

Deep Learning for Short-Term Traffic Flow Prediction in Smart Cities

Mahesh Goud¹, Swathi Yadav², Praveen Chary³, Harika Bandari⁴

ABSTRACT

Accurate short-term prediction of traffic flow is a cornerstone of intelligent transportation systems in smart cities, supporting signal control, route guidance, and congestion management. Traffic flow is nonlinear and depends on complex temporal patterns, which classical statistical models capture only partially. This paper applies deep learning, in the form of a long short-term memory network, to short-term traffic flow prediction using historical detector data. The network learns temporal dependencies from past flow observations to forecast flow over short horizons, and it was compared with historical-average, autoregressive, and support-vector-regression baselines. The deep learning model produced the most accurate forecasts, reducing the root-mean-square error to 21.4 and the mean absolute percentage error to 9.3 percent, clearly outperforming the baselines across the prediction horizons considered. The predicted flow tracked the observed diurnal pattern closely, including the morning and evening peaks. The results demonstrate that recurrent deep learning is an effective tool for short-term traffic prediction and a valuable component of smart-city traffic management.

KEYWORDS: *traffic prediction, deep learning, long short-term memory, smart cities, intelligent transportation, time-series forecasting*

INTRODUCTION

Smart cities depend on intelligent transportation systems to keep people and goods moving efficiently, and accurate prediction of traffic flow is central to these systems. Knowing the flow that is likely to occur in the next minutes or hour enables proactive signal control,

dynamic route guidance, and timely congestion warnings (1). The quality of these services depends directly on the quality of the underlying prediction (2).

Traffic flow is a nonlinear and noisy time series shaped by daily and weekly rhythms, weather, incidents, and human behaviour. Classical statistical models such as autoregressive methods capture linear temporal structure but struggle with the nonlinearity and long-range dependencies present in real traffic (3). Machine learning, and deep learning in particular, offers a way to model these complex patterns more faithfully.

This paper applies a long short-term memory network, a recurrent architecture designed to learn temporal dependencies, to short-term traffic flow prediction using historical detector data. The model is compared against statistical and machine-learning baselines using standard error metrics across several prediction horizons (4).

LITERATURE REVIEW

Traffic prediction research spans statistical time-series models, classical machine learning, and, increasingly, deep learning. Recurrent networks and their gated variants have proven well suited to sequential data, and many studies report that they outperform linear models on traffic forecasting (5).

DEEP LEARNING FOR TRAFFIC FORECASTING

Long short-term memory networks address the difficulty that simple recurrent networks have with long-range dependencies, making them effective for traffic series with strong daily structure. Studies have extended them with convolutional and graph components to capture spatial relationships between locations, but even the temporal model alone is reported to substantially improve accuracy over statistical baselines (6).

- Traffic flow exhibits nonlinear dynamics and strong daily and weekly periodicity.
- Linear statistical models capture only part of this structure.
- Long short-term memory networks learn long-range temporal dependencies effectively.
- Spatial extensions can further improve accuracy across networked locations.

Even so, many studies focus on a single horizon or omit comparison with several baseline families on the same data, leaving the practical gains of deep learning over simpler methods incompletely characterised (7).

RESEARCH GAP

Although deep learning for traffic prediction is well represented, comparisons that evaluate a recurrent model against statistical, machine-learning, and naive baselines on the same data and report multiple error metrics across horizons are less common. The degree of improvement that justifies the added model complexity is therefore not always clear (8).

This study addresses the gap by evaluating a long short-term memory network against historical-average, autoregressive, and support-vector-regression baselines on the same detector data, reporting several error metrics.

OBJECTIVES

- To prepare historical detector data for short-term traffic flow prediction.
- To develop a long short-term memory network for flow forecasting.
- To compare the model against statistical and machine-learning baselines.
- To evaluate prediction accuracy across multiple short-term horizons.

METHODOLOGY

Historical traffic data collected from loop detectors were used to train and evaluate the prediction models. The data were cleaned and formatted into input sequences, and the models were assessed on a held-out period.

1. Dataset and Preprocessing

The characteristics of the dataset are summarised in Table 1. Flow counts were aggregated at a fixed interval, and missing values arising from detector faults were imputed. The series was framed as a supervised learning problem using a sliding window of past observations to predict future flow, and the data were split chronologically into training and test periods.

Table 1: Traffic dataset description

Attribute	Description
Data source	Loop-detector traffic counts
Time span	6 months
Sampling interval	5 minutes
Features	Flow, speed, occupancy, time-of-day
Prediction horizons	15, 30, and 60 minutes

2. Model Architecture and Training

The model architecture and training settings are summarised in Table 2. A stacked long short-term memory network mapped the window of past observations to the predicted flow, and it was trained with the Adam optimizer to minimise mean-squared error. The same input framing was used to train the support-vector-regression baseline, while the autoregressive and historical-average baselines were fitted using their standard procedures.

Table 2: Model architecture and training configuration

Hyperparameter	Setting
Architecture	Stacked LSTM, 2 layers
Hidden units	128 per layer
Lookback window	12 steps (60 minutes)
Optimizer	Adam, learning rate 1×10^{-3}
Batch size	64
Training epochs	80

RESULTS AND FINDINGS

The deep learning model produced the most accurate forecasts across all horizons. Its predictions are compared with the observed flow and an autoregressive baseline over a representative day in Figure 4.

The error metrics for the deep learning model and the three baselines are summarised in Table 3.




Figure 1: Predicted and observed traffic flow over a representative day

As shown in Figure 1, the predictions of the long short-term memory network tracked the observed flow closely, capturing both the morning and evening peaks, whereas the autoregressive baseline lagged and underestimated the peaks. As reported in Table 3, the deep learning model reduced the root-mean-square error to 21.4 and the mean absolute percentage error to 9.3 percent, clearly outperforming the support-vector-regression, autoregressive, and historical-average baselines. The accuracy gain was largest around the rapidly changing peak periods, which matter most for congestion management.

Table 3: Prediction error of the evaluated models

Model	RMSE	MAE	MAPE (%)
Historical average	38.6	29.4	18.2

Model	RMSE	MAE	MAPE (%)
ARIMA	31.2	24.1	14.5
Support vector regression	27.8	21.0	12.6
LSTM (proposed)	21.4	16.2	9.3

DISCUSSION

The advantage of the long short-term memory network reflects its ability to learn the nonlinear, long-range temporal patterns that govern traffic, including the sharp transitions into and out of peak periods that linear models smooth over. By retaining relevant information across many time steps, the network anticipated changes that the autoregressive baseline could only follow with a lag, which is why its largest gains appeared around the peaks (9).

These improvements have direct operational value, because the periods that the model predicts best are exactly those in which proactive traffic management is most needed. The chief limitation of the present model is that it considers each location in isolation and does not exploit the spatial relationships between detectors, which carry additional predictive information. Incorporating spatial structure and external factors such as weather and incidents is the most promising route to further gains (10).

CONCLUSION

This paper applied a long short-term memory network to short-term traffic flow prediction and showed that it produced the most accurate forecasts, reducing the root-mean-square error to 21.4 and the mean absolute percentage error to 9.3 percent, clearly outperforming historical-average, autoregressive, and support-vector-regression baselines. The model tracked the daily flow pattern closely, including the critical peak periods, demonstrating that recurrent deep learning is an effective and valuable component of intelligent transportation systems in smart cities.

LIMITATIONS

- The model predicted flow at each location independently, ignoring spatial relationships between detectors.
- External factors such as weather and incidents were not included as inputs.
- Evaluation was limited to a single corridor over a six-month period.

FUTURE SCOPE

- Incorporating spatial structure through convolutional or graph-based components.
- Adding external inputs such as weather, events, and incident reports.
- Extending prediction to network-wide, multi-location forecasting.
- Integrating the predictions into a live adaptive signal-control system.

REFERENCES

1. Lv Y, Duan Y. Traffic flow prediction with big data: a deep learning approach. *IEEE Transactions on Intelligent Transportation Systems*. 2018; 16(2): 865-873.
2. Ma X, Tao Z. Long short-term memory networks for traffic speed prediction. *Transportation Research Part C*. 2018; 54(1): 187-197.
3. Zhang J, Zheng Y. Deep spatio-temporal residual networks for traffic prediction. *Artificial Intelligence*. 2019; 259(1): 147-166.
4. Tedjopurnomo D, Bao Z. A survey on deep learning for traffic prediction. *IEEE Transactions on Knowledge and Data Engineering*. 2021; 34(4): 1544-1561.
5. Fu R, Zhang Z. Using LSTM and GRU networks for traffic flow prediction. *Neurocomputing*. 2018; 318(1): 297-305.
6. Cui Z, Henrickson K. Traffic graph convolutional recurrent network. *IEEE Transactions on Intelligent Transportation Systems*. 2020; 21(11): 4883-4894.
7. Polson N, Sokolov V. Deep learning for short-term traffic flow prediction. *Transportation Research Part C*. 2019; 79(1): 1-17.
8. Vlahogianni E, Karlaftis M. Short-term traffic forecasting: where we are. *Transportation Research Part C*. 2018; 43(1): 3-19.

9. Wu Y, Tan H. A hybrid deep learning model for traffic flow forecasting. *Transportation Research Part C*. 2020; 90(1): 166-180.
10. Bogaerts T, Masegosa A. A graph CNN-LSTM approach for traffic forecasting. *Transportation Research Part C*. 2021; 112(1): 62-77.

***Author for Correspondence**

Mahesh Goud

E-mail: mahesh.goud88@rediffmail.com

^{1,4}Department of Computer Science and Engineering, MR Institute of Technology, Hyderabad, Telangana, India

²Department of Information Technology, Kakatiya College of Engineering, Warangal, Telangana, India

³Department of Electronics and Communication Engineering, SVU Institute, Hyderabad, Telangana, India

Received Date: 18 March 2026

Accepted Date: 14 April 2026

Published Date: 01 May 2026

Citation: Mahesh Goud, Swathi Yadav, Praveen Chary, Harika Bandari. Deep Learning for Short-Term Traffic Flow Prediction in Smart Cities. *International Journal of Smart Computing and Artificial Intelligence*. 2026; 2(1): 24-31p.