



Scripta Technica: Journal of Engineering and Applied Technology

Vol 2 No 2 December 2026, Hal. 01-15
ISSN:3110-0775(Print) ISSN: 3109-9696(Electronic)
Open Access: <https://scriptainteletektual.com/scripta-technica>

Weather-Aware Machine Learning for Urban Traffic Volume Prediction and Congestion Analysis

Dilli Ram Rai^{1*}, Maudhy Satyadharma², Edward Francis G. Gabor³

¹ Jana Chetana Secondary School and pursuing, Nepal

² Dinas Perhubungan Provinsi Sulawesi Tenggara, Indonesia

³ Pamantasan ng Cabuyao, Philippines

email: nakarakhotang@gmail.com¹

Article Info :

Received:

17-09-2026

Revised:

03-10-2026

Accepted:

14-10-2026

Abstract

Hourly traffic prediction is periodic, yet weather disturbances can alter traffic demand and prediction error, creating implications for congestion management. This study evaluates weather-aware machine-learning models using a deterministic 48,204-record benchmark calibrated to the variable definitions, coverage, and scale of the UCI Metro Interstate Traffic Volume dataset. The benchmark incorporates traffic volume, temperature, rainfall, snowfall, cloud cover, weather classes, holidays, calendar variables, lagged observations, and rolling traffic means. Chronological training, validation, testing, and walk-forward evaluation were applied to prevent future information leakage and assess temporal stability. Ridge regression, Random Forest, Extra Trees, Gradient Boosting, HistGradientBoosting, and a 24-hour persistence baseline were compared using RMSE, MAE, R^2 , and sMAPE. HistGradientBoosting achieved the lowest test error, with RMSE of 364.34 vehicles/hour, MAE of 287.64 vehicles/hour, and R^2 of 0.949. Adding weather variables to calendar and traffic-history features reduced RMSE by 3.39%. Congestion detection using the 80th-percentile threshold achieved precision of 0.898, recall of 0.807, and F1 of 0.850. Errors were highest during thunderstorms and fog, while weather-sensor perturbation increased RMSE by 3.93%. The findings indicate that weather provides incremental contextual information, whereas temporal structure and recent traffic state remain dominant predictors. The calibrated reconstruction requires validation against original source records before deployment.

Keywords: Congestion, Machine Learning, Traffic Volume Prediction, Weather, Intelligent Transportation Systems.



©2026 Authors.. This work is licensed under a Creative Commons Attribution-Non Commercial 4.0 International License.
(<https://creativecommons.org/licenses/by-nc/4.0/>)

INTRODUCTION

Traffic-volume forecasting provides a short-horizon representation of road demand for signal timing, incident response, route guidance, capacity management, and network control, making reliable prediction a central component of increasingly data-driven urban transportation systems. Recent reviews indicate that the field has moved rapidly from conventional statistical and time-series approaches toward machine-learning, recurrent, graph-based, attention-based, and transformer architectures capable of representing increasingly complex temporal and spatial dependencies (Sayed et al., 2023). The expansion is closely associated with the availability of high-frequency traffic observations and the growing integration of heterogeneous information, which allow forecasting systems to move beyond simple extrapolation of historical traffic patterns. Modern reviews further show that contemporary traffic forecasting research increasingly emphasizes spatiotemporal representation learning, adaptive architectures, and hybrid models because traffic states are jointly shaped by temporal recurrence and interactions across connected road segments (Qin et al., 2026). Yet higher architectural complexity does not automatically translate into more reliable engineering decisions, particularly when forecasting systems are evaluated primarily through aggregate error metrics under relatively stable traffic conditions. For transportation agencies, the practical value of a forecasting model depends not only on predictive accuracy but also on whether its performance remains interpretable and stable when the traffic-generating environment changes. This concern places chronological validation, transparent benchmarking, and explicit treatment of exogenous conditions at the center of an engineering-oriented assessment of traffic prediction.

Temporal structure remains one of the strongest regularities in road traffic because commuting cycles, weekdays, weekends, and recurrent demand patterns generate systematic variations that can be learned from historical observations. Conventional machine-learning approaches can exploit these regularities efficiently, while newer deep-learning architectures extend the representation of temporal dependencies through attention mechanisms and multiscale sequence learning (Gomes et al., 2023). Graph-based approaches have demonstrated that traffic forecasting can also benefit from representing spatial relationships between road segments rather than treating each observation as an independent temporal sequence. A trend graph attention network, for example, explicitly incorporates evolving spatial-temporal relationships to capture traffic dynamics that conventional temporal models may overlook (Wang et al., 2023). Other developments have extended this direction through long-short-term transformers and spatial-temporal transformer architectures, reflecting a broader movement toward models capable of integrating short-term fluctuations with longer-range dependencies (Luo et al., 2024). The methodological trajectory suggests that increasingly sophisticated architectures are valuable when network topology and complex spatial interactions are central to the forecasting problem, but it also raises an important engineering question concerning whether such complexity is necessary for individual corridors or detector-level applications. Compact machine-learning workflows using historical traffic, calendar, and operational variables may remain competitive in settings where the principal objective is accurate short-horizon prediction with manageable computational requirements.

Weather represents a persistent source of non-recurrent variation that challenges forecasting systems relying predominantly on historical traffic patterns. Rain, snow, fog, temperature extremes, and reduced visibility can alter travel demand, driving behavior, effective roadway capacity, and the stability of otherwise recurrent traffic patterns, creating conditions in which historical observations become less representative of the immediate traffic state. Empirical work using arterial-road traffic data has demonstrated the relevance of combining traffic variables with meteorological information rather than relying exclusively on traffic history (Abduljabbar et al., 2024). Graph convolutional approaches have similarly incorporated meteorological information to improve hourly traffic-volume prediction, indicating that weather can contribute useful explanatory information across temporal and spatial traffic structures (Park et al., 2024). Research focused specifically on adverse-weather conditions further suggests that the contribution of meteorological variables is not necessarily constant across operating regimes, because unusual weather may become more informative precisely when recurrent traffic relationships are disrupted (Zhang et al., 2023). This pattern creates a methodological distinction between measuring average predictive improvement and determining whether weather information changes model behavior under particular environmental conditions. A model that produces only a small improvement in overall error may still be operationally relevant if the improvement is concentrated in adverse-weather periods when forecasting uncertainty and congestion-management requirements are greater.

The challenge becomes more consequential when traffic-volume forecasting is interpreted as a congestion-management problem rather than solely as a continuous regression task. A model may achieve a favorable average RMSE while systematically underestimating high-volume periods, even though these periods are precisely those associated with greater pressure on road capacity and traffic-control operations. Studies of congestion forecasting have consequently explored multivariate and neural-network approaches that treat congestion as an outcome requiring attention to nonlinear traffic relationships rather than simple temporal extrapolation (Fahs et al., 2023). Research on traffic congestion prediction also demonstrates the growing use of deep-learning frameworks to represent complex relationships among traffic variables, reinforcing the distinction between predicting a numerical traffic state and identifying an operationally meaningful congestion condition (Qi & Cheng, 2023). Explainability research has further shifted attention toward understanding which traffic characteristics contribute to predictions, since black-box accuracy alone provides limited information for diagnosing model behavior or supporting engineering interpretation (Guo et al., 2024). These developments reveal an unresolved connection between continuous traffic-volume forecasting and event-oriented congestion assessment. Treating the two outputs as separate analytical problems may overlook whether a regression model's numerical forecasts can actually support reliable identification of high-demand periods. An integrated evaluation therefore requires both error-based regression metrics and decision-oriented measures such as precision and recall for congestion-relevant traffic states.

Despite substantial methodological progress, the literature remains fragmented across three analytical dimensions: model architecture, weather sensitivity, and operational interpretation. Many studies prioritize architectural innovation, making direct assessment of incremental feature value difficult because improvements are often attributed to increasingly complex representations rather than systematically isolated exogenous variables. Other studies demonstrate the usefulness of meteorological variables but focus on particular weather scenarios or specialized deep-learning architectures, leaving uncertainty about whether comparable benefits persist within a compact and reproducible machine-learning benchmark. The rapid emergence of transformer, attention, and other advanced architectures has also intensified the need for transparent comparisons against simpler baselines, particularly when the intended application concerns a single traffic corridor rather than a large interconnected network (Fafoutellis & Vlahogianni, 2023). At the same time, recent machine-learning studies show that traffic prediction can be approached with comparatively accessible models, suggesting that computational efficiency and reproducibility remain relevant criteria for applied transportation engineering (Berlotti et al., 2024). The literature consequently provides strong evidence that traffic is predictable, substantial evidence that weather contains additional information, and growing evidence that model interpretability matters, but these strands are less consistently integrated within a single evaluation framework. The unresolved empirical issue is not simply which algorithm obtains the lowest aggregate error, but whether weather information provides measurable and regime-specific value and whether that value translates into more reliable identification of high-congestion periods. Addressing this issue is particularly important for engineering applications in which a forecasting system must remain useful when traffic conditions deviate from their dominant historical patterns.

The present study is positioned within this gap by evaluating weather-aware machine learning as an integrated forecasting and congestion-analysis workflow rather than proposing another highly specialized predictive architecture. The analysis separates three questions that are often treated independently: whether weather improves hourly traffic-volume prediction beyond traffic history and calendar structure, whether predictive quality changes systematically under adverse weather, and whether continuous forecasts can identify high-congestion hours with useful precision and recall. The empirical benchmark is anchored to the public Metro Interstate Traffic Volume dataset, which provides traffic volume observations together with weather and temporal information suitable for examining these relationships (Hogue, 2019). Methodological transparency is maintained through a deterministic calibrated reconstruction designed to match the documented variable definitions, record count, time span, and observed scale of the source benchmark rather than presenting the reconstruction as a verbatim reanalysis of the original row-level file. Five regression models are evaluated against a persistence baseline using chronological testing, while feature ablation quantifies the incremental contribution of weather variables and weather-specific error analysis identifies conditions associated with changes in predictive performance. Permutation analysis is used to examine feature hierarchy, top-quintile congestion is framed as a decision problem, and sensor perturbation together with walk-forward testing provides complementary evidence concerning robustness. The study therefore aims to establish a reproducible engineering benchmark for weather-aware urban traffic-volume prediction and congestion analysis, with theoretical contribution arising from the explicit linkage of temporal demand regularity, meteorological disturbance, and congestion-oriented interpretation, and methodological contribution arising from a compact evaluation framework that separates average accuracy from weather-conditioned reliability and operational classification performance.

RESEARCH METHODS

This study employed an empirical machine-learning benchmarking design to evaluate weather-aware urban traffic-volume prediction and its subsequent application to congestion analysis. The reference benchmark was the Metro Interstate Traffic Volume dataset from the UCI Machine Learning Repository, comprising 48,204 hourly observations of westbound I-94 traffic recorded at Minnesota Department of Transportation ATR station 301 between 2 October 2012 and 30 September 2018, with traffic, weather, holiday, and timestamp variables (Hogue, 2019). The computational benchmark was implemented as a deterministic calibrated reconstruction generated under a fixed random seed, reproducing the documented record count, observation period, variable definitions, realistic meteorological ranges, weekday/weekend structure, peak-period behavior, holiday effects, and adverse-weather demand patterns, while explicitly distinguishing generated observations from the original row-

level UCI records. The feature-engineering pipeline encoded hour of day, day of week, and day of year using sine-cosine transformations and added weekend and holiday indicators, while traffic-history predictors comprised 1-, 2-, 3-, 24-, 48-, and 168-hour lags and 6-, 24-, and 168-hour rolling means. Weather predictors consisted of temperature, one-hour rainfall, one-hour snowfall, cloud cover, and binary indicators representing the principal weather classes, producing 48,036 usable observations after removal of records lacking sufficient warm-up history. This configuration follows the established machine-learning direction in traffic forecasting while retaining a comparatively compact feature-based architecture suitable for reproducible engineering applications rather than requiring a large spatiotemporal neural network (Abduljabbar et al., 2025).

The modeling procedure compared six forecasting configurations: a 24-hour persistence baseline, Ridge regression, Random Forest, Extra Trees, Gradient Boosting, and HistGradientBoosting, with hyperparameters fixed before final testing. The chronological partition assigned 33,935 observations to training, 8,067 to validation, and 6,034 to testing, preserving the temporal direction of forecasting and preventing future observations from entering model fitting. A walk-forward procedure was also conducted by refitting the selected model using only observations available before each evaluation period. Model performance was assessed using RMSE, MAE, R^2 , and sMAPE, with RMSE and MAE emphasized because percentage-based errors can become unstable at low nighttime traffic volumes. Chronological validation was selected because random train-test shuffling may produce overly optimistic performance estimates in temporally ordered traffic data (Liu & Shin, 2025). Feature ablation compared calendar-only, calendar-plus-traffic-history, weather-plus-calendar, and full weather-aware configurations, while congestion classification used the training-set 80th percentile threshold and was evaluated using precision, recall, and F1, complemented by weather-specific errors, permutation importance, and controlled weather-sensor perturbation tests (Zhang et al., 2024).

RESULTS AND DISCUSSION

Traffic and Weather Structure

The calibrated benchmark exhibited a pronounced temporal structure in traffic demand, with substantial dispersion around the central tendency. Mean traffic volume reached 2,139.65 vehicles/hour, while the median was 2,166 vehicles/hour, indicating that the overall distribution was centered near the typical operating level despite substantial variation between low- and high-demand periods. The 95th percentile reached 4,850 vehicles/hour and the maximum reached 6,268 vehicles/hour, confirming the presence of high-volume operating states that are materially different from the average condition. Such dispersion is consistent with the recurrent and non-recurrent components commonly identified in traffic-flow forecasting, where temporal regularity provides a strong predictive signal but does not fully describe fluctuations in observed demand (Gomes et al., 2023). The temperature distribution was comparatively broad, with a mean of 281.21 K and a maximum of 310.15 K, while cloud cover averaged 48.65% and reached 100%. Rainfall and snowfall were strongly concentrated near zero, suggesting that precipitation-related observations represented relatively infrequent deviations from predominantly dry conditions.

The descriptive profile of the principal variables is presented in Table 1, providing the numerical basis for interpreting subsequent model performance and weather-conditioned behavior. The traffic-volume standard deviation of 1,574.96 vehicles/hour was large relative to the mean, reflecting the substantial contrast between nighttime, off-peak, and commuter-period demand. The 95th-percentile traffic level was more than twice the mean, indicating that high-demand observations form a distinct operating regime rather than a minor extension of ordinary flow. This structure is particularly relevant for machine-learning models because nonlinear partitioning can distinguish combinations of temporal and meteorological conditions that would be difficult to represent through a single linear relationship. Previous traffic-prediction research similarly demonstrates that machine-learning performance depends strongly on the representation of recurrent temporal patterns and the availability of contextual predictors (Berlotti et al., 2024). The descriptive statistics therefore establish a benchmark in which traffic history and temporal variables are expected to provide substantial predictive information, while weather variables may contribute primarily to deviations from recurrent demand patterns.

Table 1. Descriptive Statistics of the 48,204-Record Calibrated Benchmark

Variable	Unit/type	Mean	SD	Median	P95	Max
Traffic volume	veh/h	2,139.65	1,574.96	2,166.00	4,850.00	6,268.00
Temperature	K	281.21	10.36	281.21	297.41	310.15
Rainfall	mm/h	0.15	0.69	0.00	1.08	13.11
Snowfall	mm/h	0.01	0.07	0.00	0.00	3.37
Cloud cover	%	48.65	26.07	49.00	95.00	100.00

Weather-class frequencies revealed a highly uneven distribution across the benchmark, which is important when interpreting both mean traffic levels and later weather-specific prediction errors. Clouds represented 61.30% of all observations, whereas clear conditions accounted for 19.36%, making these two classes the dominant operating environments. Mist contributed 8.02% of observations and rain 7.27%, while snow, fog, drizzle, and thunderstorms occurred much less frequently. The imbalance means that aggregate model performance will be influenced primarily by common weather states and may conceal behavior under comparatively rare adverse conditions. Research incorporating meteorological information into traffic prediction likewise indicates that weather effects can become more operationally relevant under particular conditions even when their contribution to aggregate accuracy remains limited (Abduljabbar et al., 2024). A meaningful interpretation therefore requires separating the frequency of each weather regime from its associated traffic-volume level.

Table 2 further shows that traffic volume declined across several adverse-weather categories, although the magnitude varied by class. Cloudy conditions produced a mean volume of 2,156.2 vehicles/hour, compared with 2,141.2 under clear conditions and 2,169.8 under mist, indicating that common weather states did not generate large differences in average demand. Rain reduced mean volume to 2,081.9 vehicles/hour, while snow, fog, and drizzle were associated with means below 2,000 vehicles/hour. Thunderstorm observations showed the lowest mean volume at 1,497.5 vehicles/hour, but the class contained only 195 records, limiting the precision with which this difference can be generalized. Comparable research has reported that meteorological variables can influence traffic volume through changes in travel behavior and roadway conditions, with the magnitude of the effect depending on weather severity and traffic context (Kendre et al., 2024). The observed pattern consequently supports treating weather as a conditional source of variation rather than assuming a uniform reduction in traffic across all adverse-weather observations.

Table 2. Traffic Volume by Weather Class

Weather	n	Share (%)	Mean Volume	Median Volume
Clouds	29,548	61.30	2156.2	2189
Clear	9,332	19.36	2141.2	2166
Mist	3,867	8.02	2169.8	2208
Rain	3,506	7.27	2081.9	2094
Snow	729	1.51	1999.9	2078
Fog	514	1.07	1970.6	2041
Drizzle	513	1.06	1938.2	1869
Thunderstorm	195	0.40	1497.5	1434

The temporal profile provides a further explanation for the large variance observed in the benchmark. As illustrated in Figure 1, traffic volume followed two pronounced commuter-related peaks, with mean volume exceeding 4,300 vehicles/hour around 07:00–08:00 and reaching 4,666.6 vehicles/hour at 16:00. Nighttime demand showed the opposite pattern, falling below 400 vehicles/hour between midnight and 03:00. The difference between these periods is substantially larger than the mean-volume differences observed across the common weather classes, indicating that time of day provides a stronger first-order description of traffic demand. Temporal periodicity is also a central component of contemporary traffic forecasting models, including approaches based on graph attention and long-short-term temporal representations (Wang et al., 2023). The strong diurnal structure provides an important

baseline against which the incremental contribution of weather can be evaluated in the subsequent forecasting experiments.




Figure 1. Mean Hourly Traffic-Volume Profile in the Test Period

The concentration of traffic demand around recurring commuting windows also clarifies why weather variables should not be interpreted independently of temporal context. A rainfall observation during a low-demand nighttime period does not represent the same operational state as rainfall occurring during a morning or afternoon peak. This interaction is consistent with research showing that meteorological information can improve traffic prediction when integrated with temporal traffic features rather than treated as an isolated explanatory block (Park et al., 2024). The benchmark structure consequently suggests a hierarchy of information in which recurrent temporal signals establish the expected traffic level and weather variables modify that expectation under specific environmental conditions. This interpretation is important for the later feature-ablation and permutation analyses because a relatively small marginal contribution from weather does not indicate that meteorological information lacks operational relevance. Its contribution may instead become more apparent when recurrent traffic patterns are disturbed or when high-demand periods coincide with adverse environmental conditions. The subsequent forecasting analysis therefore examines whether the descriptive differences identified here translate into measurable changes in predictive accuracy and whether weather-related information provides additional value beyond established temporal traffic signals.

Comparative Forecasting Performance and Feature Ablation

The comparative evaluation indicates that nonlinear ensemble models captured the traffic-volume structure more effectively than both the persistence benchmark and the regularized linear specification. HistGradientBoosting produced an RMSE of 364.34 vehicles/hour, an MAE of 287.64 vehicles/hour, and an R^2 of 0.949 on the held-out test set. Extra Trees and Random Forest generated closely comparable results, with R^2 values of 0.947 and RMSE values of 371.71 and 373.03 vehicles/hour, respectively. Gradient Boosting also achieved strong predictive performance, although its RMSE of 394.87 vehicles/hour was higher than those of the three leading tree-based models. Ridge regression performed less effectively, with an RMSE of 586.10 vehicles/hour and R^2 of 0.868, indicating that a linear mapping was less capable of representing the interactions embedded in the engineered feature space. The 24-hour persistence baseline was substantially weaker, producing an RMSE of 1,561.45 vehicles/hour and R^2 of only 0.066. This performance hierarchy is consistent with evidence that machine-learning traffic forecasting benefits from nonlinear representations capable of integrating temporal dependencies and heterogeneous predictors (Abduljabbar et al., 2025).

The complete held-out comparison is reported in Table 3, where the relatively small differences among the strongest tree ensembles are also evident. HistGradientBoosting reduced RMSE by 76.67% relative to 24-hour persistence and by 37.84% relative to Ridge regression. Its MAE of 287.64 vehicles/hour indicates that the improvement was not driven solely by a small number of extreme observations but was also reflected in typical absolute prediction errors. The comparable results of Extra Trees and Random Forest suggest that the predictive advantage is associated more broadly with

nonlinear ensemble learning than with a unique property of a single algorithm. This pattern is relevant to compact engineering applications because it indicates that substantial forecasting performance can be obtained without adopting highly complex spatiotemporal architectures. Recent reviews similarly describe tree-based machine-learning methods as useful alternatives when traffic forecasting tasks involve structured temporal and contextual predictors but do not require network-wide representation learning (Liu & Shin, 2025).

Table 3. Held-out 2018 Forecasting Performance

Model	RMSE	MAE	R ²	sMAPE (%)
HistGradient Boosting	364.34	287.64	0.949	36.15
Extra Trees	371.71	292.52	0.947	36.70
Random Forest	373.03	293.34	0.947	36.60
Gradient Boosting	394.87	310.85	0.940	36.71
Ridge	586.10	455.33	0.868	46.80
Persistence-24h	1561.45	1144.22	0.066	74.69

Figure 2 provides a visual comparison of test RMSE and reinforces the separation between the persistence baseline, Ridge regression, and the tree-based ensemble models. The three strongest ensemble models occupy a relatively narrow error range, whereas the persistence model exhibits an error more than four times that of HistGradientBoosting. The large persistence error indicates that traffic demand cannot be adequately represented by assuming that the volume at the same hour on the previous day remains unchanged. This result is particularly important because the benchmark contains strong daily recurrence, yet recurrence alone does not account for day-to-day changes in demand and environmental conditions. Advanced temporal models address related limitations through recurrent representations, attention mechanisms, and spatiotemporal transformations that learn dependencies beyond simple lag repetition (Tran et al., 2024). The present results show that explicit lag features combined with boosted-tree learning can capture a substantial proportion of this structure within a considerably simpler modeling framework.




Figure 2. Test RMSE Across Forecasting Models

Feature ablation further clarifies the sources of the observed predictive performance. The calendar-only configuration achieved an RMSE of 426.75 vehicles/hour and R² of 0.930, demonstrating that hour, day-of-week, weekend, and holiday information already captured a substantial portion of recurring traffic variation. Adding traffic-history predictors reduced RMSE to 377.14 vehicles/hour and increased R² to 0.946, confirming the additional value of recent and seasonal traffic states. The weather-

plus-calendar configuration achieved an RMSE of 409.74 vehicles/hour, while the full weather-aware model achieved 364.34 vehicles/hour, indicating that weather information was most useful when combined with traffic-history variables. The incremental improvement is therefore not simply attributable to the presence of more predictors, but to the interaction between recurrent temporal information, observed traffic history, and environmental context. This pattern corresponds with heterogeneous-data traffic forecasting research in which external variables complement, rather than replace, temporal traffic information (Wang & Susanto, 2024). The ablation results consequently provide a more direct assessment of weather contribution than the overall model comparison alone.

The numerical differences among feature configurations are summarized in Table 4, with the full model reducing RMSE by 3.39% relative to the calendar-plus-traffic-history specification. Weather plus calendar variables also reduced RMSE by 3.99% relative to calendar-only prediction, showing that meteorological information contained predictive information even without traffic-history variables. However, the strongest configuration remained the combination of calendar, traffic-history, and weather predictors, with RMSE 364.34 vehicles/hour and MAE 287.64 vehicles/hour. The relatively modest magnitude of the weather contribution should be interpreted alongside the much larger contribution of temporal and lagged traffic variables rather than as evidence that weather is uniformly unimportant. Traffic forecasting studies incorporating meteorological variables have likewise reported small aggregate gains that can coexist with greater relevance under specific environmental conditions (Kendre et al., 2024). The ablation evidence therefore supports a complementary role for weather, particularly within a model that has already established the recurrent traffic state.

Table 4. Feature-Set Ablation Using HistGradient Boosting

Feature Set	Features	RMSE	MAE	R ²
Calendar only	8	426.75	339.94	0.930
Calendar + traffic history	17	377.14	295.44	0.946
Weather + calendar only	19	409.74	327.85	0.936
Full weather-aware model	28	364.34	287.64	0.949

The temporal behavior of the selected HistGradientBoosting model is further illustrated by the representative test-week trace in Figure 3. The predicted series follows the recurring morning and afternoon peaks while avoiding the systematic one-day displacement that characterizes the persistence benchmark. Deviations remain visible around abrupt transitions into and out of peak periods, indicating that lagged traffic variables do not fully anticipate rapid changes in demand. Such residual behavior is consistent with the broader literature, where richer temporal architectures have been developed to capture short-term dynamics and multi-scale dependencies (Chu et al., 2024). Graph-based and attention-based approaches can additionally incorporate relationships among locations, which are not represented in the present single-corridor feature architecture (Tian et al., 2023). The current result therefore demonstrates the effectiveness of a compact feature-based approach within its intended single-station setting while leaving network-level spatial dependencies outside the benchmark scope.




Figure 3. Observed and Predicted Traffic Volume Over a Representative Test Week

The combined model comparison and ablation findings establish that forecasting performance arose from layered information rather than from weather variables alone. Calendar features provided the dominant recurrent structure, traffic-history variables captured short- and long-range persistence, and meteorological variables supplied additional contextual information that reduced residual error. The difference between the full model and the calendar-plus-traffic-history configuration was relatively small in aggregate terms, but its operational meaning cannot be determined from pooled RMSE alone. This issue becomes particularly relevant when adverse weather is concentrated in a small proportion of observations, because improvements or degradations within those regimes can be diluted by common-weather observations. Recent weather-aware traffic studies similarly emphasize the importance of evaluating meteorological effects beyond a single aggregate performance statistic (Park et al., 2024). The next analysis therefore disaggregates test errors by weather class to determine whether the modest overall weather contribution masks larger differences under specific environmental conditions.

Weather-Conditioned Error and Congestion Detection

Weather-conditioned evaluation revealed that predictive error varied across environmental regimes rather than remaining constant across the test set. Thunderstorm conditions produced the highest RMSE at 523.55 vehicles/hour and the highest MAE at 412.73 vehicles/hour. Fog and rain followed with RMSE values of 384.81 and 379.24 vehicles/hour, respectively, whereas mist and drizzle produced lower errors of 334.17 and 328.98 vehicles/hour. Clear and cloudy conditions generated errors close to the pooled test performance, suggesting that the model was more stable under common weather states. The pattern indicates that adverse weather can introduce deviations that are less predictable from recurrent traffic information, particularly when environmental conditions alter normal demand or traffic-flow behavior. Similar heterogeneity has been reported in weather-aware traffic forecasting, where predictive performance can deteriorate under adverse conditions even when meteorological variables improve overall forecasting accuracy (Zhang et al., 2023). The results therefore provide empirical support for evaluating weather-specific errors rather than relying exclusively on aggregate RMSE and MAE.

The weather-specific results are detailed in Table 5, which also shows substantial differences in the number of observations available for each class. Clouds contributed 3,685 test observations, whereas thunderstorms contributed only 28, making the extreme-weather estimates substantially less stable statistically. The relatively high thunderstorm error should consequently be interpreted as evidence of a difficult prediction regime within this benchmark rather than as a precise population estimate of thunderstorm behavior. Rain and snow also produced higher errors than clear and cloudy conditions, although their sample sizes were considerably larger than the thunderstorm class. This distinction is important because class frequency and predictive difficulty are separate dimensions of model evaluation. Weather-sensitive traffic studies similarly emphasize that adverse-weather observations may be less frequent while carrying disproportionate implications for forecasting reliability (Zhang et al., 2024). The concentration of errors in these less common regimes motivates a separate assessment of whether the continuous forecasts can support identification of high-volume congestion states.

Table 5. Test Error by Weather Class

Weather	n	Mean Volume	MAE	RMSE
Clouds	3685	2238.4	288.28	364.27
Clear	1193	2115.6	289.51	366.69
Mist	494	2175.2	258.80	334.17
Rain	398	2228.3	299.63	379.24
Snow	106	2048.0	297.42	372.92
Fog	68	2104.7	312.75	384.81
Drizzle	62	2116.4	265.48	328.98
Thunderstorm	28	2083.8	412.73	523.55

The congestion analysis translated continuous traffic-volume prediction into an operationally interpretable classification task. The threshold was fixed at the training-set 80th percentile of 3,591 vehicles/hour, preventing information from the held-out test observations from determining the congestion boundary. Applying this threshold to the model predictions yielded 1,089 true positives, 124 false positives, 261 false negatives, and 4,560 true negatives. The resulting precision was 0.898, recall was 0.807, and F1 was 0.850, indicating that the predicted high-volume events corresponded closely with observed congestion states while some events remained undetected. The relatively higher precision than recall reflects an asymmetric error pattern in which false congestion alarms were less frequent than missed congestion hours. Event-oriented traffic forecasting research similarly treats congestion identification as a distinct operational problem because aggregate regression accuracy does not directly indicate the reliability of traffic-state classification (Fahs et al., 2023).

The numerical classification results are presented in Table 6, where the relationship between continuous prediction and congestion detection can be examined directly. The 124 false positives indicate that the model generally avoided labeling moderate-volume observations as congestion when they did not exceed the predefined threshold. At the same time, the 261 false negatives demonstrate that some high-volume periods remained below the predicted decision boundary. This pattern is particularly relevant when the forecasting system is intended to support traffic management because the operational consequences of a false alarm and a missed congestion event may differ. A threshold calibrated for one objective may therefore not be optimal for another, especially when traffic agencies place different costs on unnecessary intervention and delayed response. Congestion forecasting literature has emphasized similar distinctions between numerical prediction quality and event-level forecasting utility (Qi & Cheng, 2023). The F1 value of 0.850 provides a compact summary of this classification trade-off without replacing the underlying precision and recall information.

Table 6. Top-Quintile Congestion Detection on the Held-Out Test Set

Threshold (veh/h)	TP	FP	FN	TN	Precision	Recall	F1
3591	1089	124	261	4560	0.898	0.807	0.850

The distribution of observed congestion across hours further demonstrates that high-volume conditions were strongly concentrated within recurring commuting periods. Congestion rates exceeded 73% at 07:00 and 08:00, while the afternoon period showed even greater concentration, reaching 79.4% at 15:00 and 90.5% at 16:00. The rate remained high at 17:00, where 83.8% of observations exceeded the congestion threshold, before declining sharply at 18:00 and reaching 0.0% at 19:00. These differences correspond closely with the diurnal traffic profile established in Section 3.1 and indicate that congestion classification is strongly conditioned by recurrent demand cycles. Machine-learning congestion studies likewise recognize temporal traffic structure as a major determinant of high-volume states, while additional contextual variables can modify the intensity and timing of those states (Qi & Cheng, 2023). Weather should therefore be interpreted within this temporal structure rather than as an independent determinant of congestion occurrence.

Table 7 summarizes the selected hourly congestion patterns and shows that the highest congestion rate coincided with the highest mean traffic volume. At 16:00, the mean volume reached 4,666.6 vehicles/hour and 90.5% of observations were classified as congestion, while the corresponding rate at 17:00 remained 83.8% with a mean volume of 4,423.0 vehicles/hour. The morning peak was also pronounced, with congestion rates of 75.7% and 73.9% at 07:00 and 08:00, respectively. In contrast, the 19:00 period recorded a mean volume of 1,908.6 vehicles/hour and no congestion observations under the selected threshold. This concentration suggests that weather-related forecasting errors may become operationally consequential when adverse conditions overlap with already elevated baseline demand. Studies of weather-aware prediction similarly indicate that meteorological effects should be evaluated in relation to traffic regimes because the same environmental disturbance may have different implications under different demand conditions (Park et al., 2024).

Table 7. Congestion Concentration During Selected Hours

Hour	n	Mean volume	Congestion Rate (%)
6	259	3028.7	21.2
7	259	4306.0	75.7
8	261	4300.4	73.9
9	257	3328.1	38.9
14	257	3271.4	26.5
15	248	4076.2	79.4
16	242	4666.6	90.5
17	253	4423.0	83.8
18	250	3310.4	38.0
19	254	1908.6	0.0

Taken together, the weather-conditioned and congestion analyses indicate that the value of weather information is concentrated in specific prediction regimes rather than uniformly distributed across the benchmark. The strongest continuous prediction errors occurred under thunderstorms, fog, and rain, while congestion was concentrated in morning and afternoon commuting windows. This combination implies that the most consequential forecasting situations may occur when environmental disturbances coincide with periods of high baseline traffic demand. The classification results also show that the selected model can identify a substantial proportion of high-volume hours, although the presence of false negatives indicates that continuous prediction error remains relevant to operational decision-making. Similar work on congestion forecasting and traffic-flow prediction supports combining numerical accuracy with event-level metrics when model outputs are intended for traffic-management applications (Fahs et al., 2023). The next analysis therefore examines which engineered features drive prediction most strongly and whether the model maintains its performance when evaluated across successive temporal windows and under controlled degradation of weather inputs.

Feature Importance, Temporal Stability, and Robustness Tests

The permutation analysis revealed a pronounced hierarchy among the predictors used by the HistGradientBoosting model. As shown in Table 10, hour_cos generated the largest increase in MAE at 1,109.26 vehicles/hour, followed by lag_1 at 258.71 vehicles/hour and hour_sin at 79.48 vehicles/hour. The magnitude of the hour_cos contribution indicates that the model relied strongly on the recurrent intraday traffic cycle represented by the sine-cosine calendar encoding. The importance of lag_1 further demonstrates that the most recent traffic condition provided substantially more predictive information than individual meteorological variables. Temporal dependence is a central component of short-term traffic forecasting because traffic volume typically exhibits recurrent patterns across adjacent time periods (Tao et al., 2022). The relatively compact feature hierarchy indicates that strong predictive performance can be achieved through explicit temporal representation without requiring a high-dimensional spatial architecture.

Table 8. Leading Permutation Importance Values

Feature	Increase in MAE	SD
hour_cos	1109.26	12.64
lag_1	258.71	6.35
hour_sin	79.48	3.63
dow_sin	48.78	3.24
weekend	15.82	1.85
is_holiday	13.98	2.53
roll_mean_6	9.04	1.99
snow_1h	6.85	0.47
rain_1h	3.64	0.75
lag_2	2.74	0.38

lag_3	1.50	0.77
wx_thunderstorm	1.30	0.41

The lower permutation contribution of meteorological variables does not indicate that weather information lacks predictive relevance. Snowfall and rainfall generated MAE increases of 6.85 and 3.64 vehicles/hour, respectively, demonstrating measurable incremental information after temporal predictors were incorporated. The relatively smaller values indicate that weather variables primarily modified an already established traffic state rather than independently determining hourly demand. This interpretation is consistent with research showing that weather information can contribute to traffic prediction while remaining secondary to traffic-history variables in aggregate forecasting performance (Abduljabbar et al., 2024). The distinction is important because permutation importance measures the marginal contribution of a predictor within a fitted feature system rather than its physical importance under every traffic condition. Explainable machine-learning approaches similarly emphasize separating dominant predictors from contextual factors when interpreting transportation models (Nesa & Yoon, 2024).

The large separation between hour_cos and the remaining predictors indicates that the recurrent daily cycle constituted the principal decision signal for traffic-volume estimation. The lag_1 variable occupied the second position, confirming that immediate traffic conditions complemented the broader calendar pattern. This structure helps explain the high predictive accuracy because the benchmark contained pronounced and recurrent morning and afternoon traffic peaks. Attention-based and transformer-based forecasting architectures have similarly been developed to capture temporal dependencies that are difficult to represent through conventional linear specifications (Luo et al., 2024). In the present single-corridor configuration, explicit temporal features provided an interpretable mechanism for representing these dependencies while maintaining a comparatively compact model structure.

Temporal stability was examined through the walk-forward procedure to determine whether model performance remained consistent across successive evaluation periods. Table 11 shows that the model produced RMSE values of 357.60 and 362.14 vehicles/hour for the two evaluation periods, while MAE values were 284.99 and 285.92 vehicles/hour. The corresponding R^2 values remained at 0.950 across both periods, indicating minimal deterioration in explanatory performance. The difference in RMSE between the two periods was only 4.54 vehicles/hour, suggesting that the engineered temporal representation remained relatively stable as the training window expanded. Chronological validation provides a more realistic assessment for sequential traffic data because random partitioning can allow temporal dependence between training and testing observations to distort estimated generalization performance (Liu & Shin, 2025). The consistency across successive periods therefore provides evidence that the forecasting configuration was not strongly dependent on one particular temporal partition.

Table 9. Walk-Forward Temporal Validation

Test year	Train n	Test n	RMSE	MAE	R^2
2017	33,935	8,067	357.60	284.99	0.950
2018	42,002	6,034	362.14	285.92	0.950

The limited variation between the two walk-forward evaluations also indicates that the model did not experience substantial temporal degradation within the evaluated benchmark. RMSE increased slightly, whereas MAE changed by less than one vehicle/hour and R^2 remained unchanged. Such stability is relevant to engineering deployment because forecasting models must operate on observations that occur after the period used for model development. Adaptive and lightweight temporal architectures have been proposed to address evolving traffic patterns while retaining computational efficiency (Zhang et al., 2025). The present results suggest that explicit calendar encoding combined with short-term traffic history can provide a stable representation in a single-corridor forecasting environment. The finding is particularly relevant because the model does not depend exclusively on meteorological variables whose availability may vary across operational settings. Broader traffic forecasting research

also emphasizes the importance of temporal adaptability when models are transferred across different observation periods (Qin et al., 2026).

Robustness testing assessed whether degradation of meteorological inputs substantially altered forecasting performance. Table 12 shows that RMSE increased from 364.34 to 378.67 vehicles/hour as the perturbation level increased from 0.00 to 1.00, while MAE increased from 287.64 to 298.31 vehicles/hour. The corresponding R^2 declined only from 0.949 to 0.945, indicating that the model retained a high proportion of its predictive performance under progressively stronger weather-data disturbance. The maximum RMSE increase was 14.33 vehicles/hour, equivalent to approximately 3.93% relative to the unperturbed condition. This limited degradation is consistent with the feature-importance results, where temporal variables contributed considerably more predictive information than individual weather predictors. Weather-aware traffic forecasting studies have nevertheless demonstrated that adverse meteorological conditions can introduce distinct prediction difficulties that justify explicit weather modeling (Zhang et al., 2023). The robustness result therefore indicates that incorporating weather information does not make the forecasting system excessively dependent on perfect meteorological observations.

Table 10. Weather-Sensor Perturbation Robustness

Perturbation Level	RMSE	MAE	R^2
0.00	364.34	287.64	0.949
0.25	374.34	295.35	0.946
0.50	377.37	297.36	0.945
1.00	378.67	298.31	0.945

The combined results establish that the forecasting system derives its primary predictive strength from temporal traffic information while retaining measurable sensitivity to meteorological conditions. This configuration differs conceptually from network-scale approaches that depend on extensive spatial representations and graph structures to model interactions among multiple road segments. Graph-based traffic forecasting can provide richer spatial representations when dense network-level observations are available (Tian et al., 2023). For a single-corridor engineering application, however, the present feature-based configuration offers a direct relationship between model inputs, predictive performance, and operational robustness. The modest performance loss under weather perturbation further indicates that the system can tolerate imperfect exogenous measurements without a major collapse in predictive accuracy. Resource-efficient traffic forecasting is increasingly relevant because model complexity must be balanced against computational and data requirements during deployment (Yaqub et al., 2025). The resulting architecture therefore provides an interpretable and relatively robust framework for combining recurrent traffic signals with weather information in hourly urban traffic-volume prediction.

CONCLUSION

The weather-aware machine-learning workflow demonstrated that a compact feature-based forecasting architecture can achieve high predictive accuracy for hourly traffic volume while retaining sensitivity to meteorological conditions. On the 48,204-record calibrated benchmark, HistGradientBoosting produced an RMSE of 364.34 vehicles/hour, an MAE of 287.64 vehicles/hour, and an R^2 of 0.949 on the held-out test period. The feature-ablation results showed that incorporating weather variables reduced RMSE by 3.39% beyond calendar and traffic-history predictors, indicating an incremental rather than dominant contribution of meteorological information. Weather-conditioned analysis further showed greater prediction errors during thunderstorms and fog, suggesting that atypical weather regimes remain more difficult to estimate from recurrent traffic patterns. Permutation analysis confirmed that hour-of-day representations and the most recent traffic observation were the principal predictive features. The resulting forecasts also provided useful congestion detection, achieving precision of 0.898, recall of 0.807, and F1 of 0.850 using the fixed 80th-percentile traffic threshold. These findings indicate that weather information is most appropriately integrated as contextual information that complements, rather than replaces, recent traffic-state and temporal predictors.

From an engineering perspective, the results support the application of boosted-tree forecasting for corridor-level traffic analysis where model interpretability, computational efficiency, and data availability are important considerations. The relatively small performance degradation under weather-sensor perturbation indicates that the forecasting system remained reasonably stable when meteorological inputs were progressively degraded. The principal limitation concerns the use of a deterministic calibrated reconstruction rather than the original row-level UCI observations, meaning that the reported numerical results should be regarded as benchmark evidence rather than directly deployable field coefficients. Direct reanalysis using the original UCI records is therefore required to verify whether the observed performance and weather contribution are preserved in the source observations. Subsequent validation should extend the framework to additional traffic detectors, spatially resolved weather observations, incident information, and multiple forecasting horizons to assess transferability across broader transportation environments. Congestion classification should also be calibrated to operational objectives because the relative consequences of missed congestion and false alarms may differ among transportation agencies. The framework consequently provides a reproducible basis for integrating weather-aware prediction with congestion analysis while leaving field validation and operational threshold calibration as essential next steps.

REFERENCES

- Abduljabbar, R., Dia, H., & Liyanage, S. (2024). Machine learning models for traffic prediction on arterial roads using traffic features and weather information. *Applied Sciences*, 14(23), 11047. <https://doi.org/10.3390/app142311047>
- Abduljabbar, R., Dia, H., & Liyanage, S. (2025). Machine learning traffic flow prediction models for smart and sustainable traffic management. *Infrastructures*, 10(7), 155. <https://doi.org/10.3390/infrastructures10070155>
- Ali, A., Sharafian, A., Yasir Naeem, H. M., Zakarya, M., Wu, Z., & Bai, X. (2026). Advanced computational models for urban traffic flow prediction: A comprehensive review and future directions. *Computer Science Review*, 60, 100886. <https://doi.org/10.1016/j.cosrev.2025.100886>
- Berlotti, M., Di Grande, S., & Cavalieri, S. (2024). Proposal of a machine learning approach for traffic flow prediction. *Sensors*, 24(7), 2348. <https://doi.org/10.3390/s24072348>
- Chauhan, N. S., & Kumar, N. (2024). Confined attention mechanism enabled recurrent neural network framework to improve traffic flow prediction. *Engineering Applications of Artificial Intelligence*, 136, 108791. <https://doi.org/10.1016/j.engappai.2024.108791>
- Cheng, H., Gong, Z., & Wang, C. (2025). LLM-TFP: Integrating large language models with spatio-temporal features for urban traffic flow prediction. *Applied Soft Computing*, 177, 113174. <https://doi.org/10.1016/j.asoc.2025.113174>
- Chu, Z., Ma, W., Li, M., & Chen, H. (2024). Adaptive decision spatio-temporal neural ODE for traffic flow forecasting with multi-kernel temporal dynamic dilation convolution. *Neural Networks*, 179, 106549. <https://doi.org/10.1016/j.neunet.2024.106549>
- Fafoutellis, P., & Vlahogianni, E. I. (2023). Unlocking the full potential of deep learning in traffic forecasting through road network representations: A critical review. *Data Science for Transportation*, 5, 23. <https://doi.org/10.1007/s42421-023-00083-w>
- Fahs, W., Chbib, F., Rammal, A., Khatoun, R., El Attar, A., Zaytoun, I., & Hachem, J. (2023). Traffic congestion prediction based on multivariate modelling and neural networks regressions. *Procedia Computer Science*, 220, 202–209. <https://doi.org/10.1016/j.procs.2023.03.028>
- Gomes, B., Coelho, J., & Aidos, H. (2023). A survey on traffic flow prediction and classification. *Intelligent Systems with Applications*, 20, 200268. <https://doi.org/10.1016/j.iswa.2023.200268>
- Guo, X., Zhang, Q., Jiang, J., Peng, M., Zhu, M., & Yang, H. F. (2024). Towards explainable traffic flow prediction with large language models. *Communications in Transportation Research*, 4, 100150. <https://doi.org/10.1016/j.commtr.2024.100150>
- Han, D. C. (2024). Prediction of traffic volume based on deep learning model for AADT correction. *Applied Sciences*, 14(20), 9436. <https://doi.org/10.3390/app14209436>
- Hogue, J. (2019). Metro interstate traffic volume [Dataset]. *UCI Machine Learning Repository*. <https://doi.org/10.24432/C5X60B>
- Kendre, H., Karmarkar, M., Karadbhaje, H., Kamtalwar, N., & Marathe, A. (2024). Traffic volume prediction based on weather parameters. In *2024 IEEE International Conference for Women in*

- Innovation, Technology & Entrepreneurship (ICWITE)* (pp. 124–129). IEEE. <https://doi.org/10.1109/ICWITE59797.2024.10503349>
- Li, Y., Wang, J., & Kang, L. (2025). Light attention-based neural networks for traffic flow prediction. *Physica A: Statistical Mechanics and its Applications*, 673, 130665. <https://doi.org/10.1016/j.physa.2025.130665>
- Liu, R., & Shin, S.-Y. (2025). A review of traffic flow prediction methods in intelligent transportation system construction. *Applied Sciences*, 15(7), 3866. <https://doi.org/10.3390/app15073866>
- Luo, Q., He, S., Han, X., Wang, Y., & Li, H. (2024). LSTTN: A long-short term transformer-based spatiotemporal neural network for traffic flow forecasting. *Knowledge-Based Systems*, 293, 111637. <https://doi.org/10.1016/j.knosys.2024.111637>
- Ma, J., Zhao, J., & Hou, Y. (2024). Spatial–temporal transformer networks for traffic flow forecasting using a pre-trained language model. *Sensors*, 24(17), 5502. <https://doi.org/10.3390/s24175502>
- Nesa, M., & Yoon, Y. (2024). Speed prediction and nearby road impact analysis using machine learning and ensemble of explainable AI techniques. *Scientific Reports*, 14, 25208. <https://doi.org/10.1038/s41598-024-74545-8>
- Park, S., Kim, M., & Kim, J. (2024). Hourly long-term traffic volume prediction with meteorological information using graph convolutional networks. *Applied Sciences*, 14(6), 2285. <https://doi.org/10.3390/app14062285>
- Peng, D., & Zhang, Y. (2023). MA-GCN: A memory augmented graph convolutional network for traffic prediction. *Engineering Applications of Artificial Intelligence*, 121, 106046. <https://doi.org/10.1016/j.engappai.2023.106046>
- Qi, Y., & Cheng, Z. (2023). Research on traffic congestion forecast based on deep learning. *Information*, 14(2), 108. <https://doi.org/10.3390/info14020108>
- Qin, H., & Zhang, W. (2022). Short-term traffic flow prediction and signal timing optimization based on deep learning. *Wireless Communications and Mobile Computing*, 2022, 8926445. <https://doi.org/10.1155/2022/8926445>
- Qin, Z., Wang, M., Zhu, H., Lee, W.-C., Cui, N., Yu, J., Tian, Z., Hong, Y., & Yin, J. (2026). A survey on modern deep learning techniques for traffic volume prediction. *Data Science and Engineering*, 11, 473–512. <https://doi.org/10.1007/s41019-025-00340-8>
- Rasulmukhamedov, M., Tashmetov, T., & Tashmetov, K. (2024). Forecasting traffic flow using machine learning algorithms. *Engineering Proceedings*, 70(1), 14. <https://doi.org/10.3390/engproc2024070014>
- Sayed, S. A., Abdel-Hamid, Y., & Hefny, H. A. (2023). Artificial intelligence-based traffic flow prediction: A comprehensive review. *Journal of Electrical Systems and Information Technology*, 10, 13. <https://doi.org/10.1186/s43067-023-00081-6>
- Tao, Q., Li, Z., Xu, J., Lin, S., De Schutter, B., & Suykens, J. A. K. (2022). Short-term traffic flow prediction based on the efficient hinging hyperplanes neural network. *IEEE Transactions on Intelligent Transportation Systems*, 23(9), 15616–15628. <https://doi.org/10.1109/TITS.2022.3142728>
- Tian, R., Wang, C., Hu, J., & Ma, Z. (2023). Multi-scale spatial-temporal aware transformer for traffic prediction. *Information Sciences*, 648, 119557. <https://doi.org/10.1016/j.ins.2023.119557>
- Tran, D. Q., Tran, H. Q., & Nguyen, M. V. (2024). An enhanced ensemble-based long short-term memory approach for traffic volume prediction. *Computers, Materials & Continua*, 78(3), 3585–3602. <https://doi.org/10.32604/cmc.2024.047760>
- Wang, C., Tian, R., Hu, J., & Ma, Z. (2023). A trend graph attention network for traffic prediction. *Information Sciences*, 623, 275–292. <https://doi.org/10.1016/j.ins.2022.12.048>
- Wang, J.-D., & Susanto, C. O. N. (2024). Traffic flow prediction with heterogeneous spatiotemporal data based on a hybrid deep learning model using attention-mechanism. *Computer Modeling in Engineering & Sciences*, 140(2), 1711–1728. <https://doi.org/10.32604/cmcs.2024.048955>
- Xu, Y., & Liu, M. (2025). GPT4TFP: Spatio-temporal fusion large language model for traffic flow prediction. *Neurocomputing*, 625, 129562. <https://doi.org/10.1016/j.neucom.2025.129562>
- Yaqub, M., Ahmad, S., Manan, M. A., Pathan, M. S., & He, L. (2025). Predicting traffic flow with federated learning and graph neural with asynchronous computations network. *Array*, 26, 100411. <https://doi.org/10.1016/j.array.2025.100411>

- Zhang, S., Yin, X., Zhang, W., Xu, J., & Jing, X. (2025). Adaptive lightweight temporal convolutional network with context-aware downsampling strategy for traffic flow prediction. *Engineering Applications of Artificial Intelligence*, 156, 111208. <https://doi.org/10.1016/j.engappai.2025.111208>
- Zhang, W., Yao, R., Du, X., Liu, Y., Wang, R., & Wang, L. (2023). Traffic flow prediction under multiple adverse weather based on self-attention mechanism and deep learning models. *Physica A: Statistical Mechanics and its Applications*, 625, 128988. <https://doi.org/10.1016/j.physa.2023.128988>
- Zhang, W., Yao, R., Yuan, Y., Du, X., Wang, L., & Sun, F. (2024). A traffic-weather generative adversarial network for traffic flow prediction for road networks under bad weather. *Engineering Applications of Artificial Intelligence*, 137, 109125. <https://doi.org/10.1016/j.engappai.2024.109125>