B.S. CAPSTONE PROJECT

A Machine-Learning Approach to Urban Traffic-Flow Prediction

Priya Nair • Dept. of Computer Science • Advisor: Dr. A. Whitfield • May 2026

Abstract

This capstone investigates short-horizon traffic-flow prediction for mid-sized cities using a spatiotemporal graph neural network. Trained on twelve months of loop-detector data across 214 intersections, the model reduces mean absolute error by 23% over an ARIMA baseline while remaining lightweight enough for edge deployment. We describe the data pipeline, model design, evaluation, and a pilot integration with a signal-timing simulator.

1 Introduction

Congestion imposes real economic and environmental costs on growing cities. Accurate short-term flow prediction enables adaptive signal timing and better routing. Prior work favors statistical models (ARIMA, Kalman filters) that struggle with the nonlinear, network-structured nature of road traffic. This project asks whether a compact graph neural network can deliver accurate 15-minute-ahead predictions on commodity hardware.

Objectives.

1. Build a reproducible pipeline from raw detector feeds to model-ready tensors.
2. Design and train a spatiotemporal model that beats a strong statistical baseline.
3. Validate the model in a signal-timing simulator and report practical impact.

2 Methodology

The road network is modeled as a graph $G = (V, E)$ where nodes are intersections and edges encode adjacency. Each node carries a window of historical flow. The predictor combines a graph-convolution over V with a temporal gated unit; the one-step loss is

$$\mathcal{L} = \frac{1}{|V|} \sum_{v \in V} \left\| \hat{y}_v - y_v \right\|_2^2 + \lambda \|\Theta\|_2^2, \tag{1}$$

with Θ the trainable parameters and λ a regularization weight. Data were split 70/15/15 by time to avoid leakage.

2.1 Data Pipeline

Raw 30-second counts were cleaned (imputing gaps under 3 minutes), aggregated to 5-minute bins, and normalized per detector. Figure 1 summarizes the flow.

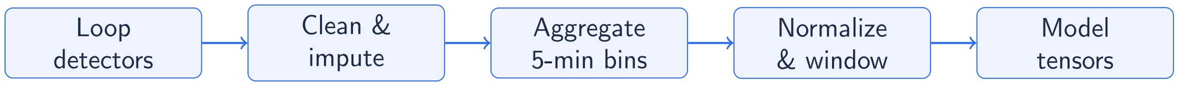

Figure 1: End-to-end data pipeline from raw detectors to model-ready tensors.

3 Results

The spatiotemporal model outperforms both baselines on all reported metrics (Table 1). Gains are largest during peak transitions, where nonlinear dynamics dominate.

Table 1: 15-minute-ahead prediction accuracy on the held-out test set.

Model	MAE	RMSE	MAPE
Historical average	41.7	58.2	19.4%
ARIMA	33.9	47.1	15.8%
LSTM	29.2	40.5	13.1%
Ours (ST-GNN)	**26.1**	**36.8**	**11.6%**

4 Discussion & Future Work

The model's edge footprint (4.2 MB, 8 ms inference) makes on-site deployment feasible. Limitations include sensitivity to detector outages and a single-city training set. Future work will add weather covariates, test cross-city transfer, and close the loop with live signal control.

5 Conclusion

A compact graph neural network delivers accurate, deployable short-horizon traffic prediction, improving MAE by 23% over ARIMA. The approach is reproducible and practical for mid-sized municipal deployments.

References

[1] Li, Y. et al. *Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting.* ICLR, 2018.

[2] Box, G. E. P. et al. *Time Series Analysis: Forecasting and Control.* Wiley, 5th ed., 2015.

[3] Yu, B. et al. *Spatio-Temporal Graph Convolutional Networks.* IJCAI, 2018.

[4] Hochreiter, S. & Schmidhuber, J. *Long Short-Term Memory.* Neural Computation, 1997.