Predictive Urban Navigation Using Hybrid TCN–STGCN Traffic Forecasting and Reinforcement Learning

Ruturaj Jadhav, Veer Jain, Aakash Hingrajiya, Soham Gore, Mrs. Manali Patekar, Dr. Megha Trivedi
Department of Computer Engineering, Vidyavardhini's College of Engineering, Vasai, India

***Abstract*—**Urban navigation systems commonly rely on static or short-term traffic predictions, limiting their ability to handle dynamic congestion caused by multiple users selecting similar routes. This project presents a hybrid traffic routing framework that integrates deep learning–based traffic prediction with graph-based path planning and reinforcement learning into a unified system. A Temporal Convolutional Network (TCN) and Spatio-Temporal Graph Convolutional Network (ST-GCN) are employed to model temporal and spatial traffic patterns and predict future edge-level travel times. These predictions are incorporated into an OpenStreetMap-based road network, where edge weights represent predicted travel time, enabling efficient route computation using the A* algorithm under various constraints such as avoiding tolls or signals. To address the limitation of deterministic routing, a Proximal Policy Optimization (PPO) framework with parameter sharing is introduced to select among multiple candidate routes generated by A*, considering both predicted and realized congestion. The system is evaluated on a Mumbai subgraph and demonstrates improved traffic distribution and more balanced routing compared to static approaches.

***Index Terms*—**Urban Traffic Forecasting, Intelligent Transportation Systems, Context-Aware LSTM, Parking Prediction, Reinforcement Learning, Smart Mobility.

I. INTRODUCTION

Rapid urbanization and the continuous growth in vehicle density have significantly increased traffic congestion in metropolitan regions such as Mumbai. Daily commuting is often characterized by unpredictable travel times, fluctuating speeds, and inefficient route selection under dynamic conditions. While modern navigation systems provide real-time traffic updates and shortest-path routing, they remain fundamentally reactive — they estimate current traffic conditions but lack the ability to anticipate future congestion states and adapt routing decisions accordingly.

Conventional routing approaches primarily depend on graph-based algorithms such as Dijkstra and A*, which compute optimal paths based on static or predicted edge weights. Although effective in isolated settings, these methods fail to account for the collective behavior of multiple agents, where identical route selection leads to congestion and degrades overall system efficiency. Similarly, existing traffic prediction models are often decoupled from routing decisions and do not incorporate feedback from real-time user interactions.

Recent advancements in deep learning have enabled more accurate modeling of traffic dynamics through spatial-temporal architectures. Models such as Temporal Convolutional Networks (TCNs) capture long-range temporal dependencies with improved parallelism, while Spatio-Temporal Graph Convolutional Networks (ST-GCNs) effectively model the structural relationships within road networks. However, these models are primarily designed for prediction tasks and do not directly address the challenges of route optimization under dynamic and multi-agent conditions.

To address these limitations, this work proposes a hybrid traffic routing framework that integrates predictive modeling with adaptive decision-making. A TCN + ST-GCN model is employed to forecast future edge-level travel times based on historical traffic patterns. These predictions are incorporated into a road network graph, enabling A* search to generate multiple candidate routes under user-defined constraints. A Proximal Policy Optimization (PPO)-based reinforcement learning layer then selects optimal routes by considering both predicted and realized congestion, allowing the system to adapt to multi-agent interactions and distribute traffic more efficiently.

II. LITERATURE REVIEW

Recent advancements in Intelligent Transportation Systems (ITS) have significantly enhanced traffic forecasting through deep learning–based methodologies. Early approaches primarily relied on recurrent neural networks such as Long Short-Term Memory (LSTM) and Gated Recurrent Units (GRU) to model temporal dependencies in traffic data. These models demonstrated strong performance in short-term prediction tasks but struggled to capture complex spatial interactions in large-scale road networks.

To address these limitations, recent research has shifted toward convolutional and graph-based models. Temporal Convolutional Networks (TCNs) have been widely adopted for capturing long-range temporal dependencies, while Spatio-Temporal Graph

Convolutional Networks (ST-GCNs) effectively model spatial correlations between interconnected road segments. Notable works such as Diffusion Convolutional Recurrent Neural Networks (DCRNN) and Graph WaveNet further improve prediction accuracy by incorporating diffusion processes and adaptive graph learning mechanisms.

More recent studies have focused on enhancing graph-based traffic prediction. Ju et al. [1] proposed a Conjoint Spatio-Temporal Graph Neural Network that captures high-order spatial and temporal relationships, improving forecasting performance. Yin et al. [2] introduced a Dynamic Diffusion STGCN (DDSTGCN), which models dynamically changing traffic patterns more effectively than static graph approaches. Li et al. [3] proposed a Spatio-Temporal Tree Attention Network incorporating hierarchical attention mechanisms. Capone et al. [4] provided a comprehensive taxonomy of graph neural network–based traffic prediction methods and highlighted key limitations such as scalability and adaptability in dynamic environments.

In parallel, classical graph-based routing algorithms such as Dijkstra and A* remain widely used due to their computational efficiency and optimality guarantees. However, in real-world scenarios, edge weights are dynamic and influenced by evolving traffic conditions. Reinforcement Learning (RL) has emerged as a promising approach for adaptive routing in dynamic environments. Khairy et al. [5] proposed an adaptive traffic prediction framework combining Graph Neural Networks with reinforcement learning, demonstrating improved adaptability in dynamic scenarios.

Overall, existing literature demonstrates significant progress in traffic prediction and adaptive decision-making. However, a clear research gap remains in integrating predictive modeling, efficient routing, and multi-agent congestion management within a unified framework. This work addresses this gap by combining TCN and ST-GCN for traffic prediction, A* for efficient candidate route generation, and PPO-based reinforcement learning for adaptive route selection.

III. METHODOLOGY

A. Data Collection

Traffic data is collected for urban road networks using OpenStreetMap (OSM) and synthetic temporal augmentation. Each road segment (edge) is associated with attributes such as length, road type, number of lanes, and estimated free-flow speed. Temporal traffic data is generated at fixed 1-hour intervals, capturing variations in congestion patterns across time. The dataset includes:

- Edge-level structural attributes (length, highway type, lanes)
- Time-dependent features (hour, day-of-week)
- Synthetic congestion indicators (traffic signals, incidents, peak-hour patterns)

All data is stored in structured parquet files and mapped to a graph representation. Historical sequences are constructed per edge to support multi-step traffic prediction.

B. Contextual Feature Engineering

To model temporal variations in traffic patterns, contextual features are encoded alongside edge attributes. Feature engineering includes time-of-day encoding:

$$h_{sin}\theta = sin(2\pi \times hour / 24)$$
$$h^{co}{}_{s}\theta = cos(2\pi \times hour / 24)$$

Traffic indicators include congestion level and incident flags. The final input feature tensor is:

$$X \in \mathbb{R}\text{\^{}}\{N \times L \times F\}$$

where N = number of edges, L = sequence length (e.g., 24), and F = number of features.

C. Traffic Forecasting Using TCN + ST-GCN

A hybrid deep learning model is used to capture both temporal and spatial dependencies.

Temporal Modeling (TCN)

TCN applies causal dilated convolutions:

$$y_t = \Sigma_k\, w_k \cdot x(t - d \cdot k)$$

where d = dilation factor.

Spatial Modeling (ST-GCN)

Graph convolution is applied over the road network graph. The model predicts travel time ratio:

$$\hat{r}_e, t+h = f(X_e)$$

The final predicted travel time is:

$$T_e = T_e\text{\^{}}\{free\} \times \hat{r}_e$$

D. A* Routing

The road network is represented as a weighted graph G = (V, E, w). The cost function for A* is:

$$f(n) = g(n) + h(n)$$

where the heuristic is:

$$h(n) = d_{have}{}^{r}{}_{s}{}^{e}{}_{ne}(n, goal) / v_{ma}{}^{x}$$

Multiple candidate routes are generated: $\{R_0, R_1, \ldots, R_k\}$. Constraints such as no toll or no signal are applied by modifying edge weights.

E. PPO-Based Route Selection

Routing is modelled as a Markov Decision Process (MDP):

- State (S): S = {R_i, predicted time, congestion, t}
- Action (A): $A \in \{0, 1, \ldots, K\}$
- Reward (R): $R = -T_{a}{}^{c}{}_{t}{}^{u}{}_{a}\ell - \lambda \cdot$ congestion

Objective:

$$\max^{L} [\sum_{t=0}^{T} \gamma^{t} R_{t}]$$

F. System Integration

The system operates as a unified pipeline:

- Input is validated and mapped to graph edges
- Historical features are extracted per edge
- TCN + ST-GCN predicts future travel times
- Graph weights are updated dynamically
- A* generates multiple candidate routes
- PPO selects optimal route based on congestion

Unlike traditional systems, prediction and routing are tightly coupled, enabling adaptive multi-agent decision-making.

G. Algorithm — Hybrid Traffic Routing Pipeline

1. Input: source, destination, departure time
2. Validate coordinates within service area
3. Map source/destination to nearest edges
4. Extract historical features per edge
5. Predict travel time ratios using TCN + ST-GCN
6. Compute edge weights: $w_e = T_e^{free} \times \hat{r}_e$
7. Generate K candidate routes using A*
8. Construct state using route features
9. Select route using PPO policy
10. Output: optimal route, ETA, congestion map

IV. RESULTS & DISCUSSION

The proposed TCN + ST-GCN model demonstrates strong predictive performance for edge-level traffic forecasting, achieving an overall Mean Absolute Error (MAE) of 0.0209 and Root Mean Squared Error (RMSE) of 0.0361 for travel time ratio (TTR). For congestion prediction, the model achieves an MAE of 0.0855 and RMSE of 0.1058, indicating stable performance across both normalized travel time and congestion estimation tasks. A detailed horizon-wise evaluation shows that prediction accuracy is highest in short-term horizons (t+1 to t+3, MAE ≈ 0.012–0.015) and gradually degrades for longer horizons (t+12 MAE = 0.0340), reflecting increasing uncertainty in long-term forecasting.

Fig. 1. Horizon-wise MAE and RMSE trends

When compared with established models such as DCRNN and Graph WaveNet evaluated on the METR-LA dataset, the proposed model demonstrates competitive performance. Prior work reports scaled MAE values typically in the range of 0.024–0.030, whereas the proposed hybrid model achieves a lower MAE of 0.0209 on the Mumbai dataset. The results suggest that combining TCN for temporal modeling and ST-GCN for spatial representation effectively captures complex traffic dynamics.

The integration of the predictive model with A* routing enables the computation of optimal paths based on dynamically updated edge weights derived from predicted travel times. Multiple candidate routes are generated under user-defined constraints, allowing the system to explore alternative paths rather than relying on a single deterministic solution.

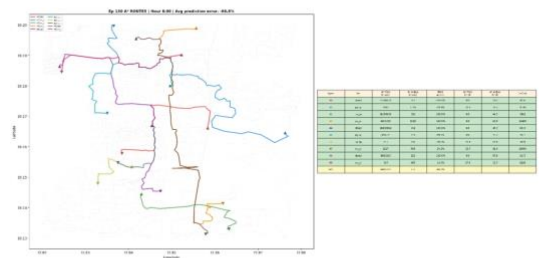

Fig. 2. Multiple routes forecasted using A for multiple agents*

To address the limitations of static routing, a PPO-based decision layer is introduced to select routes based on both predicted and realized congestion. The PPO agent learns to distribute traffic across multiple routes by avoiding over-utilized paths, thereby reducing congestion clustering and improving overall system efficiency. This results in more balanced traffic flow and improved Estimated Time of Arrival (ETA) reliability compared to traditional A*-only routing.

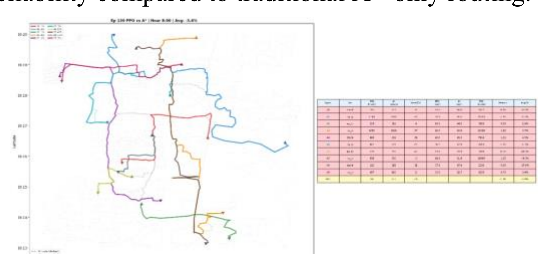

Fig. 3. Multiple routes forecasted using PPO for multiple agents

Overall, the results validate that the hybrid TCN + ST-GCN architecture provides accurate traffic predictions, while the integration of A* and PPO enables adaptive and congestion-aware routing. The combined system demonstrates improved traffic distribution and robustness in dynamic, multi-agent environments, highlighting its effectiveness for intelligent urban transportation systems.

V. CONCLUSION

This work presents a hybrid and scalable traffic routing framework that integrates deep learning–based traffic prediction with graph-based optimization and reinforcement learning for adaptive decision-making in dynamic urban environments. The proposed system leverages a combination of Temporal Convolutional Networks (TCN) and Spatio-Temporal Graph Convolutional Networks (ST-GCN) to model both temporal evolution and spatial dependencies of traffic patterns at the edge level. The model demonstrates strong predictive capability, achieving low error metrics and maintaining stable performance across multiple forecasting horizons.

The integration of predicted travel times into a graph-based routing framework enables the application of the A* algorithm for efficient path computation under various constraints. The system incorporates a Proximal Policy Optimization (PPO)–based reinforcement learning layer to address the limitations arising from multi-agent interactions. By learning from real-time congestion feedback, the PPO agent dynamically selects routes from multiple A*-generated candidates, effectively distributing traffic across the network and preventing congestion buildup caused by identical route selection.

The combined architecture transforms traffic routing from a static shortest-path problem into a dynamic, multi-agent optimization problem, leading to improved route efficiency, reduced congestion clustering, and more reliable Estimated Time of Arrival (ETA). Furthermore, the proposed framework is designed to be scalable and infrastructure-independent, relying on readily available map data and historical traffic patterns rather than dense sensor networks, making it particularly suitable for deployment in rapidly growing urban regions.

Future work can extend this framework by incorporating real-time streaming data, multi-modal transportation systems, and large-scale deployment across entire city networks, further advancing the capabilities of intelligent urban mobility systems.

REFERENCES

[1] W. Ju et al., "Conjoint Spatio-Temporal Graph Neural Network for Traffic Forecasting," 2024.

[2] X. Yin et al., "Dynamic Diffusion Spatio-Temporal Graph Convolutional Network (DDSTGCN)," Springer, 2025.

[3] H. Li et al., "Spatio-Temporal Tree Attention Network," ScienceDirect, 2025.

[4] V. Capone et al., "Spatio-Temporal Prediction using Graph Neural Networks: A Survey," 2025.

[5] M. Khairy et al., "Adaptive Traffic Prediction using Graph Neural Networks and Reinforcement Learning," 2025.