electronics

MDPI

Article

A Simulation Study of ST-GNN-Based Risk Warning and Active Visual Guidance for Low-Position Lighting on Bridge Segments

Fang Cao [1,†], Yaokun Ku [2,†], Xiaoxue Cao [2], Zhitao Chen [3], Gang Li [1], Wenyan Han [2], Bailiang Lu [2] and Zhi Chen [2,*]

[1] Guangzhou Yuexiu Group Co., Ltd., Guangzhou 510623, China
[2] School of Electrical and Control Engineering, North China University of Technology, Beijing 100144, China
[3] Yunnan Transportation Planning and Design Institute Co., Ltd., Kunming 650011, China
* Correspondence: cz86@ncut.edu.cn; Tel.: +86-13426350978
† These authors contributed equally to this work.

Abstract

Traffic risks on bridge segments exhibit spatiotemporal dependence under congestion, low visibility, and sudden traffic incidents. This study develops and evaluates a reproducible simulation framework that links short-term risk-proxy forecasting with graded low-position lighting control for a 1200 m linear bridge corridor. The framework uses six traffic and environmental scenario families, a direction-separated graph convolution and gated recurrent unit model, explicit baselines and graph ablations, and a rule-based upstream coordination controller. The model predicts a continuous conflict-exposure proxy for 20 control zones over a five-minute horizon. Across ten training seeds, the proposed DS-GCN-GRU achieved a mean RMSE of 0.1094 ± 0.0004, MAE of 0.0773 ± 0.0007, macro-F1 of 0.8339 ± 0.0082, Level 3 recall of 0.8777 ± 0.0095, and high-risk PR-AUC of 0.9292 ± 0.0010. Real-world safety benefits and driver responses to the proposed lighting commands have not been validated.

Keywords: low-position lighting; spatiotemporal graph neural network; risk warning; active visual guidance; bridge traffic safety

1. Introduction

Low-position lighting places luminaires on guardrails or along road edges to enhance road delineation and driving boundaries. Previous studies have examined its design and application on airport overpasses, interchange facilities, and suspension bridges in mountainous canyons [1–4]. However, continuous low-position lighting may also create excessive light–dark contrast and transient flicker. Chen [5] analyzed the influence of flicker from the perspectives of transient-light mechanisms and modulation, whereas An [6] and Duan et al. [7] emphasized that lighting quality, installation configuration, and drivers' visual workload should be considered jointly. The present study concerns how preset lighting patterns can be selected and coordinated from node-level forecasts, rather than establishing the prevalence or limitations of existing installations.

Bridge traffic conditions vary over time and are affected by adjacent road segments and the direction of traffic propagation, making graph-based modeling appropriate. The diffusion convolutional recurrent neural network (DCRNN), spatiotemporal graph convolutional network (STGCN), Graph WaveNet, and attention-based spatiotemporal graph convolutional network (ASTGCN) characterize spatiotemporal dependence in traffic networks through diffusion convolution, spatiotemporal convolution, adaptive adjacency, and attention mechanisms, respectively [8–11]. Related work has extended traffic flow forecasting to crash-risk prediction. Wang et al. [12] combined geographic and semantic

Check for updates

Academic Editors: Ahmed Abdulhakim Al-Absi, Mohammed Abdulhakim Al-Absi, Nadhem Ebrahim and Baseem Al-Athwari

Received: 16 September 2026
Revised: 30 October 2026
Accepted: 2 November 2026
Published: 4 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

correlations to forecast regional crash risk; Nippani et al. [13] evaluated graph neural networks for road safety modeling using large-scale road network data; and Gao et al. [14] further incorporated probabilistic distributions and uncertainty quantification. These studies demonstrate that graph neural networks provide a unified spatial representation for node-level risk prediction.

Beyond transportation, spatiotemporal graph neural networks (ST-GNNs) have been applied to structural damage identification and joint multi-parameter prediction [15,16], illustrating the use of graph structures to organize multi-source time-series signals. For the present application, a prediction model and a lighting controller address different questions: the former estimates a future risk proxy, whereas the latter converts that estimate into a discrete command sequence. The focus here is the interface between these stages, including the distinction between reference labels, predicted levels, and coordinated commands. A general claim that detection-to-warning systems are absent is not made.

Adaptive lighting and weather-responsive warning provide relevant precedents. Gibbons et al. [17] describe methods for selecting and adjusting roadway lighting levels according to the driving environment. An FHWA account of the Interstate 75 fog-warning project in Tennessee [18] documents an established detection-and-warning application. These sources place the present study within existing adaptive-control work; they do not establish the novelty or effectiveness of the proposed bridge-lighting implementation.

Driver interpretation must also be considered separately from software execution. Skinner et al. [19] investigated warning-light flash frequency as a communication cue in a service vehicle setting. Domenichini et al. [20] examined tunnel lighting and driver behavior using a driving simulator. These contexts differ from low-mounted bridge guardrail lighting. The studies motivate attention to signal comprehension and behavioral outcomes, but their results are not used as calibration coefficients for the lighting patterns considered here.

This study has three connected objectives. First, lighting-control zones and their observation areas are represented as graph nodes for short-horizon risk-proxy forecasting from traffic, visibility, and observed-event information. Second, predicted scores are converted into three command levels with explicit thresholds and a spatial coordination rule. Third, the complete pipeline is evaluated using repeat runs, non-graph baselines, graph ablations, forced Level 3 controller scenarios, and robustness checks within a controlled simulation setting.

2. Materials and Methods

2.1. Model Assumptions

The following assumptions define the scope of the case study:

(1) The bridge segment is divided at fixed spatial intervals into low-position-lighting control units, each corresponding to one guardrail light-strip node.
(2) Traffic states are updated at a fixed time step, and each node is associated with traffic, environmental, and risk features at every time step.
(3) Traffic, environmental, and observed-event variables are assumed to be available at the end of each aggregation interval and assignable to the control zones. This is an input-availability assumption, not a measured real-time acquisition capability.
(4) Local congestion, simulated traffic incidents, and reduced visibility are assumed to affect traffic conditions across adjacent control zones along the direction of travel.
(5) Low-position-lighting units and their detection zones are represented as graph nodes, while physical adjacency and directional upstream–downstream relationships are represented as graph edges.

(6) The ST-GNN is used to model spatial relationships and temporal variation and to output a short-term normalized risk-proxy score. Its predictive performance is evaluated separately from command-rule execution.
(7) Risk-classification thresholds are specified in advance for the case study and are not regarded as universal values applicable to all bridges.
(8) The guardrail light strips are assumed to support independent zonal control of color, luminance, flashing frequency, and dynamic light-flow direction.
(9) A possible effect of lighting patterns on driver perception and behavior motivates the application, but no driver-response model or behavioral benefit is validated by the reported command-level tests.

2.2. Model Constraints

To convert predictions into executable lighting commands, constraints are imposed on the normalized risk score, risk classification, node control state, and spatial coordination.

These constraints define variable domains and control logic; they do not calibrate the human-factor effectiveness of the lighting-signal parameters.

(1) Risk-calculation constraint

Bridge-node states evolve over time and are influenced by adjacent nodes. The ST-GNN maps multi-source features within a historical window to a node-level risk score at a future time:

$$\hat{P}_{i,t+1} = f_\theta\left(\mathbf{X}_{t-5:t}, A_{\text{up}}, A_{\text{down}}\right) \tag{1}$$

where $\hat{P}_{i,t+1}$ denotes the predicted normalized risk-proxy score at node i for the next interval $t+1$; $\mathbf{X}_{t-5:t}$ contains the traffic, environmental, and incident features of all nodes over six historical intervals; A_{up} and A_{down} are the upstream and downstream adjacency matrices; and f_θ is the nonlinear ST-GNN mapping with learnable parameters θ. In the classification, control, and evaluation formulas below, t indexes the target interval, so $\hat{P}_{i,t}$ uses only observations available through interval $t-1$.

The risk score is bounded by:

$$0 \le \hat{P}_{i,t} \le 1 \tag{2}$$

This constraint keeps the model output within a common normalized interval and facilitates subsequent threshold-based classification.

(2) Risk-classification constraint

The predicted continuous risk score is converted into a discrete predicted risk level using the following threshold rule:

$$\hat{L}_{i,t} = \begin{cases} 1, & 0 \le \hat{P}_{i,t} < \theta_1 \\ 2, & \theta_1 \le \hat{P}_{i,t} < \theta_2 \\ 3, & \theta_2 \le \hat{P}_{i,t} \le 1 \end{cases} \tag{3}$$

where $\hat{L}_{i,t}$ denotes the predicted risk level at node i for target interval t; θ_1 and θ_2 separate low from medium risk and medium from high risk, respectively, with $0 < \theta_1 < \theta_2 < 1$. The same thresholds are applied to the simulated reference scores to obtain the reference risk levels.

This rule assigns a predicted risk level; it does not yet determine the final coordinated lighting command.

To distinguish risk assessment from lighting control, $Y_{i,t}$ denotes the reference risk level derived from the simulated data; $\hat{L}_{i,t}$ denotes the predicted risk level; and $C_{i,t}$ denotes

the final lighting control level for the selected strategy. Prediction performance compares $Y_{i,t}$ with $\hat{L}_{i,t}$, whereas coordination and spatial smoothing act on $C_{i,t}$. For independent mapping, $C_{i,t} = \hat{L}_{i,t}$; coordinated commands may differ from the raw predictions.

(3) Low-position-lighting control constraint

Each node can independently adjust color, luminance, and flashing frequency. Only one lighting signal pattern is permitted at any time; therefore:

$$\sum_{k=1}^{3} u_{i,k,t} = 1 \tag{4}$$

$$u_{i,k,t} \in \{0, 1\} \tag{5}$$

where $u_{i,k,t}$ equals 1 when node i selects lighting pattern k for target interval t, and equals 0 otherwise. Patterns $k = 1, 2, 3$ represent normal guidance, risk warning, and incident alert, respectively.

The final lighting control level is determined by the selected binary control variable:

$$C_{i,t} = \sum_{k=1}^{3} k\, u_{i,k,t} \tag{6}$$

where $C_{i,t}$ is the final lighting control level and $u_{i,k,t}$ is its binary selection variable. Equations (4) and (5) ensure that exactly one lighting pattern is selected at each node and time step.

(4) Lighting-signal coordination constraint

Congestion and sudden incidents may affect vehicles upstream along the traffic flow direction. To convey warning information in advance, an upstream coordination constraint is applied to an incident node:

$$\hat{L}_{i,t} = 3 \Rightarrow C_{i-m,t} \geq 2, m = 1, \ldots, min(n, i-1) \tag{7}$$

where n is the selected upstream coordination range in numbers of nodes. Nodes are numbered increasingly in the direction of travel, so upstream nodes have smaller indices. The bound $min(n, i-1)$ restricts the rule to existing upstream nodes. Predicted Level 3 alerts are preserved: $\hat{L}_{i,t} = 3 \Rightarrow C_{i,t} = 3$.

For the upstream-plus-spatial strategy, the difference between final lighting levels at adjacent nodes is constrained by:

$$|C_{i,t} - C_{i-1,t}| \leq 1, i = 2, \ldots, N \tag{8}$$

The spatial constraint applies to final lighting-command levels rather than predicted risk levels. After upstream coordination, spatial smoothing removes differences greater than one while preserving existing Level 3 alerts. Independent mapping and upstream-only coordination are comparison strategies and are not assumed to satisfy Equation (8).

This constraint limits abrupt control-level changes between adjacent nodes and reduces spatial discontinuities in the lighting signal pattern.

2.3. Strategy Evaluation Functions

The control states are generated from predicted levels and a rule base rather than by optimizing a control sequence. The following quantities describe the predicted score, command changes, and notional response cost. Prediction accuracy is assessed against reference labels, while command statistics are evaluated after the selected coordination procedure. A descriptive aggregate is not treated as evidence of safety effectiveness.

(1) Risk-level evaluation term

Node weights permit a general weighted summary of predicted scores. The present example does not use measured wind exposure or sight-distance data to establish location-specific safety weights. This term describes the prediction sequence and does not represent risk reduced by a lighting action.

Let N be the number of control nodes and T the number of target intervals in one evaluation sequence. For predicted score $\hat{P}_{i,t}$ and node weight w_i, the cumulative weighted risk-proxy term is:

$$J_1 = \sum_{t=1}^{T} \sum_{i=1}^{N} w_i \hat{P}_{i,t} \tag{9}$$

where J_1 is the cumulative weighted predicted risk proxy, w_i is the weight of node i, and $\hat{P}_{i,t}$ is its predicted score.

The node weights satisfy the normalization constraint:

$$\sum_{i=1}^{N} w_i = 1, w_i \geq 0 \tag{10}$$

A larger J_1 indicates a higher cumulative weighted predicted score; a smaller value indicates a lower score. This descriptive statistic is not a measured change in traffic risk caused by lighting.

(2) Control-smoothness evaluation term

Changes in the selected command level are examined as software-output statistics. The number of changes and their level differences are distinct quantities; neither alone establishes acceptable flicker, visual workload, or driver comfort.

The cumulative magnitude of changes in the command level between consecutive target intervals is defined as:

$$J_2 = \sum_{t=2}^{T} \sum_{i=1}^{N} |C_{i,t} - C_{i,t-1}| \tag{11}$$

where J_2 is the cumulative switching magnitude, and $C_{i,t}$ and $C_{i,t-1}$ are the final lighting control levels at consecutive target intervals.

A larger J_2 indicates a greater cumulative magnitude of command changes, not necessarily more switching events. This statistic is not a direct measure of visual comfort.

Using the binary selection variables in Equation (6), the same switching magnitude can be written equivalently as:

$$J_2 = \sum_{t=2}^{T} \sum_{i=1}^{N} \left| \sum_{k=1}^{3} k(u_{i,k,t} - u_{i,k,t-1}) \right| \tag{12}$$

where $u_{i,k,t}$ is the binary selection variable. A Level 1-to-Level 3 transition contributes two units to switching magnitude but only one switching event. Switching events are counted separately as $N_{\text{switch}} = \frac{1}{2} \sum_{t=2}^{T} \sum_{i=1}^{N} \sum_{k=1}^{3} |u_{i,k,t} - u_{i,k,t-1}|$. For independent simulation runs, these definitions apply within each run; boundaries between runs are not temporal transitions.

(3) Guidance-response-cost evaluation term

The three lighting patterns represent normal guidance, warning, and incident-alert command categories. They are assigned ordered, notional response costs to summarize command use. The descriptions of steady guidance, yellow warning, or red chasing alerts identify illustrative signal concepts, not validated photometric settings or measured energy demands.

The guidance-response-cost term is defined as:

$$J_3 = \sum_{t=1}^{T} \sum_{i=1}^{N} \sum_{k=1}^{3} c_k \, u_{i,k,t} \tag{13}$$

where J_3 is the cumulative notional response cost, c_k is the cost assigned to pattern k, and $u_{i,k,t}$ indicates whether that pattern is selected for node i at interval t.

The costs of the three lighting signal patterns satisfy:

$$0 < c_1 < c_2 < c_3 \tag{14}$$

The coefficients c_1, c_2, and c_3 correspond to normal guidance, risk warning, and incident alert. A larger J_3 indicates greater notional response cost. The coefficients do not represent measured electricity consumption or monetary cost.

(4) Composite evaluation function

The cumulative terms can be normalized for descriptive comparison. Bars distinguish normalized values from the cumulative quantities in Equations (9), (11), and (13):

$$J = \alpha \bar{J}_1 + \beta \bar{J}_2 + \gamma \bar{J}_3 \tag{15}$$

The normalized terms are $\bar{J}_1 = \frac{J_1}{T}$, $\bar{J}_2 = \frac{J_2}{N(T-1)}$, and $\bar{J}_3 = \frac{J_3}{NTc_3}$. These expressions apply to a sequence with $T > 1$.

Here $\bar{J}_1$ is the time-averaged weighted predicted score, $\bar{J}_2$ is the mean switching magnitude per node-transition opportunity, and $\bar{J}_3$ is the response cost normalized by the maximum pattern cost. For three levels, $\bar{J}_2$ can range from 0 to 2; it is not a switching-event rate.

In Equation (15), α, β, and γ weight the normalized score, switching magnitude, and notional cost, respectively:

$$\alpha + \beta + \gamma = 1; \alpha, \beta, \gamma \geq 0$$

The weights express the selected descriptive priorities, not calibrated safety preferences. Because the compared controllers receive the same prediction sequence, the score term does not demonstrate a controller-induced reduction in risk. Section 3 therefore interprets the component statistics separately rather than ranking actual safety interventions using the aggregate.

A smaller J indicates a lower descriptive aggregate under the selected weights. It does not establish a reduction in actual traffic risk or the optimality of a lighting strategy.

Normal guidance, risk warning, and incident alert are determined by predefined thresholds and the rule base. The composite evaluation function summarizes the generated results and does not optimize the thresholds, weights, or control sequence.

2.4. Strategy Generation Procedure

Lighting commands are formed from a risk forecast, threshold classification, the selected coordination procedure, and signal-pattern mapping. The mapping is defined in Step 5 and applied to the finalized command levels after Step 6. This is a rule-based procedure, not a multi-objective optimization process.

Control zones are represented as nodes with fixed directional relationships. Six historical intervals of traffic, environmental, and observed-event information are used to forecast the next risk-proxy interval. Reference labels are used for prediction evaluation; they are not substituted for model outputs when evaluating prediction-driven commands.

Step 1: Data input and parameter initialization

The inputs comprise traffic flow, vehicle operation, weather conditions, traffic incidents, and node topology. The prediction horizon, historical-window length, risk thresholds θ_1

and θ_2, and lighting-signal rule-base parameters are initialized. The case study predicts one time step ahead using a six-step historical input window.

Step 2: Construction of the bridge-node graph

Low-position-lighting control zones form the graph nodes. Fixed upstream and downstream adjacency matrices, A_{up} and A_{down}, describe directional relationships between zones. The model settings are given in Section 3.2.

Step 3: Node-level risk prediction

The six-interval feature window and the two directional adjacency matrices are input to the ST-GNN to forecast the next target interval:

$$\hat{P}_{i,t+1} = f_\theta\left(\mathbf{X}_{t-5:t}, A_{\text{up}}, A_{\text{down}}\right) \tag{16}$$

where $\hat{P}_{i,t+1}$ is the predicted score for the interval following the six-step history $\mathbf{X}_{t-5:t}$, and f_θ is the parameterized ST-GNN mapping. Subsequent classification and control use the target-interval indexing defined after Equation (1).

Step 4: Risk-level classification

The predicted risk scores are classified using θ_1 and θ_2:

$$\hat{L}_{i,t} = \begin{cases} 1, & 0 \le \hat{P}_{i,t} < \theta_1 \\ 2, & \theta_1 \le \hat{P}_{i,t} < \theta_2 \\ 3, & \theta_2 \le \hat{P}_{i,t} \le 1 \end{cases} \tag{17}$$

where $\hat{L}_{i,t}$ is the raw predicted risk level. Values 1, 2, and 3 correspond to normal, warning, and incident-alert categories.

Step 5: Mapping through the lighting-signal rule base

The rule base maps a finalized command level to color, luminance, flashing frequency, and dynamic light-flow parameters. The mapping is defined here and applied after the selected coordination procedure in Step 6:

$$G_{i,t} = \Phi(C_{i,t}) \tag{18}$$

where G_t is the lighting-parameter set for node i at target interval t, M is the pattern mapping, and C_i,t is the final command level. Independent mapping uses C_i,t = L_i,t.

Specifically, G1 = (white light, 40% nominal luminance, steady output, static guidance); G2 = (yellow light, 60% nominal luminance, 0.8 Hz flashing, static warning); and G3 = (red light, 80% nominal luminance, 2.0 Hz chasing, upstream-directed incident alert).

$$G_{i,t} = \begin{cases} G_1 & C_{i,t} = 1 \\ G_2 & C_{i,t} = 2 \\ G_3 & C_{i,t} = 3 \end{cases} \tag{19}$$

G_1, G_2, and G_3 are the parameter sets for normal guidance, risk warning, and incident alert, respectively.

Step 6: Coordinated upstream control

The raw predicted levels initialize the command levels C_i,t. For strategies with upstream coordination, a predicted Level 3 alert imposes at least a Level 2 warning on each existing upstream node within the selected range. For the upstream-plus-spatial strategy, the commands are then repaired using C_i,t, not L_i,t, to satisfy the adjacent-command constraint. The repair rule only raises command levels and never lowers an existing alert; because levels are bounded by 3 and the number of nodes is finite, the procedure always terminates. The notional response costs are c1 = 1, c2 = 2, and c3 = 4.

$$\hat{L}_{i,t} = 3 \Rightarrow C_{i-m,t} \ge 2, m = 1, \ldots, min(n, i-1) \tag{20}$$

Here, n and m have the meanings in Equation (7). Existing Level 3 alerts are retained. For the upstream-plus-spatial strategy, the commands are additionally repaired to satisfy Equation (8); the reference and predicted risk levels are not changed.

Step 7: Output of the active visual-guidance strategy

The final command levels $C_{i,t}$ are mapped through Equations (18) and (19) to obtain the bridge-wide lighting commands. Prediction metrics compare $Y_{i,t}$ with $\hat{L}_{i,t}$; command switching and response-cost metrics are evaluated using $C_{i,t}$ and $u_{i,k,t}$.

3. Results

3.1. Simulation Experiment

The simulation study examines whether the prediction-control chain can forecast a simulated risk proxy and generate graded lighting commands under controlled scenarios. It includes baseline comparisons, repeated neural-network runs, classwise prediction evaluation, and separate coordination tests. The experiment does not estimate calibrated crash probability or a field safety benefit. All numerical results reported in this manuscript come from the same implementation and the same fixed simulation run set; the ten seeds are training seeds applied to that run set.

The 1200 m bridge segment was divided into 20 lighting-control zones, each 60 m long. Each zone served as a traffic-observation node, risk-prediction node, and lighting-control node. Five-minute aggregation was used. The model input consisted of six historical intervals, corresponding to 30 min, and the target was the risk proxy in the next 5 min interval.

The node features were traffic volume, mean speed, occupancy, visibility, and observed event state. Continuous variables were normalized using training-set statistics only, whereas the binary event state was kept unstandardized. The simulated scenario families included normal operation, demand increase and queue formation, local low visibility, single stopped-vehicle events, temporary capacity drops, and compound scenarios. In the DS simulation experiment, 40 independent runs were generated for each scenario family, yielding 240 simulation runs. Runs were split by complete simulation trajectories rather than by adjacent sliding windows, with a 60/20/20 train/validation/test split within each family.

3.2. Risk-Proxy Label and Model Settings

The prediction target was defined as the future 5 min car-following conflict-exposure proxy at each node. For node i and aggregation interval t, the raw exposure was computed as $e(i,t) = 0.18\,o(i,t) + 0.24\,d(i,t) + 0.18\,v(i,t) + 0.22\,q(i,t) + 0.18\,m(i,t) + \varepsilon(i,t)$, where o is occupancy, $d = \max\left(0, 1 - \frac{\text{speed}}{\text{desired}_{\text{speed}}}\right)$, v is the absolute one-step speed change divided by 100, q is the simulated incident indicator, m is the queue-memory state, and ε is zero-mean Gaussian noise with standard deviation 0.015 clipped at zero. The normalized target was $y(i,t) = \min\left(1, \frac{e(i,t)}{s_{\text{ref}}}\right)$, where s_{ref} is the 95th percentile of positive raw exposures in the training runs only. Spatial correlation is injected through upstream pressure propagation in the density update, queue-memory recursion, and a Gaussian low-visibility field over adjacent nodes. Incident scenarios inject stopped vehicles or capacity drops at randomly selected nodes and propagate their effects through the same upstream update. This score is a simulation risk proxy and is not an accident probability.

Because the incident indicator q has weight 0.22 in the raw exposure, which exceeds the training-only reference scale s_ref = 0.19683, any target cell with an active incident is normalized to 1 and is therefore labeled Level 3 regardless of the other traffic terms. The observed event state is also supplied as an input feature, so persistence of an ongoing incident and detection of a new incident onset are different tasks. On the seed-0 test output,

Level 3 recall was 0.8972 for 282 cells with an incident already present in the input window and 0.8471 for 1367 onset cells. Across the ten seeds, the corresponding mean recalls were 0.9152 ± 0.0118 and 0.8700 ± 0.0097, respectively. These results show that the Level 3 metric partly measures incident-state persistence and should not be interpreted as a stand-alone measure of advance incident prediction.

The command thresholds were fixed at $\theta_1 = 0.40$ and $\theta_2 = 0.70$. Thus, Level 1 is $y < 0.40$, Level 2 is $0.40 \leq y < 0.70$, and Level 3 is $y \geq 0.70$. Each simulation run contains six 5 min warm-up/history intervals, 24 evaluated 5 min target intervals, and one terminal interval retained for window construction, giving 31 stored 5 min bins per run. The 48 test runs therefore contain $48 \times 20 \times 24 = 23{,}040$ evaluated node-time cells. The training-only normalization value was $s_{\text{ref}} = 0.19683$. These thresholds are controlled simulation cutoffs and are not empirically calibrated traffic-safety limits.

The main model was DS-GCN-GRU, a direction-separated graph convolution and gated recurrent unit network. Its input tensor had a shape of $6 \times 20 \times 5$, representing six historical five-minute steps, 20 ordered control zones, and five features.. The model used two graph layers with 32 hidden units, a residual feature path, a single-layer GRU with 64 hidden units, and a sigmoid output for 20 continuous scores. The upstream matrix has a binary edge weight of 1 for $j = i - 1$, and the downstream matrix has a binary edge weight of 1 for $j = i + 1$. Both matrices exclude self-loops and no degree normalization is applied; the two matrices are used as separate directed channels. The model has 22,433 trainable parameters. All node indices use the zero-based convention $i = 0, \ldots, N - 1$ throughout the manuscript. The architecture and training specifications of the DS-GCN-GRU model are summarized in Table 1. The overall architecture of the proposed DS-GCN-GRU model is illustrated in Figure 1.

Table 1. DS-GCN-GRU architecture and training specification.

Item	Specification
Input	6 historical intervals × 20 nodes × 5 features
Spatial module	Two direction-separated graph layers; binary upstream/downstream edges; no self-loops; no degree normalization
Hidden dimensions	Graph layers: 32; GRU: 64
Temporal module	Single-layer GRU with residual feature fusion
Output	20 sigmoid scores for the next 5 min target interval
Trainable parameters	22,433
Optimization	Adam; learning rate 0.001; weight decay 0.0001; batch size 64; maximum 100 epochs; patience 12

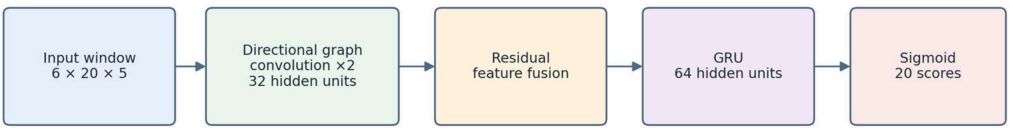


Figure 1. Architecture of the DS-GCN-GRU model.

The comparison models included a training-set node-wise mean predictor, a persistence predictor, Ridge regression using lagged features, gradient boosting regression, LSTM, GRU, a simplified STGCN baseline, and the proposed DS-GCN-GRU. The training mean is calculated independently for each node from all training target intervals and is repeated across the test horizon. The persistence predictor copies the previous target interval within a run and uses the training node mean at each run boundary. Neural-network models

were trained with Adam, learning rate 0.001, weight decay 0.0001, batch size 64, early stopping patience 12, and at most 100 epochs. The same complete-run split was used for all models: 60% training, 20% validation, and 20% testing. Ten random seeds (0 to 9) were used for repeated neural-network training. The training and validation MSE curves of the DS-GCN-GRU model with seed 0 are shown in Figure 2.

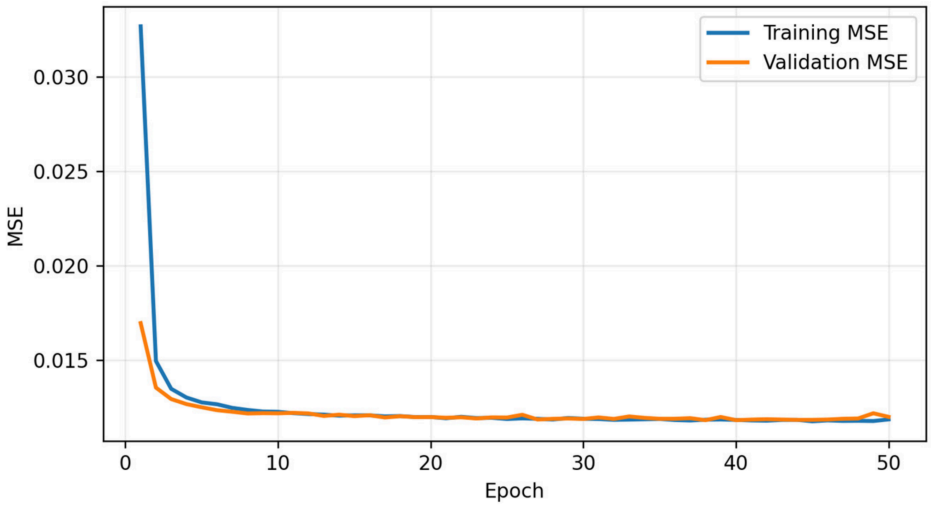


Figure 2. Training and validation MSE curves for DS-GCN-GRU with seed 0.

3.3. Prediction Performance

Table 2 reports the test-set performance. DS-GCN-GRU achieved the lowest mean RMSE (0.1094 ± 0.0004) and MAE (0.0773 ± 0.0007), and the highest macro-F1 (0.8339 ± 0.0082) and high-risk PR-AUC (0.9292 ± 0.0010) among the listed models. Its Level 3 recall was 0.8777 ± 0.0095. The persistence baseline and the two matched graph ablations make the temporal and spatial contributions directly comparable under the same run split. With a noise standard deviation of 0.015 and s_ref = 0.19683, the normalized noise scale is approximately 0.076; the remaining error also reflects model approximation and scenario variability. The test-set RMSE values of the evaluated models are compared in Figure 3.

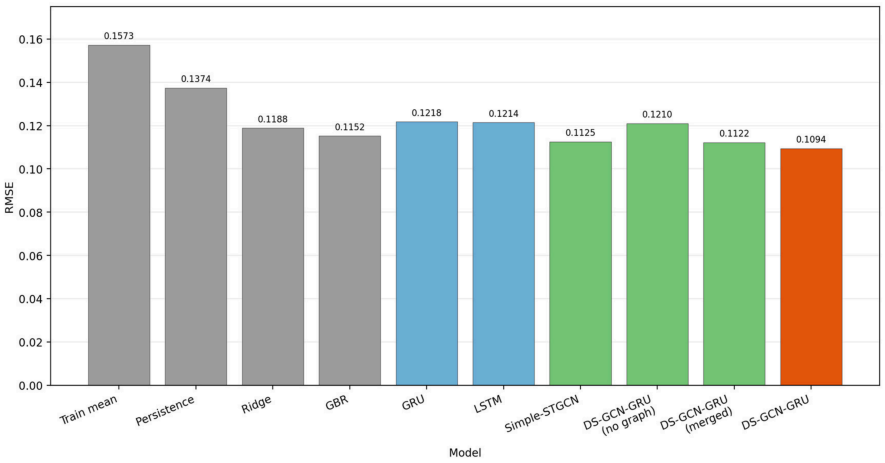


Figure 3. Test-set RMSE comparison across the evaluated models, including persistence and graph ablations.

Table 2. Test-set prediction performance in the DS simulation experiment.

Model	RMSE	MAE	Balanced Accuracy	Macro-F1	Level 3 Recall	Level 3 F1	High-Risk ROC-AUC	High-Risk PR-AUC
Train mean	0.1573	0.1026	0.7018	0.7566	0.6980	0.8218	0.8735	0.7518
Persistence	0.1374	0.0986	0.8009	0.8020	0.8969	0.8996	0.9569	0.8693
Ridge	0.1188	0.0847	0.7897	0.8197	0.8514	0.8946	0.9593	0.9098
GBR	0.1152	0.0816	0.7780	0.8144	0.8181	0.8805	0.9560	0.9146
GRU	0.1218 ± 0.0003	0.0865 ± 0.0004	0.7736 ± 0.0097	0.7929 ± 0.0047	0.8421 ± 0.0127	0.8578 ± 0.0025	0.9523 ± 0.0002	0.8839 ± 0.0038
LSTM	0.1214 ± 0.0003	0.0861 ± 0.0003	0.7750 ± 0.0051	0.7941 ± 0.0035	0.8321 ± 0.0066	0.8548 ± 0.0048	0.9524 ± 0.0004	0.8875 ± 0.0017
Simple-STGCN	0.1125 ± 0.0003	0.0801 ± 0.0006	0.7985 ± 0.0120	0.8241 ± 0.0063	0.8489 ± 0.0086	0.8926 ± 0.0045	0.9567 ± 0.0006	0.9235 ± 0.0008
DS-GCN-GRU	0.1094 ± 0.0004	0.0773 ± 0.0007	0.8039 ± 0.0155	0.8339 ± 0.0082	0.8777 ± 0.0095	0.9092 ± 0.0048	0.9557 ± 0.0006	0.9292 ± 0.0010
DS-GCN-GRU without graph	0.1210 ± 0.0005	0.0857 ± 0.0005	0.7818 ± 0.0135	0.8003 ± 0.0079	0.8546 ± 0.0101	0.8668 ± 0.0059	0.9529 ± 0.0003	0.8943 ± 0.0028
DS-GCN-GRU with merged directions	0.1122 ± 0.0003	0.0794 ± 0.0005	0.8090 ± 0.0151	0.8331 ± 0.0077	0.8842 ± 0.0051	0.9085 ± 0.0018	0.9545 ± 0.0002	0.9266 ± 0.0007

Note: Neural-network results are reported as mean ± sample standard deviation across ten training seeds (0–9; standard deviation calculated with n − 1 in the denominator). All seeds use the same simulation dataset and the same train/validation/test split. The variation therefore measures training variability, not variability across independently generated datasets, external validation, or a confidence interval. Single values are reported for baselines evaluated once on the fixed dataset.

Table 2 shows a trade-off between Level 3 sensitivity and broader predictive performance. Figure 4 illustrates an example of the reference and predicted risk-proxy scores during a simulated incident-onset episode. The persistence baseline achieved higher Level 3 recall (0.8969) than DS-GCN-GRU (0.8777 ± 0.0095), but had higher RMSE (0.1374 versus 0.1094) and MAE (0.0986 versus 0.0773), and lower macro-F1 (0.8020 versus 0.8339), Level 3 F1 (0.8996 versus 0.9092), and high-risk PR-AUC (0.8693 versus 0.9292). Thus, the proposed model improves several aggregate and precision-sensitive metrics but does not minimize Level 3 misses among the evaluated methods. This trade-off limits claims of overall warning superiority. Its ROC-AUC (0.9557) is also below that of persistence (0.9569), Ridge (0.9593), and simplified STGCN (0.9567). Removing graph information increases RMSE to 0.1210, while merging the directional channels increases RMSE to 0.1122. Direction separation improves continuous risk-proxy error, whereas its advantages over merged-direction processing are not consistent across classification metrics.

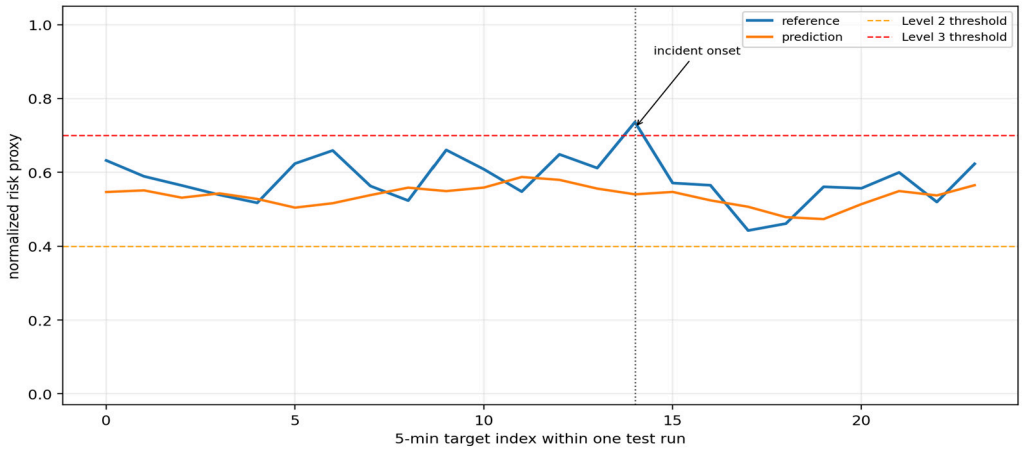


Figure 4. Example of a Level 3 episode with an incident onset at one test node.

3.4. Level 3 Risk Detection and False-Negative Analysis

Table 3 and Figure 5 compare reference and predicted levels for the seed-0 model on the test set. Rows denote reference levels and columns denote predicted levels. This single-model confusion matrix is distinct from the mean and standard deviation across ten neural-network training seeds in Table 2.

Table 3. Confusion matrix of DS-GCN-GRU on the test set.

Reference/Prediction	Level 1	Level 2	Level 3
Level 1	19,684	151	17
Level 2	745	728	66
Level 3	115	123	1411

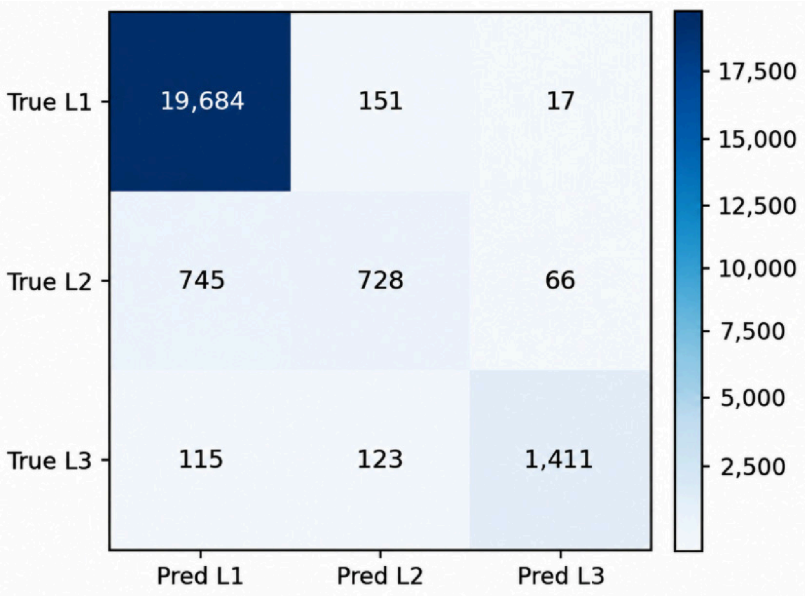


Figure 5. Seed-0 confusion matrix; rows are reference levels and columns are predicted levels.

The seed-0 confusion matrix contains 1649 reference Level 3 node-time samples. Of these, 1411 were classified as Level 3, 115 as Level 1, and 123 as Level 2, giving a Level 3 recall of 0.8557 for that seed. Across ten seeds, the mean recalls were 0.9863 ± 0.0035 for Level 1, 0.5477 ± 0.0459 for Level 2, and 0.8777 ± 0.0095 for Level 3. The intermediate class remains the main classification difficulty.

The intermediate warning class is more difficult to identify. In the seed-0 matrix, 728 of 1539 Level 2 samples were correctly classified, 745 were assigned to Level 1, and 66 to Level 3. This corresponds to a Level 2 recall of 47.30%. The result motivates reporting class-wise metrics and retaining the Level 2 limitation in the interpretation of the controller outputs.

A total of 238 reference Level 3 samples were assigned to a lower predicted level. Together with the Level 2 under-detection, this shows that graded warning errors remain despite the improved coverage of severe proxy states. The present results do not identify a unique cause of these errors. Changes to thresholds, loss weighting, or a separate event-confirmation channel are not evaluated here and are not presented as demonstrated remedies. The distributions of reference and predicted risk levels on the test set are compared in Figure 6.

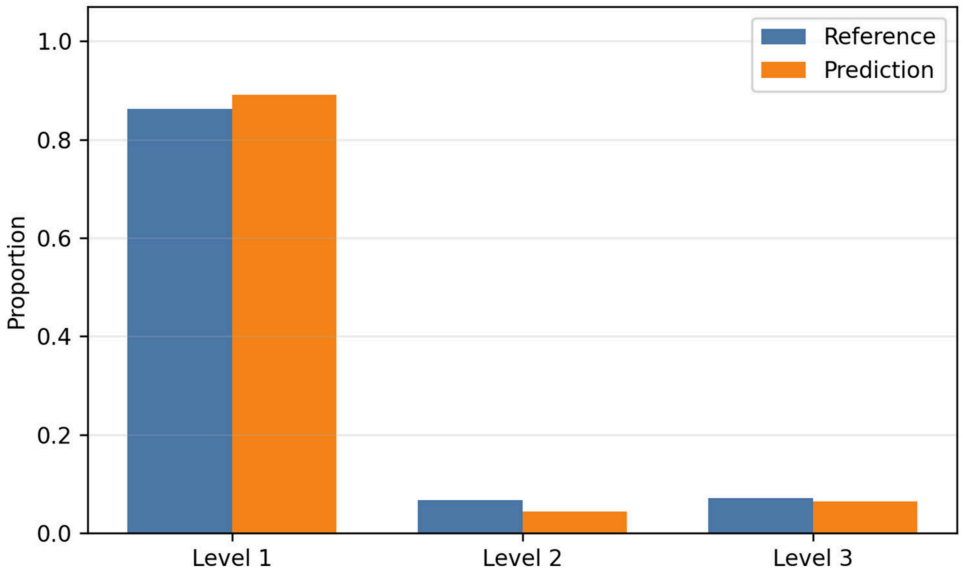


Figure 6. Reference and predicted risk-level distributions on the test set.

3.5. Active Visual-Guidance Control Results

The lighting-control comparison uses independent node-wise mapping, upstream coordination, and upstream coordination followed by spatial smoothing. Independent mapping uses each predicted level directly. Coordination raises the response on eligible upstream nodes of a predicted Level 3 alert, while the spatial stage addresses adjacent-command differences greater than one. These operations modify command levels, not reference labels or the underlying prediction scores. The spatial-constraint violations under the three lighting-control strategies are compared in Figure 7.

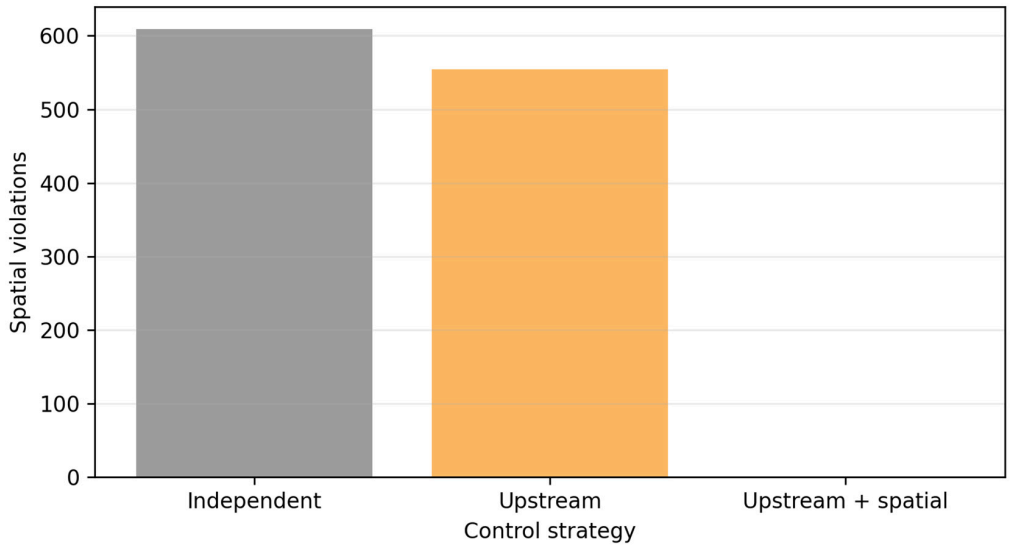


Figure 7. Spatial-constraint violations under three lighting-control strategies.

Under the reported predicted inputs, independent mapping produced 609 adjacent-command violations and upstream-only coordination produced 554. The upstream-plus-spatial strategy produced zero violations. The reported retention value of 1.0000 means that commands already triggered as Level 3 were preserved; it does not mean that all reference Level 3 samples were detected. In particular, this retention statistic cannot replace the prediction recall reported in Section 3.4. The comparative results of the three lighting-control strategies under predicted test-set inputs are summarized in Table 4.

Table 4. Control-strategy comparison under predicted test-set inputs.

Strategy	Level 1 Share	Level 2 Share	Level 3 Share	Switch Count	Switch Amplitude	Relative Cost	Spatial Violations	Level 3 Retained
Independent	0.8917	0.0435	0.0648	494	611	0.3095	609	1.0000
Upstream	0.8865	0.0487	0.0648	588	705	0.3108	554	1.0000
Upstream + spatial	0.8625	0.0727	0.0648	525	642	0.3168	0	1.0000

The spatial stage increased the Level 2 share and the reported notional cost. The controller uses an upstream range of n = 2, so each Level 3 node can impose at least Level 2 on the two valid upstream zones. The repair rule and cost coefficients are defined in Section 2.4, Step 6. These are command-allocation metrics and do not estimate driver comfort, energy use, or safety benefit. The control-strategy comparison results for the forced Level 3 scenarios are presented in Table 5.

Table 5. Forced Level 3 control scenarios.

Strategy	Level 1 Share	Level 2 Share	Level 3 Share	Switch Count	Switch Amplitude	Relative Cost	Spatial Violations	Level 3 Retained
Independent	0.9250	0.0063	0.0688	19	36	0.3032	18	1.0000
Upstream	0.8188	0.1125	0.0688	43	58	0.3297	10	1.0000
Upstream + spatial	0.7563	0.1750	0.0688	54	68	0.3453	0	1.0000

The forced-input cases cover mid-span, entry, exit, multiple and overlapping alerts, and direct Level 1-to-Level 3 adjacency. They test whether the controller preserves supplied alerts and satisfies the spatial constraint when the trigger is present. They are not additional evidence that the predictor anticipates incidents. Switching totals have a temporal interpretation only for ordered sequences within a single run; transitions between independent forced scenarios must not be interpreted as lighting behavior over time.

3.6. Summary of the Simulation Results

The available results assess two distinct parts of the workflow: prediction against a simulation proxy and execution of a discrete lighting controller. DS-GCN-GRU achieved the lowest reported RMSE and MAE and the highest macro-F1 and high-risk PR-AUC, while persistence achieved higher Level 3 recall. The confusion matrix shows substantial intermediate-level under-detection and remaining Level 3 false negatives. The control tests support the stated command-level constraint behavior for the tested inputs. They do not establish driver comprehension, visual comfort, glare acceptability, crash reduction, or field reliability.

4. Discussion

The study concerns a short-horizon prediction-to-command workflow for a linear bridge corridor. Crosswind exposure, deck geometry, vehicle composition, and sight-distance effects are not independently modeled or validated. The conclusions therefore apply to the specified simulation setting.

The same experiment also included two graph ablations. Removing both directed adjacency channels from DS-GCN-GRU tests whether the model gains from spatial message passing; merging the upstream and downstream channels tests whether direction separation matters. These ablations were trained under the same split, optimizer, batch size, and seed protocol as the main model. The ablation metrics are reported in Table 2 and are interpreted as evidence about model components within the simulation proxy.

Upstream coordination expands warning coverage, whereas the additional spatial stage enforces the tested adjacent-command constraint. Preserving a Level 3 command is distinct from identifying a reference Level 3 sample: the controller does not make the 238 seed-0 prediction false negatives disappear. The low Level 2 recall also shows that successful command execution does not establish reliable three-level warning. Notional cost and switching statistics describe these commands; they are not calibrated measures of energy, visual load, or safety benefit.

The evidence has several limitations. All traffic inputs and reference labels are simulated, and the normalized proxy is not calibrated to observed crashes. The thresholds, lighting parameters, and upstream range are therefore design values for controlled evaluation. The reported lighting controller validates command consistency, Level 3 preservation, and spatial feasibility; it does not establish crash-risk reduction, driver behavioral benefits, acceptable optical signals, or field real-time performance. Future work should connect the simulator to observed bridge data and driver-in-the-loop experiments.

Lighting levels remain logical command categories. Work on service vehicle warning flashes and tunnel lighting [19,20] provides relevant human-factor context but does not validate the proposed bridge signal meanings, brightness settings, or flash rates. Photometric output, glare, visibility in fog, possible confusion with other road signals, and abrupt braking were not measured. FHWA adaptive-lighting guidance [17] is background guidance, not approval or certification of this installation. No compliance with a particular road-lighting or traffic-control standard is asserted, and the signal descriptions are not operating recommendations for an open road.

Finally, a five-minute forecast horizon is not an end-to-end response-time measurement. Model inference, input acquisition, communication, actuator response, and operation during failures were not benchmarked in the reported study. No validated fail-safe behavior, real-time deployment capability, or independent emergency-detection performance is established. These are limits of the present evidence, not tests assumed to have been completed.

5. Conclusions

This study presents an ST-GNN-based framework that forecasts a short-term simulation risk proxy and converts predicted levels into graded low-position-lighting commands. It separates reference labels, raw predictions, and final commands so that prediction errors and coordination behavior can be interpreted independently.

Across the ten training seeds on the fixed simulation dataset, DS-GCN-GRU achieved a mean RMSE of 0.1094 $\pm$ 0.0004, MAE of 0.0773 $\pm$ 0.0007, macro-F1 of 0.8339 $\pm$ 0.0082, Level 3 recall of 0.8777 $\pm$ 0.0095, and high-risk PR-AUC of 0.9292 $\pm$ 0.0010. These values are means $\pm$ sample standard deviations across training seeds and agree with the Abstract and Table 2. The graph ablations indicate lower continuous prediction error with directed spatial processing, while classification advantages vary by metric. The persistence baseline achieved higher Level 3 recall (0.8969), whereas DS-GCN-GRU achieved lower RMSE and MAE and higher macro-F1, Level 3 F1, and high-risk PR-AUC.

For seed 0, 1411 of 1649 reference Level 3 samples were correctly identified and 238 were assigned to lower levels. The same matrix gives a Level 2 recall of 47.30%, with 745 intermediate-risk samples classified as Level 1. Both findings limit the reliability of the three-level warning interpretation and are retained alongside the aggregate performance measures.

Upstream coordination followed by spatial smoothing reduced reported adjacent-command violations to zero for the tested inputs while retaining alerts already predicted as Level 3. The forced-input cases support this command-level behavior, not incident-

prediction accuracy or safety effectiveness. The controller is a rule-based implementation rather than an optimized or human-validated lighting policy.

The findings support a bounded software-level demonstration of risk-proxy forecasting and coordinated low-position lighting commands on simulated bridge segments. The method provides an auditable mapping from simulated traffic states to risk levels and spatially feasible commands. External validation with observed traffic, calibrated safety outcomes, optical measurements, and driver response remains necessary before engineering deployment.

Author Contributions: Conceptualization, F.C. and Z.C. (Zhi Chen); methodology, F.C. and Y.K.; software, Y.K. and G.L.; validation, X.C., G.L., W.H., and B.L.; formal analysis, F.C., X.C., and B.L.; investigation, F.C. and Z.C. (Zhitao Chen); resources, Z.C. (Zhitao Chen); data curation, Y.K., X.C., and W.H.; writing—original draft preparation, F.C.; writing—review and editing, Y.K., X.C., Z.C. (Zhitao Chen), G.L., W.H., B.L., and Z.C. (Zhi Chen); visualization, Y.K. and G.L.; supervision, Z.C. (Zhi Chen); project administration, Z.C. (Zhi Chen). All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by Yunnan Transportation Planning and Design Institute Co., Ltd. through the project "Research on Key Technologies of Edge Computing Terminals for the Quasi-Unattended System of Expressway Tunnels", grant number YJSJ-ZL-2024-02, and the project "Research on Key Technologies of Digital Transformation Framework and Typical Application Scenarios of Yunnan Communications Investment & Construction Group Co., Ltd.", grant number YCIC-YF-2025-03.

Data Availability Statement: The data presented in this study are available from the corresponding author upon reasonable request.

Conflicts of Interest: Authors Fang Cao, Gang Li are employed by the company Guangzhou Yuexiu Group Co., Ltd. Author Zhitao Chen is employed by the company Yunnan Transportation Planning and Design Institute Co., Ltd. The authors declare that this study received funding from Yunnan Communications Investment & Construction Group Co., Ltd. The funder was not involved in the study design, collection, analysis, interpretation of data, in the writing of the manuscript, and in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

ST-GNN	Spatiotemporal Graph Neural Network
GNN	Graph Neural Network
RMSE	Root Mean Squared Error
MAE	Mean Absolute Error
MSE	Mean Squared Error
AUC	Area Under the Receiver Operating Characteristic Curve
LSTM	Long Short-Term Memory
GRU	Gated Recurrent Unit
STGCN	Spatiotemporal Graph Convolutional Network
DCRNN	Diffusion Convolutional Recurrent Neural Network
ASTGCN	Attention-Based Spatiotemporal Graph Convolutional Network
GCN	Graph Convolutional Network
GBR	Gradient Boosting Regression
ROC	Receiver Operating Characteristic
PR-AUC	Area Under the Precision–Recall Curve

References

1. Yan, B. Application of low-position lighting on an airport overpass. *Light Sources Light*. **2023**, *9*, 32–34. (In Chinese)
2. Lin, Q. Application and implementation of low-position lighting on an airport overpass: A case study of the Lanzhou Zhongchuan Airport Interchange. *Light Sources Light*. **2022**, *3*, 10–12. (In Chinese)

3. Guo, S. Application and implementation of low-position lighting on airport overpasses. *Urban Archit.* **2019**, *16*, 117–120. (In Chinese) [CrossRef]
4. Wang, Z.; Qi, Z.; Yu, Q.; Liu, P. Application of low-position lighting technology on a suspension bridge in a mountainous canyon. *Highway* **2020**, *65*, 204–207. (In Chinese)
5. Chen, S. Transient-Light Mechanism and Modulation Technology for low-Position Lighting. Master's Thesis, Shenzhen University, Shenzhen, China, 2020. (In Chinese) [CrossRef]
6. An, R. Application of low-position continuous lighting in urban-tunnel approach sections. *China Light Light.* **2021**, *3*, 10–13. (In Chinese)
7. Duan, Y.; Sun, Q. Application of low-position approach lighting in highway tunnels. *Highway* **2019**, *64*, 333–337. (In Chinese)
8. Li, Y.; Yu, R.; Shahabi, C.; Liu, Y. Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. In Proceedings of the International Conference on Learning Representations, Vancouver, BC, Canada, 30 April–3 May 2018.
9. Yu, B.; Yin, H.; Zhu, Z. Spatio-temporal graph convolutional networks: A deep learning framework for traffic forecasting. In Proceedings of the 27th International Joint Conference on Artificial Intelligence, Stockholm, Sweden, 13–19 July 2018; pp. 3634–3640.
10. Wu, Z.; Pan, S.; Long, G.; Jiang, J.; Zhang, C. Graph WaveNet for deep spatial-temporal graph modeling. In Proceedings of the 28th International Joint Conference on Artificial Intelligence, Macao, China, 10–16 August 2019; pp. 1907–1913.
11. Guo, S.; Lin, Y.; Feng, N.; Song, C.; Wan, H. Attention-based spatial-temporal graph convolutional networks for traffic flow forecasting. *Proc. AAAI Conf. Artif. Intell.* **2019**, *33*, 922–929. [CrossRef]
12. Wang, B.; Lin, Y.; Guo, S.; Wan, H. GSNet: Learning spatial-temporal correlations from geographical and semantic aspects for traffic accident risk forecasting. *Proc. AAAI Conf. Artif. Intell.* **2021**, *35*, 4402–4409. [CrossRef]
13. Nippani, A.; Li, D.; Ju, H.; Koutsopoulos, H.N.; Zhang, H.R. Graph neural networks for road safety modeling: Datasets and evaluations for accident analysis. *Adv. Neural Inf. Process. Syst.* **2023**, *36*, 52009–52032. [CrossRef]
14. Gao, X.; Jiang, X.; Haworth, J.; Zhuang, D.; Wang, S.; Chen, H.; Law, S. Uncertainty-aware probabilistic graph neural networks for road-level traffic crash prediction. *Accid. Anal. Prev.* **2024**, *208*, 107801. [CrossRef] [PubMed]
15. Yang, H.; Yan, B.; Gao, Y.; Deng, H.; Wu, X.; Wu, C. Global damage identification method based on multi-source signals and ST-GNN for frame structures. *Int. J. Struct. Stab. Dyn.* **2026**, 2750131. [CrossRef]
16. Ding, J. Joint prediction of mine-ventilation parameters using a spatiotemporal graph neural network. *Autom. Appl.* **2026**, *67*, 213–215. (In Chinese) [CrossRef]
17. Gibbons, R.; Guo, F.; Medina, A.; Terry, T.; Du, J.; Lutkevich, P.; Corkum, D.; Vetere, P. *Guidelines for the Implementation of Reduced Lighting on Roadways*; FHWA-HRT-14-050; Federal Highway Administration: Washington, DC, USA, 2014. Available online: https://www.fhwa.dot.gov/publications/research/safety/14050/ (accessed on 13 October 2026).
18. Federal Highway Administration. ITS in Tennessee: Fog Project on I-75 between Chattanooga and Knoxville. ITS Deployment Evaluation Source Record. 2000. Available online: https://www.itskrs.its.dot.gov/SRC-2000-c4fb4cedf18c4c3e852570ba0075d820 (accessed on 13 October 2026).
19. Skinner, N.P.; LaPlumm, T.T.; Bullough, J.D. Warning light flash frequency as a method for visual communication to drivers. *Transp. Res. Rec.* **2021**, *2675*, 88–93. [CrossRef]
20. Domenichini, L.; La Torre, F.; Vangi, D.; Virga, A.; Branzi, V. Influence of the lighting system on the driver's behavior in road tunnels: A driving simulator study. *J. Transp. Saf. Secur.* **2017**, *9*, 216–238. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.