 99.5% still induces no accuracy drop against the non-localised baselines.
- **Full localisation of AGCRNs is still practical.** Even when we fully localise the AGCRNs, which in effect turn them into independent RNNs ignoring all spatial dependencies, they can still provide decent test accuracies. As shown in Figure. 6, on transportation datasets (PeMSD3, PeMSD4, PeMSD7 and PeMSD8), only minor drops are observed. On blockchain datasets (Bytom, Decentraland and Golem) and biosurveillance datasets (CA&TX), we can observe that the test accuracy is no worse at 100% sparsity compared with the non-localised baselines.
- **Localised AGCRNs cannot be relearned without the dense spatial graphs.** As shown in Figure. 6, when we reinitialise the partially or fully localised AGCRNs and then

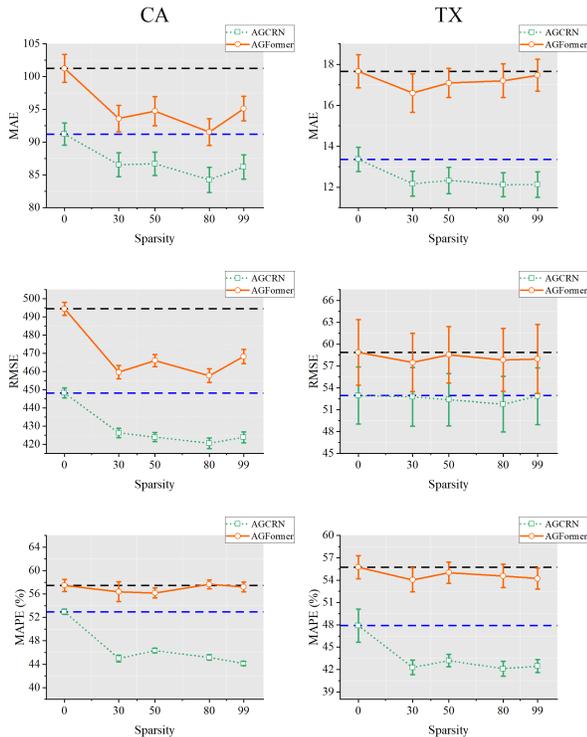


Figure 4: Test accuracies of original and localised (up to 99%) AGCRNs and AGFormers, tested on biosurveillance datasets (CA and TX). Horizontal dash lines represent the baselines of non-localised AGCRNs and AGFormers.

retrain them on the nine datasets, we can observe a consistent and considerable drop in inference accuracy.

Upon these observations, we suggest the following hypotheses:

- In many spatial-temporal datasets, the information provided by the spatial dependencies is primarily included in the information provided by the temporal dependencies. Therefore, the spatial dependencies can be safely ignored for inference without a noteworthy loss of accuracy.
- Although the information contained in the spatial and temporal dependencies overlaps, such overlapping provides the vital redundancy necessary for properly training a spatial-temporal graph model. Thus, the spatial dependencies cannot be ignored during training.

6 ABLATION STUDIES

6.1 Impact on Resource Efficiency

As mentioned in Sec. 1, one of the reasons that we are particularly interested in the localisation of ASTGNNs is its resource efficiency. Non-localised ASTGNNs usually learn spatial graphs that are complete. Therefore, the number of edges and consequently the computation overhead grows quadratically with the number of vertices.

Localisation of ASTGNNs is equivalent to pruning edges in the learned spatial graph, which could dramatically reduce the computation overhead associated with the spatial dependencies and thus improve resource efficiency. To this end, we calculated the amount of computation during inference of 99%-localised AGCRNs and AGFormers, measured in FLOPs, and summarised the results in Table 2. We can see that the localisation of AGCRNs and AGFormers effectively reduces the amount of computation required for inference. The acceleration is more prominent on AGCRNs against AGFormers because a larger portion of the total computation required by AGFormers is used on their temporal module (transformer layers), whereas AGCRNs use much lighter temporal modules (RNN layers).

6.2 Localised AGCRNs vs. Other Non-Localised ASTGNNs

In Figure. 3, Figure. 4 and Figure. 5, we can clearly see that the localisation up to 99% is able to improve the test accuracy slightly. For example, 99%-localised AGCRNs outperform non-localised AGCRNs by decreasing the RMSE/MAE/MAPE by 3.6%/3.7%/2.0% on PeMSD3. Such improvement is consistently observed across both AGCRNs and AGFormers and all tested datasets. We reckon that this improvement is caused by the regularisation effect of sparsifying the spatial graph, which may suggest that the non-localised AGCRNs and AGFormers all suffer from overfitting to a certain degree.

Recent works on ASTGNNs proposed improved architectures of AGCRN, including Z-GCNETs [7], STG-NCDE [8], and TAMP-S2GCNets [6] for different applications. We are therefore curious about how our localised AGCRNs compare to these variations. Hence we compare the test accuracy of 99%-localised AGCRNs with these architectures. Results are shown in Table 3, Table 4 and Table 5. We can see that our localised AGCRNs can generally provide competitive inference performance even against those delivered by state-of-the-art architectures. This observation also agrees with our first hypothesis mentioned in Sec. 5.3: in many spatial-temporal datasets, the information provided by the spatial dependencies are primarily included in the information provided by the temporal dependencies. Therefore, different spatial modules, given that the temporal modules are properly trained, may not make a significant difference in inference performance.

6.3 Localisation of Non-temporal Graphs

To further investigate spatial dependencies and indirectly test our hypotheses, we conduct additional experiments and extend AGS to non-temporal graphs. We attempt to sparsify the spatial graph pre-given by non-temporal datasets, including Cora, CiteSeer, and Pubmed [20]. On these datasets, we train two non-temporal graph neural network architectures, GCN[14] and GAT[23]. On GCN and GAT, as they don't own node embeddings E , we can use the representation H learned by pre-trained GCN and GAT to replace E in (6), then prune the edges as in Alg. 1, where the weighting factor λ takes control of graph sparsity.

Table 6 shows the accuracy of localised GCN and GAT non-temporal datasets. The pre-defined graphs of Cora, CiteSeer, and PubMed are sparsified to 30%, 50%, 80% and 100%, respectively. We can observe a significant drop in the test accuracy among all localised non-temporal graph models. This indicates that, in the

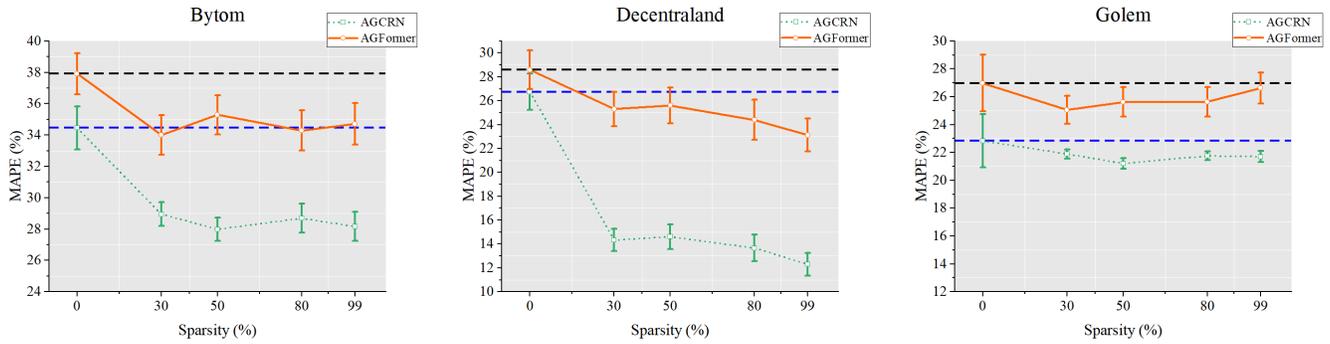


Figure 5: Test accuracies of original and localised (up to 99%) AGCRNs and AGFormers, tested on blockchain datasets (Bytom, Decentraland and Golem). Horizontal dash lines represent the baselines of non-localised AGCRNs and AGFormers.

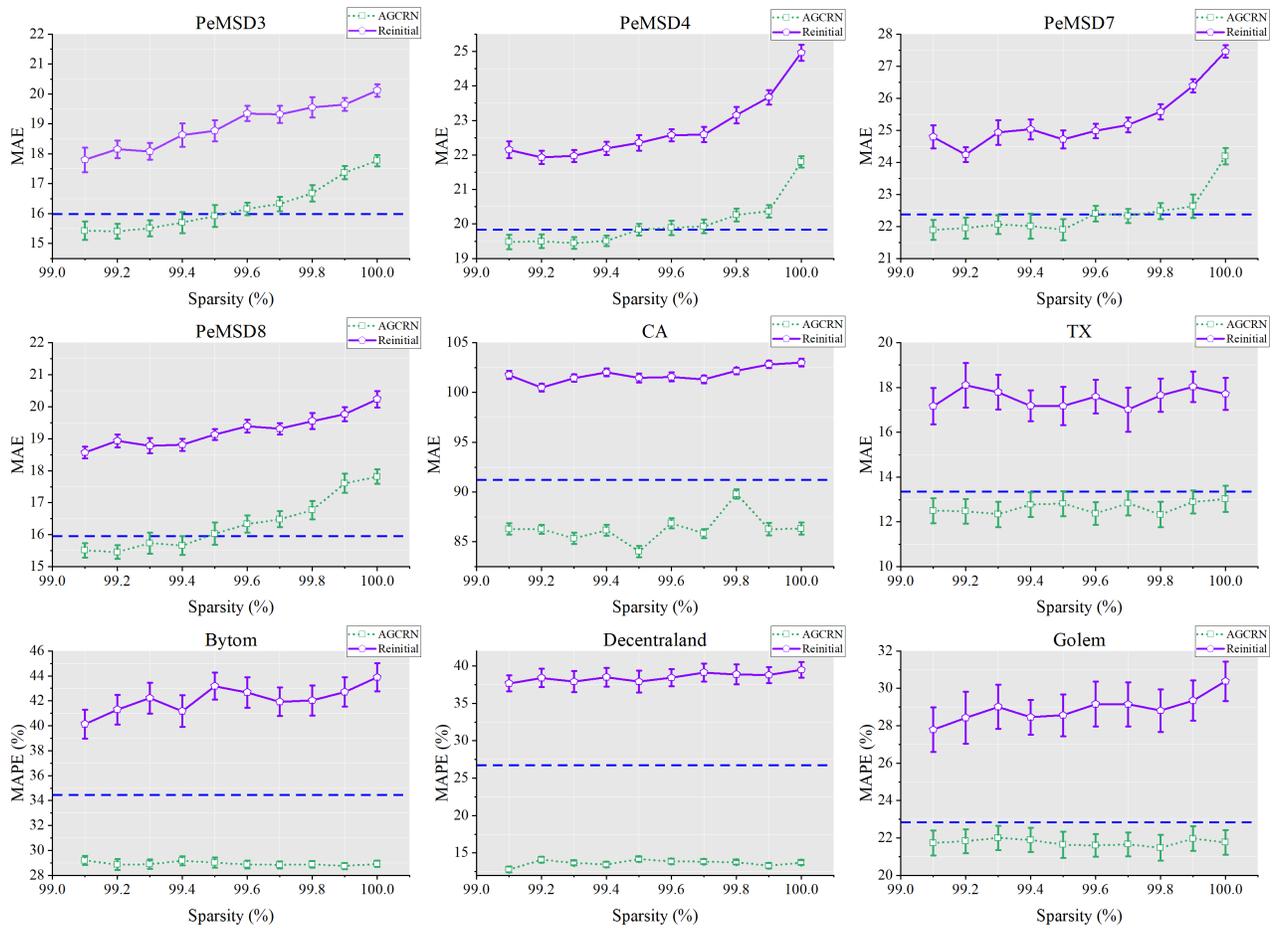


Figure 6: Test accuracies of localised (99.1% to 100%) AGCRNs and their reinitialised&retrained counterparts, tested on all datasets. Horizontal dash lines represent the baselines of non-localised AGCRNs.

absence of temporal dependencies, the information provided by spatial dependencies plays a major role during inference and thus can not be ignored via localisation.

7 CONCLUSION

In this paper, we ask the following question: whether and to what extent can we localise spatial-temporal graph models? To facilitate

Table 2: Computation cost during inference on original and 99%-localised AGCRNs and AGformers. The amount of computation is measured in MFLOPs, and acceleration factors are calculated in the round brackets.

Methods	Computation Cost for Inference (MFLOPs)								
	PeMSD3	PeMSD4	PeMSD7	PeMSD8	Decentraland	Bytom	Golem	CA	TX
Original AGCRN	400.26	188.59	1131.41	153.06	8.58	8.58	8.58	161.42	850.29
Original AGFormer	122.01	99.45	453.89	47.39	15.02	15.02	15.02	21.50	266.97
Localised AGCRN	253.33(\uparrow 1.6 \times)	80.55(\uparrow 2.3 \times)	237.56(\uparrow 4.8 \times)	119.93(\uparrow 1.3 \times)	2.82(\uparrow 3.0 \times)	2.82(\uparrow 3.0 \times)	2.82(\uparrow 3.0 \times)	145.50(\uparrow 1.1 \times)	706.21(\uparrow 1.2 \times)
Localised AGFormer	80.13(\uparrow 1.5 \times)	68.64(\uparrow 1.4 \times)	199.17(\uparrow 2.3 \times)	37.95(\uparrow 1.3 \times)	11.75(\uparrow 1.3 \times)	1.75(\uparrow 1.3 \times)	11.75(\uparrow 1.3 \times)	19.56(\uparrow 1.1 \times)	226.43(\uparrow 1.1 \times)

Table 3: Performance of 99%-localised AGCRNs compared with other non-localised ASTGNN architectures on transportation datasets.

Methods	Datasets	PeMSD3			PeMSD4			PeMSD7			PeMSD8			Average		
		MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE	MAE	RMSE	MAPE
AGCRN		15.98	28.25	15.23%	19.83	32.30	12.97%	22.37	36.55	9.12%	15.95	25.22	10.09%	18.53	30.58	11.85%
Z-GCNETs		16.64	28.15	16.39%	19.50	31.61	12.78%	21.77	35.17	9.25%	15.76	25.11	10.01%	18.42	30.01	12.11%
STG-NCDE		15.57	27.09	15.06%	19.21	31.09	12.76%	20.53	33.84	8.80%	15.45	24.81	9.92%	17.69	29.21	11.64%
TAMP-S2GCNets		16.03	28.28	15.37%	19.58	31.64	13.22%	22.16	36.24	9.20%	16.17	25.75	10.18%	18.49	30.48	11.99%
Localised AGCRN		15.41	27.21	14.93%	19.55	31.88	12.70%	21.03	34.56	8.53%	15.63	24.78	9.78%	17.91	29.61	11.49%

Table 4: Performance of 99%-localised AGCRNs compared with other non-localised ASTGNN architectures on biosurveillance datasets.

Methods	Datasets	CA			TX		
		MAE	RMSE	MAPE	MAE	RMSE	MAPE
AGCRN		91.23 \pm 1.69	448.27 \pm 2.78	53.67 \pm 0.42	13.36 \pm 0.58	52.96 \pm 3.92	47.89% \pm 2.22
TAMP-S2GCNets		76.53\pm0.87	371.60\pm2.68	92.90 \pm 1.57	11.29\pm1.05	48.21\pm3.17	52.34% \pm 3.87
Localised AGCRN		86.22 \pm 1.43	423.91 \pm 3.01	44.12\pm0.35	12.31 \pm 0.59	52.88 \pm 3.73	42.47%\pm1.41

Table 5: Performance of 99%-localised AGCRNs compared with other non-localised ASTGNN architectures on blockchain datasets.

Model	MAPE in %		
	Bytom	Decentraland	Golem
AGCRN	34.46 \pm 1.37	26.75 \pm 1.51	22.83 \pm 1.91
Z-GCNETs	31.04 \pm 0.78	23.81 \pm 2.43	22.32 \pm 1.42
STG-NCDE	29.65 \pm 0.63	24.13 \pm 1.07	22.24 \pm 1.53
TAMP-S2GCNets	29.26 \pm 1.06	19.89 \pm 1.49	20.10\pm2.30
Localised AGCRN	28.17\pm0.93	13.33\pm0.31	21.71 \pm 0.49

Table 6: Classification accuracy (%) of localised GCN and GAT on citation graph datasets.

Sparsity(%)	Cora		Citeseer		PubMed	
	GCN	GAT	GCN	GAT	GCN	GAT
0%	80.20	82.10	69.40	72.52	78.90	79.00
30%	80.35	83.17	69.23	72.31	79.14	79.23
50%	72.73	75.40	69.37	72.70	78.82	79.31
80%	65.19	70.81	58.47	63.18	68.37	77.03
100%	56.22	63.29	53.13	57.50	61.02	64.25

our investigation, we propose AGS, a novel algorithm dedicated to the sparsification of adjacency matrices in ASTGNNs. We use AGS to localise two ASTGNN architectures: AGCRN and AGFormer, and conduct extensive experiments on nine different spatial-temporal

datasets. Primary experiment results showed that The spatial adjacency matrices could be sparsified to over 99.5% without deterioration in test accuracy on all datasets. Furthermore, when the ASTGNNs are fully localised, we still observe no accuracy drop on the majority of the tested datasets, while only minor accuracy deterioration happened on the rest datasets. Based on these observations, we suggest two hypotheses regarding spatial and temporal dependencies: (i) in the tested data, the information provided by the spatial dependencies is primarily included in the information provided by the temporal dependencies and, thus, can be essentially ignored for inference; and (ii) although the spatial dependencies provide redundant information, it is vital for effective training of ASTGNNs and thus cannot be ignored during training. Last but not least, we conduct additional ablation studies to show the practical impact of ASTGNN’s localisation on resource efficiency and to verify our hypotheses from different angles further.

8 ACKNOWLEDGEMENT

The authors are grateful to the KDD anonymous reviewers for many insightful suggestions and engaging discussion which improved the quality of the manuscript.

REFERENCES

- [1] Gowtham Atluri, Anuj Karpatne, and Vipin Kumar. 2018. Spatio-temporal data mining: A survey of problems and methods. *Comput. Surveys* 51, 4 (2018), 1–41.
- [2] Lei Bai, Lina Yao, Salil S. Kanhere, Xianzhi Wang, and Quan Z. Sheng. 2019. STG2Seq: Spatial-Temporal Graph to Sequence Model for Multi-step Passenger Demand Forecasting. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, Sarit Kraus (Ed.). ijcai.org, 1981–1987.
- [3] Lei Bai, Lina Yao, Can Li, Xianzhi Wang, and Can Wang. 2020. Adaptive graph convolutional recurrent network for traffic forecasting. In *Advances in Neural Information Processing Systems*. 17804–17815.
- [4] Chao Chen, Karl Petty, Alexander Skabardonis, Pravin Varaiya, and Zhanfeng Jia. 2022. Freeway Performance Measurement System: Mining Loop Detector Data. *Transportation Research Record* 1748, 1 (2001), 96–102. <https://doi.org/10.3141/1748-12>
- [5] Tianlong Chen, Yongduo Sui, Xuxi Chen, Aston Zhang, and Zhangyang Wang. 2021. A Unified Lottery Ticket Hypothesis for Graph Neural Networks. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 1695–1706.
- [6] Yuzhou Chen, Ignacio Segovia-Dominguez, Baris Coskunuzer, and Yulia R. Gel. 2022. TAMP-S2GCNets: Coupling Time-Aware Multipersistance Knowledge Representation with Spatio-Supra Graph Convolutional Networks for Time-Series Forecasting. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*.
- [7] Yuzhou Chen, Ignacio Segovia-Dominguez, and Yulia R. Gel. 2021. Z-GCNets: Time Zigzags at Graph Convolutional Networks for Time Series Forecasting. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event (Proceedings of Machine Learning Research, Vol. 139)*, Marina Meila and Tong Zhang (Eds.). PMLR, 1684–1694.
- [8] Jeongwhan Choi, Hwangyong Choi, Jeehyun Hwang, and Noseong Park. 2022. Graph Neural Controlled Differential Equations for Traffic Forecasting. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*. AAAI Press, 6367–6374.
- [9] Andriy Mnih Chris J. Maddison and Yee Whye Teh. 2018. Learning Sparse Neural Networks through L_0 Regularization. In *The Tenth International Conference on Learning Representations, ICLR 2018*.
- [10] Wenyang Duan, Xiaoxi He, Lu Zhou, Lothar Thiele, and Hong Rao. 2022. Combating Distribution Shift for Accurate Time Series Forecasting via Hypernetworks. In *28th IEEE International Conference on Parallel and Distributed Systems, ICPADS 2022, Nanjing, China, January 10-12, 2023*. IEEE, 900–907.
- [11] Shengnan Guo, Youfang Lin, Ning Feng, and Chao Song. 2019. Attention Based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting. In *Thirty-third AAAI Conference on Artificial Intelligence, AAAI AAAI Press, 922–929*.
- [12] Rongzhou Huang, Chuyin Huang, Yubao Liu, Genan Dai, and Weiyang Kong. 2020. LSGCN: Long Short-Term Traffic Prediction with Graph Convolutional Networks. In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020*, Christian Bessiere (Ed.). ijcai.org, 2355–2361.
- [13] Diederik Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. *International Conference on Learning Representations* (12 2015).
- [14] Thomas N. Kipf and Max Welling. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- [15] Xiangyuan Kong, Jian Zhang, Xiang Wei, Weiwei Xing, and Wei Lu. 2022. Adaptive spatial-temporal graph attention networks for traffic flow forecasting. *Applied Intelligence* (2022), 1–17.
- [16] Jiayu Li, Tianyun Zhang, Hao Tian, Shengmin Jin, Makan Fardad, and Reza Zafarani. 2020. SGCN: A Graph Sparsifier Based on Graph Convolutional Networks. In *Advances in Knowledge Discovery and Data Mining - 24th Pacific-Asia Conference, PAKDD 2020, Singapore, May 11-14, 2020, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 12084)*, Hady W. Lauw, Raymond Chi-Wing Wong, Alexandros Ntoulas, Ee-Peng Lim, See-Kiong Ng, and Sinno Jialin Pan (Eds.). Springer, 275–287. https://doi.org/10.1007/978-3-030-47426-3_22
- [17] Shiyang Li, Xiaoyong Jin, Yao Xuan, Xiyu Zhou, Wenhui Chen, Yu-Xiang Wang, and Xifeng Yan. 2019. Enhancing the locality and breaking the memory bottleneck of transformer on time series forecasting. In *Advances in Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, 5243–5253.
- [18] Yitao Li, Umar Isambekov, Cuneyt Gurcan Akcora, Ekaterina Smirnova, Yulia R. Gel, and Murat Kantarcioglu. 2020. Dissecting Ethereum Blockchain Analytics: What We Learn from Topology and Geometry of the Ethereum Graph?. In *Proceedings of the 2020 SIAM International Conference on Data Mining, SDM 2020, Cincinnati, Ohio, USA, May 7-9, 2020*, Carlotta Demeniconi and Nitesh V. Chawla (Eds.). SIAM, 523–531. <https://doi.org/10.1137/1.9781611976236.59>
- [19] Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. 2018. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In *Proceedings of International Conference on Learning Representations*.
- [20] Prithviraj Sen, Galileo Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. 2008. Collective Classification in Network Data. *AI Mag.* 29, 3 (2008), 93–106.
- [21] Youngjoo Seo, Michaël Defferrard, Pierre Vandergheynst, and Xavier Bresson. 2018. Structured Sequence Modeling with Graph Convolutional Recurrent Networks. In *Neural Information Processing - 25th International Conference, ICONIP 2018, Siem Reap, Cambodia, December 13-16, 2018, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 11301)*, Long Cheng, Andrew Chi-Sing Leung, and Seiichi Ozawa (Eds.). Springer, 362–373. https://doi.org/10.1007/978-3-030-04167-0_33
- [22] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*. Curran Associates Inc., Red Hook, NY, USA, 6000–6010.
- [23] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- [24] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE Transactions on Neural Networks and Learning Systems* 32, 1 (2020), 4–24.
- [25] Zonghan Wu, Shirui Pan, Guodong Long, Jing Jiang, and Chengqi Zhang. 2019. Graph Wavenet for Deep Spatial-Temporal Graph Modeling. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (Macao, China) (IJCAI'19)*. AAAI Press, 1907–1913.
- [26] Sijie Yan, Yuanjun Xiong, and Dahua Lin. 2018. Spatial Temporal Graph Convolutional Networks for Skeleton-Based Action Recognition. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*, Sheila A. McIlraith and Kilian Q. Weinberger (Eds.). AAAI Press, 7444–7452.
- [27] Haoran You, Zhihan Lu, Zijian Zhou, Yonggan Fu, and Yingyan Lin. 2022. Early-Bird GCNs: Graph-Network Co-optimization towards More Efficient GCN Training and Inference via Drawing Early-Bird Lottery Tickets. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022, Thirty-Fourth Conference on Innovative Applications of Artificial Intelligence, IAAI 2022, The Twelveth Symposium on Educational Advances in Artificial Intelligence, EAAI 2022 Virtual Event, February 22 - March 1, 2022*. AAAI Press, 8910–8918. <https://ojs.aaai.org/index.php/AAAI/article/view/20873>
- [28] Bing Yu, Haoteng Yin, and Zhanxing Zhu. 2018. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden, Jérôme Lang (Ed.)*. ijcai.org, 3634–3640.
- [29] Cheng Zheng, Bo Zong, Wei Cheng, Dongjin Song, Jingchao Ni, Wenchao Yu, Haifeng Chen, and Wei Wang. 2020. Robust Graph Representation Learning via Neural Sparsification. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event (Proceedings of Machine Learning Research, Vol. 119)*. PMLR, 11458–11468.

A APPENDIX

A.1 Additional Dataset Details

The detailed datasets and configurations are as follows:

- **Transportation:** We use four widely-studied traffic forecasting datasets from Caltrans Performance Measure System (PeMS): **PeMSD3**, **PeMSD4**, **PeMSD7** and **PeMSD8** [4]. Following [3], PeMSD3, PeMSD4, PeMSD7, and PeMSD8 are split with a ratio of 6:2:2 for training, validation and testing. The traffic flows are aggregated into 5-minute intervals. We conduct 12-sequence-to-12-sequence forecasting, the standard setting in this domain. The accuracy is measured in Mean Absolute Error (MAE), Root Mean Square Error (RMSE) and Mean Absolute Percentage Error (MAPE).
- **Blockchain:** We use three Ethereum price prediction datasets: **Bytom**, **Decentral** and **Golem**[18]. These data are represented in graphs, with nodes and edges being the addresses of users and digital transactions, respectively. The

Table 7: Dataset-specific hyperparameter setup for AGCRN and AGFormer.

Task type	Datasets	hidden dimension		embedding dimension		batch size		learning rate	
		AGCRN	AGFormer	AGCRN	AGFormer	AGCRN	AGFormer	AGCRN	AGFormer
Traffic	PeMSD3	64	32	5	2	64	128	3e-3	5e-4
Traffic	PeMSD4	64	32	2	5	64	128	3e-3	5e-4
Traffic	PeMSD7	64	32	2	2	64	128	3e-3	5e-4
Traffic	PeMSD8	64	32	5	5	64	128	3e-3	5e-4
Ethereum	Decentraland	32	32	1	2	8	32	1e-2	1e-3
Ethereum	Bytom	32	32	1	2	8	32	1e-2	1e-3
Ethereum	Golem	32	32	1	2	8	32	1e-2	1e-3
COVID-19	CA	128	64	5	2	16	64	6e-3	2e-3
COVID-19	TX	128	64	5	5	16	64	6e-3	2e-3

interval between two consecutive time points is one day. Following [6], Bytom, Decentraland, and Golem token networks are split with a ratio of 8:2 for training and testing. We use 7 days of historical data to predict a future of 7 days. MAE, RMSE, and MAPE are also used as accuracy metrics.

- **Biosurveillance:** We use California (CA) and Texas (TX), COVID-19 biosurveillance datasets[6] used to forecast the number of hospitalized patients. The time interval of these datasets is one day. Following [6], CA and TX are split with a ratio of 8:2 for training and testing. Three days of historical data are used to predict a future of fifteen days. With these two datasets, MAE and RMSE are not used for measuring the inference accuracy, as their values are extremely small and, thus, difficult to reflect the performances realistically. Following [6], we only use MAPE.

A.2 Experiment Details

The dataset-specific hyperparameters chosen for AGCRNs and AGFormers are summarised in Table 7. All experiments are implemented in Python with Pytorch 1.8.2 and executed on a server with one NVIDIA RTX3090 GPU. We optimize all the models using the Adam optimizer[13]. For transportation data, we set the number of pre-training iterations $N_1 = 150$, and the number of sparsification iterations $N_2 = 400$. We use an early stop strategy with a patience of 15 for pre-training. Parameters are chosen through a parameter-tuning process on the validation set. For blockchain data, we set the number of pre-training iterations $N_1 = 100$, and the number of sparsification iterations $N_2 = 200$. We use an early stop strategy with a patience of 30 for pre-training. Parameters are chosen through a parameter-tuning process on the training set. For biosurveillance data, we set the number of pre-training iterations $N_1 = 100$, and the number of sparsification iterations $N_2 = 200$. We use an early stop strategy with a patience of 15 for pre-training. Parameters are chosen through a parameter-tuning process on the training set.

A.3 Notation

Frequently used notations are summarized in Table 8.

Table 8: The main symbols and definitions in this paper.

Notation	Definition
\mathcal{G}	the spatial network
N	the number of nodes
\mathcal{X}	a sequence of discrete frames
\mathcal{T}	length of historical observations
\mathcal{H}	length of future observations
X_t	node feature matrix at timestep t
E	the learnable node embedding
d	the node embedding dimension
C	the node feature dimension
F	the feature dimension
A_{adp}	the normalized adaptive adjacency matrix
M_A	mask to prune A_{adp}
\mathcal{B}	the Bernoulli distribution
$\mathcal{F}(\cdot)$	the ASTGNN model
N_1	number of pre-training iterations
N_2	number of sparsification iterations