JOURNAL DIGITAL TECHNOLOGY TREND
E-ISSN: 2963-8143, P-ISSN: 2963-9190
Volume 5 Number 2, December 2026, pp 232-241.
DOI: https://doi.org/10.56347/jdtt.v5i2.489.

Short-Term Traffic Condition Prediction for the Next 15 Minutes Using XGBoost Based on Temporal Features

Sartika Lina Mulani Sitio [1*], Nur Rofiq [2], Nardiono [3]
[1*,2,3] Department of Informatics Engineering, Faculty of Computer Science, Universitas Pamulang, South Tangerang City, Banten Province, Indonesia.

ABSTRACT
Short-term traffic condition prediction is essential for supporting responsive monitoring and proactive traffic management within intelligent transportation systems. This study aims to predict multiclass traffic conditions across four categories—normal, heavy, high, and low—for the subsequent 15-minute horizon using an eXtreme Gradient Boosting (XGBoost) model driven by temporal features. The experiment utilized a Kaggle traffic dataset comprising 2,976 chronological observations encompassing timestamp records, categorical vehicle counts, and traffic situation labels. The methodology involved data preprocessing, temporal and cyclical feature engineering, time-ordered data partitioning (2,296 training observations across the first 24 days and 671 testing observations across the final seven days), and hyperparameter optimization. Validation was performed using a three-fold TimeSeriesSplit with an observation gap to prevent temporal data leakage. Experimental results on the holdout test set demonstrated that the proposed XGBoost model achieved an overall accuracy of 75.71%, a weighted F1-score of 73.37%, a macro F1-score of 58.31%, and a balanced accuracy of 56.33%. This performance outperforms the majority baseline (56.18%) and represents an increase of 7.75 percentage points in accuracy over the persistence baseline (67.96%). Class-level analysis revealed strong predictive stability for heavy (F1-score 86.50%) and normal conditions (F1-score 82.21%), although minority classes exhibited lower sensitivity. These findings confirm that integrating XGBoost with engineered temporal features yields an effective approach for 15-minute traffic condition forecasting while rigorously maintaining chronological dependencies.

ARTICLE HISTORY
Received: 18 September 2026.
Revised: 19 September 2026.
Accepted: 10 October 2026.
Available Online: 20 October 2026.
Published: 1 December 2026.

KEYWORDS
Traffic Prediction; XGBoost; Temporal Features; Traffic Conditions; Short-Term Predictions.

CONTACT Sartika Lina Mulani Sitio dosen00847@unpam.ac.id Department of Informatics Engineering, Faculty of Computer Science, Universitas Pamulang, South Tangerang City, Banten Province, Indonesia.
© 2026 Information KAWANAD Institute

1. Introduction

The rapid expansion of urban mobility and vehicle ownership has rendered roadway traffic conditions exceedingly dynamic, non-linear, and difficult to anticipate over short temporal intervals. Abrupt fluctuations in vehicular volume, caused by bottlenecks, localized bottlenecks, and intermittent traffic signals, can rapidly trigger severe congestion across arterial corridors. Consequently, responsive traffic management demands robust forecasting mechanisms capable of projecting upcoming traffic conditions within immediate forward intervals, allowing operators to execute proactive interventions rather than reactive responses (Sun *et al*., 2025). Within modern intelligent transportation systems (ITS), short-term traffic prediction serves as a foundational component for automated signal coordination, dynamic route guidance, and situational incident detection (Soon *et al*., 2023). However, achieving reliable forecasts at granular intervals—such as a 15-minute horizon—remains challenging due to pronounced traffic volatility, recurring daily peaks, and non-stationary traffic dynamics that conventional statistical models often fail to capture adequately.

To address these complexities, recent transportation research has increasingly shifted toward machine learning architectures capable of extracting informative patterns from historical and temporal traffic records. Among ensemble learning methods, extreme gradient boosting (XGBoost) has emerged as an exceptionally effective framework for structured tabular data, demonstrating a superior capability to capture intricate, non-linear dependencies between traffic variables without incurring the heavy computational overhead typically associated with deep neural networks (Wang & Fang, 2024). Previous investigations have confirmed that tree-based gradient boosting models can effectively model short-term traffic flow variations, particularly when historical traffic observations are augmented with temporal and spatial indicators (L. Liu *et al*., 2024). Furthermore, extensive benchmark studies across related data-intensive domains have substantiated the robustness of XGBoost in tabular classification tasks, showing performance competitive with deep neural models such as Transformers (Sitio *et al*., 2024) while offering enhanced variance reduction, faster training convergence, and superior stability compared to conventional bagging and single-tree algorithms (Sitio *et al*., 2026).

Despite these documented advantages, several critical methodological challenges persist in the practical deployment of tree-based models for short-term traffic condition forecasting. First, many existing short-term traffic forecasting studies concentrate primarily on continuous flow regression rather than multiclass state classification, whereas traffic operators often require discrete categorical states (such as normal, heavy, high, or low conditions) to enact standardized operational protocols. Second, traffic state categories in real-world environments are inherently imbalanced, which complicates classification performance across less frequent yet operationally critical congestion regimes. Third, a prevalent shortcoming in machine learning workflows for transportation is the improper implementation of randomized data splitting and conventional k-fold cross-validation. Applying randomized partitioning to sequential time series data introduces temporal data leakage—where past instances are predicted using future information—leading to artificially inflated performance metrics that fail to generalize under real-world, forward-facing operational settings.

To resolve these limitations, this study presents a dedicated short-term traffic condition prediction framework for a forward 15-minute horizon using XGBoost structured on comprehensive temporal feature engineering. The methodological design strictly enforces chronological integrity by utilizing time-ordered data partitioning alongside forward-chaining TimeSeriesSplit validation, thereby completely preventing temporal data leakage. Engineered inputs incorporate vehicle categorical volumes, cyclical temporal transformations, and historical indicators to represent changing traffic states. The resulting classification model is evaluated against four distinct multiclass metrics—overall accuracy, balanced accuracy, weighted F1-score, and macro F1-score—and rigorously benchmarked against both majority class and persistence baselines. Through this design, the research provides a transparent, computationally lightweight, and reproducible framework for short-term traffic state forecasting that directly supports real-time decision-making in intelligent traffic surveillance environments.

2. Methodology

This research applies a supervised machine learning framework based on extreme gradient boosting (XGBoost) to predict multiclass traffic conditions across a forward 15-minute horizon using engineered temporal features. The overall research pipeline comprises six integrated stages: data acquisition, data preprocessing, temporal feature engineering, chronological data partitioning, model optimization via time-series cross-validation, and multi-metric performance evaluation. Figure 1 illustrates the end-to-end procedural workflow adopted in this study.

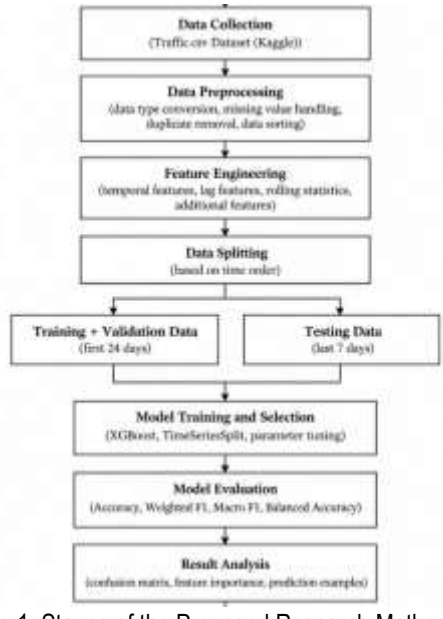


Figure 1. Stages of the Proposed Research Methodology

2.1 Data Acquisition and Preprocessing

The experimental data were obtained from a publicly accessible traffic repository hosted on Kaggle (Traffic.csv). The original dataset comprises 2,976 chronological observations recorded at 15-minute intervals across 31 continuous days, structured along nine attributes: Time, Date, Day of the week, CarCount, BikeCount, BusCount, TruckCount, Total, and Traffic Situation. The target variable, Traffic Situation, categorizes road occupancy into four distinct classes: normal (1,669 instances; 56.08%), heavy (682 instances; 22.92%), high (321 instances; 10.79%), and low (304 instances; 10.22%). Because this benchmark dataset represents an urban roadway corridor rather than a localized municipal network, it provides an objective baseline for examining short-term temporal dynamics in intelligent transportation systems (Abduljabbar *et al*., 2025; Rasaizadi & Seyedabrishami, 2022). Initial data validation confirmed that the dataset contained zero duplicate records and zero missing values across all feature columns (Guo *et al*., 2024). Data preprocessing involved datetime type casting of the temporal fields, standardizing string variables, and verifying the strict ascending chronological order of the time records. Explicit sorting by date and time was conducted to ensure that historical sequences were preserved prior to feature extraction, which is essential for preventing future-to-past data contamination (Zhu *et al*., 2022).

2.2 Feature Engineering and Chronological Data Partitioning

Feature engineering was conducted to capture multi-scale temporal dependencies, cyclic rhythms, and vehicular volume dynamics without incorporating future observations. The predictor feature set encompasses four primary components: 1) granular vehicle counts comprising individual volumes for cars, motorcycles, buses, and trucks, as well as the total vehicular volume; 2) cyclical temporal features generated through sine and cosine trigonometric transformations of the hour and day-of-week variables to preserve cyclic boundary continuity (e.g., the transition between 23:45 and 00:00); 3) historical lag features and

volume differencing across prior time steps to reflect instantaneous traffic acceleration or deceleration; and 4) rolling window statistics, vehicular composition ratios, and binary indicators denoting peak operational periods (Rasaizadi & Seyedabrishami, 2022; Sun *et al*., 2021). The target variable was constructed by shifting the traffic condition category forward by one 15-minute interval (t+1). Following the generation of lag predictors, the dataset yielded 2,967 complete observations. To prevent temporal data leakage and replicate real-world deployment conditions, data partitioning strictly adhered to chronological order rather than random sampling (Weng *et al*., 2023). The initial 24 days (comprising 2,296 observations) were allocated for model training, feature selection, and cross-validation, whereas the final 7 consecutive days (comprising 671 observations) were reserved as an untouched holdout test set. Preserving temporal boundaries in this manner ensures that the predictive model is evaluated solely on unseen future periods (Chen *et al*., 2025).

2.3 Model Training, Hyperparameter Tuning, and Evaluation

Model training was conducted using the XGBoost classification algorithm, which utilizes gradient-boosted decision trees capable of resolving complex non-linear interactions within structured tabular features. Model selection and hyperparameter optimization were executed on the training split using a three-fold TimeSeriesSplit with an observation gap of one time step between the training and validation windows. This walk-forward validation strategy prevents boundary leakage between adjacent temporal folds (Deepika & Pandove, 2025; Lam, 2025; B. Liu *et al*., 2026). Multiple model candidates were evaluated across distinct feature subsets (basic, basic with lags, basic with lag-trends, and full feature sets) and class weighting strategies. The optimal model configuration was selected using a composite objective score calculated as:

$$Selection\ Score = 0.65 \times Accuracy + 0.25 \times Macro\ F1 + 0.10 \times Balanced\ Accuracy \quad (1)$$

Final performance assessment was executed exclusively on the 671 holdout test observations. Given the pronounced class imbalance across the four traffic conditions, model efficacy was evaluated using a comprehensive suite of metrics: overall classification accuracy, balanced accuracy, weighted F1-score, and macro-averaged F1-score (Berlotti *et al*., 2024). The tuned XGBoost model was systematically benchmarked against two baseline strategies: a majority-class baseline (which persistently predicts the dominant class) and a persistence baseline (which assumes the traffic condition at $t+1$ remains identical to time t). Furthermore, model behavior was characterized through confusion matrix analysis, class-level precision and recall metrics, prediction probability distributions, and Gini-based feature importance rankings (Moumen *et al*., 2023).

3. Results

This section presents the empirical outcomes of the study, covering dataset quality, temporal traffic dynamics, chronological data partitioning, XGBoost hyperparameter selection, classification performance across evaluation metrics, confusion matrix diagnostics, and feature importance rankings. All experimental evaluations strictly adhered to the chronological sequence of observations.

3.1 Dataset Quality and Temporal Pattern Exploration

The experimental dataset comprised 2,976 continuous observations and 9 attributes: Time, Date, Day of the week, CarCount, BikeCount, BusCount, TruckCount, Total, and Traffic Situation. Preliminary data validation confirmed that all attributes were complete, with zero missing values and zero duplicate records, as summarized in Table 1.

Table 1. Dataset Characteristics

Features	Results
Observations	2,976
Attributes	9
Missing values	0
Duplicate records	0

Target classes	4

The target variable, Traffic Situation, exhibits a pronounced multiclass imbalance across its four categories. As detailed in Table 2, the normal class constitutes more than half of the entire dataset (56.08%), followed by the heavy class (22.92%), while the high and low traffic conditions represent the minority categories, accounting for 10.79% and 10.22% of the observations, respectively.

Table 2. Traffic Class Distribution

Conditions	Quantity	Percentage
Normal	1,669	56.08%
Heavy	682	22.92%
High	321	10.79%
Low	304	10.22%
Total	2,976	100.00%

To examine the cyclic behavior of traffic conditions, temporal pattern exploration was conducted by calculating the mean total vehicular volume for each hour of the day. The resulting hourly volume trajectory is illustrated in Figure 2.

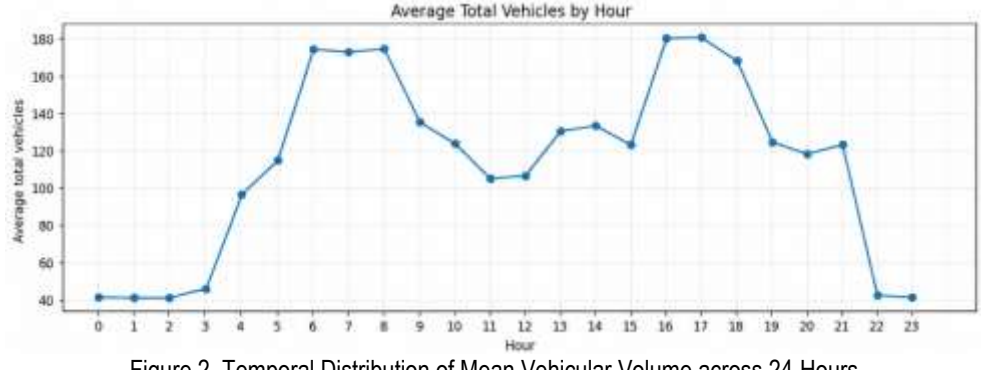


Figure 2. Temporal Distribution of Mean Vehicular Volume across 24 Hours

As depicted in Figure 2, vehicular volume exhibits pronounced diurnal variations. Traffic density remains lowest during the early morning hours (01:00 to 05:00), begins an upward trajectory during the morning commute, and peaks during late afternoon and early evening hours, specifically reaching maximum values between 16:00 and 17:00. This diurnal trajectory validates the inclusion of cyclic hour-of-day features to capture temporal variations in traffic congestion states.

3.2 Chronological Partitioning and Model Optimization

Following temporal feature extraction and target variable lagging, the final modeling dataset yielded 2,967 observations. To preserve chronological dependencies, the observations were divided into a training partition covering the initial 24 days (2,296 observations) and a holdout test partition comprising the final 7 consecutive days (671 observations). The class distributions across both partitions are detailed in Table 3.

Table 3. Chronological Data Partitioning Across Traffic Classes

Traffic Situation	Training	Testing	Total
Normal	1,291	377	1,668
Heavy	519	163	682
High	250	71	321
Low	236	60	296
Total	2,296	671	2,967

The model selection phase evaluated multiple feature configurations, class-weighting options, and hyperparameter sets using a three-fold TimeSeriesSplit on the training partition. Table 4 summarizes the top-performing candidate configurations ranked by composite selection score.

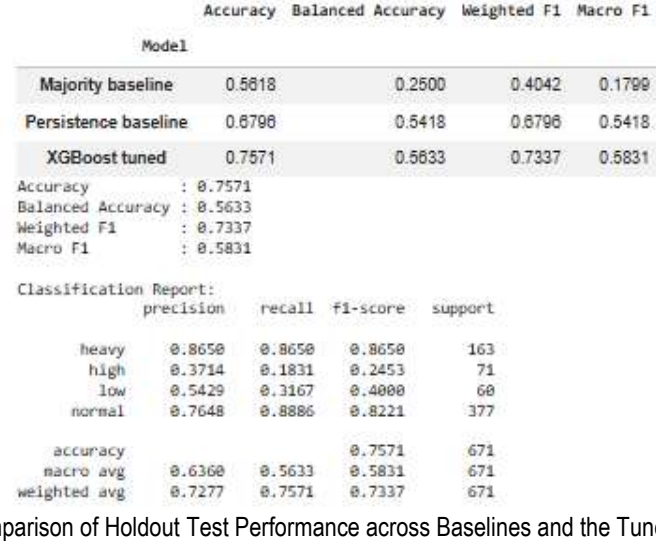
Table 4. Model and Hyperparameter Selection Results

Feature Set	Weighting	Candidate	Selection Score	Mean Accuracy	Mean Macro F1	Mean Balanced Accuracy
basic	none	2	0.687509	0.736355	0.539968	0.538870
full	none	3	0.686350	0.739257	0.530813	0.531452
basic	none	1	0.686308	0.735192	0.538474	0.531847
basic_lag	none	3	0.685284	0.732869	0.539714	0.539095
basic_lag	none	1	0.684220	0.738095	0.527178	0.526834
basic_lag	none	3	0.682937	0.737515	0.524002	0.525519
basic_lag_trend	none	3	0.682889	0.735772	0.528494	0.528131
full	none	1	0.661128	0.736934	0.519032	0.523627
full	balanced	3	0.660471	0.718351	0.552328	0.554813

The basic feature set without class weighting (Candidate 2) attained the highest composite selection score of 0.687509 and was designated as the final operational model. This configuration utilizes 14 distinct features. The selected optimal hyperparameters comprise n_estimators = 450, max_depth = 4, learning_rate = 0.035, min_child_weight = 2, subsample = 0.95, and colsample_bytree = 0.90.

3.3 Overall Model Performance and Error Analysis

The finalized XGBoost model was evaluated against the 671 unseen holdout test observations. Performance was benchmarked against a majority-class baseline and a persistence baseline across four evaluation metrics, as illustrated in Figure 3.



Model	Accuracy	Balanced Accuracy	Weighted F1	Macro F1
Majority baseline	0.5618	0.2500	0.4042	0.1799
Persistence baseline	0.6796	0.5418	0.6796	0.5418
XGBoost tuned	0.7571	0.5633	0.7337	0.5831

```
Accuracy          : 0.7571
Balanced Accuracy : 0.5633
Weighted F1       : 0.7337
Macro F1          : 0.5831

Classification Report:
              precision    recall  f1-score   support

       heavy     0.8650    0.8650    0.8650       163
        high     0.3714    0.1831    0.2453        71
         low     0.5429    0.3167    0.4000        60
      normal     0.7648    0.8886    0.8221       377

    accuracy                         0.7571       671
   macro avg     0.6360    0.5633    0.5831       671
weighted avg     0.7277    0.7571    0.7337       671
```

Figure 3. Comparison of Holdout Test Performance across Baselines and the Tuned XGBoost Model

The tuned XGBoost model achieved an overall accuracy of 75.71%, a balanced accuracy of 56.33%, a weighted F1-score of 73.37%, and a macro F1-score of 58.31%. This performance significantly outperforms the majority baseline (accuracy 56.18%) and provides a 7.75 percentage point improvement over the persistence baseline (accuracy 67.96%). At the individual class level, the model exhibited the strongest discrimination in the heavy category, achieving 86.50% across precision, recall, and F1-score. The normal class achieved a precision of 76.48%, a recall of 88.86%, and an F1-score of 82.21%. Conversely, the minority classes demonstrated lower predictive performance, with the low and high categories recording F1-

scores of 40.00% and 24.53%, respectively. To examine specific classification trajectories and error distributions, the confusion matrix on the test partition is presented in Figure 4.

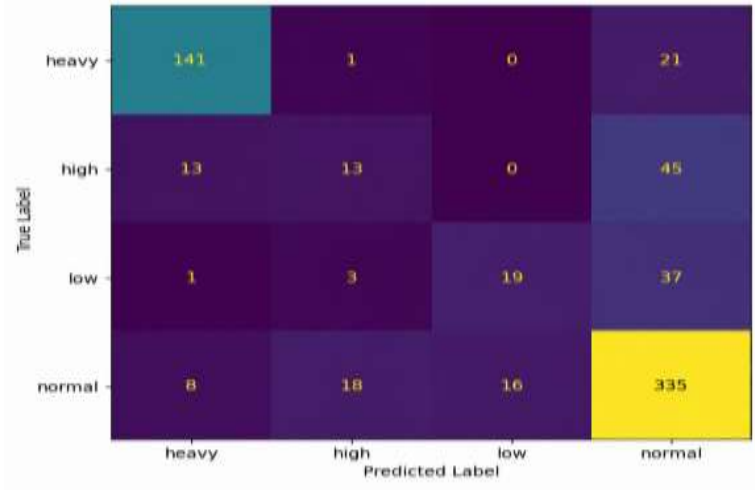


Figure 4. Confusion Matrix of the Optimized XGBoost Model on the Test Partition

As shown in the confusion matrix in Figure 4, the diagonal cells represent correct predictions across all categories. The model correctly classified 141 of 163 heavy instances, 13 of 71 high instances, 19 of 60 low instances, and 335 of 377 normal instances, yielding 508 correct predictions out of 671 total observations (75.71% overall accuracy). The primary source of classification error occurred in the high and low categories, where 45 high instances and 37 low instances were erroneously predicted as normal. For the heavy category, 21 instances were misclassified as normal. Within the normal category, 8 instances were predicted as heavy, 18 as high, and 16 as low. These error patterns demonstrate that while the model separates the dominant normal and heavy regimes effectively, distinguishing the boundary transitions for low and high traffic states remains more prone to misclassification into the majority class.

3.4 Relative Feature Importance

Feature importance analysis was conducted to quantify the relative contribution of each input variable to the decision tree splits within the final XGBoost model. Figure 5 depicts the relative importance scores of all 14 features.

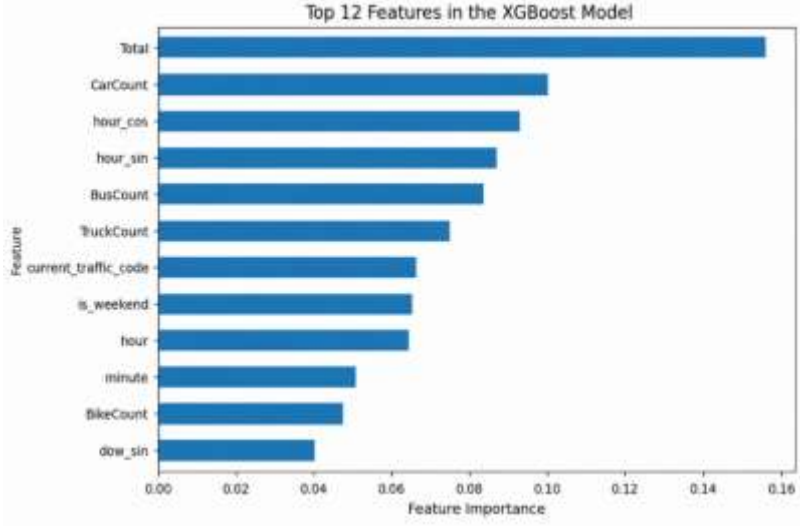


Figure 5. Relative Feature Importance Scores of the Optimized XGBoost Model

As illustrated in Figure 5, the Total vehicular count feature contributed the highest relative importance score at 0.154969, followed by CarCount (0.099857), hour_cos (0.092565), hour_sin (0.086755), and BusCount (0.083442). Secondary contributions were provided by TruckCount, current_traffic_code, is_weekend, hour, minute, BikeCount, and dow_sin. These results demonstrate that aggregate and passenger vehicular volumes, combined with cyclic harmonic time representations, provide the primary predictive signal for 15-minute forward traffic state estimation. These feature importance values represent algorithmic decision contributions and do not imply causal transport mechanisms.

4. Discussion

The empirical findings demonstrate that an XGBoost classifier integrated with engineered temporal features can effectively capture short-term traffic transitions within a forward 15-minute horizon. Achieving an overall accuracy of 75.71% and a weighted F1-score of 73.37% on the holdout test set confirms the model's capacity to recognize prevailing traffic states across sequential observation windows. Notably, the model achieved a 7.75 percentage point improvement in accuracy over the persistence baseline (67.96%). In short-term traffic monitoring, persistence represents a strong benchmark because traffic conditions naturally exhibit high temporal autocorrelation over narrow intervals (Soon *et al*., 2023). Surpassing this threshold indicates that the engineered feature set provides actionable predictive signals that anticipate impending state shifts rather than merely echoing current road occupancy. Furthermore, maintaining strict chronological ordering during cross-validation ensured that these performance gains reflect genuine prospective forecasting capability rather than optimistic artifacts caused by temporal data leakage (Chen *et al*., 2025; Weng *et al*., 2023).

Despite strong aggregate accuracy, granular class-level diagnostics reveal pronounced performance disparities across traffic categories. The model exhibited excellent discriminatory stability on the dominant heavy (F1-score 86.50%) and normal (F1-score 82.21%) regimes, but displayed markedly lower sensitivity toward the minority low (F1-score 40.00%) and high (F1-score 24.53%) classes. As evidenced by the confusion matrix, this disparity stems primarily from asymmetric boundary errors, wherein 45 high instances and 37 low instances were erroneously categorized as normal. This misclassification pattern explains the substantial divergence between the weighted F1-score (73.37%) and the unweighted macro F1-score (58.31%). Because the normal regime comprises 56.08% of the dataset, the decision trees inherently favor splits that maximize majority purity when borderline vehicular counts occur. In operational traffic environments, rapid transitions from normal to acute congestion (high) represent critical inflection points; hence, the tendency to misclassify these transitional states into the modal category highlights the fundamental challenge of managing severe class imbalance in tabular traffic streams without distorting broader classification precision (Berlotti *et al*., 2024).

The feature importance distribution provides structural insight into the predictive mechanism governing the tree ensembles. Aggregate vehicular count (Total) and passenger vehicle volume (CarCount) emerged as the dominant predictive drivers, collectively accounting for more than 25% of split importance. Concurrently, harmonic cyclical transformations (hour_cos and hour_sin) contributed substantially more predictive power than linear time indicators, confirming that continuous trigonometric representations effectively model recurring diurnal congestion waves across day-night transitions (Rasaizadi & Seyedabrishami, 2022; Sun *et al*., 2021). Modal composition variables, particularly commercial bus and truck volumes, also exerted measurable influence, likely reflecting the disproportionate impact of heavy vehicles on traffic stream impedance. While these importance rankings illustrate the internal logic of the model, they denote statistical associations rather than direct causal transport mechanisms. From an operational perspective, the lightweight nature of this feature-engineered XGBoost pipeline makes it highly amenable to edge deployment in intelligent transportation systems, provided that future iterations incorporate targeted minority-class boundary refinement.

5. Conclusion

This study demonstrates that an extreme gradient boosting (XGBoost) model structured on engineered temporal features provides an effective, lightweight approach for short-term multiclass traffic condition forecasting across a forward 15-minute horizon. By enforcing chronological data partitioning and forward-chaining TimeSeriesSplit validation, the experimental framework prevented temporal data leakage and simulated realistic operational conditions. Evaluated on a seven-day holdout test set comprising 671 observations, the proposed model achieved an overall accuracy of 75.71%, a weighted F1-score of 73.37%, a macro F1-score of 58.31%, and a balanced accuracy of 56.33%, surpassing the majority baseline (56.18%) and providing a 7.75 percentage point improvement over the persistence baseline (67.96%). While the framework exhibited robust classification stability for the dominant heavy and normal traffic regimes, its discriminatory power was constrained on the minority low and high congestion categories due to substantial class imbalance. For practical application in intelligent transportation systems, this model offers a transparent and computationally efficient baseline for automated traffic monitoring and early warning systems. Future research should expand upon these findings by extending observation windows across longer multi-seasonal periods and incorporating multi-sensor spatial dependencies alongside external environmental factors, such as adverse weather and localized incident reports. Furthermore, investigating specialized class-imbalance remedies, including cost-sensitive loss formulations and boundary-focused resampling techniques, is recommended to enhance model sensitivity during abrupt transitions into severe congestion regimes.

References

Abduljabbar, R., Dia, H., & Liyanage, S. (2025). Machine learning traffic flow prediction models for smart and sustainable traffic management. *Infrastructures*, *10*(7), Article 155. https://doi.org/10.3390/infrastructures10070155

Berlotti, M., Di Grande, S., & Cavalieri, S. (2024). Proposal of a machine learning approach for traffic flow prediction. *Sensors*, *24*(7), Article 2348. https://doi.org/10.3390/s24072348

Chen, F., Liu, X. Q., Yang, J. J., Liu, X. K., Ma, J. H., Chen, J., & Xiao, H. Y. (2025). Traffic accident severity prediction based on an enhanced MSCPO-XGBoost hybrid model. *Scientific Reports*, *15*(1), Article 25729. https://doi.org/10.1038/s41598-025-00797-7

Deepika, & Pandove, G. (2025). Prediction of traffic time using XGBoost model with hyperparameter optimization. *Multimedia Tools and Applications*, *84*(30), 101895–101918. https://doi.org/10.1007/s11042-025-20646-z

Guo, K., Shen, C., Zhou, X., Ren, S., Hu, M., Shen, M., Chen, X., & Guo, H. (2024). Traffic data-empowered XGBoost-LSTM framework for infectious disease prediction. *IEEE Transactions on Intelligent Transportation Systems*, *25*(1), 774–785. https://doi.org/10.1109/TITS.2022.3172206

Lam, K. N. (2025). Traffic prediction using LSTM, RF and XGBoost. In *Proceedings of the 6th International Conference on Data Analytics and Machine Learning (DAML 2024)* (pp. 267–274). SCITEPRESS. https://doi.org/10.5220/0013515600004619

Liu, B., Zhang, X., Liang, C., Zhang, T., & Xiang, K. (2026). XGBoost–LSTM regional traffic congestion ratio prediction integrating spatio-temporal and weather features. *IET Intelligent Transport Systems*, *20*(1), Article e70145. https://doi.org/10.1049/itr2.70145

Liu, L., Li, C., Yang, Y., & Wang, J. (2024). Short-term traffic flow forecasting based on a novel combined model. *Sustainability*, *16*(23), Article 10216. https://doi.org/10.3390/su162310216

Moumen, I., Abouchabaka, J., & Rafalia, N. (2023). Adaptive traffic lights based on traffic flow prediction using machine learning models. *International Journal of Electrical and Computer Engineering*, *13*(5), 5813–5823. https://doi.org/10.11591/ijece.v13i5.pp5813-5823

Rasaizadi, A., & Seyedabrishami, S. (2022). Stacking ensemble learning process to predict rural road traffic flow. *Journal of Advanced Transportation*, *2022*, Article 3198636. https://doi.org/10.1155/2022/3198636

Sitio, S. L. M., Darmawati, S., & Samudra, Y. (2026). Comparative analysis of bagging and boosting models in ensemble learning for graduation prediction. *JITK (Jurnal Ilmu Pengetahuan dan Teknologi Komputer)*, *11*(3), 979–986. https://doi.org/10.33480/jitk.v11i3.7579

Sitio, S. L. M., Thoyyibah, Maryani, Ardhiansyah, M., Rosmawarni, N., Djaksana, Y. M., & Arianti, N. D. (2024). Comparison of the ensemble XGBoost and Transformer models with machine learning for classification of Indonesian music moods of the 70s and 80s era. *Journal of Theoretical and Applied Information Technology*, *102*(24), 9157–9165.

Soon, K. L., Chan, R. K. C., Lim, J. M.-Y., & Parthiban, R. (2023). Short-term traffic forecasting model: Prevailing trends and guidelines. *Transportation Safety and Environment*, *5*(3), Article tdac058. https://doi.org/10.1093/tse/tdac058

Sun, B., Sun, T., & Jiao, P. (2021). Spatio-temporal segmented traffic flow prediction with ANPRS data based on improved XGBoost. *Journal of Advanced Transportation*, *2021*, Article 5559562. https://doi.org/10.1155/2021/5559562

Sun, B., Wang, Y., Shen, T., Zhang, L., & Geng, R. (2025). Heuristic feature engineering for enhancing neural network performance in spatiotemporal traffic prediction. *Computers, Materials & Continua*, *82*(3), 4219–4236. https://doi.org/10.32604/cmc.2025.060567

Wang, X., & Fang, F. (2024). Short-term traffic flow prediction based on wavelet analysis and XGBoost. *International Journal of Transportation Engineering and Technology*, *10*(1), 15–24. https://doi.org/10.11648/j.ijtet.20241001.12

Weng, J., Feng, K., Fu, Y., Wang, J., & Mao, L. (2023). Extreme gradient boosting algorithm based urban daily traffic index prediction model: A case study of Beijing, China. *Digital Transportation and Safety*, *2*(3), 220–228. https://doi.org/10.48130/DTS-2023-0018

Zhu, L., Shu, S., & Zou, L. (2022). XGBoost-based travel time prediction between bus stations and analysis of influencing factors. *Wireless Communications and Mobile Computing*, *2022*, Article 3504704. https://doi.org/10.1155/2022/3504704