Online ISSN: 3104-9338
Print ISSN: 3104-932X

Assignment-Informed Link-Flow Forecasting: A Case Study in Wuxi

Tianbin Wu[1,2], Tong Liu[1], Xihua Liu[2], Kun Jin[3*]

[1] Wuxi Public Security Bureau Traffic Management Division, Wuxi, China

[2] National Engineering Research Center of Road Safety Control Technology, China

[3] School of Transportation, Southeast University, Nanjing, China

Received: Oct 22, 2026

Revised: Oct 27, 2026

Accepted: Oct 30, 2026

Published online: November 2, 2026

To appear in: International Journal of Advanced AI Applications, Vol. 2, No. 12 (December 2026)

* Corresponding Author:Kun Jin (jinkun@seu.edu.cn)

Abstract. Road restrictions and changing travel demand can alter link flows before recent detector counts describe the new pattern. This study develops a selective assignment correction for traffic flow forecasting in Wuxi. A relative cost Logit model translates OD demand and road costs into changes in route use. Link specific calibration adjusts the response scale, and an assignment state rule directs the correction toward substantial departures from recurring conditions. The study combines observed detector counts, OD demand and the Wuxi road network. Comparisons cover supplied road cost changes, several forecasting backbones, prediction horizons and persistent flow drops. Assignment improves forecasts in strong path states and adds a useful correction to models that already use recent detector histories. The benefit increases with assigned state strength and remains visible across test dates. Calibrating service ratios also improves forecasts during persistent drops. These findings indicate how route information can supplement an established detector forecast. Broader observation periods, independent demand sources and documented disruptions would support assessment across a wider range of operating conditions.

Keywords: *route choice; origin destination demand; detector; flow forecasting*

1. Introduction

1.1. Research Background

Traffic forecasts help road agencies anticipate demand on individual links and respond to changes across an urban network. Recent detector counts are a useful starting point because

they describe what vehicles are doing now. A forecast based on these counts can follow recurring commuting patterns and short fluctuations in traffic volume. When a road restriction raises the cost of using a route, drivers may begin to use other roads before the resulting pattern appears clearly in the measured history. The agency then needs a forecast that can use its knowledge of the changed road conditions.

Graph models such as STGCN and Graph WaveNet learn temporal patterns together with relationships between measured links [1,2]. These relationships are valuable under the traffic conditions represented in the training data. A newly supplied change in travel cost adds information about route choice that may arrive earlier than the next set of useful detector observations. A forecast can benefit from that information if the expected response is translated to the scale of the measured flow. The aim is to update the forecast at the point when knowledge of the network changes.

Traffic assignment provides a connection between travel demand, road costs, and the routes used by travelers. The relative cost Logit model represents variation in route choice through a cost sensitive loading process [3]. The wider assignment literature explains how travel demand and network conditions jointly determine link flows [4]. Repeating an assignment with a different road cost estimates where flow is expected to increase or decrease. Assignment can also describe an unusual current distribution of trips when no particular road restriction has been specified.

Assigned volumes and detector counts describe related quantities at different scales. The origin destination (OD) data represent a set of trips, the network model represents possible routes, and each detector samples vehicles at a particular location. These representations can produce different volumes on the same link. Directly adding an assigned flow to a detector forecast would have little meaning. We use the change between assigned states and estimate how strongly that change should affect the forecast. This keeps the observed traffic pattern as the reference while allowing route information to influence the prediction.

1.2. Related Work

Graph based traffic forecasting has developed several ways to represent spatial and temporal dependence. A survey by Jiang and Luo [5] describes the use of graph neural networks for this purpose. Recent models incorporate trends over several prediction steps [6], combine information across spatial and temporal scales [7], or simplify graph computation while maintaining predictive accuracy [8]. These studies establish the value of learning from measured traffic streams. Their relevance here is the strength of the detector forecast that an

additional source of network information must improve.

Incomplete observations have also received substantial attention. Joint matrix factorization uses shared structure to recover missing traffic values [9]. A graph recurrent imputation model combines spatial relationships with temporal information [10]. Other methods address missing information and incomplete features more directly [11,12], while a recent survey organizes the broader literature on traffic data imputation [13]. These approaches help recover a measured traffic record. A supplied road restriction carries a different kind of information because it describes a change in travel conditions before its full effect has been observed.

Reliable observations remain essential in either setting. Fault detection methods identify unusual sensor behavior [14], and sensor reconciliation can improve the consistency of a traffic sensor network [15]. Traffic state estimation also combines several sources of measurements [16] and can accommodate errors in automated sensor data [17]. Roving sensor networks extend the spatial coverage available for city scale estimation [18]. These studies show why the quality and coverage of observations matter when interpreting an additional model output. Assignment contributes a representation of possible routes, with its own dependence on demand and network inputs.

1.3. Study Objectives

The study considers two main uses of this approach. The first starts with a normal network forecast and a supplied change in road cost. Assignment estimates the redistribution caused by that change. The second starts with a detector based forecast and asks whether a strong departure in the current assigned path state helps predict the future count. Both uses begin with a forecast that can already be formed from the available inputs. The assigned response supplies a correction whose strength is selected before the final evaluation.

We study these uses with March 2025 vehicle passage data from Wuxi. The observed part of the study contains 548 links with counts aggregated at 5 min intervals. A road cost experiment applies the specified changes to the same Wuxi network and OD inputs, with independently solved MSA BPR flows serving as the reference. Persistent flow drop episodes provide another comparison in which the count at forecast issue time is available. Routine predictions and forecasts one to seven days ahead show how much the same assignment information contributes when the usual measured pattern is already informative.

The paper contributes a straightforward way to combine these sources without requiring a new forecasting architecture. It describes the inputs available in each setting, calibrates the

assigned response, and reports the resulting accuracy under a common error measure. The strongest improvements occur in the controlled cost change and in the small set of direct predictions with strong assigned path states. The remaining comparisons show where an agency can retain a simpler forecast. These results support a selective role for assignment in a broader traffic forecasting system.

The present study brings the route response into an existing flow forecast through a calibrated correction. It examines a supplied change in cost separately from a strong current path state because the information arrives in different forms. A common sample comparison then adds the same update to three models that already use recent detector histories. Detector based and constant bias corrections show how much improvement can arise from a small additional calibration. The resulting comparison establishes the circumstances in which an assigned departure adds useful predictive information.

2. Methodology

2.1. Forecast Inputs and Assignment

Let y denote the detector count on a directed road link during a 5 min interval. A forecast issued at time t predicts the count at a later interval t + h. Every empirical prediction uses inputs available when it is issued, with the future count reserved for evaluation. The controlled road cost scenario starts from a normal network flow estimate and receives the specified change in cost. The flow drop comparison also receives the count observed on the affected link at the time of prediction. These input conditions determine which reference forecast can be formed.

The assignment model loads each time slice of OD demand through the road network using a destination based relative cost Logit formulation [3]. Demand is divided into five increments. Link costs increase with the accumulated volume through a Bureau of Public Roads (BPR) cost function. For each destination, a shortest cost search identifies outgoing links that reduce the remaining cost of reaching that destination. The Logit probabilities distribute available demand among those links. Updating costs between increments allows later demand to respond to the volume already loaded onto the network.

The five loading shares are 5/15, 4/15, 3/15, 2/15, and 1/15. The relative cost formulation scales outgoing route costs by their mean, and the Logit dispersion parameter is 7.2. Each demand slice is converted to an equivalent hourly rate before loading. The same network representation and loading procedure are used for the normal and changed cost assignments. Their difference describes the response of the model to the supplied cost change, while the

absolute assigned volume remains dependent on the represented demand.

For forecasts based on detector observations, we also calculate the usual assigned flow for each link and time of week. A current assignment can then be expressed as a departure from that usual value. This departure identifies a path state that differs from the recurring pattern. It can be large when the location or intensity of OD demand changes, even when no specific road event has been supplied. The forecast uses the departure only after its relationship with earlier detector forecast errors has been fitted.

2.2. Forecast Update for a Supplied Cost Change

The controlled scenario begins with the normal flow on each link. We run the assignment with the normal road cost and again with a supplied cost multiplier κ. The difference between those assigned volumes gives the expected flow change. A positive difference indicates an expected increase in the use of a link; a negative difference indicates an expected decrease. The corrected forecast is given by Equation (1), where λ controls the size of the response.

$$\hat{f}_l^{\mathrm{E}} = f_l^0 + \frac{\lambda}{\kappa - 1} \Delta q_l^{\mathrm{A}}(\kappa) \tag{1}$$

The multiplier κ is part of the scenario input. The response strength λ is selected using a fivefold cost change and then held fixed when a twofold cost change is evaluated. The denominator in Equation (1) follows the response scaling used in the original experiment. After the correction is added, a nonnegative bound is applied to the forecast. The normal flow remains the starting point for every link, and the assigned difference provides the direction and size of the update.

The reference flow for each supplied road cost setting is calculated using the method of successive averages with BPR costs (MSA BPR) on the Wuxi network and the same OD demand. The Logit update is evaluated against this independently solved assignment. The comparison examines how well the route response transfers the specified cost change into the direction and size of the resulting flow redistribution.

We report the result for all links with normal reference flow of at least 20 vehicles per hour. We also examine links whose normal assigned flow reaches the same threshold and changes by at least 10% under the supplied restriction. This second group focuses on links for which the assignment expects a substantial response. Group membership is determined from the normal flow and assignment outputs before the changed reference flow is used to calculate errors. Both groups are retained because an average over the whole network can be dominated by roads with very little response.

2.3. Direct Forecasts under Strong Path States

For direct forecasts, the starting point is a detector based estimate of the future link count. We calculate the difference between the recent assigned flow and its usual value for the corresponding time of week. The assignment time is offset by a horizon specific lag measured in 5 min bins. The departure is $r^{A}_{l,t-\ell_h} = q^{A}_{l,t-\ell_h} - \bar{q}^{A}_{l,k(t-\ell_h)}$, where the bar denotes the usual weekly assignment. The response coefficient β for each link and horizon is fitted to earlier errors in the detector forecast using the ridge penalty ξ. Equation (2) gives this fitted coefficient. The training sum uses only intervals in the fitting period.

$$\beta_{l,h} = \frac{\sum_{r \in \mathcal{T}_{tr}} [y_{l,r+h} - \hat{y}^{B}_{l,r+h}] r^{A}_{l,r-\ell_h}}{\sum_{r \in \mathcal{T}_{tr}} (r^{A}_{l,r-\ell_h})^2 + \xi_{l,h}} \tag{2}$$

The fitted coefficient translates the assignment departure into an expected correction to the detector forecast. Its sign allows the relationship to differ across links. The ridge term reduces the size of a response when the available variation gives weak support for a large coefficient. The lag allows a path state observed earlier to relate to flow at the prediction horizon. These quantities are estimated before the final test period and remain fixed when the future count is withheld.

The correction is activated when the assignment departure is large relative to the usual assigned flow at the forecast target time. Equation (3) defines the state strength for each prediction horizon; k indexes the weekly time bin. The reference floor ε is 25 vehicles per 5 min, which limits unstable ratios when the usual assigned volume is low. A threshold η identifies the path states used for the correction. Expressing strength relative to the usual assigned flow makes the rule depend on the departure from each link's own pattern.

$$u_{l,t,h} = \frac{|r^{A}_{l,t-\ell_h}|}{\max\left(\bar{q}^{A}_{l,k(t+h)}, \varepsilon\right)} \tag{3}$$

Equation (4) adds the fitted assignment response to the detector forecast when the threshold is met and bounds the prediction below by zero. The indicator equals one for a sufficiently strong state and zero otherwise. A horizon specific gain γ controls the amount of the fitted response that enters the forecast. The threshold and gain are selected on validation data. The same settings are then applied throughout the test period, including predictions for which the indicator remains zero.

$$\hat{y}^{H}_{l,t+h} = \max\left\{0, \hat{y}^{B}_{l,t+h} + \mathbb{I}(u_{l,t,h} \geq \eta)\, \gamma_h \beta_{l,h} r^{A}_{l,t-\ell_h}\right\} \tag{4}$$

This rule gives assignment a selective role. Ordinary path states retain the detector forecast, while strong states receive a correction based on the relationship learned earlier. The final

evaluation reports both the performance under routine conditions and the performance in the activated group. The number of activated predictions is needed to interpret the result: an improvement in a small group can be useful for particular situations while contributing little to the error averaged over all traffic.

2.4. Comparison with Forecasts Using Recent Counts

The profile comparison uses a usual weekly count and omits recent observations from the target link. We also evaluate three 30 min models that use those observations. Temporal autoregression combines the target time profile with count departures at issue time and 5, 10, 15 and 30 min earlier, together with the counts one day and one week before the target. Spatial temporal autoregression (STAR) adds departures on ten links selected by training correlation. Both regressions standardize their predictors, fit an intercept and use ridge strength 100. A week of count history makes their first fitting target March 8 at 00:30. Their coefficients are fixed after March 20.

The neural comparison is a compact STGCN style implementation with two temporal convolutions and graph aggregation. It uses the preceding twelve 5 min bins, eight hidden channels and ten neighbors selected from training correlations. Training uses Adam with learning rate 0.001 for fifteen epochs and random seed 2026; validation selects the retained epoch. This implementation provides a graph based reference for the assignment update. It is smaller than many recent forecasting networks and represents one training run.

All three models receive the same additional correction. For each link, one slope relates training forecast errors to the assignment departure, scaled by its training root mean square. The ridge penalty is 10, with no added intercept. Two controls use the same fitting period and activation rule: one replaces the assignment departure with the target detector departure at issue time, and the other estimates a constant link bias. Each control adds one fitted coefficient per link. The controls retain the assignment gate, allowing the comparison to isolate the value of the correction magnitude after the strong state has been identified.

Correction fitting uses the common March 8 to 20 targets and each model's training predictions. These predictions are from the model fitting period. The gains 0, 0.25, 0.5, 0.75, 1 and 1.25 are compared using activated validation WMAPE; ties select the smaller gain. Zero retains the original forecast. The original 30 min lag of one bin, denominator floor of 25 and threshold of 0.5 stay fixed. The 3,990 validation cases select the gains, which are then applied to the same 5,267 test cases for every model. Target times, link identifiers and observed counts are matched exactly across the saved predictions and assignment data. This comparison changes

the forecast being corrected while retaining the state definition.

2.5. Forecasts during Persistent Flow Drops

The episode comparison uses persistent drops identified in detector counts. The screen examines the period from 06:00 to 22:00 and compares each link with its usual count for the same weekday and time. The usual count must be at least 20 vehicles per 5 min. A low state has a count no more than 45% of that reference and a standardized departure no greater than minus 2.5. The preceding count must be at least 80% of its reference to identify an abrupt onset. Consecutive low states form an episode, and at least two intervals are required. The reference statistics use March 1 to 10; later dates supply the calibration, validation, and test episodes.

The target is the count in the interval following the first detected drop, giving a 5 min forecast. The first low count is available at issue time. Episode duration is established from the subsequent record, so the evaluation is a retrospective comparison conditional on a persistent drop. The sample describes observed traffic states. Their causes are not externally verified. The assignment branch receives a supplied fivefold road cost multiplier on the affected link, with capacity unchanged, and a historical OD slice for the matching weekday and time. This multiplier represents a scenario assumption applied consistently across episodes.

A weekly profile formed without the selected episode intervals gives the normal count. Persistence carries the issue count forward. The fixed service forecast multiplies the normal profile by one fifth. Its capped version takes the smaller of that result and the issue count, with a minimum cap of 20 vehicles per 5 min. The raw assignment forecast multiplies the normal profile by the changed to normal assigned flow ratio. This ratio is bounded between zero and one, and the normal assigned flow has a denominator floor of one.

We also fit a simple service ratio using three assignment summaries. The first is the raw ratio on the target link. The second is the ratio of total outgoing assigned flow at the target link's source node. The third is the sum of absolute flow changes over the network divided by total normal assigned flow. Together they describe local route use and the size of the network redistribution. Each feature is centered and scaled using the calibration episodes. A ridge regression relates these features to the observed count divided by the normal profile. Its intercept is the calibration mean of that count ratio.

The fitted service ratio is bounded between 0.05 and one before multiplying the normal profile. The same issue count cap can then be applied. Ridge strengths of 0.1, 1, 10, and 100 are compared on validation episodes, and the selected model remains fixed in testing. The

method adds only three fitted slopes and an intercept. This limited calibration allows assignment to contribute variation across episodes while keeping the forecast easy to estimate with a small sample.

2.6. Forecasts at Longer Leads

The longer lead comparison predicts one to seven days ahead with supplied future OD demand. A weekly detector profile provides the principal reference. It carries the usual count for the target time and allows the current departure from that pattern to decay with lead time. Persistence and a recursive autoregression provide two additional count based comparisons. The autoregression moves the current count toward the training mean. Both decay coefficients are fitted on training intervals and bounded in absolute value by 0.95.

The future OD demand is assigned to the network, and its departure from the usual weekly assignment is fitted to earlier errors in the weekly forecast. The response is estimated for each link and 10 min time bin using ridge regression. Validation selects the ridge strength and a cap on the correction magnitude. A relative correction score determines activation, with the threshold set to the 90th percentile of positive validation scores. A validation selected gain controls the final update. These settings are shared across all seven leads.

All leads use the same target intervals, while their issue dates differ. The target time weekly profile and supplied assignment are shared across those forecasts. Only the issue observation and its decaying residual depend on the lead. The comparison describes forecasts conditional on the supplied demand scenario; future OD demand is not estimated by this procedure. The shared inputs make the recurring forecast stable across long leads.

3. Results

3.1. Study Area and Evaluation

The Wuxi case uses March 2025 vehicle passage records to derive detector counts and OD demand. The observed area contains 548 directed links and 4,892,544 records at 5 min intervals. The assignment network covers 57,161 nodes and 116,696 directed links (Figure 1), allowing routes to extend beyond the observed area. Calibration translates the assigned response to the count scale at each detector.

The main forecasting models use March 1 to 20 for fitting, March 21 to 24 for validation, and March 25 to 31 for testing. The episode comparison uses March 1 to 10 to define the usual flow and March 11 to 20 for service ratio calibration. All model selection uses validation data.

The four validation days run from Friday through Monday and include two weekend days. Testing covers a later week within the same month.

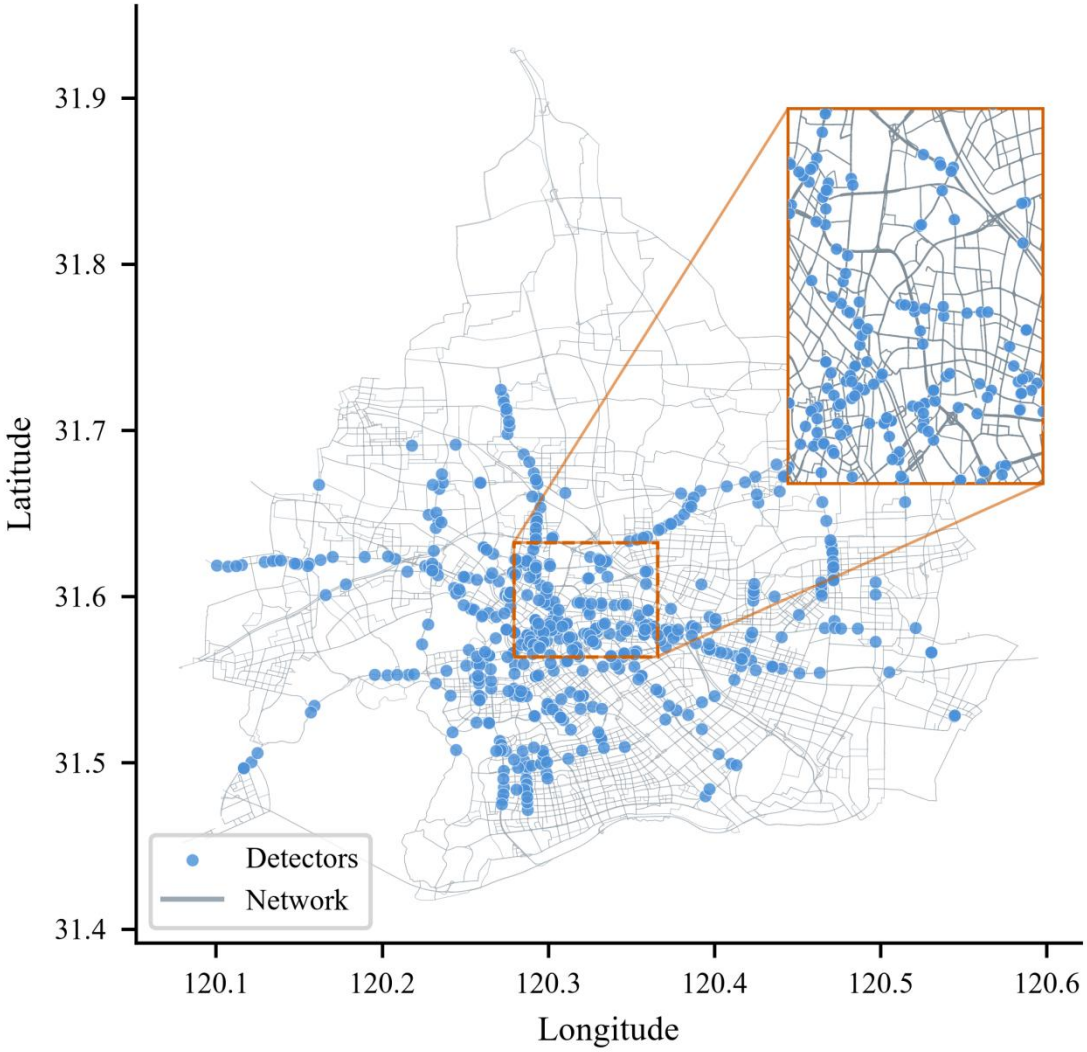


Figure 1. Observed links in Wuxi within the larger assignment network.

Time sliced OD demand is derived from passage trajectories in the same underlying dataset as the detector counts. The two representations capture different aspects of those records. Their joint use shows what route information can contribute within this data source. A separately collected OD feed could represent another set of trips and would require its own calibration against detector flow.

Accuracy is measured by weighted mean absolute percentage error. WMAPE is the sum of absolute forecast errors divided by the sum of observed or reference flow over the evaluated cases. Empirical comparisons use detector counts in the denominator. The controlled cost scenario uses a separate MSA BPR flow solution. Mean absolute error (MAE) and root mean squared error (RMSE) describe the magnitude of errors in vehicles per interval, or vehicles per

hour for the static controlled scenario. The different units follow the corresponding target variable.

The main direct forecast intervals use 2,000 paired resamples of the seven complete test calendar days, with seed 2026. All predictions from a sampled day remain together, preserving changes shared by different roads on that date. We also resample link by day blocks to describe sensitivity to the resampling unit. Daily and leave one day out results show how the pooled estimate depends on particular dates. Episode intervals use 500 paired resamples of the 14 test episodes. The longer lead comparison also resamples complete test days. These intervals describe variation within the available period; seven days provide limited information about transfer to another week.

3.2. Flow Redistribution under a Supplied Cost Change

The supplied road cost change gives assignment a specific task: update the normal flow pattern before observations from the changed state are available. A response strength of 2.0 is selected using the fivefold validation scenario and applied to the separate twofold scenario. The changed flow is evaluated against an independently solved MSA BPR state with 12 loading iterations. Link selection uses normal flows and the assigned response, before reading the changed reference flow.

Table 1. Controlled twofold cost change against the independent MSA BPR reference. The responsive cohort is a subset of the wider network.

Cohort	Links	Normal WMAPE	Corrected WMAPE	MAE (vehicles/h) normal → corrected
Normal flow ≥ 20 vehicles/h	9,495	0.0083	0.0025	0.98 → 0.30
Assigned response ≥ 10%	45	0.2309	0.0997	61.94 → 26.76

Table 1 separates the wider network from the roads with a substantial response. Across 9,495 links carrying at least 20 vehicles/h in the normal reference, WMAPE decreases from 0.0083 to 0.0025. This wider group contains many roads with little change. The responsive cohort contains 45 links with normal reference and assigned flows of at least 20 vehicles/h and an assigned change of at least 10%. Its WMAPE falls from 0.2309 to 0.0997, a relative reduction of 56.80%. These cohorts overlap and describe different parts of the same scenario.

Figure 2 compares the correction on every responsive link with its change in the independent reference. All 45 assigned responses are positive in this particular scenario, corresponding to roads receiving diverted traffic. Most large reference increases receive a correction in the same range. A smaller group receives an increase even though its reference flow changes little. Points

coincide where successive links share the same route volume, so the number of visible points is smaller than the number of evaluated links.

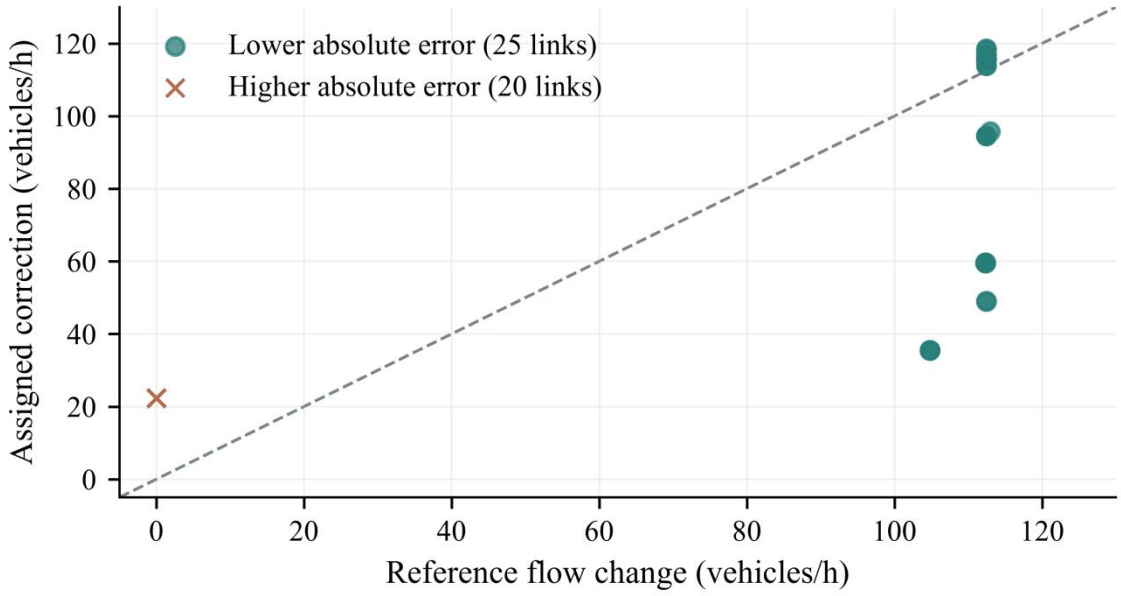


Figure 2. Assigned correction and independent reference change for all 45 responsive links in the twofold cost scenario. The dashed line indicates equal changes. Points can coincide when links share flow values.

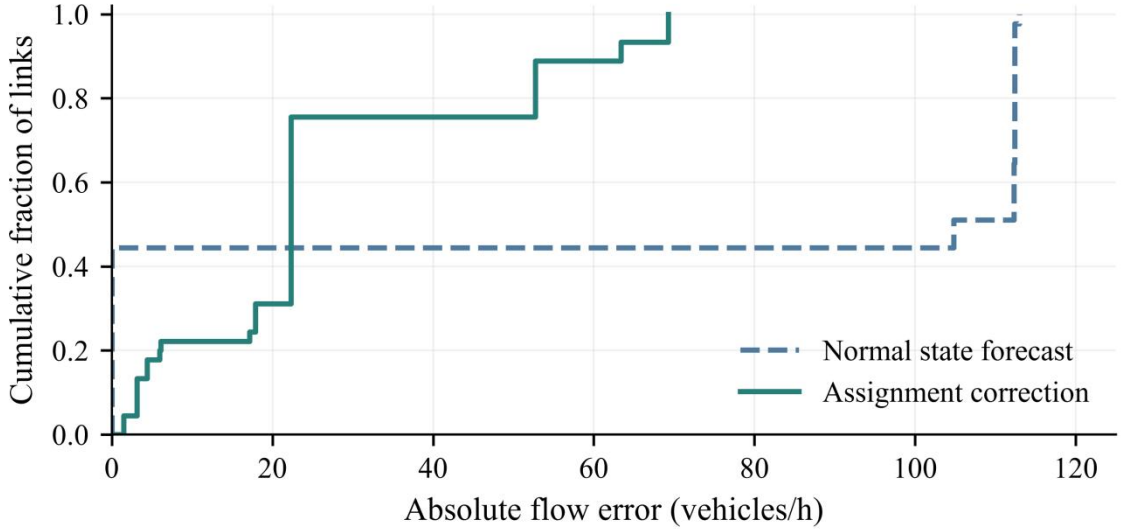


Figure 3. Absolute error distributions across the same 45 links, with each link weighted equally. The correction reduces large errors while increasing errors on some initially accurate links.

The aggregate reduction combines improvements on 25 links with increases in error on 20 links. Figure 3 shows the full absolute error distributions with each link weighted equally. The normal forecast has near zero error on some roads and errors above 100 vehicles/h on many others. The correction reduces the larger errors while moving some initially accurate estimates away from the reference. Median absolute error decreases from 104.78 to 22.30 vehicles/h.

MAE falls from 61.94 to 26.76 vehicles/h, and RMSE falls from 83.12 to 33.48 vehicles/h.

The controlled result shows how an assigned response can transfer a supplied change in road conditions into a flow update. It also identifies a local weakness: an assigned diversion can be too broad for the independently solved state. For an application focused on a particular alternative route, the link level response in Figure 2 is relevant alongside the pooled WMAPE. The road cost comparison evaluates redistribution on the Wuxi network against the MSA BPR reference. The detector forecasting comparisons use the observed future counts as their targets.

3.3. Direct Forecast Accuracy across Horizons

The direct forecast comparison retains the state strength threshold of 0.5 selected on validation data. Selection requires at least 500 validation cases. The response gains are 0.375 for 30 and 60 min and 0.450 for 120 min, with assignment lags of one, zero and two 5 min bins. These settings define a common operating rule: a routine state retains the detector forecast, and a strong assigned departure receives the calibrated route response.

Table 2. Direct forecasts within strong assigned states, with 1,104,768 predictions per horizon. MAE is in vehicles per 5 min.

Horizon	Active cases	Active share	Profile WMAPE	Corrected WMAPE	MAE (veh/5 min) profile → corrected
30 min	5,267	0.477%	0.2814	0.2018	17.97 → 12.89
60 min	5,384	0.487%	0.2873	0.2087	17.17 → 12.48
120 min	5,597	0.507%	0.2944	0.2173	14.32 → 10.57

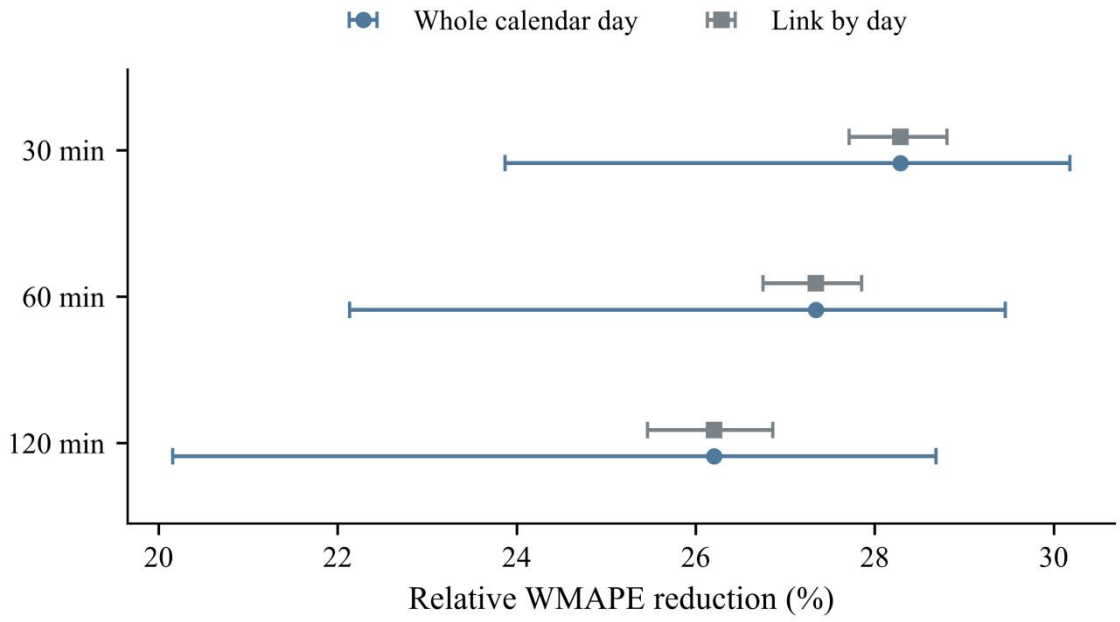


Figure 4. Relative WMAPE reductions for the weekly profile within strong assigned states: 5,267, 5,384 and 5,597 cases at 30, 60 and 120 min. Points are pooled estimates; bars are 95% intervals from 2,000 paired resamples. Whole calendar days are the principal unit.

Table 2 reports the activated predictions at each horizon. WMAPE decreases from 0.2814 to 0.2018 at 30 min, from 0.2873 to 0.2087 at 60 min, and from 0.2944 to 0.2173 at 120 min. MAE decreases at every horizon. The activated groups contain 5,267, 5,384 and 5,597 predictions out of 1,104,768 test predictions per horizon. Their frequencies remain 0.477% to 0.507%, so the result describes infrequent departures from the weekly pattern.

The relative WMAPE reductions are 28.29%, 27.34% and 26.21%. The corresponding whole day bootstrap intervals are 23.87% to 30.18%, 22.13% to 29.46% and 20.16% to 28.68%. Figure 4 also reports intervals obtained by resampling link by day blocks. Whole days preserve the traffic variation shared by roads on the same date, and provide the principal temporal comparison. The two resampling units describe different aspects of variation within the test week.

The relative benefit decreases gradually with prediction horizon. A current path departure carries less information about the count as the target moves further ahead. Route information still improves the forecast at every evaluated lead. The groups differ slightly because the horizon specific lag affects which assigned departures cross the threshold. Their observed count totals and mean flow levels also differ. WMAPE and MAE should consequently be compared between forecasts within each horizon, with sample size reported alongside accuracy.

Across every 30 min test prediction, applying the activation rule changes profile WMAPE from 0.18835 to 0.18781. This smaller network average follows from retaining the original profile on most predictions. The activated result describes where the route response contributes, while the complete sample describes its effect on the forecasting service as a whole. Both are needed when deciding how much additional calculation to introduce into an existing system.

3.4. Accuracy with Recent Detector Histories

The common 30 min group contains 5,267 predictions from 413 links across seven days, with an observed total of 336,386 vehicles. The weekly profile gives WMAPE 0.2814 before correction. Recent detector histories provide a stronger starting estimate: WMAPE is 0.1916 for STGCN style, 0.1641 for temporal autoregression and 0.1608 for STAR. Keeping these target counts and the activation rule identical allows the added assignment response to be compared across forecasting backbones. Table 3 summarizes this comparison.

Assignment reduces STGCN style WMAPE to 0.1489, a relative reduction of 22.28%, with a whole day interval of 17.83% to 24.97%. MAE decreases from 12.24 to 9.51 vehicles per 5 min, and RMSE from 19.28 to 15.11. The update improves both the typical count error and the

contribution of larger errors. The selected gain remains 1.0 for this backbone.

Table 3. Forecasts on the same 5,267 strong state cases at 30 min. The weekly profile omits recent target counts. The other models use them. All controls share the assignment gate. Intervals use 2,000 paired resamples of seven complete test days.

Forecast	Original WMAPE	Assignment WMAPE	Detector WMAPE	Bias WMAPE	Assignment reduction % [95% day interval]
Weekly profile	0.2814	0.2018	—	—	28.29 [23.87, 30.18]
STGCN style	0.1916	0.1489	0.1726	0.1863	22.28 [17.83, 24.97]
Temporal AR	0.1641	0.1342	0.1625	0.1637	18.19 [13.15, 20.72]
STAR	0.1608	0.1342	0.1592	0.1602	16.50 [11.46, 19.26]

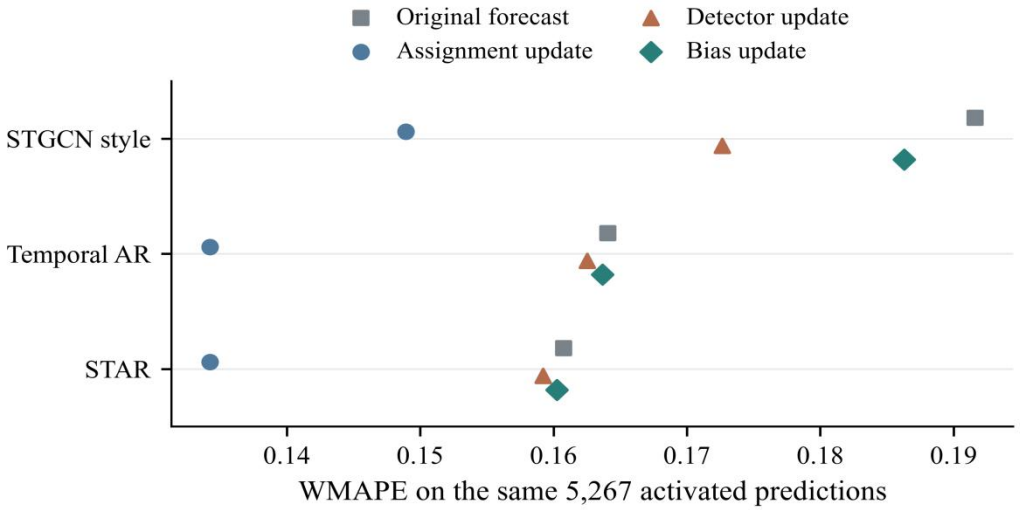


Figure 5. Accuracy of the original models and three corrections on the same 5,267 strong state cases from 413 links. Detector and bias controls retain the assignment gate. Table 3 reports the paired intervals.

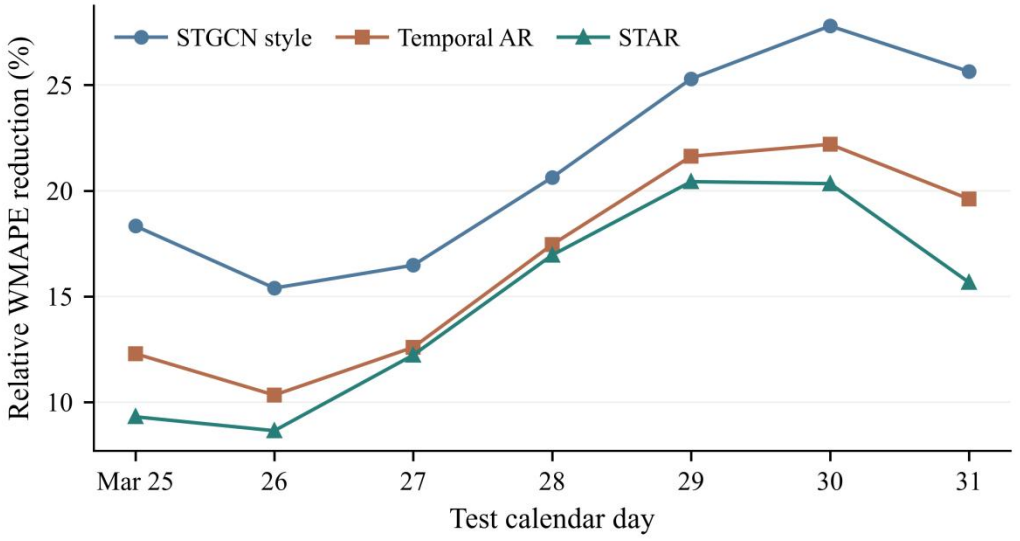


Figure 6. Assignment reductions against each original forecast for all seven test dates. Daily sample sizes are 460, 577, 588, 711, 2,079, 492 and 360. Lines connect consecutive daily summaries, with no additional smoothing.

Temporal autoregression and STAR retain gains of 1.25 and both reach WMAPE close to 0.1342. Their relative reductions are 18.19% and 16.50%, with corresponding intervals of 13.15% to 20.72% and 11.46% to 19.26%. Strong recent count forecasts continue to benefit from the assigned departure, with a smaller relative response than the weekly profile. This ordering is consistent with the information already represented in their input histories.

Figure 5 compares assignment with detector departure and constant bias corrections. The detector correction reaches WMAPE 0.1726 for STGCN style, 0.1625 for temporal autoregression and 0.1592 for STAR. The assignment correction has lower WMAPE in each comparison. Relative to the detector corrected forecast, its reductions are 13.74%, 17.40% and 15.69%. These comparisons use the same count history and the same assigned state gate.

The bias correction gives WMAPE 0.1863, 0.1637 and 0.1602 for the three backbones. A constant link offset captures part of the recurring calibration error. Assignment describes a departure that varies with the represented travel demand and route loading, allowing the response magnitude to change between targets. The comparison with the detector correction further indicates that the route departure adds a useful description of the network state within the activated group.

Figure 6 shows every test date. STGCN style daily reductions range from 15.39% to 27.79%; temporal autoregression ranges from 10.33% to 22.20%, and STAR from 8.64% to 20.43%. Each model improves on all seven dates, while the size of the benefit changes with the day's count and assigned state distribution. March 29 contains 2,079 activated predictions, or 39.5% of the group, and is retained together with the other dates.

Leaving out one test date at a time gives STGCN style reductions from 20.34% to 23.21%, temporal autoregressive reductions from 15.61% to 19.02%, and STAR reductions from 13.94% to 17.43%. The direction of the result is preserved across these comparisons. Variation in the estimate still matters because the test period contains only one week. Extending the comparison across additional months would describe a broader range of demand and road conditions.

Across all 1,104,768 predictions, STGCN style WMAPE changes from 0.168136 to 0.167846. Temporal autoregression changes from 0.159043 to 0.158841, and STAR from 0.158768 to 0.158588. The complete sample gains are smaller because the correction is restricted to 0.477% of predictions. The comparison supports a targeted update while retaining the original forecast during routine conditions.

3.5. Accuracy across Assigned State Strengths

Figure 7 separates predictions into five disjoint groups of assigned state strength. The threshold and fitted response settings are retained. Below the deployment threshold, the plot evaluates the additional response on the same test targets, while the forecast rule itself keeps the original prediction. This comparison describes how the value of the route departure changes with its magnitude, without selecting another threshold from the test week.

At 30 min, the groups below 0.10 and from 0.10 to 0.25 contain 795,406 and 245,245 predictions. Their relative WMAPE changes are -0.59% and 0.08%. The group from 0.25 to 0.50 contains 58,850 predictions and improves by 7.53%. The group from 0.50 to 1.00 contains 5,153 predictions and improves by 27.93%. The strongest group improves by 41.41%, with 114 predictions. Reporting every group makes the concentration of benefit visible alongside its frequency.

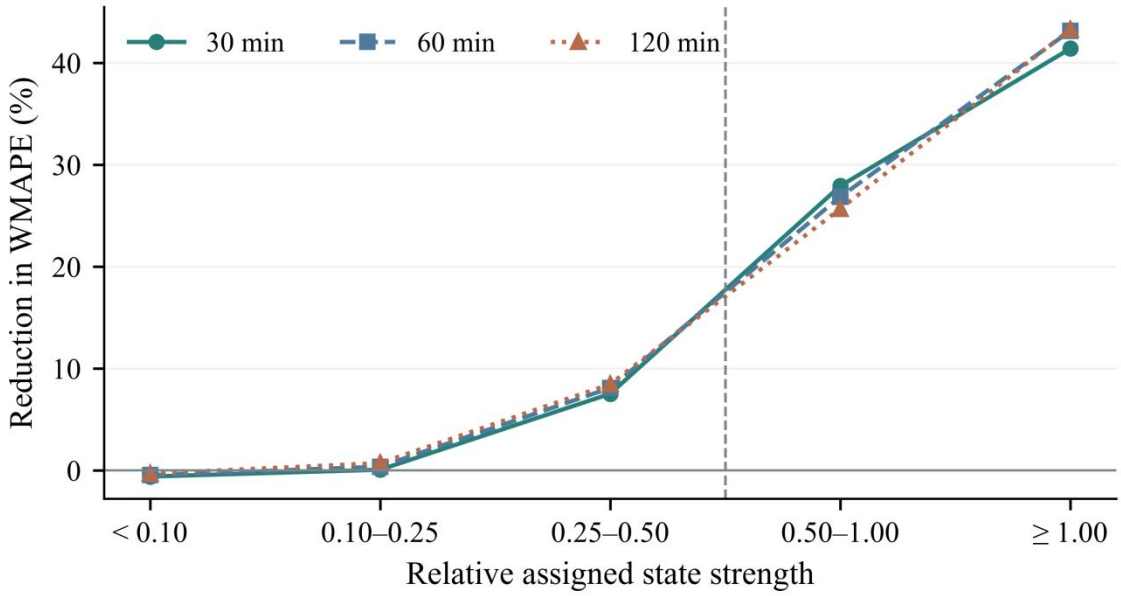


Figure 7. Correction performance in disjoint test state strength groups. The dashed boundary marks the activation threshold. Plotted corrections below it are excluded from the deployed prediction. The strongest groups contain only 114 to 139 cases.

The same ordering appears at 60 and 120 min: small assigned departures provide little additional value, intermediate departures give a modest correction, and large departures yield the largest relative reductions. The strongest groups contain 128 and 139 predictions at those horizons. Their limited size calls for caution when transferring their percentages to a whole network. The concentration of improvement above the activation threshold gives a practical reason for using the correction selectively.

3.6. Forecasts during Persistent Flow Drop Episodes

The detector screen and assignment inclusion rules produce 55 candidate episodes, of which 53 have sufficient assigned flow for comparison. Their dates place 26 in calibration, 13 in validation and 14 in testing. The test episodes occur on March 29 to 31 and include all qualifying episodes in the reserved test period. The first low count is available at issue time, and the target is the following 5 min interval under the supplied road cost setting.

Table 4. Forecasts during persistent flow drops: 13 validation and 14 test episodes. MAE and RMSE are in vehicles per 5 min.

Forecast	Validation WMAPE	Test WMAPE	Test MAE	Test RMSE
Current count persistence	0.5207	0.8552	8.86	12.54
Fixed service ratio	0.3901	0.5375	5.57	8.11
Fixed service ratio + count cap	0.3876	0.4414	4.57	6.78
Raw assignment ratio	0.8552	0.9651	10.00	12.88
Calibrated assignment ratio	0.2907	0.3979	4.12	6.07
Calibrated assignment + count cap	0.2907	0.3294	3.41	5.07

The fixed service forecast and its capped form provide compact references. A raw assigned flow ratio can predict an excessively large reduction because the represented OD trips and the count observed at a detector have different scales. The calibrated ratio adjusts this response using the earlier episodes. The selected ridge strength remains 1.0, and the cap carries forward the upper bound supplied by the current count. These components address response scale and the amount of traffic already observed after the drop begins.

Table 4 reports the episode comparison. Persistence gives test WMAPE 0.8552. The fixed service ratio lowers it to 0.5375, and adding the issue count cap lowers it to 0.4414. The raw assignment ratio gives 0.9651. Calibrating the ratio gives WMAPE 0.3979 without the cap and 0.3294 with it. The calibrated capped result has MAE 3.41 and RMSE 5.07, compared with 4.57 and 6.78 for the fixed capped forecast. The lower RMSE indicates that the additional correction also improves the scale of the larger errors.

Figure 8 shows all 14 test episodes in chronological order. Calibrated service ratios range from 0.155 to 0.273. The capped forecast improves absolute error on 12 episodes, increases it on 1 and leaves 1 unchanged. It follows count variation more closely than the fixed ratio, with remaining errors reflecting local variation.

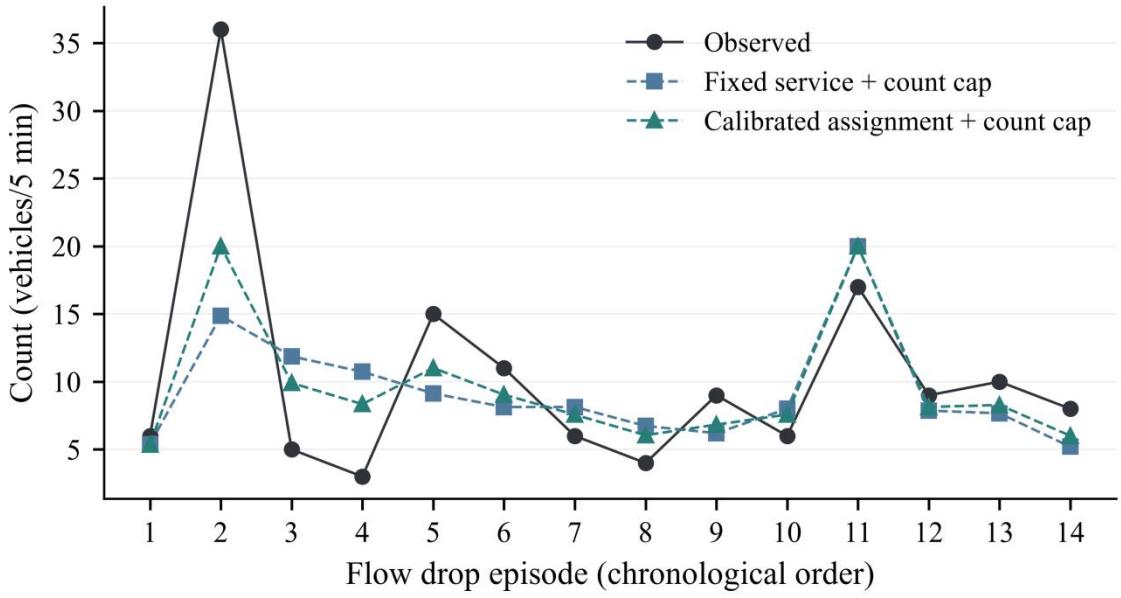


Figure 8. Observed and predicted counts for all 14 test flow drop episodes in chronological order. Lines connect separate episodes to aid comparison. Both service forecasts use the same issue count cap.

The WMAPE reduction against the fixed capped forecast is 25.37%, with a paired episode interval of 21.65% to 28.62%. This is an additional contribution after the simple service rule has already captured much of the decrease. Calibration permits the route response to vary between episodes instead of reducing all counts by the same proportion. A larger collection of independently documented disruptions would provide a broader basis for estimating that response and assessing its transfer across locations.

3.7. Forecasts One to Seven Days Ahead

The longer lead comparison uses the same 552,384 target link time cases at each lead. The validation settings remain a ridge multiplier of 2.0, a correction threshold of 0.126 and a gain of 0.25. Persistence WMAPE is 0.1826 at one day, 0.2171 at three days and 0.1586 at seven days. Recursive autoregression is close to 0.595. The weekly profile gives WMAPE 0.1556, and assignment alone gives 0.4266. These references show the value of retaining a stable observed pattern as the starting point.

The combined forecast gives WMAPE 0.1455 at one day, 0.1512 at three days and 0.1527 at seven days. The corresponding complete sample reductions are 6.47%, 2.82% and 1.83%. The targets and weekly profile are shared across leads. Their corrections differ because the issue state and its distance from the target affect which supplied departures receive an update. The relative benefit decreases as the forecast reaches further ahead.

At seven days, the correction acts on 7.4% of predictions. Within this group, WMAPE

decreases from 0.2667 to 0.2062; its relative reduction interval is 18.11% to 25.44%. The complete sample benefit is smaller because the weekly profile is retained on the remaining targets. These results support a calibrated contribution from the supplied OD scenario while showing why the recurring pattern continues to determine much of a long lead forecast.

4. Discussion

4.1. When the Assigned Response Helps

Assignment adds useful information when travel demand and route use depart from the conditions represented by the starting forecast. A supplied restriction changes the available route costs. Recomputing the route loading then describes where flow can move across the network before the full response appears in recent counts. In the Wuxi cost comparison, this response reduces the larger differences from the independently solved flow state. The roads with remaining local errors indicate where route alternatives and loading assumptions deserve closer attention.

The detector comparisons extend this role to changing demand patterns. An assigned departure describes a distribution of represented trips across the network. It can carry information about approaching traffic and the interaction among origins, destinations and available paths. Recent detector histories describe the traffic already observed at particular locations. Their combination supports the larger improvements within strong assigned states and the retained benefit across several forecasting backbones.

The correction remains useful after a detector departure or a constant link bias is allowed to improve the same starting forecast. This comparison relates the benefit to the changing route state and its response amplitude. Each control uses the same assigned state gate, so the result concerns the magnitude of the correction after that state has been identified. A broader comparison could evaluate alternative state selection rules with the same range of demand inputs and observation histories.

The state strength comparison explains why the correction is selective. Weak departures account for most predictions and supply little additional information. Large departures occur less often, yet coincide with larger errors in the recurring profile. The activation rule concentrates the additional response within those conditions. A useful local correction can consequently coexist with a smaller change in the complete sample score. This relationship should guide how an agency interprets the accuracy gain and the calculation needed to obtain it.

Calibration gives the assigned response an appropriate count scale. OD loading and detector measurements represent different trip populations and spatial coverage. The controlled response strength, link coefficients and episode service ratio each translate a route response into the corresponding forecast quantity. The episode results show the value of adjusting the raw ratio before using it to estimate a local count. The remaining episode variation supports further evaluation across different road conditions and observation locations.

4.2. Implications for Traffic Forecasting Practice

An agency can introduce the route correction alongside an existing detector forecast. The count model supplies the recurring traffic estimate, while assignment supplies a response to changing road costs or represented demand. A consistent mapping between modeled links and detector locations allows the response to be interpreted at the point where a forecast is required. Updating the usual assigned pattern is also important when the demand distribution changes over a longer period.

Input timing determines the operational use. A published restriction supplies information about a future road condition and permits assignment before the restriction begins. A current assigned departure uses a recent demand estimate to describe the developing network state. The persistent drop comparison uses the first low count and predicts the following interval. These applications should be evaluated with the information available at their respective issue times so that their accuracy corresponds to the decisions they support.

The forecasting backbone can remain compact. Each link requires one response coefficient, and validation determines the global gain. Detector departure and constant bias corrections provide useful continuing references. Comparing these alternatives helps establish whether route loading adds enough information to justify its maintenance and calculation. The assignment update can be retained where its contribution is consistent across dates and operating conditions.

The long lead results support keeping the weekly profile as an anchor. Current count departures lose influence as the target moves further ahead. A supplied future demand scenario can still direct an additional route response toward targets that differ from their recurring pattern. Its benefit should be assessed across several dates, demand levels and seasons, with attention to the accuracy and availability of that demand feed.

Link level performance also matters for diversion decisions. A lower network average can accompany errors on individual alternative roads. The predicted redistribution should be

considered together with road capacity, local access and incoming observations. The present comparisons evaluate flow accuracy. Operational route management would require these local constraints and the uncertainty in restriction duration to be represented alongside the forecast.

5. Conclusion

This study developed an assignment based update for link flow forecasting in Wuxi. The method uses changes in assigned flows to adjust an existing detector forecast, calibrates the response scale and applies the correction during substantial assigned departures. The evaluation combines Wuxi detector counts, OD demand and the available road network through road cost comparisons, several forecasting backbones, persistent flow drops and longer prediction horizons.

The findings indicate that route information is most useful when it describes a substantial change beyond the recurring traffic pattern. Calibration allows that response to improve a detector forecast while accommodating differences between represented trips and observed counts. Strong assigned states give the largest benefit, and recent count models retain a useful additional correction. The consistent direction across dates supports selective use of the route response. The complete sample results show that its contribution depends on how frequently those states occur.

The study covers one city over a short observation period, with OD demand and detector counts derived from the same passage dataset. Further work should evaluate independent demand sources across cities and seasons. Documented disruptions with measured severity and duration would support a closer assessment of redistribution. Repeated neural model training and evaluation with delayed or incomplete demand inputs would establish forecast stability across a broader range of operating conditions.

Acknowledgements

This study was supported by the Open Research Project of the National Engineering Research Center of Road Safety Control Technology (grant 2024GCZXKFKT08B).

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author Contributions

Tianbin Wu: Conceptualization, Methodology, Writing – original draft. Tong Liu: Software, Writing – review & editing. Xihua Liu: Software, Writing – review & editing. Kun Jin: Supervision, Conceptualization, Writing – review & editing.

References

[1] Yu, B., Yin, H., & Zhu, Z. (2018). Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (pp. 3634–3640). International Joint Conferences on Artificial Intelligence Organization. https://doi.org/10.24963/ijcai.2018/505

[2] Wu, Z., Pan, S., Long, G., Jiang, J., & Zhang, C. (2019). Graph WaveNet for Deep Spatial-Temporal Graph Modeling. In Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence (pp. 1907–1913). International Joint Conferences on Artificial Intelligence Organization. https://doi.org/10.24963/ijcai.2019/264

[3] Wang, W., Wang, J., Zhao, D., & Jin, K. (2022). A Revised Logit Model for Stochastic Traffic Assignment With a Relatively Stable Dispersion Parameter. IEEE Intelligent Transportation Systems Magazine, 14(2), 92–104. https://doi.org/10.1109/MITS.2021.3131887

[4] Peeta, S., & Ziliaskopoulos, A. K. (2001). Foundations of Dynamic Traffic Assignment: The Past, the Present and the Future. Networks and Spatial Economics, 1(3), 233–265. https://doi.org/10.1023/A:1012827724856

[5] Jiang, W., & Luo, J. (2022). Graph neural network for traffic forecasting: A survey. Expert Systems with Applications, 207, 117921. https://doi.org/10.1016/j.eswa.2022.117921

[6] Zhao, L., Guo, B., Dai, C., Shen, Y., Chen, F., Zhao, M., & Hu, Y. (2024). Multi-step trend aware graph neural network for traffic flow forecasting. Big Data Research, 38, 100482. https://doi.org/10.1016/j.bdr.2024.100482

[7] Wang, Z., Wang, Y., Jia, F., Zhang, F., Klimenko, N., Wang, L., He, Z., Huang, Z., & Liu, Y. (2024). Spatiotemporal Fusion Transformer for large-scale traffic forecasting. Information Fusion, 107, 102293. https://doi.org/10.1016/j.inffus.2024.102293

[8] Hu, J., Zheng, T., Peng, L., Teng, F., Du, S., & Li, T. (2025). LightST: A Simplifying Spatio-Temporal Graph Neural Network for Traffic Flow Forecasting. IEEE Transactions on Big Data, 11(5), 2517–2528. https://doi.org/10.1109/tbdata.2025.3544131

[9] Jia, X., Dong, X., Chen, M., & Yu, X. (2021). Missing data imputation for traffic congestion data based on joint matrix factorization. Knowledge-Based Systems, 225, 107114. https://doi.org/10.1016/j.knosys.2021.107114

[10] Kong, X., Zhou, W., Shen, G., Zhang, W., Liu, N., & Yang, Y. (2023). Dynamic graph convolutional recurrent imputation network for spatiotemporal traffic missing data. Knowledge-Based Systems, 261, 110188. https://doi.org/10.1016/j.knosys.2022.110188

[11] Hou, M., Tang, T., Xia, F., Sultan, I., Kaur, R., & Kong, X. (2024). MissII: Missing Information Imputation for Traffic Data. IEEE Transactions on Emerging Topics in Computing, 12(3), 752–765. https://doi.org/10.1109/tetc.2023.3280481

[12] Li, H., Han, S., Yang, M., Liu, J., Zhou, J., Zhang, T., & Philip Chen, C. L. (2025). Dynamic Spatial-Temporal Imputation Network With Missing Features for Traffic Data Imputation. IEEE Internet of Things Journal, 12(14), 29005–29017. https://doi.org/10.1109/jiot.2025.3567697

[13] Zhang, Y., Kong, X., Zhou, W., Liu, J., Fu, Y., & Shen, G. (2024). A Comprehensive

Survey on Traffic Missing Data Imputation. IEEE Transactions on Intelligent Transportation Systems, 25(12), 19252–19275. https://doi.org/10.1109/tits.2024.3478816

[14] Huang, Y., & Yang, J. J. (2024). Symmetric contrastive learning for robust fault detection in time-series traffic sensor data. International Journal of Data Science and Analytics, 20(2), 1199–1213. https://doi.org/10.1007/s41060-024-00521-0

[15] Gagliardi, G., Casavola, A., & D'Angelo, V. (2023). Reliable Traffic Sensor Networks Via Fault-Tolerant Sensor Reconciliation Schemes. IFAC-PapersOnLine, 56(2), 10558–10563. https://doi.org/10.1016/j.ifacol.2023.10.312

[16] Mishra, M., Mishra, S., & Har, D. (2024). Integrating Multisourced Sensor Data for Enhanced Traffic State Estimation. IEEE Sensors Journal, 24(12), 19614–19625. https://doi.org/10.1109/jsen.2024.3397534

[17] Fulari, S., Vanajakshi, L., & Subramanian, S. C. (2017). Artificial Neural Network–Based Traffic State Estimation Using Erroneous Automated Sensor Data. Journal of Transportation Engineering, Part A: Systems, 143(8), 05017003. https://doi.org/10.1061/jtepbs.0000058

[18] Aslam, J., Lim, S., Pan, X., & Rus, D. (2012). City-scale traffic estimation from a roving sensor network. In Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems (pp. 141–154). ACM. https://doi.org/10.1145/2426656.2426671