Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

# Corridor-Wide Freeway Speed Field Reconstruction from Sparse Detector Data: A Deep Learning Framework Supervised by Connected Vehicle Data

**Md Mushtaque Tahmid**

md891073@ucf.edu

University of Central Florida

**Zubayer Islam**

University of Central Florida

**Mohamed Abdel-Aty**

University of Central Florida

**Research Article**

**Keywords:** Freeway speed, Microwave vehicle detection system, Graph neural network and Connected Vehicle Data

**Posted Date:** November 24th, 2026

**DOI:** https://doi.org/10.21203/rs.3.rs-11623524/v1

**License:** This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

**Additional Declarations:** No competing interests reported.

# Corridor-Wide Freeway Speed Field Reconstruction from Sparse Detector Data: A Deep Learning Framework Supervised by Connected Vehicle Data

Md Mushtaque Tahmid[1*], Zubayer Islam, Ph.D. [2], and Mohamed Abdel-Aty, Ph.D. [3]

[1] PhD Student, Department of Civil, Environmental and Construction Engineering, University of Central Florida, Orlando, FL, USA, 32816; email: md891073@ucf.edu (Corresponding Author)

[2] Assistant Professor, Department of Civil, Environmental and Construction Engineering, University of Central Florida, Orlando, FL, USA, 32816; email: Zubayer.Islam@ucf.edu

[3] Pegasus Professor, Department of Civil, Environmental and Construction Engineering, University of Central Florida, Orlando, FL, USA, 32816; email: M.Aty@ucf.edu

*Corresponding author

**Abstract:** Accurate corridor-wide traffic speed information is essential for freeway operations, incident management, and proactive traffic control. However, conventional Microwave Vehicle Detection Systems (MVDS) provide traffic measurements only at discrete sensor locations, leaving large portions of freeway corridors unobserved. This study formulates corridor-wide speed field reconstruction as a cross-modal supervised learning problem: sparse MVDS speeds are the only input at deployment, and Connected Vehicle (CV)-derived speeds on a 0.05-mile grid serve as supervision during training only. A two-stage spatiotemporal deep learning framework was developed for the State Road SR-417, Florida; unlike fusion methods, it requires no probe data at deployment, and unlike imputation methods, it learns from a sensing modality independent of the detectors. Stage 1 predicted speed at detector-adjacent locations; twelve models spanning recurrent, transformer, and graph architectures were evaluated. Stage 2 propagated Stage 1 estimates through a gap-filling framework integrating graph-based spatial learning with temporal modelling to reconstruct speed across unobserved segments. An Attention-Based Spatiotemporal Graph Convolutional Network (ASTGCN) provided superior performance under congested conditions and was selected as the Stage 1 backbone, achieving a Mean Absolute Error (MAE) of 4.450 mph. The full pipeline reconstructed a continuous speed field at 0.05-mile resolution, achieving an MAE of 5.072 mph. Detector-only observations underestimated the spatiotemporal extent of congestion, whereas the reconstructed field recovered the congestion footprint with close agreement to connected vehicle ground truth. The framework supports shockwave localization, congestion severity quantification, and Active Traffic Management evaluation for agencies operating sparse fixed-sensor infrastructure.

**Keywords:** Freeway speed, Microwave vehicle detection system; Graph neural network and Connected Vehicle Data

## Statements and Declarations

No funding was received to assist with the preparation of this manuscript. The data that support the findings of this study are available from the corresponding author upon reasonable request. 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.

## Introduction

Traffic speed is a fundamental freeway-management parameter, supporting real-time traffic control, incident detection, travel time estimation, and route choice (Abdelraouf et al. 2022b; Qian et al. 2024; Zhang et al. 2025a). Prior studies have substantially advanced traffic speed estimation and prediction methods (Abdelraouf et al. 2022b; Yu et al. 2017; Zhao et al. 2018; Zou et al. 2023; Zou et al. 2024; Zhang et al. 2025c). Existing studies have relied on aggregated macro-level traffic features such as average speed and volume collected at discrete fixed sensor locations to model spatiotemporal dependencies and estimate traffic conditions (Yu et al. 2017; Zhao et al. 2018; Zou et al. 2023; Zhang et al. 2025c). Microwave Vehicle Detection Systems (MVDS) and inductive loop detectors remain the principal freeway monitoring infrastructure, providing high-fidelity measurements at their installation points but no coverage between them. On SR-417 in Central Florida, detectors are spaced on average 0.7 miles apart, leaving the majority of the spatial domain unobserved. The continuous distribution of speed over space and time along the corridor, referred to as the speed field, cannot be directly measured between adjacent detectors. Reconstructing this field, particularly during non-recurrent congestion events, remains an operational challenge for traffic management.

Connected vehicle (CV) technology offers to complement fixed-sensor infrastructure with continuous, trajectory-based traffic observations. Unlike fixed sensors that provide traffic information at sparse and discrete points, CV trajectories span full road segments and capture continuous traffic dynamics along the corridor (Han et al. 2024a; Wang et al. 2024). CV data offer high temporal resolution at 1–3 second update intervals, precise spatial positioning at decimeter level, and rich kinematic information including instantaneous speed and acceleration (Gu et al. 2023; Han et al. 2024b; Zhang and Abdel-Aty 2022). These characteristics enable the extraction

of speed volatility which serve as leading indicators of traffic fluctuation and imminent congestion onset, capturing traffic dynamics that are entirely invisible to aggregated fixed-sensor data.

CV data platforms have extended these capabilities to freeway-network scale. At an estimated penetration rate of 4–5% on Florida freeways (Han et al. 2026b), StreetLight CV data provides individual probe trajectories into segment level speed estimates that capture both recurrent daily patterns and non-recurrent congestion events. The spatial coverage across the SR-417 corridor exceeds that of the fixed MVDS detector network, which provides measurements only at discrete locations separated by more than half a mile on average. This corridor-wide spatial continuity is the property that makes StreetLight CV speed a uniquely suitable candidate as a spatial ground truth for supervised freeway speed field reconstruction.

The problem of reconstructing a complete, spatially continuous speed field along a freeway corridor from sparse fixed sensor observations remains unaddressed. While recent studies have validated StreetLight CV data for freeway speed prediction (Han et al. 2026b) and city-scale arterial traffic state estimation (Han et al. 2026a), neither addresses the gap-filling formulation in which CV-derived speed serves as the spatial training target for reconstruction between sparse detectors. Deep learning approaches have focused on forecasting at instrumented locations rather than on the space between them. Furthermore, the potential of high-resolution CV probe data as a supervised training signal for corridor-scale gap filling has not yet been explored. The setting in which a spatially dense, independent modality supervises training while a sparse modality is the only input at deployment has not been examined for freeway speed field reconstruction. These gaps motivate the present study.

This paper proposes a two-stage spatiotemporal pipeline for speed field reconstruction on SR-417. In Stage 1, a comprehensive comparative evaluation of twelve spatiotemporal models was

first conducted to identify the most suitable model for estimating StreetLight CV-consistent speed at detector-adjacent (near-detector) locations from sparse MVDS real-time and baseline speed. In Stage 2, a corridor-wide gap filler model propagated the Stage 1 speed signal across all unobserved spatial locations between detectors, producing a complete speed field validated against CV ground truth at 0.05-mile spatial resolution.

The main contributions of this paper are as follows. First, corridor-wide speed field reconstruction was formulated as a cross-modal supervised learning problem, in which StreetLight CV-derived speeds at 0.05-mile resolution serve as dense spatial supervision during training only, while sparse MVDS speeds are the sole input at deployment. Unlike fusion methods, this formulation requires no probe data at deployment, and unlike imputation methods, it learns from a sensing modality independent of the detectors. Second, a two-stage spatiotemporal deep learning framework was developed that first estimates CV-consistent speeds at detector-adjacent locations from MVDS inputs and then reconstructs speeds at all unobserved locations between detectors through a Graph–LSTM gap filler. Within this framework, congestion-aware loss weighting emphasizes abnormal traffic conditions, the most operationally critical regime for freeway management, and teacher forcing with CV ground-truth anchors decouples gap filler training from Stage 1 estimation errors. Third, the reconstructed speed field was evaluated beyond pointwise accuracy: in representative congestion events, detector-only observations substantially underestimated the spatiotemporal extent of congestion, whereas the reconstructed field recovered the congestion footprint in close agreement with CV ground truth.

## Literature Review

### *Multi-Source Traffic Data Fusion*

Historically, freeway speed monitoring has relied on fixed MVDS installed at discrete locations. These provide high-fidelity point measurements, but installation and maintenance costs limit their spatial coverage. On SR-417 in Central Florida, detectors are spaced more than half a mile apart on average, leaving large gaps in the observed speed field. Reconstructing a continuous spatiotemporal speed field across these gaps is the central challenge addressed in this paper.

The earliest alternative to fixed sensors was the floating car, in which a GPS-equipped vehicle reports position and instantaneous speed at regular intervals as it traverses the network. Even a small number of probe vehicles improves freeway speed estimation when fused with loop detector data through Kalman filtering (Nanthawichit et al. 2003). An early systematic comparison evaluated seven fusion techniques for combining loop detector and Bluetooth probe data on a micro simulated stretch of a major freeway in Toronto (Bachmann et al. 2013a). Most improved on single-source accuracy, but the benefit depended heavily on probe penetration rate and traffic state, with loop detectors proving unreliable under congestion where point measurements fail to capture the true average link speed. A follow-up study on real-world Highway 401 data confirmed these findings. Fused estimates were at least as accurate as the better of the two sources, even though which source was more reliable changed with conditions (Bachmann et al. 2013b). Both studies aimed at better speed estimates at detector locations, not at reconstructing the unobserved space between them.

A Progressive Extended Kalman Filter assimilated cellphone activity records alongside microwave sensor readings within a second-order macroscopic flow model, producing speed and density estimates across all links even where no fixed sensor exists (Liu et al. 2018). The method

remains tightly coupled to a calibrated flow model, requires manual specification of physical parameters and link boundaries, and depends on cellular infrastructure with sufficient spatial resolution; these are conditions that differ substantially from the connected vehicle trajectory setting. On the German A9, a comparison of adaptive and phase-based smoothing methods showed that adding Bluetooth data does not always improve reconstruction quality and can degrade space-time speed accuracy when detector spacing is large relative to congestion feature scales, while floating car data combined with loop detectors consistently outperformed all other sensor combinations (Kessler et al. 2021). A multisensor system for an urban Beijing corridor fused coil detectors, microwave sensors, and GPS probes, but targeted discrete traffic state classification rather than continuous speed field estimation (Zeng et al. 2014).

Taken together, these studies established that fixed-sensor and probe fusion is well-motivated, that no single fusion technique dominates universally across traffic conditions and sensor configurations, and that probe data carries information that sparse fixed sensors simply cannot replicate across spatial gaps. Discussed methods used probe data as an input at estimation; none used probe data solely as a training signal for a model that afterwards operates on detector data alone.

***Spatiotemporal Deep Learning for Traffic Speed Modeling***

Existing traffic speed prediction methods can broadly be categorized into three groups: statistical approaches, machine learning algorithms, and deep learning models. Statistical methods such as Kalman filtering (Liu et al. 2018; Tampere and Immers 2007), historical averaging (Yao et al. 2017), and autoregressive integrated moving average models along with their variants (Yao et al. 2017) have been widely applied, but their reliance on strict theoretical assumptions limits their ability to capture the complex and nonlinear dynamics of real-world traffic conditions. Machine

learning approaches including gradient boosting decision trees (Wu et al. 2021), support vector machines (Yao et al. 2017), hidden Markov models (Jiang and Fei 2017), and artificial neural networks (Li et al. 2017; Wu et al. 2019) improve upon these limitations by extracting nonlinear relationships from high-dimensional traffic data. However, they predominantly model temporal dependencies while failing to account for the spatial structure inherent in road networks, which is particularly critical for corridor-level speed reconstruction.

Recent studies have primarily focused on modeling spatial-temporal dependencies of macro traffic flow features with deep learning (DL) models (e.g., Spatial-Temporal Graph Neural Networks (STGNNs) (Cheng et al. 2025; Yu et al. 2017) and Transformer models (Ouyang et al. 2025; Zhang et al. 2025c). Deep learning models have emerged as the dominant paradigm for spatiotemporal modeling. Spatiotemporal graph neural networks (STGNNs) were proposed to simultaneously capture spatial and temporal dependencies in traffic data, achieving substantially improved prediction performance (Guo et al. 2019; Yu et al. 2017). A variety of STGNN-based architectures including Graph WaveNet (Wu et al. 2019), ASTGCN (Guo et al. 2019), and T-GCN (Zhao et al. 2018) have been developed to explore spatiotemporal relationships from different perspectives, each offering distinct advantages in modeling the complex interdependencies of traffic networks. Transformer-based models have recently demonstrated strong performance through attention mechanisms that learn important spatiotemporal traffic dynamics (Ouyang et al. 2025; Zhang et al. 2025c).

Data from five toll collection gates and a microwave sensor were fused in a hybrid stacked autoencoder–LSTM (Long Short-Term Memory) model to forecast traffic flow, and higher accuracy was achieved than with single-model or single-source alternatives (Jin et al. 2018). An attention-based multi-encoder-decoder predicted 60-minute-ahead speeds on State Road 408 in

Orlando from short-term, daily, and weekly traffic sequences (Abdelraouf et al. 2022b), and MMCAformer integrated macro flow features with micro connected vehicle behavior across four Florida freeways, with gains most pronounced under congestion (Han et al. 2026b). CrossGraphNet, a bidirectional spatiotemporal graph neural network, shares the reconstruction objective. It was pre-trained on 0.5-m satellite imagery and sparse ground sensors covering about 6% of road segments in the Yuen Long District of Hong Kong, after which street-level speeds were estimated from satellite-detected vehicle densities alone (Zhang et al. 2025b). Satellite revisit cycles of several days, however, preclude intraday operations.

None of these works reconstructs a complete, continuous spatiotemporal speed field across a freeway corridor where fixed sensors exist only at discrete, widely spaced locations and where the reconstruction target is not a point estimate or a state label but a full speed field resolving fine-scale congestion dynamics. This is precisely the gap the present study fills. In this study, StreetLight connected vehicle GPS trajectories, with their corridor-wide spatial coverage, serve as ground truth to supervise a model that learns to propagate speed information across the unobserved space between MVDS loop detectors, producing a corridor-wide speed field.

**Data Collection and Preparation**

Road network geometry was obtained from a shapefile of the SR-417 corridor containing individual link segments with unique identifiers. Start and end coordinates were extracted for each segment to reflect travel direction, northbound or southbound. Segment lengths were computed using the haversine formula, and spatial bins were defined at 0.05-mile resolution spanning the full corridor length. Each MVDS detector was mapped to five spatial consecutive bins centered on its milepost location. **Fig. 1** shows the study corridor and MVDS detector locations, with an average inter-detector spacing of approximately 0.7 miles.

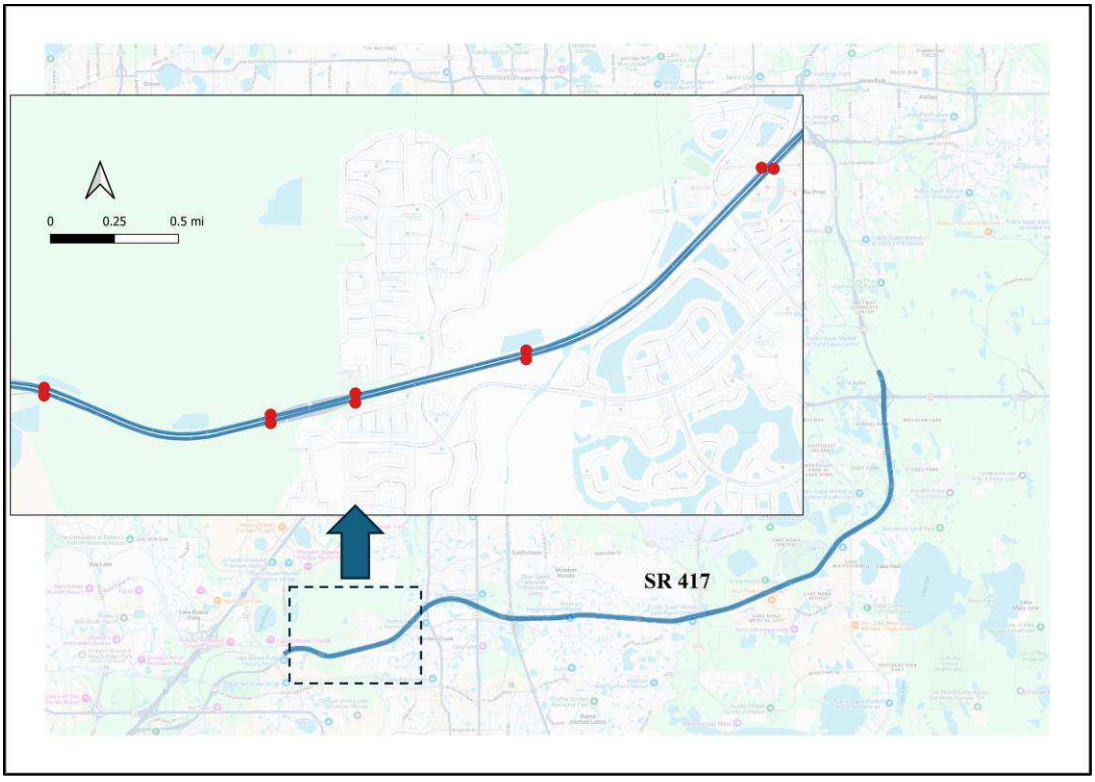



**Fig. 1 Study Corridor with Locations of MVDS Detectors (red points)**

MVDS detector data were loaded at per-minute speed. For each detector and each calendar day in the study period, a dense array of 1,440 minute-resolution speed entries was pre-constructed, allowing any 30-minute analysis window to be extracted as a contiguous array slice.

StreetLight probe vehicle trajectories were matched to corridor segments, projected onto the road geometry, and validated at the journey level through the three-step procedure as shown in **Fig. 2**. Through Geofence Filtering, GPS points falling within a segment's spatial buffer were retained as candidates for that segment. Then points were grouped by journey identifier, and any journey containing a temporal gap of 10 seconds or more between consecutive observations was discarded, as such gaps indicate a discontinuous or unreliable trace.

The study period spanned September 1–15, 2025, in the US/Eastern time zone. Analysis windows of 30 minutes duration were extracted using a sliding window approach, advancing in

one-minute steps across all daytime hours (06:00–21:00). For each window and each aggregation granularity (1, 2, 3, and 5 minutes), MVDS speed readings were aggregated into temporal slots by computing the mean across constituent minutes. Probe vehicle observations within the window were aggregated by computing the mean speed of unique journeys per spatial bin per time slot.

A baseline speed array was constructed for each window by identifying all other windows in the dataset that shared the same time-of-day and day-of-week type (weekday or weekend). Candidate windows from other dates were stacked, and the median value was taken across entries. This provides a reference representation of typical conditions, which serves as a second input feature alongside the MVDS readings.

Each window was labeled as normal, or abnormal based on the MVDS and probe speed signals. Windows in which both signals indicated free-flow conditions (minimum speed above 57 mph) were labeled normal. The 57 mph threshold was used because speeds below 57 mph on a 70 mph posted speed facility are indicative of at least moderate congestion onset, as reflected in the Highway Capacity Manual (HCM). Windows in which both signals indicated congestion were labeled abnormal. Windows with missing values were discarded. For normal windows, September 1–11 was designated as the training period and September 12–15 as the test period, following a strict chronological partition to prevent data leakage. Because abnormal windows are rare and concentrated around specific congestion events, a strict date cut might yield too few abnormal samples for training. A separate rank-based split was therefore applied to this subset: the earliest 70 percent of abnormal observations sorted by timestamp were assigned to training and the remainder to testing, ensuring sufficient representation of congested conditions during model learning.

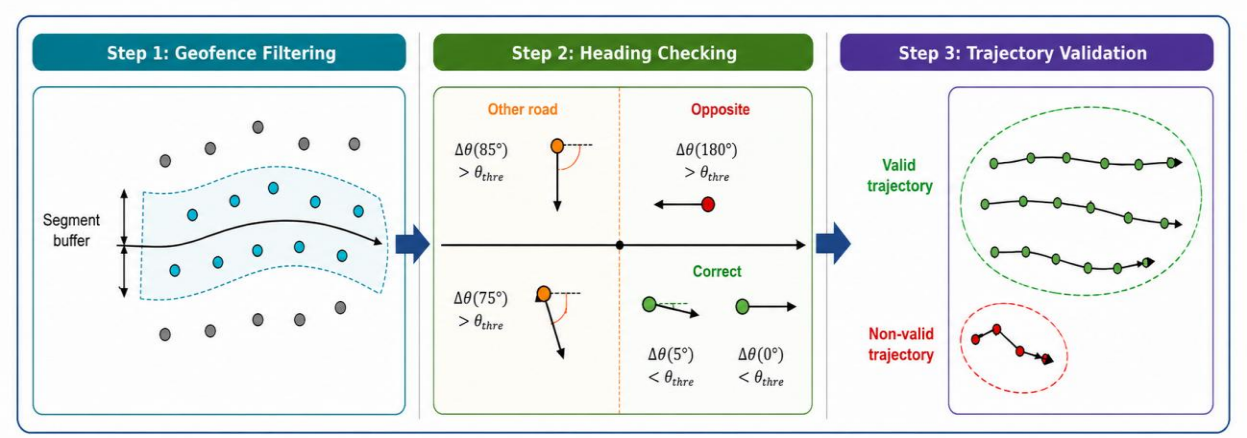



**Fig. 2 CV Data Spatial Matching Steps**

A scalar market penetration estimate was computed for each window as the ratio of unique probe vehicle journeys observed within the detector's spatial bin group to the total MVDS vehicle volume recorded over the same 30-minute period. This quantity provides a window-level indication of the fraction of traffic represented in the probe data and is stored alongside the speed arrays for use in subsequent modeling. **Fig. 3** shows the data preparation flowchart.

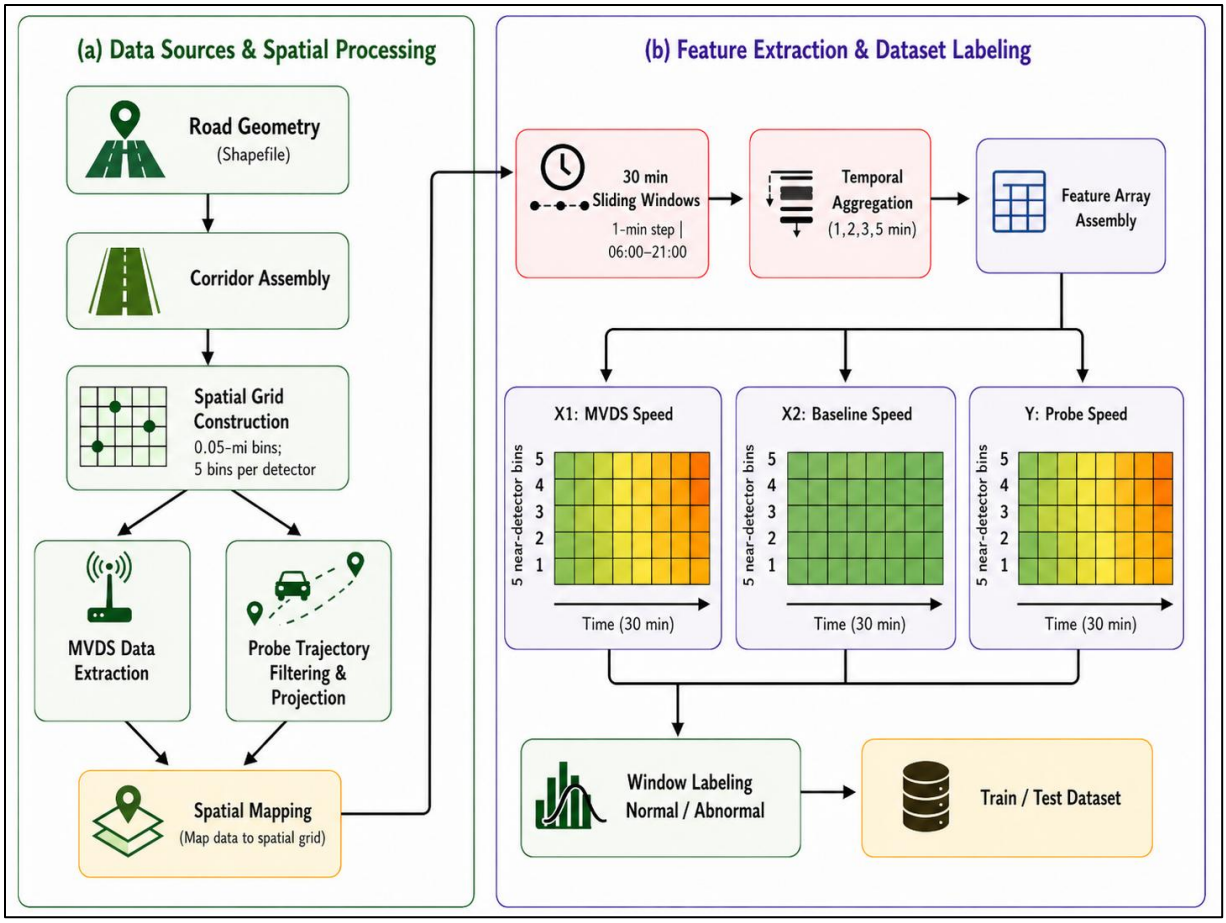



**Fig. 3 Data Preparation Flowchart**

## Methodology

### *Stage 1: Comparative Model Evaluation for Detector-adjacent Locations Speed Prediction*

The first stage of the proposed pipeline addresses the prediction of speed field at detector-adjacent (near-detector) spatial bin locations from sparse MVDS real-time and baseline speed inputs. A systematic comparative evaluation of twelve spatiotemporal deep learning models was conducted on the dataset of five-bin near-detector groups.

All twelve models share a common input formulation and a unified gated dual-branch architecture, enabling a fair comparison that isolates the effect of the spatial-temporal encoder design from input and fusion differences. The input to each model consists of three components:

$X_1 \in \mathbb{R}^{T \times N_a}$ : Real-time MVDS speed at the five active bins ($N_a$= 5) surrounding each detector, across $T$ = 15 temporal slots of 2-minute resolution within the 30-minute analysis window,

$X_2 \in \mathbb{R}^{T \times N_a}$ : Baseline MVDS speed at the same five active bins ($N_a$= 5) surrounding each detector, across $T$ = 15 temporal slots of 2-minute resolution within the 30-minute analysis window,

$abn \in \{0,1\}$ : Binary abnormality flag: 1 if the minimum MVDS speed across active bins falls below 57 mph, indicating congested conditions; 0 otherwise

The target output is $Y \in \mathbb{R}^{T \times N_a}$ , the CV probe speeds at the five near-detector bins, which serve as the ground truth.

Each model processes $X_1$ and $X_2$ through two independent encoding branches, a real-time branch (Branch-RT) and a baseline branch (Branch-Base) before fusing their representations through a learned sigmoid gate. The gate is computed as:

$$G = \sigma\left(W_g\left[H_{rt} \,\|\, H_{base} \,\|\, dev_{proj} \,\|\, abn_{tiled}\right] + b_g\right) \quad (1)$$

$$H_{fused} = G \odot H_{rt} + (1 - G) \odot H_{base} \quad (2)$$

Here, $H_{rt}$ and $H_{base}$ are the encoded representations from each branch, $dev_{proj}$ is a linear projection of the absolute difference $|X_1 - X_2|$ to the hidden dimension, providing the gate with an explicit signal of how much the current traffic state deviates from its historical baseline, and $abn_{tiled}$ is the abnormality flag broadcast and tiled to match the spatial and temporal dimensions of $H_{rt}$. Gate $G$ allows the model to weigh the real-time branch more heavily when $abn$ = 1 (congested conditions with unreliable baseline patterns) and lean toward the baseline branch when

$abn$ = 0 (normal recurrent conditions). The fused representation $H_{fused}$ was passed to a TimeDistributed Dense output layer with sigmoid activation to produce the predicted speed at all five active bins across all T time slots.

All twelve models were trained using the same asymmetric weighted Huber loss, separating the contribution of normal and abnormal training windows:

$$L_{total} = 0.3\, L_{Huber}^{normal} + 0.7\, L_{Huber}^{abnormal} \quad (3)$$

The 0.7 weight on abnormal windows ensures that the model allocates disproportionate learning capacity to congested conditions, preventing the class imbalance between normal and abnormal windows from biasing training entirely toward free-flow accuracy. This weighting addresses the class imbalance between normal and abnormal training windows, ensuring abnormal cases are not underweighted during training (Han et al. 2026a).

The twelve models span three architectural families: recurrent, transformer-based, and graph neural networks. They differ in how they encode spatial and temporal relationships among the five active bins surrounding each detector. All were embedded in the same dual-branch framework with their original encoders retained, so that performance differences reflect encoder design rather than framework configuration.

Recurrent baselines comprised Bidirectional LSTM, Sequence-to-Sequence with Attention, and ConvLSTM, each receiving the five-bin spatial input jointly at every time step (Abdelraouf et al. 2022a; Abdelraouf et al. 2022b). Transformer models process the full analysis window in parallel rather than sequentially (Ouyang et al. 2025; Zou et al. 2024; Zhang et al. 2025c). Five were evaluated: a standard Transformer Encoder; Patch Time Series Transformer (PatchTST) (Nie et al. 2023), which groups consecutive observations into patch-level tokens; Temporal Fusion Transformer (TFT) (Lim et al. 2021), which adds gated residual networks and

recurrent encoding; a Chronos-style encoder with bidirectional attention (Ansari et al. 2024); and STAEformer (Liu et al. 2023), which adds learnable spatial embeddings that distinguish bins without an explicit graph.

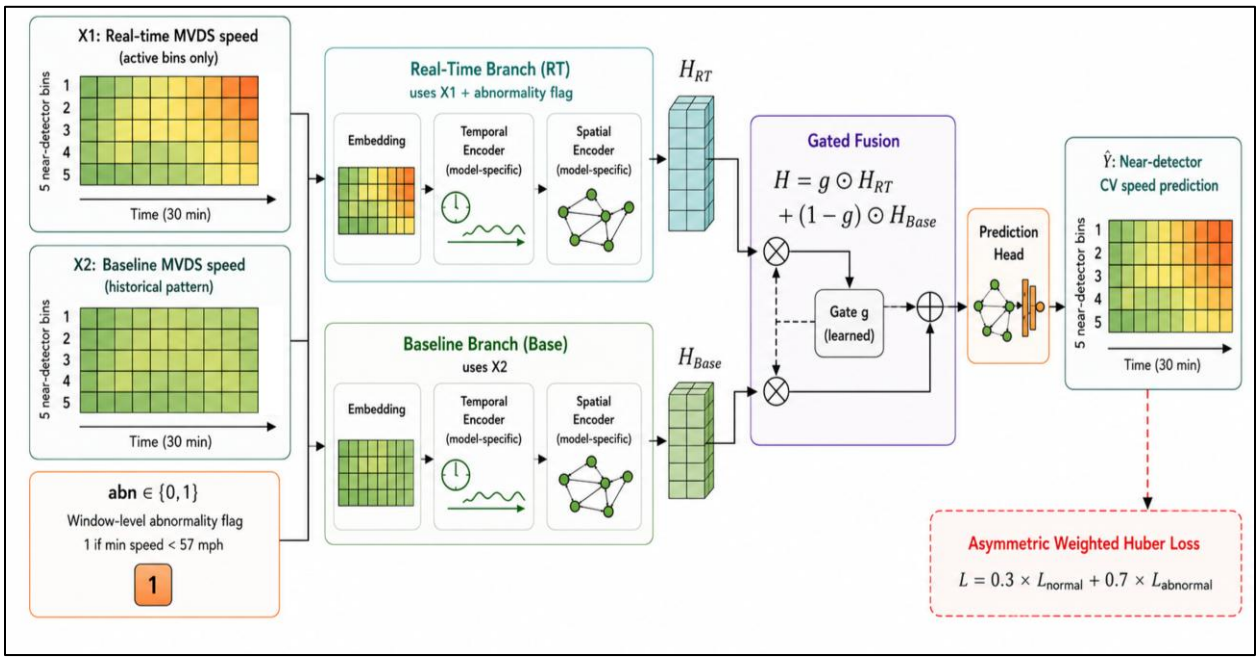



**Fig. 4 Stage 1 Flowchart**

Graph Neural Network Family incorporates the spatial relationships among the five active bins through graph representation. Unlike recurrent and transformer architectures that primarily learn spatial interactions from the input features, graph neural networks used a predefined spatial structure to exchange information between neighboring locations. This allowed traffic conditions at one bin to directly influence the representation learned at adjacent bins. Attention-Based Spatiotemporal Graph Convolutional Network (ASTGCN) (Guo et al. 2019; Han et al. 2024b) combined graph convolution with spatial and temporal attention mechanisms. The model dynamically determines the relative importance of different spatial locations and temporal positions before performing graph-based feature aggregation. Diffusion Convolutional Recurrent

Neural Network (DCRNN) (Han et al. 2024a; Li et al. 2017) represented traffic evolution as a diffusion process over the graph structure. Information is propagated across neighboring bins through diffusion operations, while temporal dependencies are incorporated through sequential temporal processing. Multi-Scale Temporal Graph Neural Network (MTGNN) (Wu et al. 2020) extended graph-based learning with multi-scale temporal modelling. GraphWaveNet (Han et al. 2026a; Han et al. 2026b; Wu et al. 2019) combined adaptive graph learning with dilated temporal convolutions. The model jointly learns spatial relationships from the data while expanding its temporal receptive field through hierarchical convolutional operations. **Fig. 4** shows the flowchart of the stage 1 framework.

***Stage 2: Corridor-wide Speed Field Reconstruction***

This section presents second stage of the proposed pipeline: a corridor-wide gap filler that propagates first stage near-detector speed estimates across all unobserved spatial locations between detectors. The gap filler is a graph-based spatiotemporal model, referred to as the Graph–LSTM gap filler, operating on a uniform spatial grid of 0.05-mile bins spanning the full corridor. Within each ramp-to-ramp segment, MVDS detectors were mapped to a set of five consecutive spatial bins centered on the detector's known milepost location; these are designated as active bins. All remaining bins between detectors were designated gap bins. The train and test windows followed the same chronological partition as Stage 1.

At test time, the Stage 1 ASTGCN was first applied to each five-bin group associated with the segment's detectors to produce predictions independently over its five active bins. The gap filler incorporated an abnormality indicator representing the prevailing traffic condition for each analysis window. This information was integrated into the model prior to spatial and temporal feature extraction, allowing the reconstruction process to adapt to differing traffic conditions. As

a result, the model could utilize different representations for normal and abnormal traffic states when estimating speeds at unobserved locations. Following abnormality conditioning, the input features were transformed into an internal representation suitable for graph-based processing. This step combined the available traffic, spatial, and contextual information into a unified feature space that served as the input to the graph convolution layers. A series of graph convolution layers were then applied to exchange information between adjacent locations and propagate traffic information across the corridor for GraphConv Gap Filler. Residual connections and normalization layers were incorporated to stabilize training and preserve information learned in earlier layers.

Following spatial feature extraction, a Long Short-Term Memory (LSTM) network was used to model the temporal evolution of traffic conditions within each analysis window. While the graph convolution layers captured spatial interactions among segment locations, the LSTM learned how traffic conditions change over time. The resulting spatiotemporal representation was subsequently used to reconstruct the corridor-wide speed field.

During training, the gap filler used the ground-truth StreetLight CV speeds at active bin locations rather than the Stage 1 ASTGCN predictions. This teacher-forcing strategy provided the model with reliable anchor information while it learned to reconstruct speeds at unobserved locations between detectors. At inference time, the ground-truth speeds were no longer available and were replaced by the Stage 1 ASTGCN predictions. The gap filler then used these predicted anchor values, together with the spatial and temporal information available within the segment, to estimate speeds at gap locations. By separating the spatial reconstruction task from the prediction errors of Stage 1 during training, teacher forcing enabled the gap filler to focus on learning the spatial and temporal relationships required for corridor-wide reconstruction. This resulted in a

more stable training process while preserving the operational deployment framework in which Stage 1 predictions served as the anchor inputs.

The Gap filler was trained using prediction errors computed only at gap bins, which represent locations without direct detector coverage. Active bins were excluded from the loss calculation because their final values were provided by the Stage 1 ASTGCN predictions during inference. A Huber loss function was employed to measure reconstruction error. Because abnormal traffic conditions occur less frequently but are of greater operational importance, a congestion-aware weighting strategy was incorporated into the training loss. Reconstruction errors associated with abnormal windows received a higher weight than those associated with normal traffic conditions, encouraging the model to place greater emphasis on accurately reconstructing speeds during congestion. This weighting strategy helped address the imbalance between normal and abnormal traffic conditions within the training data.

To evaluate the contribution of the spatial operator within Stage 2, three variants of the Graph–LSTM gap filler were compared. All three configurations shared an identical architecture differing only in the spatial operator applied within each of the three stacked spatial processing blocks. The second configuration, DiffusionConv Gap Filler, replaced the fixed symmetric chain convolution with a bidirectional diffusion convolution operator. This operator separately computed forward and backward random-walk diffusion, concatenating the resulting representations before projection. The third configuration, AdaptiveGraphConv Gap Filler, replaced the fixed adjacency with a learned adjacency matrix. The adjacency was computed as a softmax-normalized product of two embedding matrices, blended with the fixed chain adjacency during training. This allowed the model to discover spatial dependencies beyond the immediate

geographic neighborhood if warranted by the training data. **Fig. 5** shows the flowchart of the stage 2 framework.

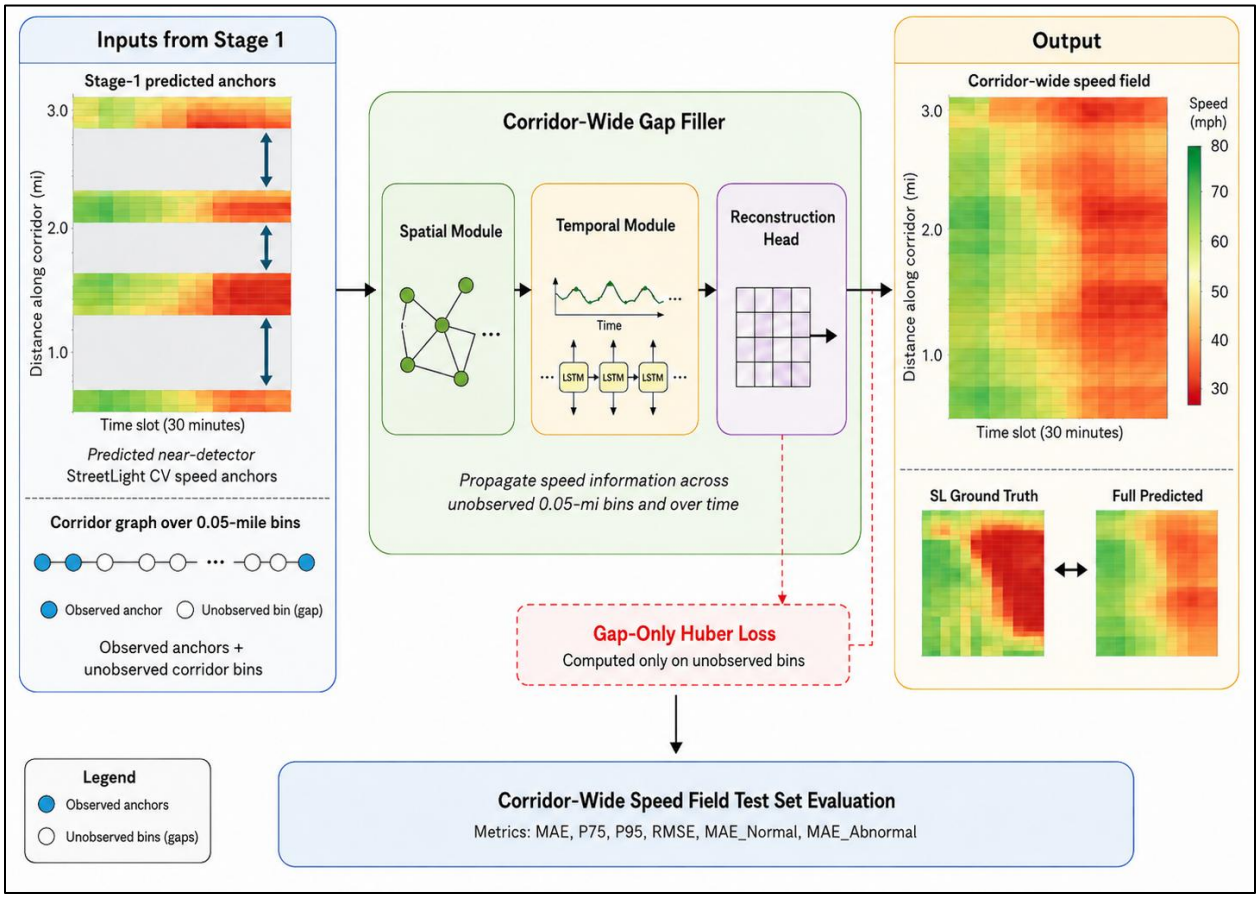



**Fig. 5 Stage 2 Flowchart**

## Results and Discussion

**Table 1** presents test-set performance for the twelve Stage 1 models across six metrics. The mean absolute error (MAE) is the average absolute difference between the estimated and CV ground-truth speeds and indicates typical accuracy. The root mean square error (RMSE) squares each error before averaging, so it penalizes large errors more heavily than MAE. P75 and P95 are the 75th and 95th percentiles of the absolute error distribution. MAE_Normal and MAE_Abnormal are the

MAE computed separately over normal and abnormal windows. The test set contains 149,938 normal and 3,145 abnormal windows, from totals of 531,860 and 10,382 respectively.

Across all twelve models, overall MAE ranged from 4.339 mph (TFT) to 4.536 mph (MTGNN), representing a relatively narrow spread. TFT achieved the lowest overall MAE (4.339 mph), the lowest normal-window MAE (4.280 mph), and the lowest RMSE (5.331 mph), demonstrating strong performance under the prevailing traffic conditions represented in the dataset. While recurrent and transformer-based models generally produced abnormal-window MAE values between approximately 6.5 and 6.9 mph, several graph-based models achieved substantially lower errors. Among all models, ASTGCN produced the lowest abnormal-window MAE (5.778 mph). This result indicates that architectures that explicitly incorporate spatial relationships are better able to represent the traffic patterns observed during non-recurrent congestion. Although TFT achieved the best overall MAE, its abnormal-window MAE (6.722 mph) was considerably higher than that of ASTGCN. Because accurate reconstruction during congestion is the primary operational objective of the proposed framework, performance under abnormal traffic conditions was considered more important than marginal differences in overall MAE. In addition, ASTGCN explicitly models spatial relationships among adjacent locations through graph-based processing, making it conceptually consistent with the corridor-wide spatial propagation strategy adopted in Stage 2. Therefore, ASTGCN was selected as the backbone for the corridor-wide reconstruction framework.

Stage 2 performed corridor-wide reconstruction at locations without direct detector coverage. The model must infer speeds at gap bins located between detectors using spatial propagation and temporal relationships learned during training. These locations do not possess direct traffic speed observations and therefore contain greater uncertainty than detector-adjacent

bins. Furthermore, Stage 2 inherited any residual error produced by the Stage 1 backbone. During inference, ASTGCN predictions at active bins served as anchor values for the reconstruction model. Any inaccuracies in these anchor predictions could propagate into neighboring gap locations, creating an additional source of reconstruction error that is absent from the Stage 1 evaluation. Abnormal traffic conditions further amplified this challenge. Congestion often produces sharper spatial gradients and more localized speed variations between detectors. Reconstructing these patterns throughout the corridor is more difficult than predicting speeds at detector-adjacent locations, which explains why the abnormal-window MAE increases from 5.778 mph for ASTGCN in Stage 1 to 9.674 mph for the corridor-wide reconstruction task in Stage 2.

GraphConv Gap Filler achieved the lowest gap-bin MAE (5.072 mph) and abnormal-window MAE (9.674 mph) and was therefore adopted as the final gap filler. To assess whether more expressive spatial operators could improve corridor-wide reconstruction, DiffusionConv Gap Filler and AdaptiveGraphConv Gap Filler were evaluated under identical training conditions. DiffusionConv Gap Filler, which incorporated bidirectional diffusion convolution to represent asymmetric upstream and downstream propagation, produced MAE of 5.629 mph and abnormal-window MAE of 11.588 mph, representing a degradation relative to GraphConv Gap Filler across all reported metrics. AdaptiveGraphConv Gap Filler, which employed a learned adjacency matrix, produced a gap-bin MAE of 5.719 mph and abnormal-window MAE of 13.386 mph, the highest among the three configurations. These results indicate that increased spatial complexity did not improve reconstruction accuracy for the SR-417 corridor. The superior performance of GraphConv Gap Filler suggests that the geographic chain topology in which each spatial bin interacts symmetrically with its immediate neighbors adequately captures the spatial structure of freeway speed fields in this corridor.

**Table 1. Model Performance**

| Stage 1 Model Performance | | | | | | |
|---|---|---|---|---|---|---|
| **Model** | **MAE (mph)** | **P75 (mph)** | **P95 (mph)** | **RMSE (mph)** | **MAE_Normal (mph)** | **MAE_Abnormal (mph)** |
| BiLSTM | 4.395 | 6.071 | 9.946 | 5.386 | 4.341 | 6.621 |
| Seq2Seq_Attn | 4.352 | 6.011 | 9.884 | 5.344 | 4.301 | 6.524 |
| ConvLSTM | 4.354 | 6.017 | 9.903 | 5.345 | 4.296 | 6.685 |
| Transformer | 4.465 | 6.176 | 10.101 | 5.467 | 4.408 | 6.779 |
| PatchTST | 4.478 | 6.199 | 10.199 | 5.491 | 4.425 | 6.621 |
| TFT | 4.339 | 5.992 | 9.883 | 5.331 | 4.280 | 6.722 |
| Chronos | 4.454 | 6.151 | 10.062 | 5.453 | 4.394 | 6.861 |
| STAEformer | 4.458 | 6.191 | 10.166 | 5.482 | 4.421 | 5.958 |
| **ASTGCN** | **4.450** | **6.167** | **10.067** | **5.465** | **4.417** | **5.778** |
| DCRNN | 4.457 | 6.172 | 10.128 | 5.467 | 4.412 | 6.298 |
| MTGNN | 4.536 | 6.281 | 10.221 | 5.546 | 4.489 | 6.467 |
| GraphWaveNet | 4.445 | 6.152 | 10.095 | 5.454 | 4.404 | 6.092 |
| **Stage 2 Corridor-Wide Reconstruction Performance (Gap Bins)** | | | | | | |
| GraphConv Gap Filler | **5.072** | **7.001** | **13.719** | **6.707** | **5.048** | **9.674** |
| DiffusionConv Gap Filler | 5.629 | 7.789 | 14.534 | 7.280 | 5.598 | 11.588 |
| AdaptiveGraphConv Gap Filler | 5.719 | 7.991 | 15.732 | 7.596 | 5.679 | 13.386 |

To assess the spatial consistency of Stage 1 ASTGCN predictions across the SR-417 corridor, MAE was computed independently for each near-detector location and plotted as a function of detector position along the northbound and southbound directions. As shown in **Fig. 6**, prediction errors were generally distributed within a narrow band across both directions. MAE values across detector locations ranged from 3.4 to 5.8 mph, indicating spatially consistent performance with no detector exhibiting error substantially outside this narrow band. Overall, the spatial distribution of errors confirms that the Stage 1 model provided a reliable anchor signal across the full SR-417 corridor extent.

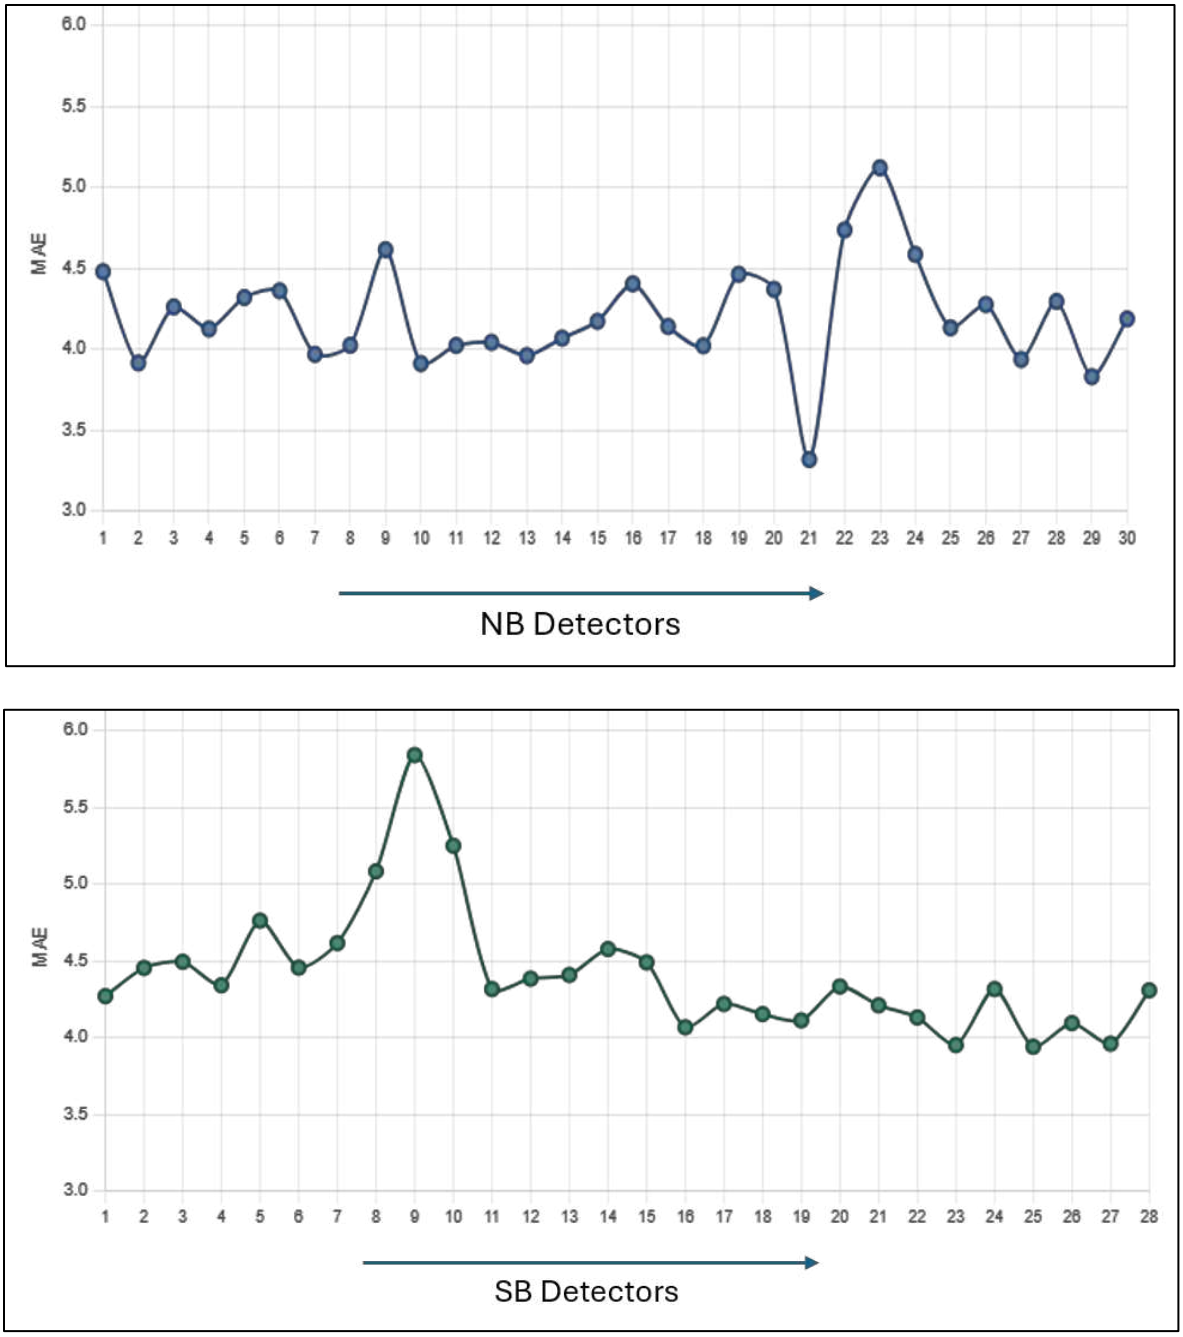



**Fig. 6 MAE (mph) Distribution across Northbound and Southbound Detectors**

The effect of probe data characteristics on Stage 1 prediction accuracy was further examined through a sensitivity analysis across market penetration rate and temporal aggregation interval, presented in **Fig. 7**. Results reveal a clear and consistent improvement in MAE as both penetration rate and aggregation interval increase. At the lowest penetration band (below 0.01)

with 1-minute aggregation, MAE reached 5.17 mph, reflecting the high noise level inherent in very sparse probe observations at fine temporal resolution. Increasing the aggregation interval to 5 minutes reduced MAE to 5.13 mph even at low penetration, as averaging over longer intervals individual trajectory noise. The most substantial gains were observed when penetration exceeded 0.03, where 5-minute aggregated data yield the minimum MAE of 3.22 mph. These findings confirm that the 2-minute aggregation interval adopted in the present study represents a reasonable operating point because it provided sufficient temporal granularity to capture congestion onset dynamics within the 30-minute analysis window.

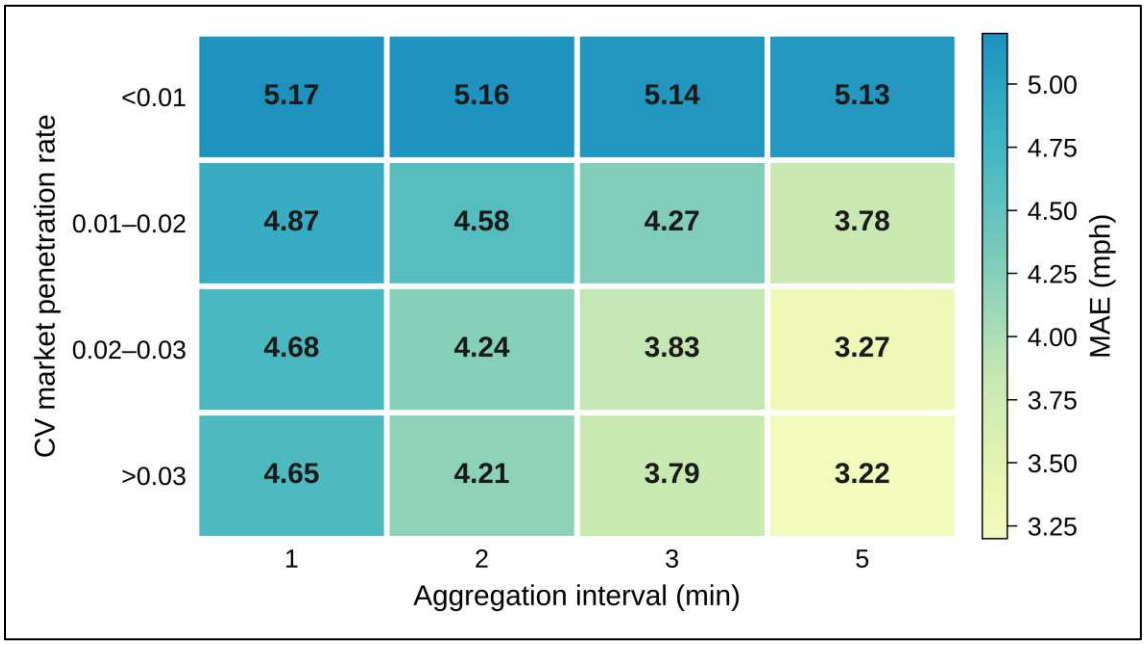



**Fig. 7 MAE Sensitivity to Penetration Rate and Data Aggregation**

The qualitative behavior of the Stage 1 ASTGCN model under congested traffic condition is illustrated in **Fig. 8**, which presents side-by-side heatmaps of the MVDS real-time input (X1), StreetLight CV ground truth (Y), and ASTGCN prediction ($\hat{Y}$). From the figure, the model prediction captured the onset and approximate spatial extent of the congestion.

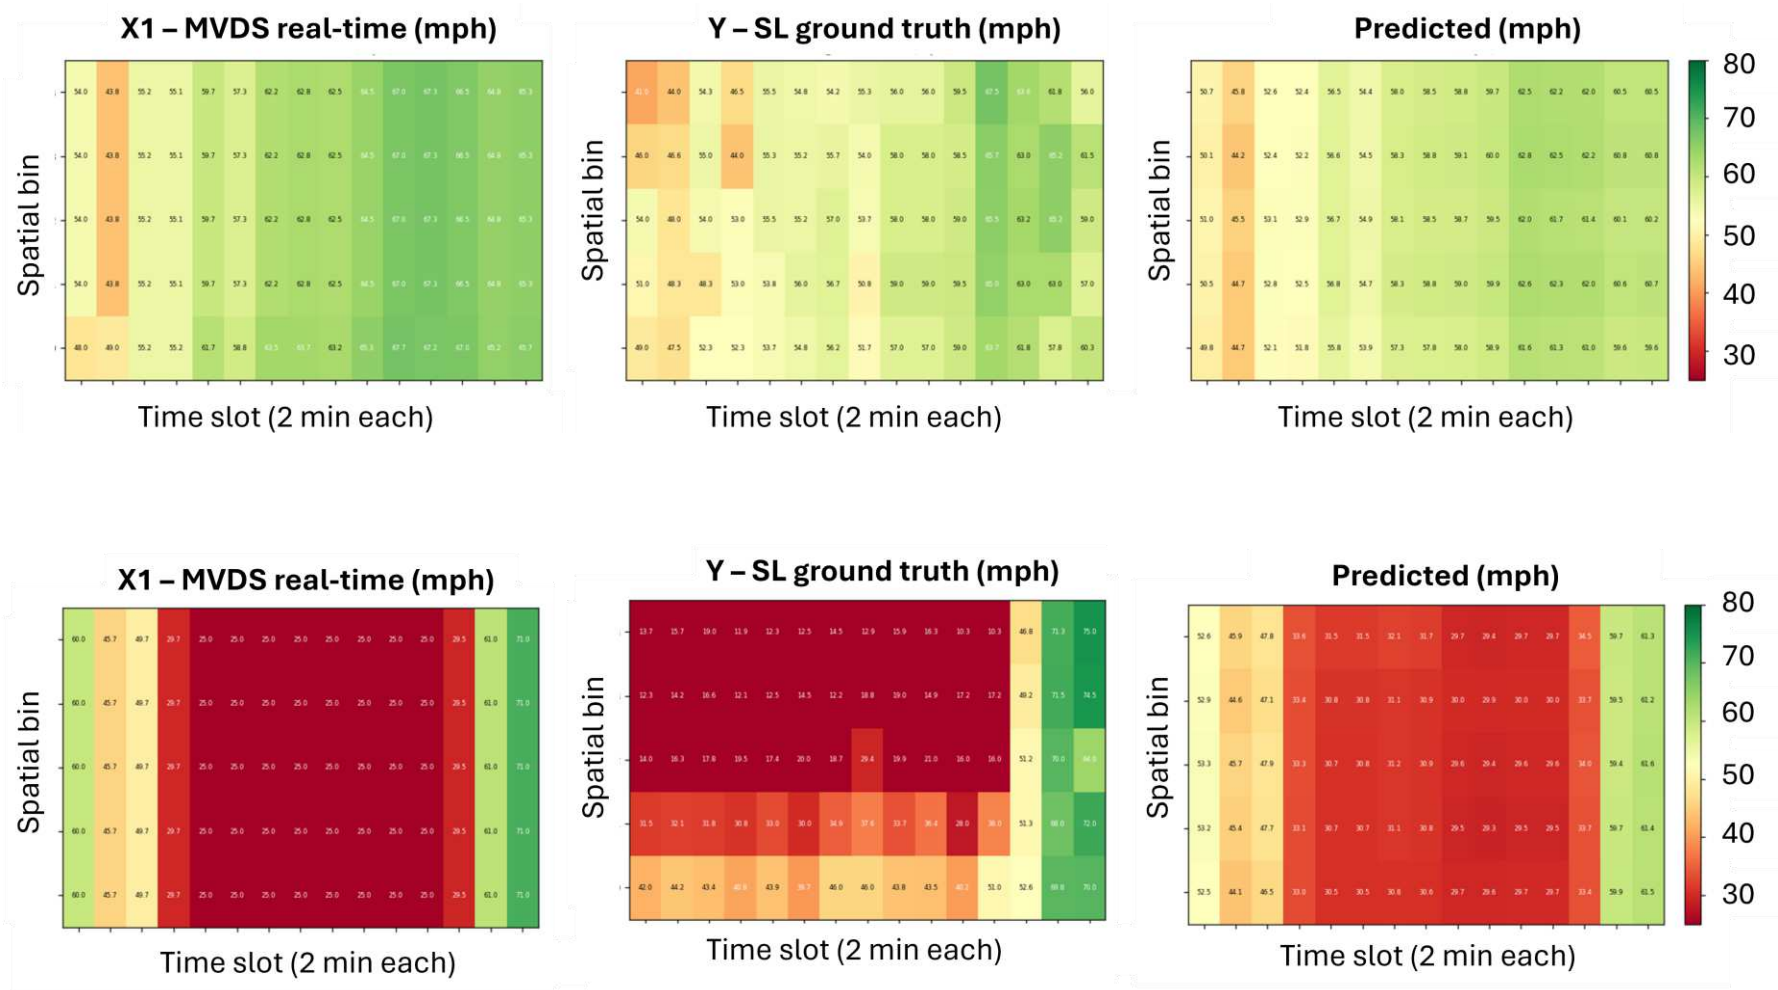



**Fig. 8 Near Detector Heatmaps**

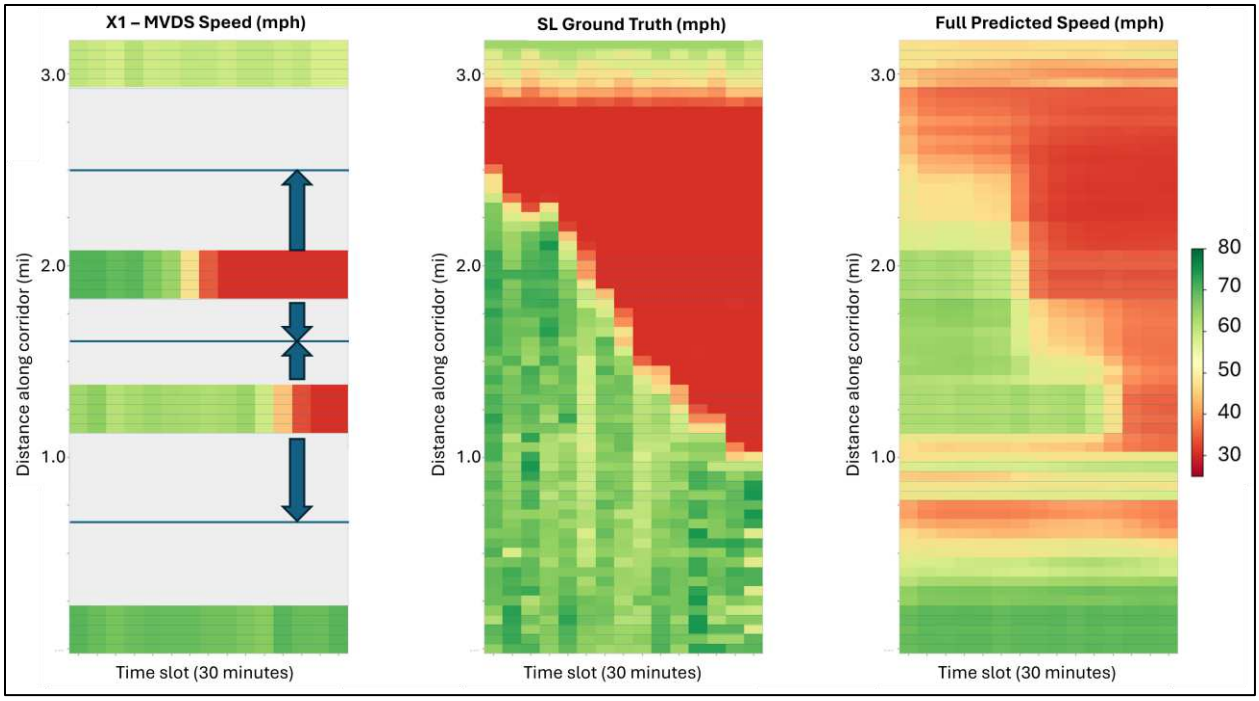



**Fig. 9 Example of Corridor-wide Speed Field Reconstruction Under Congested Condition**

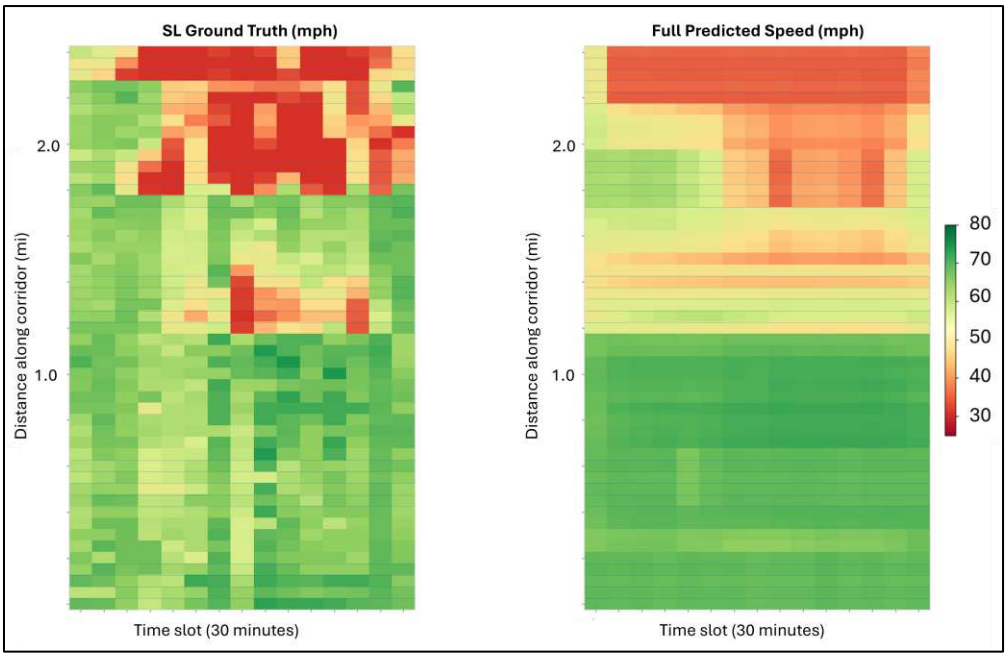



**Fig. 10 Additional Reconstruction Example**

**Fig. 9** presents an example of segment-wide reconstruction produced by the proposed two-stage framework. The left panel shows the detector-adjacent observations available from the sparse MVDS network, the center panel shows the StreetLight CV ground-truth speed field, and the right panel shows the reconstructed speed field generated by the gap filler. Visual inspection indicates that the reconstructed speed field successfully reproduced the major congestion pattern observed in the ground truth, including both the spatial extent and temporal evolution of the speed reduction. To quantify this improvement, congestion severity was measured as the total space–time area with speeds below 57 mph, computed as the number of spatiotemporal cells classified as congested and converted to mile-minutes. Using only detector observations and assuming conditions between detectors were represented by the detector locations alone, the estimated congestion area was 22.9 mile-minutes, substantially underestimating the true congestion extent. The StreetLight CV ground truth yielded a congestion area of 55.6 mile-minutes, while the reconstructed speed field produced an estimate of 61.1 mile-minutes. The reconstructed value differed from the ground truth by only 5.5 mile-minutes (9.9%), demonstrating that the Stage 2 model is able to recover the majority of the unobserved congestion footprint between detectors. These results highlight the importance of corridor-wide reconstruction, as detector-only observations substantially underestimate the spatial and temporal extent of congestion.

**Fig. 10** also illustrates a representative reconstruction result. Based on a 57 mph congestion threshold, the ground-truth congestion area was 22.2 mile-minutes compared with 28.4 mile-minutes in the reconstructed speed field. Despite a modest overestimation of congestion extent, the reconstructed field preserved the spatiotemporal characteristics of the observed congestion pattern.

## Conclusions

This study proposed and validated a two-stage spatiotemporal pipeline for corridor-scale freeway speed field reconstruction on the SR-417 corridor in Central Florida. In Stage 1, a comparative evaluation of twelve spatiotemporal deep learning models demonstrated that graph neural network architectures outperform recurrent and transformer-based models under congested traffic conditions, with ASTGCN achieving the lowest abnormal-window MAE of 5.778 mph. In Stage 2, the corridor-wide gap filler, trained with teacher forcing and a congestion-aware loss function, successfully reconstructed the full corridor speed field at 0.05-mile spatial resolution. Together, these results demonstrate that commercial connected vehicle probe data, when used as a dense spatial ground truth to supervise a graph-based spatial model, provides a practically deployable pathway to continuous freeway speed field monitoring without additional roadside hardware investment. Case-study analyses further showed that detector-only observations can substantially underestimate congestion footprints, whereas the reconstructed speed field closely approximated the congestion extent observed in the connected vehicle ground truth. In a representative congestion event, detector-only monitoring underestimated the congestion footprint by 58.8%, whereas the reconstructed field was within 9.9% of the CV ground truth.

The framework requires ground-truth connected vehicle probe data only for the training period. Full-year probe data coverage is not required; once trained, the model reconstructs corridor-wide speed fields using only the fixed MVDS detector readings that are continuously available in real time. This makes the approach operationally feasible for agencies that collect probe data periodically or through targeted deployments rather than continuously.

The proposed framework provides corridor-wide speed coverage at 0.05-mile resolution, which enables traffic analyses that are not possible with sparse detector data alone. One direct

application is shockwave identification. A shockwave is the boundary between a congested queue and the free-flow traffic upstream of it, which moves in the opposite direction of travel as the queue grows. Accurately locating this boundary requires speed observations at fine spatial intervals; with only detector-spaced measurements, the shockwave position can only be bracketed between two detectors, introducing a positional uncertainty equal to the inter-detector spacing. The reconstructed speed field resolves this boundary at 0.05-mile intervals, substantially reducing positional uncertainty.

A related benefit concerns congestion severity assessment. Congestion severity depends not only on the minimum speed reached but on the spatial extent and duration of the speed reduction below threshold. With detector-only data, the congested area between two detectors is estimated by assumption rather than measurement. The reconstructed field provides a spatially continuous speed map from which congested area can be computed directly as the product of the number of bins below the threshold speed and the bin size, aggregated over time. This produces a more defensible measure of congestion severity than interpolation between isolated detector readings.

Finally, fixed MVDS detectors cover only the specific locations where sensors are installed, leaving the majority of the corridor spatially unobserved. The proposed method transforms a sparse, detector-constrained snapshot of corridor conditions into a complete spatiotemporal speed field suitable for corridor-level operational analysis.

## Statements and Declarations


### *Funding*

No funding was received to assist with the preparation of this manuscript.


### *Data Availability*

The data that support the findings of this study are available from the corresponding author upon reasonable request.

### *Competing Interests*

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

### *Author Contributions*

The authors confirm contribution to the paper as follows: study conceptualization and design: Tahmid, Islam and Abdel-Aty; data collection: Tahmid; formal analysis and interpretation of results: Tahmid, Islam and Abdel-Aty; draft manuscript preparation: Tahmid, Islam and Abdel-Aty. All authors reviewed the results and approved the final version of the manuscript.

## References


Abdelraouf A, Abdel-Aty M, Mahmoud N (2022a) Sequence-to-sequence recurrent graph convolutional networks for traffic estimation and prediction using connected probe vehicle data. IEEE Transactions on Intelligent Transportation Systems 24:1395–1405

Abdelraouf A, Abdel-Aty M, Yuan J (2022b) Utilizing Attention-Based Multi-Encoder-Decoder Neural Networks for Freeway Traffic Speed Prediction. IEEE Transactions on Intelligent Transportation Systems 23:11960–11969. https://doi.org/10.1109/TITS.2021.3108939

Ansari AF, Stella L, Turkmen C, Zhang X, Mercado P, Shen H, Shchur O, Rangapuram SS, Arango SP, Kapoor S, Zschiegner J, Maddix DC, Wang H, Mahoney MW, Torkkola K, Wilson AG, Bohlke-Schneider M, Wang Y (2024) Chronos: Learning the Language of Time Series

Bachmann C, Abdulhai B, Roorda MJ, Moshiri B (2013a) A comparative assessment of multi-sensor data fusion techniques for freeway traffic speed estimation using microsimulation modeling. Transportation Research Part C: Emerging Technologies 26:33–48. https://doi.org/10.1016/j.trc.2012.07.003

Bachmann C, Roorda MJ, Abdulhai B, Moshiri B (2013b) Fusing a Bluetooth Traffic Monitoring System With Loop Detector Data for Improved Freeway Traffic Speed Estimation. Journal of Intelligent Transportation Systems 17:152–164. https://doi.org/10.1080/15472450.2012.696449

Cheng R, Liu M, Xu Y (2025) ST_AGCNT: Traffic Speed Forecasting Based on Spatial–Temporal Adaptive Graph Convolutional Network with Transformer. Sustainability 17:1829. https://doi.org/10.3390/su17051829

Gu Y, Liu D, Arvin R, Khattak AJ, Han LD (2023) Predicting intersection crash frequency using connected vehicle data: A framework for geographical random forest. Accident Analysis & Prevention 179:106880. https://doi.org/10.1016/j.aap.2022.106880

Guo S, Lin Y, Feng N, Song C, Wan H (2019) Attention Based Spatial-Temporal Graph Convolutional Networks for Traffic Flow Forecasting. Proceedings of the AAAI Conference on Artificial Intelligence 33:922–929. https://doi.org/10.1609/aaai.v33i01.3301922

Han L, Abdel-Aty M, Joo Y-J (2026a) Arterial Network Traffic State Prediction with Connected Vehicle Data: An Abnormality-Aware Spatiotemporal Network. arXiv preprint arXiv:260304432

Han L, Abdel-Aty M, Kim Y, Joo Y-J, Islam Z (2026b) MMCAformer: Macro-Micro Cross-Attention Transformer for Traffic Speed Prediction with Microscopic Connected Vehicle Driving Behavior. arXiv preprint arXiv:260216730

Han L, Abdel-Aty M, Yu R, Wang C (2024a) LSTM + Transformer Real-Time Crash Risk Evaluation Using Traffic Flow and Risky Driving Behavior Data. Trans Intell Transport Sys 25:18383–18395. https://doi.org/10.1109/TITS.2024.3438616

Han L, Yu R, Wang C, Abdel-Aty M (2024b) Transformer-based modeling of abnormal driving events for freeway crash risk evaluation. Transportation Research Part C: Emerging Technologies 165:104727. https://doi.org/10.1016/j.trc.2024.104727

Jiang B, Fei Y (2017) Vehicle Speed Prediction by Two-Level Data Driven Models in Vehicular Networks. IEEE Transactions on Intelligent Transportation Systems 18:1793–1801. https://doi.org/10.1109/TITS.2016.2620498

Jin Y, Tan E, Li L, Wang G, Wang J, Wang P (2018) Hybrid Traffic Forecasting Model With Fusion of Multiple Spatial Toll Collection Data and Remote Microwave Sensor Data. IEEE Access 6:79211–79221. https://doi.org/10.1109/ACCESS.2018.2884933

Kessler L, Rempe F, Bogenberger K (2021) Multi-Sensor Data Fusion for Accurate Traffic Speed and Travel Time Reconstruction. Front Future Transp 2. https://doi.org/10.3389/ffutr.2021.766951

Li Y, Yu R, Shahabi C, Liu Y (2017) Diffusion convolutional recurrent neural network: Data-driven traffic forecasting. arXiv preprint arXiv:170701926

Lim B, Arık SÖ, Loeff N, Pfister T (2021) Temporal Fusion Transformers for interpretable multi-horizon time series forecasting. International Journal of Forecasting 37:1748–1764. https://doi.org/10.1016/j.ijforecast.2021.03.012

Liu H, Dong Z, Jiang R, Deng J, Deng J, Chen Q, Song X (2023) STAEformer: Spatio-Temporal Adaptive Embedding Makes Vanilla Transformer SOTA for Traffic Forecasting

Liu Y, He S, Ran B, Cheng Y (2018) A Progressive Extended Kalman Filter Method for Freeway Traffic State Estimation Integrating Multisource Data. Wireless Communications and Mobile Computing 2018:6745726. https://doi.org/10.1155/2018/6745726

Nanthawichit C, Nakatsuji T, Suzuki H (2003) Application of Probe-Vehicle Data for Real-Time Traffic-State Estimation and Short-Term Travel-Time Prediction on a Freeway. Transportation Research Record: Journal of the Transportation Research Board 1855:49–59. https://doi.org/10.3141/1855-06

Nie Y, Nguyen NH, Sinthong P, Kalagnanam J (2023) A Time Series is Worth 64 Words: Long-term Forecasting with Transformers

Ouyang N, Ao L, Cai Q, Wan W, Ren X, He X, Sheng K (2025) Graph Transformer-Based Dynamic Edge Interaction Encoding for Traffic Prediction. IEEE Transactions on Intelligent Transportation Systems 26:4066–4079. https://doi.org/10.1109/TITS.2024.3513325

Qian W, Nielsen TD, Zhao Y, Larsen KG, Yu JJ (2024) Uncertainty-Aware Temporal Graph Convolutional Network for Traffic Speed Forecasting. IEEE Trans Intell Transport Syst 25:8578–8590. https://doi.org/10.1109/TITS.2024.3365721

Tampere CMJ, Immers LH (2007) An Extended Kalman Filter Application for Traffic State Estimation Using CTM with Implicit Mode Switching and Dynamic Parameters. In: 2007 IEEE Intelligent Transportation Systems Conference. pp 209–216

Wang X, Jerome Z, Wang Z, Zhang C, Shen S, Kumar VV, Bai F, Krajewski P, Deneau D, Jawad A, Jones R, Piotrowicz G, Liu HX (2024) Traffic light optimization with low penetration rate vehicle trajectory data. Nat Commun 15:1306. https://doi.org/10.1038/s41467-024-45427-4

Wu W, Xia Y, Jin W (2021) Predicting Bus Passenger Flow and Prioritizing Influential Factors Using Multi-Source Data: Scaled Stacking Gradient Boosting Decision Trees. IEEE Transactions on Intelligent Transportation Systems 22:2510–2523. https://doi.org/10.1109/TITS.2020.3035647

Wu Z, Pan S, Long G, Jiang J, Chang X, Zhang C (2020) Connecting the Dots: Multivariate Time Series Forecasting with Graph Neural Networks. In: arXiv.org. https://arxiv.org/abs/2005.11650v1. Accessed 3 June 2026

Wu Z, Pan S, Long G, Jiang J, Zhang C (2019) Graph WaveNet for Deep Spatial-Temporal Graph Modeling. 1907–1913

Yao B, Chen C, Cao Q, Jin L, Zhang M, Zhu H, Yu B (2017) Short-Term Traffic Speed Prediction for an Urban Corridor. Computer-Aided Civil and Infrastructure Engineering 32:154–169. https://doi.org/10.1111/mice.12221

Yu B, Yin H, Zhu Z (2017) Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In: arXiv.org. https://arxiv.org/abs/1709.04875v4. Accessed 3 June 2026

Zeng Y, Lan J, Ran B, Jiang Y (2014) A Novel Multisensor Traffic State Assessment System Based on Incomplete Data. ScientificWorldJournal 2014:532602. https://doi.org/10.1155/2014/532602

Zhang H, Dong H, Yang Z (2025a) TSGDiff: Traffic state generative diffusion model using multi-source information fusion. Transportation Research Part C: Emerging Technologies 174:105081. https://doi.org/10.1016/j.trc.2025.105081

Zhang S, Abdel-Aty M (2022) Real-time crash potential prediction on freeways using connected vehicle data. Analytic Methods in Accident Research 36:100239. https://doi.org/10.1016/j.amar.2022.100239

Zhang Y, Kwan M-P, Cai J, Wang J, Ma P (2025b) CrossGraphNet: a cross-spatiotemporal graph-based method for traffic speed reconstruction using remote sensing vehicle

detection. International Journal of Digital Earth 18:2528640. https://doi.org/10.1080/17538947.2025.2528640

Zhang Y, Zhou Q, Wang J, Kouvelas A, Makridis MA (2025c) CASAformer: Congestion-aware sparse attention transformer for traffic speed prediction. Communications in Transportation Research 5:100174. https://doi.org/10.1016/j.commtr.2025.100174

Zhao L, Song Y, Zhang C, Liu Y, Wang P, Lin T, Deng M, Li H (2018) T-GCN: A Temporal Graph ConvolutionalNetwork for Traffic Prediction. In: arXiv.org. https://arxiv.org/abs/1811.05320v3. Accessed 3 June 2026

Zou G, Lai Z, Ma C, Li Y, Wang T (2023) A novel spatio-temporal generative inference network for predicting the long-term highway traffic speed. Transportation Research Part C: Emerging Technologies 154:104263. https://doi.org/10.1016/j.trc.2023.104263

Zou Y, Chen Y, Xu Y, Zhang H, Zhang S (2024) Short-term freeway traffic speed multistep prediction using an iTransformer model. Physica A: Statistical Mechanics and its Applications 655:130185. https://doi.org/10.1016/j.physa.2024.130185