TU Delft
Danae Savvidi
Email: d.n.savvidi@student.tudelft.nl
Supervisor: Elena Congeduti

Scalability of Graph Neural Networks in Traffic Forecasting

Assessing Accuracy and Computational Efficiency in Varying Road Network Sizes and Complexities

1. Introduction

Traffic Forecasting: predict future traffic conditions based on historical data, structure of the network and more.

Scalability: effectively handle increasing amounts of data without significant loss of performance or high increases in computational resources.

Graph Neural Networks (GNNs) in traffic forecasting: represent sensors as nodes, and edges as connections between them.

GNN Scalability problems: GPU memory constraints, reliability issues in subsampling [1].

Figure 1: Sensors in PEMS-BAY dataset.

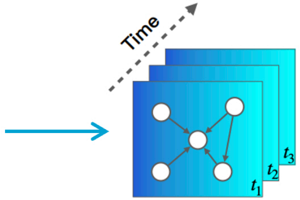

Figure 2: Graph-structured traffic data [2].

2. Problem Description

Given a set of historical traffic speed observations from the past T_h time steps $\mathcal{X} = [X_{t-T_h+1}, \ldots, X_{t-1}, X_t] \in \mathbb{R}^{T_h \times N}$ with $X_t \in \mathbb{R}^N$ at timestep t over N sensors in a traffic network G, predict the future traffic speed observations $\mathcal{Y} = [X_{t+1}, X_{t+2}, \ldots, X_{t+T_f}]$. Adapted from [2].

3. Research Question

How does the accuracy and computational efficiency of Graph Neural Networks in traffic forecasting vary with the size and complexity of road networks?

References

[1] W. Jiang and J. Luo, "Graph neural network for traffic forecasting: A survey," Expert Systems with Applications, vol. 207, p. 117921, Nov. 2022. [Online]. Available: http://dx.doi.org/10.1016/j.eswa.2022.117921
[2] Z. Shao, Z. Zhang, W. Wei, F. Wang, Y. Xu, X. Cao, and C. S. Jensen, "Decoupled dynamic spatial-temporal graph neural network for traffic forecasting," arXiv preprint arXiv:2206.09112, 2022

4. Methodology

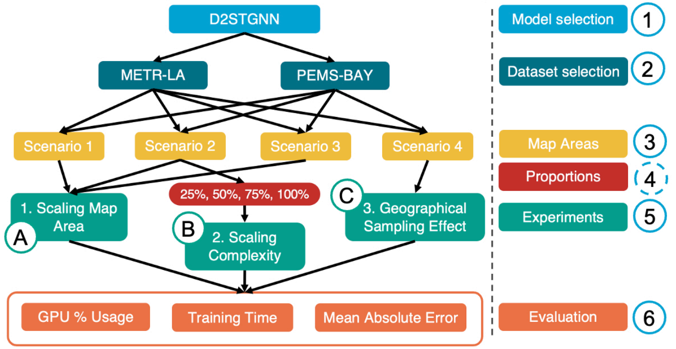

Figure 3: Overview of the methodology used in this study. The model is the Decoupled Dynamic Spatio-Temporal Graph Neural Network (D2STGNN) [2].

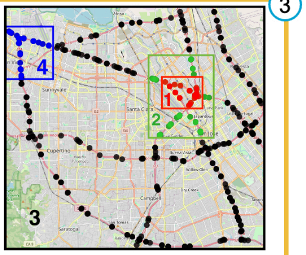

Figure 4: Areas selected for each scenario. The areas are non-mutually exclusive (PEMS-BAY).

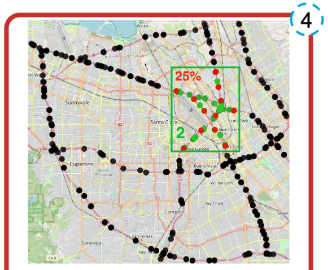

Figure 5: Example of sampling with 25% proportion in Scenario 2 (PEMS-BAY).

6

$$MAE = \frac{1}{T_f \cdot N} \sum_{t=1}^{T_f} \sum_{i=1}^{N} |\hat{\mathcal{Y}}_{ti} - \mathcal{Y}_{ti}| \qquad Training\ Time\ Per\ Node = \frac{Total\ Training\ Time}{N}$$

5. Results

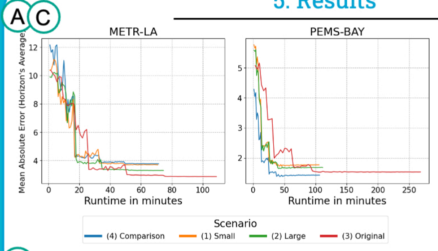

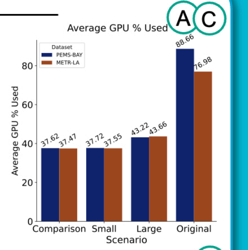

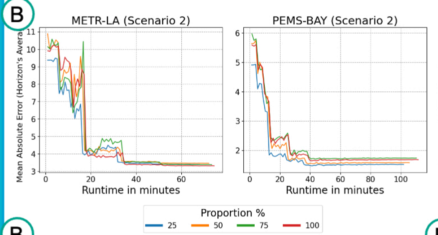

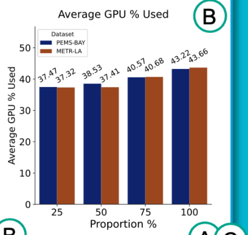

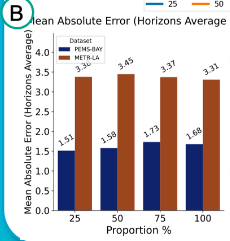

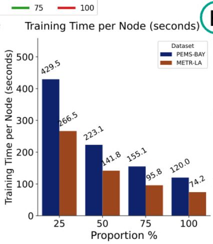

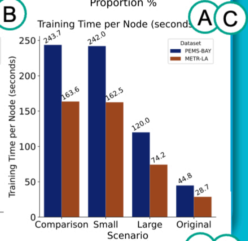

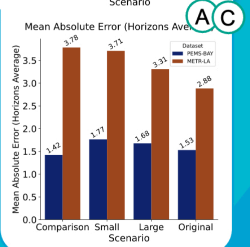

6. Limitations & Future Work

Limitations: (i) Geographic and dataset size limitations restrict the generalizability of the findings. (ii) Results may not fully generalize to other models.

Future Work: (i) Using larger more diverse datasets. (ii) Exploring other dimensions of scalability.

7. Conclusions

A Larger graphs result in shorter training times per node with higher GPU utilization and can generally improve accuracy.

B The model maintains robust accuracy with increased graph complexity, though GPU usage slightly rises.

C Sensor geographic location impacts accuracy but minimally affects computational resources.