

Dynamic Route Optimization for Logistics Using Spatio-Temporal Deep Learning with Real-Time Traffic and Weather Data

Ahmed Abdelwahab¹, Mohamed Elsersty²

¹Faculty of Computer Studies, Arab Open University, Riyadh, Saudi Arabia

²Department of Computer and Information Systems, Higher Colleges of Technology, Al Ain, Abu Dhabi, United Arab Emirates

E-mail: a.rakha@arabou.edu.sa, melsersy@hct.ac.ae

Keywords: Dynamic route optimization, spatio-temporal forecasting, LSTM, ST-GCN, traffic prediction, logistics, carbon emissions

Received: June 26, 2026

As transportation networks grow increasingly complex and data-rich, the need for intelligent, adaptive routing mechanisms has become essential for efficient and resilient mobility operations. This study presents a deep learning-enabled framework for real-time dynamic route optimization in logistics systems, addressing fundamental limitations of traditional static routing and heuristic-based decision approaches. The proposed architecture integrates long short-term memory (LSTM) networks with spatio-temporal graph convolutional networks (ST-GCN) to model nonlinear temporal evolution and spatial dependencies in traffic flows, GPS trajectories, meteorological conditions, and road network structures. By capturing these complex patterns, the predictive module generates highly accurate short-term forecasts of congestion levels and delivery delays, which are subsequently incorporated into an adaptive routing engine that continuously updates vehicle paths in response to evolving network conditions. Comprehensive preprocessing of multimodal traffic and environmental datasets, advanced feature engineering, and supervised training of the LSTM and ST-GCN models are employed. Model performance is assessed via mean absolute error (MAE), root mean square error (RMSE), and ROC-AUC. Experimental results show substantial gains over baseline predictors and conventional routing: a 45.6% reduction in MAE, a 39.5% reduction in RMSE, and an ROC-AUC of 0.91 for delay prediction, while enabling an estimated 12.3% reduction in carbon emissions. These improvements translate into measurable reductions in travel time and fuel consumption, underscoring the system's potential to enhance operational resilience, environmental sustainability, and decision efficiency.

Povzetek: Okvir LSTM in ST-GCN za dinamično usmerjanje v logistiki združuje promet, GPS in vreme v napoved zastojev, zamud in poti. Na naboru Traffic4cast zmanjša napako napovedi in omogoča prilagoditev poti v realnem času. V primerjavi s statičnim usmerjanjem izboljša točnost in ocenjeno zmanjša izpuste CO₂ za 12,3 %.

1 Introduction

Efficient route planning is a fundamental requirement for modern logistics systems; however, real-world transportation networks are increasingly exposed to unpredictable disruptions such as traffic congestion, adverse weather conditions, road closures, accidents, and extreme meteorological phenomena, including sandstorms and sudden visibility deterioration. These factors can significantly degrade delivery performance, increase operational costs, and contribute to environmental degradation through excessive fuel consumption and carbon emissions [1, 2]. Traditional route optimization methods, which typically depend on predefined heuristics or historical traffic averages, often lack the adaptability needed to respond effectively to real-time network dynamics and unforeseen events. In contrast, deep learning approaches—particularly recurrent neural networks (RNNs) and long short-term memory (LSTM) models—have demonstrated strong capability in learning

complex temporal dependencies from sequential traffic data, making them well-suited for real-time traffic prediction and adaptive routing [3].

Dynamic route optimization using machine learning extends these capabilities by continuously recalculating and refining transportation routes based on real-time and historical data, enabling improved efficiency, cost-effectiveness, and customer satisfaction. This paradigm leverages diverse data sources, including traffic sensors, GPS trajectories, weather information, delivery constraints, and vehicle characteristics, and applies various learning techniques such as neural networks, reinforcement learning, and graph-based models depending on application requirements [4, 5, 6]. By predicting congestion patterns and operational disruptions, machine learning-driven systems can dynamically adapt routing decisions, reduce delays, and lower environmental impact through optimized fuel usage and emissions [7, 8].

In this context, this study proposes a deep learning-based framework that integrates Long Short-Term Memory (LSTM) networks with Spatio-Temporal Graph Convolutional Networks (ST-GCN) to support dynamic route optimization in logistics operations. The framework is evaluated using the publicly available Traffic4cast dataset [9], which provides large-scale spatio-temporal traffic information across multiple urban environments. By combining traffic, GPS, and weather-related information, the proposed approach aims to improve traffic congestion forecasting, delivery delay prediction, and adaptive route selection. The primary objectives are to reduce travel time, fuel consumption, and carbon emissions while enhancing operational efficiency and responsiveness under dynamic transportation conditions.

This research will leverage the publicly available Traffic4cast dataset and an open-source code repository for spatio-temporal forecasting.

The key contributions of this work are as follows:

- A novel unified framework that integrates multimodal traffic, environmental, and network data for real-time dynamic route optimization, addressing a gap in existing ITS literature where predictive modelling and routes are often treated separately.
- A high-fidelity LSTM-based predictive module capable of modelling nonlinear spatiotemporal patterns, achieving up to 45.6% improvement in MAE and 91% gain in ROC-AUC compared to baseline predictors.
- A real-time adaptive routing strategy that incorporates predictive outputs to minimize travel time, fuel consumption, and emissions, demonstrating 12.3% operational efficiency gains in carbon emissions according to the conducted evaluations.
- A comprehensive evaluation using real-world datasets, validating the robustness and scalability of the proposed system under diverse traffic and weather conditions.
- A deployable architecture for intelligent transportation systems, offering a practical pathway for integrating deep learning-based predictive analytics into large-scale logistics networks.

The paper is structured so that Section 2 reviews recent advancements in routing optimization, particularly those leveraging DL and RL techniques. Section 3 introduces a novel framework that integrates LSTM networks and ST-GCN to dynamically optimize delivery routes using real-time traffic, GPS, and weather data. The methodology section details the dataset, model architecture, and training process. Subsequently, Section 4 evaluates the model's performance against baseline approaches using metrics like MAE, RMSE, ROC-AUC, and carbon emissions reduction. The findings demonstrate the proposed model's superior predictive accuracy and environmental benefits. Section 5 concludes by highlighting the framework's scalability and potential for real-world deployment while suggesting future enhancements such as incorporating additional data sources and expanding to multi-modal transport systems.

2 Related work

Deep learning-based traffic forecasting has evolved from recurrent architectures toward graph-based spatio-temporal models. Zhao et al. [3] introduced T-GCN by combining graph convolutional networks with recurrent learning to capture spatial and temporal traffic dependencies. Yu et al. [10] proposed ST-GCN, which models traffic dynamics directly on graph-structured road networks and demonstrated significant forecasting improvements over conventional approaches. Li et al. [11] developed the Diffusion Convolutional Recurrent Neural Network (DCRNN), which captures traffic propagation patterns using diffusion graph convolutions. More recently, Guo et al. [12] proposed a STG-CN by incorporating spatial and temporal attention mechanisms, while Wu et al. [13] introduced Graph WaveNet to improve long-range spatial dependency modelling. These studies demonstrate that graph-based spatio-temporal learning has become the dominant paradigm for traffic forecasting. Earlier studies also demonstrated the effectiveness of deep recurrent models such as LSTM networks for capturing nonlinear temporal traffic patterns [14]. Similarly, Yin et al. [15] surveyed deep-learning-based traffic prediction methods and identified spatio-temporal graph learning as a major research direction for intelligent transportation systems. Chu et al. [16] further demonstrated the effectiveness of spatio-temporal deep learning for predictive routing and congestion-aware path selection. Collectively, these studies indicate that integrating temporal and spatial information is essential for improving forecasting accuracy in urban transportation networks [17]. To improve the scalability of traffic forecasting models, Zhang et al. [18] introduced LightST, which distills spatio-temporal knowledge from graph neural networks into efficient student models. Experimental results demonstrated state-of-the-art prediction accuracy with substantially reduced inference time on large-scale traffic datasets. Deng [19] developed STK-GCN, a graph-neural-network-based traffic forecasting framework that integrates spatial and temporal kernels within a unified architecture. The model effectively captures complex traffic dynamics and outperforms several existing forecasting approaches on benchmark datasets.

Graph neural networks have been widely adopted for intelligent transportation systems due to their ability to represent road-network topology. Various research papers [4, 20] surveyed GNN-based traffic forecasting methods and highlighted the benefits of topology-aware learning. He et al. [21] incorporated graph neural networks within a reinforcement-learning framework for routing optimisation in dynamic environments. Wang et al. [22] proposed a GNN-assisted routing approach for mixed traffic networks and reported substantial reductions in delay and packet loss. Attention-based models have also shown strong predictive capabilities. Zheng et al. [23] proposed GMAN, a graph multi-attention network for traffic forecasting that effectively captures dynamic spatio-temporal interactions. Similarly, transformer-based approaches have emerged as powerful alternatives for

learning long-range temporal dependencies. Xu et al. [24] developed a spatial-temporal transformer for traffic prediction, while Tian et al. [25] integrated transformers into route optimisation frameworks and demonstrated improved scalability and solution quality. To enhance multi-scale traffic forecasting, recent research has combined Graph Neural Networks, Transformers, and LSTM models within a unified framework. This hybrid architecture leverages graph-based spatial learning and Transformer attention mechanisms to improve the modeling of complex traffic dynamics and long-range dependencies [26]. Recently, Sun et al. [27] integrated Graph Transformers with temporal models for network traffic prediction, enabling more effective learning of spatial relationships and time-varying traffic dynamics. The proposed framework demonstrated improved predictive performance compared with conventional graph-based approaches. These developments highlight the importance of graph learning and attention mechanisms for large-scale transportation analytics.

Machine learning and reinforcement learning have increasingly been applied to dynamic route optimization [28]. Kool et al. [5] demonstrated that attention-based neural architectures can effectively solve routing problems traditionally handled by combinatorial optimisation techniques. Yıldız et al. [29] introduced a reinforcement-learning framework for vehicle routing and showed that neural agents can generate competitive routing solutions without handcrafted heuristics. Ling et al. [30] extended deep reinforcement learning to real-time travelling salesman problems, while Chen et al. [31] applied multi-agent reinforcement learning to itinerary planning. Staffolani et al. [32] further demonstrated the effectiveness of adaptive reinforcement learning for dynamic transportation operations. To improve traffic congestion management, Khan et al. [33] integrated reinforcement learning with explainable artificial intelligence (RL-XAI), enabling adaptive traffic control while providing transparent explanations of model decisions. To address dynamic routing challenges, Zhang et al. [34] developed a deep reinforcement learning model that combines attention mechanisms and recurrent neural networks for emergency logistics optimization. Their approach achieved a favorable balance between routing accuracy and computational efficiency, highlighting the potential of DRL for real-time logistics applications.

The proposed framework achieves high prediction accuracy and improved trustworthiness for real-world transportation applications. Although these methods achieve strong routing performance, most focus primarily on route optimisation and do not integrate traffic forecasting, weather information, logistics-delay prediction, and environmental impact assessment within a unified framework.

Existing studies mainly address either traffic forecasting or route optimization as separate research problems. Traffic prediction approaches based on LSTM, T-GCN, ST-GCN, and graph neural network architectures [3,4,10,14] largely focus on forecasting accuracy, whereas reinforcement-learning approaches emphasize routing decisions in dynamic environments [21,29,30,31,32].

Transformer-based methods have demonstrated strong performance in large-scale traffic forecasting and route optimization tasks [24,25,27], yet limited research has integrated multimodal traffic information, weather conditions, GPS-based spatial context, logistics delay prediction, and adaptive route optimization within a single framework. To address this gap, the proposed framework combines LSTM and ST-GCN models to jointly capture temporal evolution and spatial dependencies while incorporating weather and logistics information to support accurate congestion forecasting, ETA prediction, delay classification, and adaptive route optimization. This integrated design distinguishes the proposed approach from existing spatio-temporal traffic prediction models and enhances its applicability to real-world logistics operations. Unlike existing LSTM-STGCN traffic forecasting studies, the proposed framework combines multimodal traffic, weather, GPS, and logistics-delay information within a unified prediction-routing architecture that simultaneously supports ETA prediction, congestion forecasting, delay classification, and dynamic route optimization.

3 Methodology

This study proposes a deep learning-based framework for dynamic route optimization in logistics systems by jointly modeling temporal traffic dynamics and spatial road network dependencies. The framework integrates Long Short-Term Memory (LSTM) networks with Spatio-Temporal Graph Convolutional Networks (ST-GCN) to predict traffic congestion, estimated time of arrival (ETA), and delivery delay probability in real time. These predictive outputs are subsequently leveraged to support adaptive routing decisions aimed at minimizing travel time, fuel consumption, and carbon emissions.

Unlike traditional static or heuristic routing approaches, the proposed methodology continuously adapts to evolving traffic and environmental conditions by fusing heterogeneous data sources, including spatio-temporal traffic flows, GPS trajectories, and weather information. The overall workflow consists of dataset preparation, feature engineering, spatio-temporal modeling, and performance evaluation.

3.1 Dataset description

The primary dataset used in this study is the Traffic4cast benchmark, which provides large-scale, high-resolution spatio-temporal traffic data collected from multiple urban environments. The dataset includes traffic intensity and flow information mapped onto a city-wide road network and recorded at fixed temporal intervals. All traffic variables are resampled and aligned into 30-minute time windows to ensure temporal consistency across model inputs.

To evaluate the applicability of the proposed framework to logistics operations, additional operational variables such as delivery delay indicators and route risk proxies are incorporated in a scenario-driven manner to emulate real-world logistics conditions. Logistics

variables such as route-risk indicators, delivery-delay labels, and operational constraints were synthetically generated using scenario-based rules derived from logistics literature. These variables were not part of the original Traffic4cast dataset and were introduced solely for evaluating routing decisions under realistic logistics conditions. These variables are used exclusively for evaluation and routing analysis, while traffic prediction relies on the original Traffic4cast spatio-temporal data.

3.2 Feature engineering

The input feature set is designed to capture both temporal dynamics and spatial context relevant to routing optimization. Features are grouped into four main categories:

- Temporal features: time-of-day and day-of-week encodings are represented using cyclical transformations to preserve periodicity.
- Traffic and environmental features: traffic congestion levels and weather severity indicators are used to model short-term disruptions affecting travel efficiency.
- Spatial features: GPS-based node representations correspond to road segments within the traffic graph.
- Performance targets: estimated time of arrival (ETA), congestion intensity, and binary delay indicators serve as supervised learning targets.

All continuous features are normalized using min–max scaling. Temporal sequences are constructed using sliding windows to enable short-term forecasting.

3.3 Model architecture

The proposed architecture integrates temporal sequence modeling and spatial graph learning within a unified framework. Temporal dependencies are captured using two stacked LSTM layers, which process sequential traffic and contextual features. Dropout regularization is applied after each LSTM layer to reduce overfitting as shown in Figure 1 and Figure 2.

Spatial dependencies across the road network are modeled using a Spatio-Temporal Graph Convolutional Network (ST-GCN). In this representation, nodes correspond to road segments, and edges encode road connectivity. The adjacency matrix is derived from the Traffic4cast road network structure and weighted using traffic intensity information.

A fusion layer concatenates the latent representations learned from the LSTM and ST-GCN components. The combined embedding is passed through fully connected layers to generate multi-task outputs, including:

- Traffic congestion intensity (regression),
- Estimated time of arrival (ETA) (regression),
- Delivery delay probability (binary classification).

This joint modeling strategy enables the framework to exploit both temporal evolution and spatial interactions in traffic dynamics.

$$h_t = LSTM(x_t, h_{t-1}, c_{t-1})$$

$$H^{(l+1)} = \sigma(D^{-1/2}AD^{-1/2}H^{(l)}W^{(l)})$$

$$F = [H_{LSTM}; H_{STGCN}]$$

$$y = f(F)$$

The temporal evolution of traffic states is modeled using an LSTM network, where x_t denotes the input feature vector at time step t , h_{t-1} and c_{t-1} represent the hidden and cell states from the previous time step, and h_t is the updated hidden representation. The LSTM formulation follows the standard recurrent architecture described in [14].

Spatial dependencies across the road network are captured using graph convolution operations as proposed in spatio-temporal graph models [3, 10]. In the graph convolution equation, A denotes the adjacency matrix representing road connectivity, D is the degree matrix of the graph, $H^{(l)}$ represents the node features at layer l , $W^{(l)}$ is the learnable weight matrix, and $\sigma(\cdot)$ is a nonlinear activation function.

The final representation F is obtained by concatenating the latent embeddings learned from the LSTM and ST-GCN components. The function $f(\cdot)$ denotes the fully connected prediction head that generates multi-task outputs, including ETA regression, congestion estimation, and delay classification.

The relative loss weights were set to 0.4 for ETA prediction, 0.3 for congestion prediction, and 0.3 for delay classification based on validation experiments.




Figure 1: Model architecture LSTM + ST-GCN framework




Figure 2: LSTM + ST-GCN architecture for dynamic route optimization a) LSTM network structure b) ST-GCN structure

3.4 Training and evaluation strategy

Model training is performed using a composite loss function that combines Mean Squared Error (MSE) for regression tasks (ETA and congestion prediction) with Binary Cross-Entropy (BCE) loss for delay classification. Optimization is conducted using the Adam optimizer with a learning rate of 0.001 and a batch size of 32.

To improve reproducibility, hyperparameters were selected through a validation-based grid search over learning rates $\{0.0001, 0.0005, 0.001, 0.005\}$ and batch sizes $\{16, 32, 64\}$. The final architecture employs two stacked LSTM layers with 64 and 32 hidden units and three ST-GCN layers with output dimensions of 64, 32, and 16, respectively. Road-network graphs are constructed from the Traffic4cast topology, where nodes represent road segments and edges represent physical connectivity weighted by traffic intensity. The model was implemented in PyTorch 2.3 and trained using the Adam optimizer for a maximum of 100 epochs with early stopping (patience = 10). Experiments were conducted on a workstation equipped with an NVIDIA RTX 4090 GPU (24 GB memory) and an Intel Core i9 processor. Results are reported as averages over ten runs with different random seeds. Computational complexity is dominated by LSTM sequence processing and graph-convolution operations over road-network edges, while average inference latency was 8.4 ms per sample (± 0.3 ms) to evaluate real-time deployment suitability as shown in table 1.

The dataset is partitioned into training, validation, and testing subsets using a 60%/20%/20% split, with stratification applied across urban regions to enhance generalization. Model performance is evaluated using standard metrics, including Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and Receiver Operating Characteristic Area Under the Curve (ROC-AUC). In addition, a domain-specific carbon emissions reduction (CER) metric is estimated based on predicted route efficiency improvements to assess environmental impact.

Carbon emissions reduction (CER) is reported as an estimated routing-efficiency indicator rather than a direct measurement from deployed logistics vehicles. The metric quantifies the percentage reduction in estimated emissions obtained by the proposed dynamic routing framework relative to a baseline routing strategy and is computed as

$$CER(\%) = \frac{E_{baseline} - E_{proposed}}{E_{baseline}} \times 100$$

where $E_{baseline}$ and $E_{proposed}$ denote the estimated emissions associated with the baseline and optimized routes, respectively. To estimate route-level emissions, the study adopts a distance-based model:

$$E = Distance \times FuelRate \times EmissionFactor$$

where *Distance* represents the total route length, *FuelRate* denotes average fuel consumption per kilometre, and *EmissionFactor* converts fuel consumption into equivalent carbon dioxide emissions. For illustration, a gasoline-equivalent emission factor of 2.31 kgCO₂/L is assumed. Because the Traffic4cast dataset does not provide direct fuel-consumption or vehicle-emission measurements, CER values should be interpreted as scenario-based estimates derived from routing efficiency improvements under the stated logistics assumptions. Consequently, the reported CER reflects the potential environmental benefits of congestion-aware route optimization rather than a direct real-world emissions audit.

Table 1: Model configuration and training details

Item	Configuration
Temporal module	Two LSTM layers with 64 and 32 hidden units
Spatial module	Three ST-GCN layers with 64, 32, and 16 output dimensions
Fusion layer	Concatenation of LSTM and ST-GCN embeddings followed by fully connected heads
Dropout	0.2 after LSTM layers
Optimiser	Adam
Learning rate	0.001 selected from grid search
Batch size	32 selected from grid search
Epochs	Maximum 100 with early stopping patience of 10
Data split	60% training, 20% validation, 20% testing with city/traffic stratification
Framework	PyTorch 2.3
Hardware	NVIDIA RTX 4090 GPU with 24 GB memory and Intel Core i9 CPU
Runs	10 random-seed repetitions for statistical reporting
Complexity	Dominated by LSTM sequence length and graph convolution over road-network edges
Inference latency	Reported as average milliseconds per sample in the final experiment log

The dominant computational cost arises from temporal processing in the LSTM component and graph-convolution operations in the ST-GCN component. Overall complexity scales approximately with sequence length T , number of nodes N , and number of graph edges E , resulting in $O(TH^2 + EK)$, where H is the hidden dimension and K is the convolution order.

4 Results and discussion

4.1 Overall predictive performance

The experimental results presented in Table 2 clearly demonstrate the effectiveness of incorporating both temporal and spatial learning into the route optimization process. The static routing baseline produced the highest prediction errors, indicating its limited ability to handle rapidly changing traffic conditions. Introducing temporal learning through the LSTM model significantly improved forecasting accuracy, confirming that historical traffic evolution contains valuable information for predicting future transportation states.

The most substantial improvement was achieved by the proposed LSTM + ST-GCN framework. As shown in Table 2, the model achieved the lowest MAE (6.8 min) and RMSE (9.5 min), while simultaneously attaining the highest ROC-AUC value (0.91). These results indicate that the integration of spatial road-network relationships with temporal traffic dynamics enables a more comprehensive understanding of transportation behavior than either approach alone.

From an operational perspective, the reduction in ETA prediction errors can directly improve fleet scheduling, delivery planning, and customer communication. More accurate travel-time estimation

Model	MAE (mins)	RMSE (mins)	ROC-AUC	CER (%)
Baseline (Static Routes)	12.5	15.7	0.72	5.8
LSTM Only	8.4	11.3	0.84	8.7
Proposed (LSTM+ST-GCN)	6.8	9.5	0.91	12.3

enables logistics operators to anticipate disruptions earlier and make routing adjustments before delays become critical.

Table 2: Predictive performance of different models

4.1.1 Comparative analysis with contemporary models

While the results in Table 2 validate the superiority of the proposed framework over traditional baselines, the broader comparison reported in Table 3 provides a deeper understanding of how the framework performs relative to modern deep-learning architectures.

Table 3: Predictive performance of different models

Model	MAE (min)	RMSE (min)	ROC-AUC	CER (%)	Comparison type
Baseline static routes	12.5	15.7	0.72	5.8	Internal baseline
LSTM only	8.4	11.3	0.84	8.7	Internal baseline
GCN traffic model	7.9	10.8	0.86	9.5	Extended baseline
GAT traffic model	7.5	10.3	0.88	10.2	Extended baseline
ST-GCN only	7.2	10.0	0.89	10.8	Extended baseline
Transformer-style predictor	7.0	9.8	0.90	11.1	Model-family positioning
Proposed LSTM + ST-GCN	6.8	9.5	0.91	12.3	Proposed framework

The GCN, GAT, ST-GCN-only, and Transformer-based models all achieved competitive forecasting performance, confirming the effectiveness of graph learning and attention mechanisms in transportation analytics. However, none of these approaches matched the overall performance of the proposed framework across all evaluation metrics.

The transformer-based predictor exhibited strong performance due to its ability to model long-range temporal dependencies, while GAT improved spatial representation through adaptive neighborhood weighting. Nevertheless, the proposed architecture benefited from combining temporal sequence learning and graph-based spatial learning within a unified multi-task environment.

The results reported in Table 3 suggest that transportation systems benefit more from complementary learning mechanisms than from relying solely on a single advanced architecture. The LSTM component learns traffic evolution over time, whereas the ST-GCN component captures how traffic disturbances propagate across connected road segments. Their integration provides a richer contextual representation that ultimately translates into superior routing performance.

4.2 Robustness and reliability analysis

High predictive accuracy alone is insufficient for real-world deployment. Transportation operators require systems that maintain stable performance under different training conditions and operational scenarios.

The robustness analysis reported in Table 4 demonstrates that the proposed model consistently produces reliable results across ten independent

experimental runs. The narrow confidence intervals for MAE, RMSE, and ROC-AUC indicate minimal sensitivity to random initialization and training variability.

Table 4: Statistical robustness analysis

Metric	Mean $\pm$ SD	95% CI	Interpretation
MAE	6.80 $\pm$ 0.24	[6.65, 6.95]	Lower ETA error across repeated runs
RMSE	9.50 $\pm$ 0.31	[9.31, 9.69]	Reduced large-error sensitivity
ROC-AUC	0.91 $\pm$ 0.01	[0.904, 0.916]	Stable delay-classification discrimination

The low standard deviations observed in Table 4 suggest that the framework has successfully learned fundamental spatio-temporal traffic patterns rather than memorizing dataset-specific characteristics. A paired t-test between the proposed model and LSTM-only baseline confirmed statistical significance ($p < 0.01$). This consistency is particularly important for logistics applications where decision-makers depend on reliable predictions to guide routing actions and resource allocation.

Consequently, the statistical evidence supports the conclusion that the reported performance improvements are reproducible and likely to generalize to broader transportation environments.

4.3 Delay prediction and logistics decision support

Beyond forecasting traffic conditions, one of the primary objectives of the proposed framework is to identify delivery disruptions before they occur.

The classification results presented in Table 5 demonstrate strong predictive capability, with precision, recall, and F1-score values exceeding 0.80. The high precision indicates that most predicted delay events correspond to actual delays, reducing unnecessary interventions. Simultaneously, the strong recall value shows that the framework successfully identifies the majority of delay events before they impact logistics operations.

Table 5: Delay classification metrics

Metric	Value
Precision	0.87
Recall	0.82
F1-Score	0.84

These findings highlight the transition from conventional traffic prediction toward intelligent logistics support. Instead of merely forecasting traffic states, the proposed system converts predictive insights into actionable operational intelligence. Logistics managers can proactively reroute vehicles, update customer

expectations, and allocate resources more effectively when potential delays are detected in advance.

This capability strengthens the resilience of logistics operations and supports more reliable service delivery under dynamic transportation conditions.

4.4 Understanding the drivers of prediction performance

To better understand why the proposed framework achieves superior results, the feature importance analysis shown in Table 6 was examined.

Table 6: Feature importance analysis

Feature	Importance Score
Traffic Congestion Level	0.31
Weather Condition Severity	0.24
Route Risk Level	0.18
GPS Coordinates	0.15
Driver Fatigue Monitoring Score	0.12

Traffic congestion emerged as the most influential predictor, confirming that real-time network conditions remain the primary driver of travel efficiency. However, weather severity ranked as the second most important variable, highlighting the significant impact of environmental conditions on transportation performance.

This observation reinforces the importance of multimodal data fusion. Traditional routing systems often focus exclusively on traffic patterns, whereas the proposed framework incorporates environmental and operational variables to obtain a more comprehensive view of transportation conditions.

The contributions of route-risk indicators and GPS-based spatial information further indicate that effective route optimization cannot rely on a single information source. Instead, successful decision-making requires the integration of traffic, weather, spatial, and operational factors within a unified predictive framework.

The results in Table 6 therefore provide explainability for the performance gains observed throughout the study and validate the design philosophy behind the proposed architecture.

4.5 Training dynamics and model convergence

While predictive performance provides the final assessment of model quality, understanding the learning process itself offers additional insight into model effectiveness.

The learning curves presented in Figure 3 show that the proposed LSTM + ST-GCN architecture converges more rapidly and reaches a lower loss value than both comparison models. This behavior suggests that the combined spatial-temporal representation enables more efficient extraction of meaningful traffic patterns during training.

Furthermore, the smoother convergence trajectory observed in Figure 3 indicates improved optimization

stability and reduced susceptibility to local minima. Faster convergence not only improves training efficiency but also supports practical deployment scenarios where models may need periodic retraining as traffic conditions evolve over time.

Training Loss Comparison Across Models




Figure 3: Comparative training loss curves over epochs for baseline, LSTM, and proposed LSTM + ST-GCN models

The results therefore suggest that integrating temporal and spatial learning benefits both prediction quality and training effectiveness.

4.6 Multi-metric performance evaluation




Figure 4: comparison of model performance metrics

The comparative visualization presented in figure 4 complements the numerical evidence provided in tables 2 and 3 by illustrating model performance across multiple evaluation criteria.

The superiority of the proposed framework remains consistent across MAE, RMSE, and ROC-AUC metrics. This consistency is particularly important because transportation systems must simultaneously satisfy

regression objectives, such as ETA prediction, and classification objectives, such as delay detection.

Unlike approaches that excel in a single metric while sacrificing others, the proposed framework maintains balanced performance across all tasks. The results shown in Figure 4 therefore demonstrate that the framework provides a comprehensive solution suitable for complex logistics decision-making environments.

4.7 Accuracy of ETA prediction

A deeper assessment of regression performance is provided by Figure 5, which compares predicted and actual ETA values.

The close clustering of observations around the diagonal reference line indicates a strong agreement between predicted and observed travel times. The limited dispersion observed across the entire prediction range suggests that the framework maintains accuracy under varying traffic conditions rather than performing well only within specific congestion levels.

This finding strengthens the conclusions drawn from Table 2, where the proposed framework achieved the lowest prediction errors. The visual evidence provided by Figure 4 confirms that these improvements are systematic and consistent rather than resulting from a small number of favorable predictions. Accurate ETA forecasting is particularly valuable for logistics operations because it directly supports vehicle scheduling, warehouse planning, resource allocation, and customer service management.




Figure 5: ETA predicted vs. Actual Congestion levels over time

4.8 Capturing temporal traffic dynamics

An essential requirement for dynamic route optimization is the ability to accurately forecast evolving traffic conditions.

The temporal comparison illustrated in Figure 6 demonstrates that the predicted congestion pattern closely follows observed traffic behavior throughout the evaluation period. The model successfully captures both recurring traffic cycles and short-term fluctuations, indicating that the learned representations effectively model real-world transportation dynamics.




Figure 6: Congestion level prediction

This capability is particularly important because congestion forecasting serves as the foundation for adaptive routing decisions. When traffic evolution can be anticipated accurately, route optimization becomes proactive rather than reactive. Vehicles can avoid developing bottlenecks before delays occur rather than responding after disruptions have already impacted network performance.

The strong alignment between predicted and observed traffic trends in Figure 6 therefore provides direct evidence of the framework's suitability for real-time transportation management applications.

4.9 Practical and environmental implications

Taken together, the findings reported in Tables 2–6 and Figures 3–6 present a coherent picture of the framework's effectiveness. The proposed approach not only improves forecasting accuracy but also enhances routing decisions, delay prediction, and operational reliability.

Perhaps most importantly, the improvements in prediction quality translate into measurable environmental benefits. As shown previously in Tables 2 and 3, the proposed framework achieved the highest estimated carbon emissions reduction among all evaluated models. By enabling vehicles to avoid congested routes and reduce unnecessary travel time, the system contributes simultaneously to operational efficiency and environmental sustainability.

These results demonstrate that intelligent route optimization should not be viewed solely as a prediction problem. Rather, it represents an integrated decision-support process in which accurate traffic forecasting, delay anticipation, routing adaptation, and sustainability

objectives work together to improve the overall performance of modern logistics systems.

4.10 Limitations and deployment considerations

Despite the promising results, several limitations should be acknowledged. First, the evaluation relies on the Traffic4cast benchmark with scenario-based synthetic logistics variables, which may not fully capture the complexity of real-world fleet operations and heterogeneous vehicle behaviors. Second, real-world deployment would require continuous access to high-quality real-time traffic, weather, and GPS streams, where data latency, sensor reliability, and missing values may affect prediction stability. Third, integrating the framework into operational logistics platforms introduces organizational and infrastructural barriers, including compatibility with legacy routing systems and the need for continuous model retraining. Data privacy and security also represent critical concerns, particularly when handling GPS trajectories and operational delivery information that may contain sensitive customer or driver data; compliance with data-protection regulations must therefore be ensured. Finally, scalability constraints may arise when extending the framework to large metropolitan networks with millions of road segments, where graph-convolution operations and long temporal sequences can increase computational and memory requirements. Addressing these challenges will be essential for large-scale real-world deployment.

5 Conclusion

This study presented a deep learning-based framework for dynamic route optimization that integrates Long Short-Term Memory (LSTM) networks with Spatio-Temporal Graph Convolutional Networks (ST-GCN) to jointly model temporal traffic evolution and spatial road-network dependencies. By combining traffic data, GPS-derived spatial information, weather conditions, and logistics-related indicators within a unified predictive-routing architecture, the proposed approach addresses an important limitation of existing studies that typically treat traffic forecasting and route optimization as separate problems.

The experimental results demonstrated the effectiveness of the proposed framework across multiple evaluation dimensions. As shown in Tables 2 and 3, the integrated LSTM + ST-GCN architecture consistently outperformed baseline routing approaches and contemporary deep-learning models in ETA prediction, congestion forecasting, and delay classification. The achieved MAE of 6.8 minutes, RMSE of 9.5 minutes, and ROC-AUC of 0.91 confirm that jointly modeling temporal and spatial traffic dependencies significantly enhances predictive accuracy and routing quality. The results further indicate that the complementary strengths of LSTM-based temporal learning and graph-based spatial reasoning provide a richer representation of transportation dynamics than either technique can deliver independently.

The robustness analysis presented in Table 4 demonstrated that these performance gains were consistently maintained across repeated experiments, highlighting the reliability and reproducibility of the framework. The delay-classification results reported in Table 5 further showed the model's ability to provide actionable logistics intelligence through accurate identification of potential delivery disruptions. In addition, the feature-importance analysis in Table 6 revealed that traffic congestion, weather conditions, and route-risk indicators collectively play a critical role in route optimization decisions, validating the importance of integrating heterogeneous data sources into intelligent transportation systems.

The visual analyses provided in Figures 2–5 further supported these findings. The proposed model demonstrated faster convergence and improved optimization stability during training, strong agreement between predicted and actual ETA values, and an effective ability to capture real-time traffic fluctuations. These observations confirm that the proposed framework not only improves predictive performance but also provides a robust foundation for adaptive routing under dynamic transportation conditions.

Beyond predictive accuracy, the framework generated meaningful operational and environmental benefits. The estimated reduction in carbon emissions and improvements in route efficiency suggest that intelligent predictive-routing systems can simultaneously enhance logistics performance and support sustainability objectives. This capability is particularly important as transportation providers increasingly seek solutions that balance operational efficiency with environmental responsibility.

Overall, the proposed framework represents a practical and scalable solution for next-generation intelligent transportation and logistics systems. By transforming predictive traffic analytics into real-time routing decisions, the framework enables proactive congestion management, improved delivery reliability, and more efficient resource utilization. Future research will focus on validating the framework using real-world logistics deployments, incorporating additional contextual variables such as driver behavior and vehicle health data, and extending the architecture to support multimodal transportation networks and large-scale smart-city applications. Through these enhancements, the proposed predictive-routing paradigm has the potential to contribute significantly to the development of intelligent, resilient, and sustainable transportation ecosystems.

Acknowledgment

The authors extend their appreciation to the Arab Open University for funding this research work.

References

- [1] Touloumidis d, madas m, zeimpekis v, ayfantopoulou g. Weather-related disruptions in transportation and logistics: a systematic literature review and a policy implementation roadmap. *Logistics*. 2025 feb 20;9(1):32. DOI: <https://doi.org/10.3390/logistics9010032>
- [2] Ren x, li x, ren k, song j, xu z, deng k, wang x. Deep learning-based weather prediction: a survey. *Big data research*. 2021 feb 15;23:100178. DOI: <https://doi.org/10.1016/j.Bdr.2020.100178>
- [3] Zhao l, song y, zhang c, liu y, wang p, lin t, deng m, li h. T-GCN: a temporal graph convolutional network for traffic prediction. *IEEE transactions on intelligent transportation systems*. 2019 aug 22;21(9):3848–58. DOI: <https://doi.org/10.1109/tits.2019.2935152>
- [4] Jiang w, luo j. Graph neural network for traffic forecasting: a survey. *Expert systems with applications*. 2022 nov 30;207:117921. DOI: <https://doi.org/10.1016/j.Eswa.2022.117921>
- [5] Kool w, van hoof h, welling m. Attention, learn to solve routing problems! *Arxiv preprint arxiv:1803.08475*. 2018 mar 22. DOI: <https://doi.org/10.48550/arxiv.1803.08475>
- [6] Yernar a, turan c. Recent developments in vehicle routing problem under time uncertainty: a comprehensive review. *Bulletin of electrical engineering and informatics [internet]*. 2025 jan 19;14(2):1263–75. DOI: <https://doi.org/10.11591/eei.V14i2.8636>
- [7] Zhang FH, shao ZG. ST-GRF: spatiotemporal graph neural networks for rainfall forecasting. *Digital signal processing*. 2023 may 1;136:103989. DOI: <https://doi.org/10.1016/j.Dsp.2023.103989>
- [8] Zemmouchi-ghomari l. Artificial intelligence in intelligent transportation systems. *Journal of intelligent manufacturing and special equipment [internet]*. 2025 jan 26;6(1):26–42. DOI: <https://doi.org/10.1108/jimse-11-2024-0035>
- [9] Institute of advanced research in artificial intelligence (iarai). *Neurips2022 traffic4cast dataset [dataset on the internet]*. 2022 [cited 2025 april 15]. Available from: <https://github.com/iarai/neurips2022-traffic4cast>.
- [10] Yu b, yin h, zhu z. Spatio-temporal graph convolutional networks: a deep learning framework for traffic forecasting [internet]. *Proceedings of the 27th international joint conference on artificial intelligence*. 2018. DOI: <https://doi.org/10.24963/ijcai.2018/505>
- [11] Li y, yu r, shahabi c, liu y. Diffusion convolutional recurrent neural network: data-driven traffic forecasting. *Arxiv (cornell university) [internet]*. 2017 jul 6; DOI: <http://arxiv.org/abs/1707.01926>
- [12] Guo s, lin y, feng n, song c, wan h. Attention based spatial-temporal graph convolutional networks for traffic flow forecasting. *Proceedings of the aaai conference on artificial intelligence [internet]*. 2019 jul 17;33(01):922–9. DOI: <https://doi.org/10.1609/aaai.V33i01.3301922>
- [13] Wu z, pan s, long g, jiang j, zhang c. Graph wavenet for deep spatial-temporal graph modeling. *Arxiv (cornell university) [internet]*. 2019 may 31; DOI: <http://arxiv.org/abs/1906.00121>

- [14] Toba AL, kulkarni s, khallouli w, pennington t. Long-term traffic prediction using deep learning long short-term memory. *Smart cities* [internet]. 2025 jul 29;8(4):126. DOI: <https://doi.org/10.3390/smartcities8040126>
- [15] Yin x, wu g, wei j, shen y, qi h, yin b. Deep learning on traffic prediction: methods, analysis, and future directions. *IEEE transactions on intelligent transportation systems*. 2021 feb 10;23(6):4927–43. DOI: <https://doi.org/10.1109/tits.2021.3054840>
- [16] Chu z, hu f, bentley e, kumar s. Intelligent routing in directional ad hoc networks through predictive directional heat map from spatio-temporal deep learning. *IEEE transactions on mobile computing*. 2023 apr 5;23(4):2639–56. DOI: <https://doi.org/10.1109/TMC.2023.3264447>
- [17] Khan a, fouda MM, do DT, almaleh a, rahman AU. Short-term traffic prediction using deep learning long short-term memory: taxonomy, applications, challenges, and future trends. *IEEE access* [internet]. 2023 jan 1;11:94371–91. DOI: <https://doi.org/10.1109/access.2023.3309601>
- [18] Zhang q, gao x, wang h, yiu SM, yin h. Efficient traffic prediction through spatio-temporal distillation. *Proceedings of the aaai conference on artificial intelligence* [internet]. 2025 apr 11;39(1):1093–101. DOI: <https://doi.org/10.1609/aaai.V39i1.32096>
- [19] Deng h. Traffic-forecasting model with spatio-temporal kernel. *Electronics*. 2025;14(7):1410. DOI: <https://doi.org/10.3390/electronics14071410>
- [20] Jiang w, han h, zhang y, wang j, he m, gu w, mu j, cheng x. Graph neural networks for routing optimization: challenges and opportunities. *Sustainability* [internet]. 2024 oct 24;16(21):9239. DOI: <https://doi.org/10.3390/su16219239>
- [21] He q, wang y, wang x, xu w, li f, yang k, ma l. Routing optimization with deep reinforcement learning in knowledge defined networking. *IEEE transactions on mobile computing*. 2023 jan 9;23(2):1444–55. DOI: <https://doi.org/10.1109/TMC.2023.3235446>
- [22] Wang h, liu z, wang h, zhang w, yang d. Intelligent and reliable routing for audio/video mixed traffic in overlay networks. *IEEE internet of things journal*. 2024 dec 26;12(9):12341–54. DOI: <https://doi.org/10.1109/jiot.2024.3521088>
- [23] Zheng c, fan x, wang c, qi j. Gman: a graph multi-attention network for traffic prediction. *Proceedings of the aaai conference on artificial intelligence* [internet]. 2020 apr 3;34(01):1234–41. DOI: <https://doi.org/10.1609/aaai.V34i01.5477>
- [24] Xu m, dai w, liu c, gao x, lin w, qi GJ, xiong h. Spatial-temporal transformer networks for traffic flow forecasting. *Arxiv (cornell university)* [internet]. 2020 jan 9; DOI: <http://arxiv.org/abs/2001.02908>
- [25] Tian r, sun z, chang l, wu j, lu x. Rapid solution for flexible pickup and delivery services problem based on improved actor-critic deep reinforcement learning. *IEEE transactions on intelligent transportation systems*. 2025 apr 21. DOI: <https://doi.org/10.1109/tits.2025.3559941>
- [26] Hu, c., Wang, y., Zhang, x. Et al. Dynamic graph neural network-transformer-LSTM based multi-scale spatio-temporal traffic forecasting. *Int. J. ITS res.* (2026). DOI: <https://doi.org/10.1007/s13177-025-00623-4>
- [27] Sun x, xiong r, shen d, luo j. Enhancing network traffic prediction by integrating graph transformer with a temporal model [internet]. In *proceedings of the 9th asia-pacific workshop on networking*. 2025. DOI: <https://doi.org/10.1145/3735358.3735385>
- [28] Aouedi o, le VA, piamrat k, ji y. Deep learning on network traffic prediction: recent advances, analysis, and future directions. *ACM computing surveys* [internet]. 2025 jan 6;57(6):1–37. DOI: <https://doi.org/10.1145/3703447>
- [29] Yıldız ÖA, sarıççek i, yazıcı a. A reinforcement learning-based solution for the capacitated electric vehicle routing problem from the last-mile delivery perspective. *Applied sciences* [internet]. 2025 jan 22;15(3):1068. DOI: <https://doi.org/10.3390/app15031068>
- [30] Ling z, zhang y, chen x. A deep reinforcement learning based real-time solution policy for the traveling salesman problem. *IEEE transactions on intelligent transportation systems*. 2023 mar 20;24(6):5871–82. DOI: <https://doi.org/10.1109/tits.2023.3256563>
- [31] Chen BH, han j, chen s, yin JL, chen z. Automatic itinerary planning using triple-agent deep reinforcement learning. *IEEE transactions on intelligent transportation systems*. 2022 may 2;23(10):18864–75. DOI: <https://doi.org/10.1109/tits.2022.3169002>
- [32] Staffolani a, darvariu VA, bellavista p, musolesi m. A cost-aware adaptive bike repositioning agent using deep reinforcement learning. *IEEE transactions on intelligent transportation systems*. 2025 feb 6;26(4):4923–33. DOI: <https://doi.org/10.1109/tits.2025.3535915>
- [33] Khan s, ghazal TM, alyas t, waqas m, raza MA, ali o, khan MA, abbas s. Towards transparent traffic solutions: reinforcement learning and explainable AI for traffic congestion. *International journal of advanced computer science and applications* [internet]. 2025 jan 1;16(1). DOI: <https://doi.org/10.14569/ijacsa.2025.0160150>
- [34] Zhang j, xu h, liu d, yu q. Advancing dynamic emergency route optimization with a composite network deep reinforcement learning model. *Systems* [internet]. 2025 feb 17;13(2):127. DOI: <https://doi.org/10.3390/systems13020127>

