Reliability Engineering and System Safety 278 (2027) 113443



Contents lists available at ScienceDirect

Reliability Engineering and System Safety

journal homepage: www.elsevier.com/locate/ress



A reliability-informed stochastic hybrid graph attention network (SHGAT) framework for uncertainty-aware road operation-maintenance risk inference

Dong Feng

Faculty of Civil Engineering, RWTH Aachen University, Mies-van-der-Rohe-Straße 1, 52074, Aachen, Germany

ARTICLE INFO

Keywords:
Road network reliability
System safety
Risk-informed maintenance
Stochastic hybrid graph attention network (SHGAT)
Uncertainty quantification
Camera-sparse infrastructure systems

ABSTRACT

Urban road networks are complex infrastructure systems whose reliable operation depends on timely maintenance decisions under incomplete and heterogeneous observations. Direct maintenance records and condition measurements are often sparse, whereas traffic trajectories, road topology, camera deployment, and geometric attributes are continuously available but unevenly reliable. This study proposes a reliability-informed stochastic hybrid graph attention network (SHGAT) framework for inferring latent road operation-maintenance pressure as a reliability-risk proxy. Vehicle trajectories are transformed into segment-time evidence, including flow exposure, traversal-time burden, inverse-speed pressure, temporal irregularity, recurrence, bottleneck contrast, camera-derived reliability, and observation uncertainty. A hybrid segment graph represents local propagation, while stacked SHGAT layers learn reliability-aware neighborhood interactions considering observation confidence, road-class compatibility, length contrast, and propagation structure. A reliability-gated fusion module integrates geometry, flow, temporal, and camera evidence before stochastic pressure inference. Feature noise, dropout, camera degradation, trajectory degradation, geometry perturbation, and Monte Carlo sampling are introduced for robustness and uncertainty evaluation. In the Shenzhen case, SHGAT achieved a predictive correlation of $r = 0.992$. Cross-city validation on the Jinan network retained a correlation of $r = 0.978$ without retraining. The framework produces pressure maps, uncertainty fields, source sensitivities, propagation gaps, intervention values, and maintenance priorities, supporting risk-informed decision-making for sensor-sparse urban road networks.

1. Introduction

Urban road networks are critical infrastructure systems whose operational reliability directly affects mobility, service continuity, emergency accessibility, and maintenance resource allocation [1–3]. In practice, road agencies must identify high-pressure or potentially vulnerable segments before direct condition evidence becomes complete or before service degradation becomes severe [4,5]. Traditional operation and maintenance strategies often rely on periodic inspection, engineering experience, and delayed maintenance records, which are difficult to update continuously at the scale of a dense urban network [6,7]. With the development of automated sensing, traffic cameras, geographic information systems, connected mobility records, and large-scale vehicle trajectories, road agencies now have access to richer descriptions of network operation [8–10]. Nevertheless, the availability of more data does not automatically lead to more reliable maintenance decisions, because the data sources are heterogeneous, unevenly observable, and affected by measurement noise, missing coverage, and temporal variability [11,12].

From the perspective of reliability engineering and system safety, road operation-maintenance inference is a problem of decision making under uncertain evidence [13,14]. The state of a road segment is not determined by traffic volume alone. It is shaped by repeated service exposure, traversal-time burden, speed deterioration, temporal irregularity, local bottleneck behavior, road-class heterogeneity, geometric constraints, camera observability, and pressure propagation from neighboring segments [15–17]. In many real-world networks, direct maintenance labels or pavement-condition measurements are sparse, delayed, or unavailable [18,19]. Therefore, it is necessary to construct an operational proxy that can represent latent maintenance pressure while explicitly accounting for observation reliability and uncertainty. Information fusion provides a natural basis for this problem, because no single data source is sufficient to support reliable network-scale

E-mail address: dong.feng@rwth-aachen.de.

https://doi.org/10.1016/j.ress.2026.113443
Received 21 June 2026; Received in revised form 29 August 2026; Accepted 2 September 2026
Available online 5 September 2026
0951-8320/© 2026 The Author(s). Published by Elsevier Ltd. This is an open access article under the CC BY license (http://creativecommons.org/licenses/by/4.0/).

Nomenclature

Abbreviations

BA	Balanced accuracy
FN	False negative
FP	False positive
GAT	Graph attention network
GNN	Graph neural network
KS	Kolmogorov-Smirnov
MAE	Mean absolute error
MCC	Matthews correlation coefficient
MLP	Multilayer perceptron
MSE	Mean squared error
ReLU	Rectified linear unit
RMSE	Root mean squared error
SHGAT	Stochastic hybrid graph attention network
TN	True negative
TP	True positive

Symbols

$\mathbf{A}^H$	Hybrid segment adjacency matrix
$A^H_{e,e'}$	Hybrid adjacency entry between road segments e and e'
C_e	Binary camera-coverage indicator of road segment e
$F_{e,\tau}$	Trajectory-derived flow count of segment e in temporal window τ
$\widetilde{F}_{e,\tau}$	Normalized flow evidence
$\mathscr{G}$	Directed road graph
G^P_e	Local pressure propagation gap between segment e and its hybrid-neighborhood average
$\mathbf{H}^G$	Final matrix of graph-attention embeddings for all road segments
$\mathbf{h}^G_e$	Final graph-attention embedding vector of segment e
$\mathscr{N}^H$	Hybrid-neighbor set of road segment e
$R_{e,\tau}$	Recurrence-enhanced flow descriptor
S^G_e, S^F_e, S^T_e, S^C_e	Geometry-, flow-, temporal-, and camera-modality evidence scores of segment e
$S_{e,j}$	Finite-difference sensitivity of segment e to feature channel j
$\overline{T}_{e,\tau}$	Mean traversal-time burden of segment e in temporal window τ
$U_{e,\tau}$	Empirical segment speed
$U^O_{e,\tau}$	Observation-uncertainty evidence
$\mathscr{U}_n$	Stochastic uncertainty score of sample n
V^I_e	Local intervention value of road segment e
$\mathbf{W}^{(0)}$	Initial projection matrix from hybrid static features to the SHGAT hidden space
$\mathbf{W}^{Q,l}$, $\mathbf{W}^{K,l}$, $\mathbf{W}^{V,l}$	Query, key, and value transformation matrices at SHGAT layer l
$\mathbf{x}^H_e$	Static hybrid feature vector used as SHGAT input
$\mathbf{X}^H$	Matrix collecting hybrid static feature vectors of all road segments
$\mathbf{x}_{e,\tau}$	Final segment-time inference feature vector
$\mathbf{X}$	Segment-time sample feature matrix used by the pressure-inference head
$Y_{e,\tau}$	Weak operation-maintenance pressure target of segment e in temporal window τ
$\widehat{Y}_{e,\tau}$	Predicted operation-maintenance pressure of segment e in temporal window τ

Greek symbols and parameters

$\alpha^{(l)}_{e,e'}$	Normalized graph-attention coefficient from segment e to neighboring segment e' at layer l
$\boldsymbol{\beta}_e$	Modality-weight vector of segment e over geometry, flow, temporal, and camera channels
ε	Small positive stabilizing constant used to prevent division by zero or numerical instability
η_l	Self-retention coefficient controlling the residual mixture of local and neighboring SHGAT information at layer l
$\theta_{a,\rho}$, $\theta_{a,\ell}$, $\theta_{a,\kappa}$	Reliability preference, length-contrast penalty, and road-class mismatch penalty in the attention score
$\theta_{\rho,0}$, $\theta_{\rho,C}$, $\theta_{\rho,F}$, $\theta_{\rho,\Omega}$	Coefficients defining the observation-reliability model
$\theta_{Y,0}$, $\theta_{Y,F}$, $\theta_{Y,T}$, $\theta_{Y,U}$, $\theta_{Y,\Omega}$, $\theta_{Y,R}$, $\theta_{Y,B}$, $\theta_{Y,O}$	Coefficients defining the weak pressure target
κ_e	Categorical road-class label of segment e
$\rho_{e,\tau}$	Observation reliability of segment e in temporal window τ
$\sigma(\cdot)$	Sigmoid activation function
τ	Temporal-window index
τ^*	Representative peak-demand temporal window
χ_e	Lightweight road-geometry curvature descriptor of segment e
$\Omega_{e,\tau}$	Temporal irregularity index of segment e in temporal window τ
Π_e	Fused maintenance-priority score of road segment e
Θ	Collection of trainable parameters of the pressure-inference head

maintenance inference [20,21]. Static road attributes describe the physical and functional context of each segment, including length, curvature, road class, and topological position [22,23]. Dynamic trajectory evidence reflects how segments are repeatedly used over time, including flow exposure, traversal-time burden, empirical speed, recurrent demand, and temporal instability [24,25]. Camera deployment provides direct observation support, but it also reveals the spatially uneven reliability of the sensing system [26,27]. A camera-covered segment may provide more confident evidence, whereas a camera-sparse segment may still experience high operation-maintenance pressure that can only be inferred indirectly through trajectories, topology, and neighboring segments [28,29]. Therefore, the key challenge is not simply to concatenate these heterogeneous features, but to learn how they should be weighted, propagated, and interpreted under nonuniform reliability conditions.

With the advancement of artificial intelligence (AI), deep learning is playing an increasingly significant role in the field of engineering [30–32]. More broadly, advanced computational and data-driven modeling approaches have demonstrated considerable potential for performance evaluation, spatiotemporal diagnosis, and risk assessment across a wide range of complex engineering systems [33–35]. Existing studies in intelligent transportation systems have made substantial progress in traffic prediction, network state estimation, data augmentation, spatio-temporal graph learning, and multi-source traffic data fusion [36–38]. Graph neural networks (GNNs) [39,40] and graph attention networks (GATs) [41] are particularly useful for capturing spatial dependencies among road links or intersections, while trajectory-based studies have improved travel-time estimation, mobility pattern mining, and demand reconstruction [42,43]. However, several limitations remain when these methods are transferred to reliability-oriented road operation-maintenance inference. First, many graph-based transportation models are designed to predict traffic states such as speed, flow, or congestion, whereas maintenance pressure is a broader latent reliability-risk state associated with repeated service burden, instability, uncertainty, and decision consequences [44,45]. Second, camera sparsity and observation uncertainty are often treated as

missing-data issues, rather than as reliability indicators that should actively guide attention and fusion [46,47]. Third, neighboring road segments are commonly aggregated using fixed topological relations, while road-class compatibility, length contrast, local recurrence, and reliability-aware propagation are less explicitly represented [48,49]. Fourth, many models are evaluated mainly under ordinary prediction settings, whereas maintenance decision support requires robustness under camera dropout, trajectory degradation, feature perturbation, geometry uncertainty, and source-dependent failure modes [50,51].

To address these gaps, this study proposes a reliability-informed stochastic hybrid graph attention network (SHGAT) framework for uncertainty-aware road operation-maintenance risk inference. The proposed framework makes four main contributions (see Fig. 1). First, it converts multi-source road-network observations into maintenance-oriented reliability evidence, including flow exposure, traversal-time burden, inverse-speed pressure, temporal irregularity, neighbor-induced recurrence, bottleneck contrast, camera-derived reliability, and observation uncertainty. Second, it constructs an edge-level hybrid segment graph in which road segments are treated as inference entities, allowing local pressure propagation to be modeled through shared network context. Third, it develops a reliability-aware SHGAT module that incorporates observation confidence, road-class compatibility, length contrast, neighbor propagation, and layer normalization to learn interpretable graph interactions. Fourth, it introduces reliability-gated modality fusion and stochastic robustness evaluation to quantify source contribution, uncertainty, degradation sensitivity, counterfactual intervention effects, and maintenance-priority stability. In this way, the framework connects automated sensing, graph-based reliability inference, uncertainty quantification, and risk-informed road maintenance decision support within a unified modeling pipeline.

2. Related work

Research on intelligent road-network modeling has developed rapidly with the expansion of automated sensing, connected mobility, traffic cameras, and data-driven transportation management [52,53]. A large body of work has focused on traffic-state prediction, where historical speed, volume, or occupancy measurements are used to forecast future network conditions [54]. GNNs are especially suitable for this task because road networks have strong spatial dependencies, directed connectivity, and localized propagation. Early graph-based methods mainly relied on predefined adjacency matrices [55], while later studies introduced adaptive graphs [56], attention mechanisms [57], spatio-temporal convolution [58], and recurrent structures [59] to capture more flexible traffic dependencies. These studies provide an important methodological basis, but their primary target is usually traffic prediction rather than maintenance-oriented pressure inference [60].

Multi-source information fusion has also become an important direction in intelligent transportation systems. Data from cameras, loop detectors, GPS trajectories, probe vehicles, road maps, and weather conditions can be fused to improve traffic monitoring, incident detection, travel-time estimation, and decision support [61]. Existing fusion strategies can be roughly divided into data-level fusion, feature-level fusion, model-level fusion, and decision-level fusion [62,63]. However, the quality and reliability of different sources are often uneven. Camera observations may be spatially sparse [64], trajectories may be incomplete [65], and map attributes may not fully describe operational stress [66]. Therefore, an effective fusion model should not only concatenate multi-source features, but also learn how much confidence should be assigned to each source under different road and traffic contexts [67].

GATs provide a useful mechanism for learning heterogeneous neighborhood influence. In transportation networks, attention can assign different weights to upstream, downstream, adjacent, or functionally related road segments [68]. Compared with fixed-weight graph convolution, attention-based methods can better reflect dynamic or context-dependent dependencies [69]. Nevertheless, most existing attention mechanisms are designed for predictive accuracy and do not explicitly encode observation reliability, road-class compatibility, length contrast, or camera sparsity [70,71]. For road operation-maintenance pressure inference, these factors are essential because the same amount of traffic exposure may imply different pressure levels depending on road type, structural position, and data reliability.

Vehicle trajectories offer another important source of operational

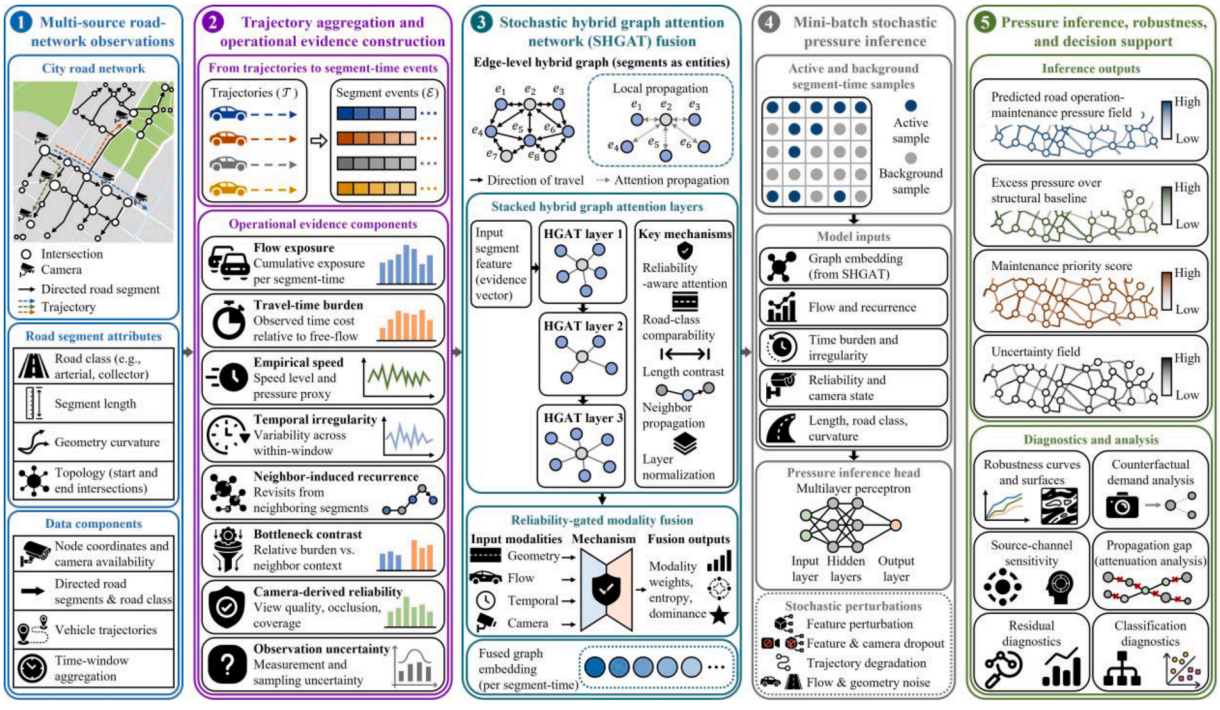


Fig. 1. Conceptual architecture of the proposed reliability-gated SHGAT fusion framework for road operation-maintenance pressure inference.

information. Unlike fixed-point traffic measurements, trajectory data describe movement sequences and can reveal route concentration, traversal time, speed variation, and recurrent use of specific network elements [72,73]. These properties are highly relevant to operation-maintenance pressure because repeated traversal, unstable travel time, and bottleneck-like local concentration may indicate hidden service burden. However, trajectory data are rarely transformed into explicit maintenance-oriented evidence. Most trajectory-based studies emphasize map matching [74], demand estimation [75], travel-time prediction [76], or mobility-pattern mining [77]. The present study instead uses trajectories to construct a weak pressure proxy and multiple operational evidence components for maintenance-oriented inference.

Road maintenance and infrastructure management studies have traditionally relied on pavement-condition surveys, inspection records, distress indices, and engineering assessment [78,79]. These data are valuable but often costly, sparse, and delayed. Recent data-driven studies have explored automated condition assessment using images [80], sensors [81], or crowdsourced information [82]. Yet these approaches usually require direct condition observations. In contrast, the proposed framework focuses on the inference of latent operation-maintenance pressure from indirect operational evidence. This distinction is important because pressure inference can support early screening, prioritization, and risk-aware monitoring even when direct condition labels are unavailable.

Table 1 summarizes the relationship between the main research streams and the proposed framework. The table shows that existing work has made substantial progress in traffic prediction, graph learning, trajectory analysis, and infrastructure monitoring, but these streams are usually developed separately. The proposed framework integrates them through reliability-gated multi-source fusion, stochastic hybrid graph attention, trajectory-derived operational evidence, and maintenance-oriented pressure inference.

3. Methodology

This section presents the proposed reliability-gated SHGAT fusion methodology for road operation-maintenance pressure inference. The central objective is to transform heterogeneous road-network observations into a latent pressure representation that reflects traffic exposure, temporal instability, camera observability, structural context, and neighborhood propagation. The methodology is designed as a general modeling framework rather than a dataset-specific procedure. It first defines the directed road network and its multi-source attributes, then converts trajectory records into segment-time operational evidence and a weak pressure field. Based on these representations, a SHGAT fusion architecture is developed to learn reliability-aware interactions among road segments. A mini-batch stochastic learning scheme is then used to infer pressure from fused graph embeddings and dynamic evidence. Finally, robustness, counterfactual inference, and decision-oriented diagnostics are introduced to evaluate the stability, interpretability, and practical value of the inferred pressure field.

3.1. Road network representation and multi-source operational evidence

Road operation-maintenance pressure emerges from the interaction between physical road structure, traffic usage, and observation conditions. Therefore, the first step is to construct a directed segment-level representation that can accommodate both static infrastructure attributes and observation-related information. In this representation, intersections provide spatial and camera-related context, while directed road segments serve as the primary inference entities. Road class, length, geometric curvature, endpoint topology, and camera coverage are encoded as multi-source operational descriptors. These variables do not directly represent maintenance damage, but they provide the structural and observational background needed for later pressure inference.

Let the urban road system be represented by a directed road graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is the set of intersections and $\mathcal{E}$ is the set of directed road segments. The numbers of nodes and edges are denoted by $N_V = |\mathcal{V}|$ and $N_E = |\mathcal{E}|$, respectively. Each node $i \in \mathcal{V}$ is associated with a spatial coordinate and a camera-observation indicator, while each directed segment $e \in \mathcal{E}$ connects an origin node o_e to a destination node d_e. The directed graph is written as:

$$\mathcal{G} = (\mathcal{V}, \mathcal{E}, \{\mathbf{s}_i\}_{i\in\mathcal{V}}, \{\mathbf{r}_e\}_{e\in\mathcal{E}}), \tag{1}$$

Table 1
Summary of related research streams and positioning of the proposed study.

Research stream	Main focus	Typical data sources	Common modeling strategy	Main limitation for this study	Role in the proposed framework
Traffic-state prediction	Forecasting speed, flow, congestion, or travel time	Sensors, cameras, GPS, historical traffic records	Spatio-temporal models, graph convolution, recurrent networks, attention networks	Usually predict traffic states rather than maintenance pressure	Provides graph-based and temporal modeling foundations
Multi-source transportation fusion	Integrating heterogeneous traffic and infrastructure data	Cameras, trajectories, maps, detectors, probe vehicles	Data-level, feature-level, model-level, or decision-level fusion	Reliability differences across sources are not always explicitly modeled	Motivates reliability-gated fusion across geometry, flow, temporal, and camera evidence
GNNs for road networks	Learning spatial dependencies over network topology	Road graphs, adjacency matrices, traffic states	Graph convolutional networks (GCNs), GATs, adaptive graph learning, spatio-temporal GNNs	Often rely on fixed or prediction-oriented graph relations	Supports hybrid segment graph construction and attention-based propagation
Graph attention mechanisms	Assigning unequal importance to neighboring nodes or segments	Node features, edge features, dynamic traffic states	Attention coefficients, multi-head attention, adaptive weighting	Reliability, road-class compatibility, and length contrast are rarely embedded together	Forms the SHGAT core for reliability-aware neighborhood aggregation
Trajectory-based mobility analysis	Extracting route, travel-time, and movement patterns	Vehicle trajectories, GPS traces, camera-recovered paths	Map matching, demand estimation, travel-time estimation, mobility mining	Maintenance-oriented operational evidence is usually not the primary target	Provides segment-time flow, burden, irregularity, recurrence, and bottleneck evidence
Camera-sparse traffic monitoring	Estimating traffic states under incomplete observation	Sparse cameras, detectors, probe trajectories	Reconstruction, imputation, sensor placement, state estimation	Camera sparsity is often handled as missing data rather than fusion reliability	Becomes a reliability and uncertainty channel in pressure inference
Road maintenance analytics	Assessing pavement condition, deterioration, or maintenance priority	Inspection records, pavement images, distress surveys, sensors	Condition indices, statistical models, machine learning, vision models	Often requires direct condition labels or inspection data	Inspires the pressure-inference objective under limited direct labels
Robust and stochastic learning	Improving model stability under noise or missing data	Perturbed features, dropout, noisy observations	Data augmentation, dropout, adversarial or stochastic training	Less frequently connected to road maintenance decision support	Supports feature perturbation, camera dropout, trajectory degradation, and robustness diagnostics

where $\mathbf{s}_i$ denotes the node attribute vector of intersection i, and $\mathbf{r}_e$ denotes the static road-segment attribute vector of edge e. The node-level attribute vector is defined as:

$$\mathbf{s}_i = [\lambda_i, \phi_i, c_i]^{\mathrm{T}}, \tag{2}$$

where λ_i and ϕ_i are the longitude and latitude of node i, and c_i is a binary camera-availability indicator. The edge-level static attribute vector is expressed as:

$$\mathbf{r}_e = [o_e, d_e, \kappa_e, g_e, \ell_e]^{\mathrm{T}}, \tag{3}$$

where κ_e is the categorical road-class label, g_e is the polyline geometry of the segment, and ℓ_e is the segment length.

Because the framework operates at the road-segment level, the camera state of a segment is induced by the camera states of its two endpoint intersections. The edge is considered camera-covered when at least one endpoint has camera support, while the camera-asymmetry variable measures whether the two endpoints have different observation states. These two variables are denoted by C_e and A_e^C, respectively. They encode complementary information: C_e reflects whether the segment is directly supported by camera infrastructure, whereas A_e^C captures directional or boundary-like observability imbalance.

The road geometry is summarized through a lightweight curvature descriptor. Let $g_e = \left\{\mathbf{p}_{e,a}\right\}_{a\in\mathscr{I}_e}$ be the ordered coordinate sequence of segment e, where $\mathbf{p}_{e,a}$ is the a-th geometry point and $\mathscr{I}_e$ is the corresponding index set. The local turning angle at an interior point is computed from two adjacent geometry vectors. The mean absolute angular change is:

$$\chi_e = \frac{1}{|\mathscr{I}_e^\circ|}\sum_{a\in\mathscr{I}_e^\circ}\left|\arccos\left(\frac{(\mathbf{p}_{e,a}-\mathbf{p}_{e,a^-})^{\mathrm{T}}(\mathbf{p}_{e,a^+}-\mathbf{p}_{e,a})}{\|\mathbf{p}_{e,a}-\mathbf{p}_{e,a^-}\|_2\|\mathbf{p}_{e,a^+}-\mathbf{p}_{e,a}\|_2+\varepsilon}\right)\right|, \tag{4}$$

where $\mathscr{I}_e^\circ$ is the set of valid interior geometry-point indices, a^- and a^+ are the predecessor and successor indices of a, and ε is a small stabilizer.

For heterogeneous quantities with different units and distributions, a robust normalization operator is introduced. Given a scalar feature x and lower and upper empirical reference quantiles $Q_L(x)$ and $Q_U(x)$, the normalized value is:

$$\mathscr{N}(x) = \mathrm{clip}_{[0,1]}\left(\frac{x-Q_L(x)}{Q_U(x)-Q_L(x)+\varepsilon}\right), \tag{5}$$

where ε is a small positive stabilizing constant introduced to prevent numerical instability when the difference between $Q_U(x)$ and $Q_L(x)$ is zero or very small.

This operator is applied to length, road-class index, curvature, flow evidence, time burden, inverse speed, irregularity, recurrence, and other continuous descriptors. Compared with ordinary min-max normalization, this robust operator reduces the influence of extreme observations while preserving a bounded scale suitable for sigmoid-based pressure inference.

The topological exposure of a segment is derived from the endpoint degrees. Let d_i denote the total incident edge count of node i. The segment-level structural degree is:

$$d_e^E = d_{o_e} + d_{d_e}. \tag{6}$$

The normalized degree $\widetilde{d}_e^E = \mathscr{N}(d_e^E)$ is later used as a structural indicator of local graph complexity. A segment connected to high-degree intersections may experience stronger routing interactions, higher conflict exposure, and greater propagation of neighboring traffic disturbances.

The static evidence vector of edge e is therefore constructed as:

$$\mathbf{x}_e^S = \left[\widetilde{\ell}_e, \widetilde{\kappa}_e, C_e, A_e^C, \widetilde{d}_e^E, \widetilde{\chi}_e\right]^{\mathrm{T}}, \tag{7}$$

where $\widetilde{\ell}_e = \mathscr{N}(\ell_e)$, $\widetilde{\kappa}_e = \mathscr{N}(\kappa_e)$, and $\widetilde{\chi}_e = \mathscr{N}(\chi_e)$. This vector defines the static physical, categorical, observational, and topological background of the segment.

3.2. Trajectory aggregation and weak pressure field construction

Vehicle trajectories contain rich information about how road segments are repeatedly used over time, but they must be transformed into segment-level evidence before they can support maintenance-oriented inference. This subsection describes how ordered trajectory points are converted into directed segment-time traversal events and then aggregated into traffic exposure, traversal-time burden, empirical speed, temporal irregularity, recurrence, bottleneck contrast, and observation uncertainty (see Fig. 2). Since direct maintenance labels are not assumed to be available, these evidence components are further fused into a weak pressure field. The resulting target is not intended to replace engineering inspection, but to provide a data-driven operational proxy for identifying road elements with persistent service burden and uncertain observability.

The selection of the operational indicators is guided by their complementary representation of road service exposure, operating efficiency, temporal stability, local traffic interaction, and observation confidence. Flow exposure reflects the accumulated utilization intensity of a road segment and is closely associated with repeated traffic demand and service loading [83]. Traversal-time burden and inverse-speed evidence describe prolonged occupancy and reduced operating efficiency, while temporal irregularity characterizes fluctuations in segment performance and travel-time reliability [84]. Recurrence captures the persistence of demand across neighboring and successive observations, and bottleneck contrast measures the degree to which traffic pressure is locally concentrated relative to surrounding road segments, which is consistent with the spatial organization and propagation of recurrent congestion in urban road networks [85]. Observation reliability further characterizes the confidence associated with the available sensing evidence. These variables jointly describe cumulative exposure, operational inefficiency, temporal instability, recurrent demand, localized bottleneck effects, and observation quality, providing a physically interpretable basis for constructing the weak operation-maintenance pressure field under incomplete direct condition information.

To represent local propagation, a hybrid segment-neighborhood graph is constructed. Two segments are treated as neighbors when they share at least one endpoint. Let $A_{e,e'}^H$ be the hybrid adjacency entry between segments e and e'. It is defined as:

$$A_{e,e'}^H = \mathbb{I}[e \neq e', \{o_e, d_e\} \cap \{o_{e'}, d_{e'}\} \neq \varnothing]. \tag{8}$$

Observation reliability is modeled as a bounded function of camera support, flow evidence, and temporal irregularity. Let $\rho_{e,\tau} \in [0,1]$ denote the reliability of segment-time evidence. It is given by:

$$\rho_{e,\tau} = \sigma\left(\theta_{\rho,0} + \theta_{\rho,C}C_e + \theta_{\rho,F}\widetilde{F}_{e,\tau} - \theta_{\rho,\Omega}\widetilde{\Omega}_{e,\tau}\right), \tag{9}$$

where $\sigma(\cdot)$ is the sigmoid function and $\theta_{\rho,0}$, $\theta_{\rho,C}$, $\theta_{\rho,F}$, and $\theta_{\rho,\Omega}$ are learnable or user-specified reliability parameters. The reliability decreases when observations are camera-sparse and dynamically irregular, and increases when camera support and flow evidence are stronger. The three variables (C_e, $\widetilde{F}_{e,\tau}$, $\widetilde{\Omega}_{e,\tau}$) characterize complementary aspects of observation quality: camera support represents direct sensing availability, flow evidence reflects the amount of trajectory information

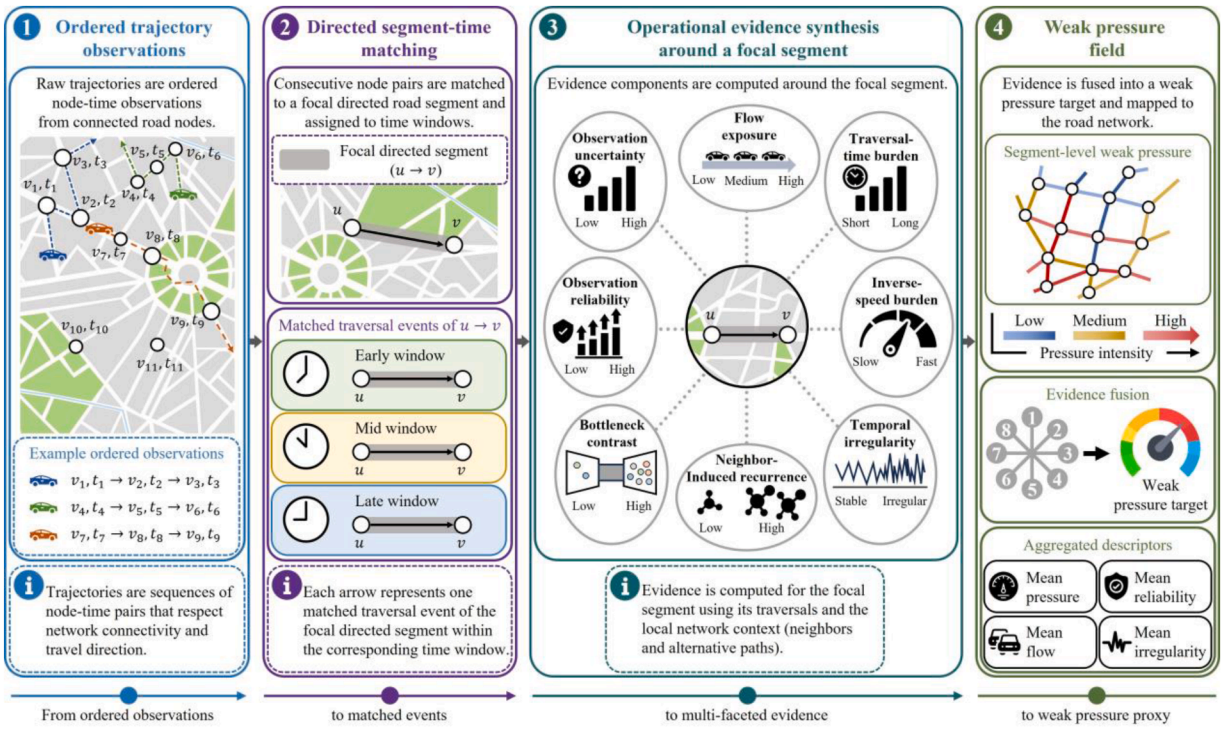


Fig. 2. Conceptual construction of trajectory-derived operational evidence and weak pressure field.

supporting a segment-time state, and temporal irregularity indicates instability in the observed operating condition. The sigmoid function provides a smooth and monotonic mapping of these heterogeneous factors to a bounded confidence scale, allowing reliability to be consistently incorporated into subsequent attention weighting, uncertainty characterization, and information fusion. Accordingly, $\rho_{e,\tau}$ is interpreted as a relative observation-confidence proxy whose magnitude reflects the combined strength and stability of the available evidence.

The observation-uncertainty evidence is:

$$U^{O}_{e,\tau} = \eta_O\left(1-\rho_{e,\tau}\right)\left(1+\widetilde{\Omega}_{e,\tau}\right), \tag{10}$$

where η_O is a positive scaling parameter. The normalized uncertainty descriptor is denoted by $\widetilde{U}^{O}_{e,\tau} = \mathscr{N}\left(U^{O}_{e,\tau}\right)$.

Because direct maintenance labels are not assumed to be available, the pressure field is constructed as a weak, multi-source operational proxy. Let $Y_{e,\tau} \in [0,1]$ be the weak pressure target of segment e at time window τ. It is defined as:

$$\begin{aligned} Y_{e,\tau} = \sigma\Big(&\theta_{Y,0} + \theta_{Y,F}\widetilde{F}_{e,\tau} + \theta_{Y,T}\widetilde{T}_{e,\tau} + \theta_{Y,U}\widetilde{U}^{-1}_{e,\tau} + \theta_{Y,\Omega}\widetilde{\Omega}_{e,\tau} + \theta_{Y,R}\widetilde{R}_{e,\tau} \\ &+ \theta_{Y,B}\widetilde{B}_{e,\tau} + \theta_{Y,O}\widetilde{U}^{O}_{e,\tau}\Big), \end{aligned} \tag{11}$$

where the θ_Y terms are non-negative fusion coefficients except for the intercept $\theta_{Y,0}$. This proxy combines traffic exposure, traversal burden, low-speed tendency, temporal instability, neighbor-induced recurrence, bottleneck contrast, and observation uncertainty. It should be interpreted as an operation-maintenance pressure signal rather than a direct pavement-damage measurement.

Because several operational variables contributing to Eq. (11) also enter the subsequent inference stage, the weak pressure field is treated as a supervisory proxy that provides an initial operational reference for learning. The SHGAT model combines this supervision with road geometry, road class, camera state, observation reliability, hybrid-neighborhood propagation, and reliability-gated modality fusion, enabling the inferred pressure to capture spatial and contextual dependencies beyond the local weak-label construction. The stability of this learned representation is further examined through holdout evaluation, architectural ablation, stochastic robustness analysis, and frozen cross-city validation.

The static hybrid feature vector used for graph attention embedding is:

$$\mathbf{x}^{H}_{e} = \left[\widetilde{\ell}_e, \widetilde{\kappa}_e, C_e, A^{C}_{e}, \widetilde{d}^{E}_{e}, \widetilde{\chi}_e, \mathscr{N}(\overline{F}_e), \mathscr{N}(\overline{\Omega}_e), \overline{\rho}_e, \overline{Y}_e\right]^{\mathrm{T}}. \tag{12}$$

By including $\overline{Y}_e$ and $\overline{\rho}_e$ in the embedding input, the graph representation is guided not only by physical structure but also by aggregate pressure tendency and observation confidence.

3.3. Stochastic hybrid graph attention fusion architecture

Once segment-level evidence has been obtained, the analysis turns to the relationships among neighboring road segments and the contribution of different information sources to the resulting representation. This subsection introduces the SHGAT fusion architecture, in which road segments are treated as graph entities and hybrid adjacency is induced by shared network context (see Fig. 3). The attention mechanism is designed to be reliability-aware and structure-sensitive, so that neighboring segments contribute differently according to observation confidence, road-class compatibility, length contrast, and local propagation relationships. In parallel, a reliability-gated modality fusion mechanism allocates importance across geometry, flow, temporal, and camera channels, allowing the model to represent both source complementarity and source uncertainty.

Let $\mathbf{X}^{H} \in \mathbb{R}^{N_E \times D_H}$ be the matrix collecting $\mathbf{x}^{H}_{e}$ for all segments, where D_H is the hybrid static feature dimension. The initial hidden representation of segment e is:

$$\mathbf{h}^{(0)}_{e} = \tanh\left(\mathbf{W}^{(0)}\mathbf{x}^{H}_{e} + \mathbf{b}^{(0)}\right), \tag{13}$$

where $\mathbf{h}^{(0)}_{e} \in \mathbb{R}^{D_Z}$, D_Z is the embedding dimension, $\mathbf{W}^{(0)}$ is a projection matrix, and $\mathbf{b}^{(0)}$ is a bias vector. The complete hidden matrix at layer l is denoted by $\mathbf{H}^{(l)} \in \mathbb{R}^{N_E \times D_Z}$.

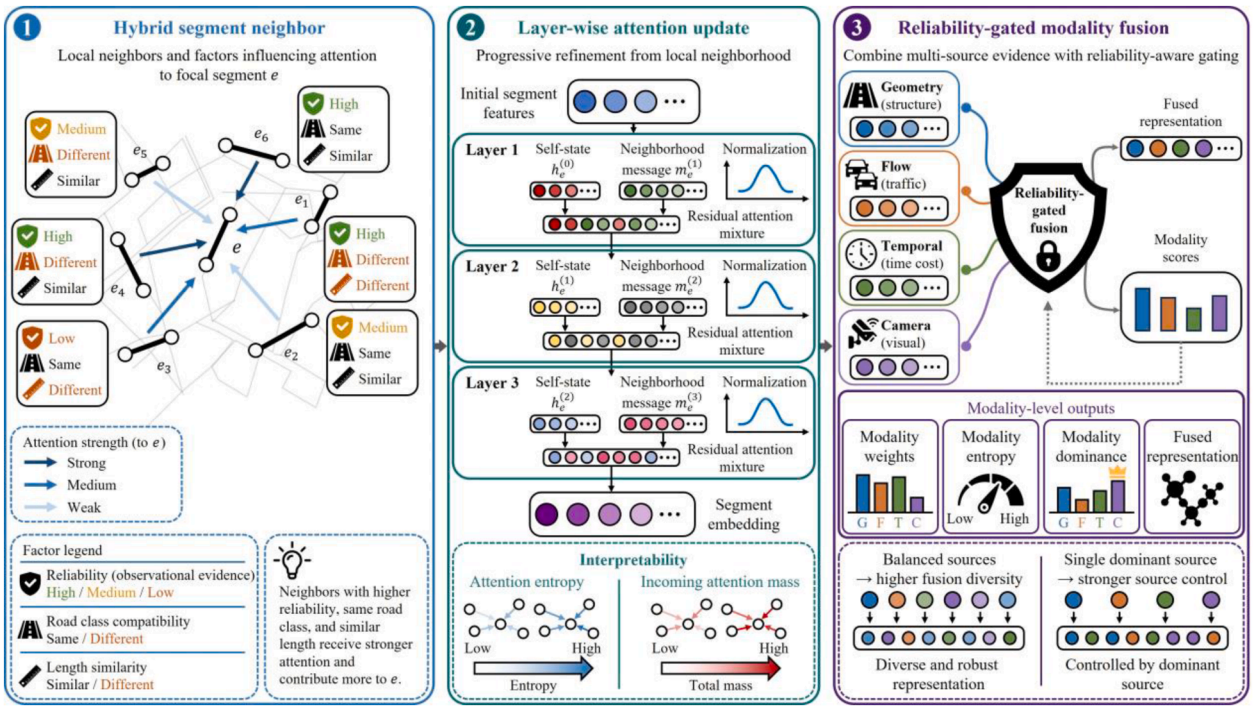


Fig. 3. Reliability-aware hybrid attention and modality gating in SHGAT.

For each graph attention layer l, query $\mathbf{q}_e^{(l)}$, key $\mathbf{k}_e^{(l)}$, and value $\mathbf{v}_e^{(l)}$ representations are computed as:

$$\begin{aligned}\mathbf{q}_e^{(l)} &= \mathbf{W}^{Q,l}\mathbf{h}_e^{(l)},\\ \mathbf{k}_e^{(l)} &= \mathbf{W}^{K,l}\mathbf{h}_e^{(l)},\\ \mathbf{v}_e^{(l)} &= \mathbf{W}^{V,l}\mathbf{h}_e^{(l)},\end{aligned} \tag{14}$$

where $\mathbf{W}^{Q,l}$, $\mathbf{W}^{K,l}$, and $\mathbf{W}^{V,l}$ are layer-specific transformation matrices. For a directed attention pair (e, e') satisfying $A^H_{e,e'} = 1$, the raw attention score is defined by:

$$s^{(l)}_{e,e'} = \frac{(q^{(l)}_e)^T k^{(l)}_{e'}}{\sqrt{D_Z}} + \theta_{a,\rho}\bar{\rho}_{e'} - \theta_{a,\ell}|\tilde{\ell}_e - \tilde{\ell}_{e'}| - \theta_{a,\kappa} II[\kappa_e \neq \kappa_{e'}], \tag{15}$$

where $\theta_{a,\rho}$, $\theta_{a,\ell}$, and $\theta_{a,\kappa}$ control reliability preference, length-contrast penalty, and road-class mismatch penalty, respectively. This design makes attention source-aware and road-structure-aware. A reliable neighboring segment with compatible length and road class receives stronger influence, whereas structurally dissimilar neighbors are suppressed. The attention formulation combines learned neighborhood similarity with explicit reliability and structural compatibility. Following the local adaptive weighting principle of graph attention networks [86], observation reliability increases the contribution of better-supported neighboring segments, while length contrast and road-class mismatch suppress propagation between less compatible road elements. This compact additive form preserves interpretability and sparse local propagation, while more general graph-Transformer formulations can represent richer global structural interactions through explicit structural encodings [87].

The normalized attention coefficient is:

$$\alpha^{(l)}_{e,e'} = \frac{\exp\left(s^{(l)}_{e,e'}\right)}{\sum_{r\in\mathcal{N}^H(e)}\exp\left(s^{(l)}_{e,r}\right) + \varepsilon}, \tag{16}$$

where $\mathcal{N}^H(e) = \left\{e' \in \mathcal{E} : A^H_{e,e'} = 1\right\}$ is the hybrid neighbor set of segment e. The attention-weighted message is:

$$\mathbf{m}_e^{(l)} = \sum_{e'\in\mathcal{N}^H(e)} \alpha^{(l)}_{e,e'}\mathbf{v}^{(l)}_{e'}. \tag{17}$$

The hidden state is updated through a residual attention mixture:

$$\mathbf{z}_e^{(l+1)} = \tanh\left(\eta_l \mathbf{h}_e^{(l)} + (1-\eta_l)\mathbf{m}_e^{(l)}\right), \tag{18}$$

where $\eta_l \in [0,1]$ controls the balance between self-retention and neighboring message aggregation. To stabilize multi-layer propagation, layer normalization is applied row-wise:

$$\mathbf{h}_e^{(l+1)} = \frac{\mathbf{z}_e^{(l+1)} - \mu_e^{(l+1)}\mathbf{1}}{\sigma_e^{(l+1)} + \varepsilon}, \tag{19}$$

where $\mu_e^{(l+1)}$ and $\sigma_e^{(l+1)}$ are the mean and standard deviation of the components of $\mathbf{z}_e^{(l+1)}$, and 1 is an all-ones vector. After L layers, the graph-attention embedding of segment e is:

$$\mathbf{h}_e^G = \mathbf{h}_e^{(L)}. \tag{20}$$

The final embedding matrix is $\mathbf{H}^G \in \mathbb{R}^{N_E \times D_Z}$. The attention entropy of segment e is used as an interpretability measure of whether attention is concentrated or dispersed:

$$E^A_e = -\frac{\sum_{e'\in\mathcal{N}^H(e)}\alpha^{(L)}_{e,e'}\log\left(\alpha^{(L)}_{e,e'} + \varepsilon\right)}{\log\left(D^H_e + \varepsilon\right)}. \tag{21}$$

The incoming attention mass of segment e is:

$$M^A_e = \sum_{r\in\mathcal{E}} A^H_{r,e}\alpha^{(L)}_{r,e}. \tag{22}$$

While E^A_e reflects local attention dispersion, M^A_e reflects how strongly a segment is attended by its neighbors. These two quantities are later useful for explaining graph-based pressure propagation.

The framework further constructs modality-level evidence scores. Let S^G_e, S^F_e, S^T_e, and S^C_e denote geometry, flow, temporal, and camera-source scores, respectively. They are represented as:

$$S_e^G = \boldsymbol{\omega}_G^{\mathrm{T}}\left[\widetilde{\ell}_e, \widetilde{\chi}_e, \widetilde{d}_e^E\right]^{\mathrm{T}}, \tag{23}$$

$$S_e^F = \boldsymbol{\omega}_F^{\mathrm{T}}[\mathcal{N}(\overline{F}_e), \mathcal{N}(\overline{R}_e)]^{\mathrm{T}}, \tag{24}$$

$$S_e^T = \boldsymbol{\omega}_T^{\mathrm{T}}\left[\overline{T}_e^*, \overline{\Omega}_e\right]^{\mathrm{T}}, \tag{25}$$

$$S_e^C = \boldsymbol{\omega}_C^{\mathrm{T}}\left[C_e, \overline{\rho}_e, A_e^C\right]^{\mathrm{T}}, \tag{26}$$

where $\boldsymbol{\omega}_G$, $\boldsymbol{\omega}_F$, $\boldsymbol{\omega}_T$, and $\boldsymbol{\omega}_C$ are source-specific weight vectors, $\overline{R}_e$ is the temporal mean of $R_{e,\tau}$, and $\overline{T}_e^*$ is a representative normalized traversal-time burden. The four scores are stacked into a modality score vector:

$$\mathbf{s}_e^M = \left[S_e^G, S_e^F, S_e^T, S_e^C\right]^{\mathrm{T}}. \tag{27}$$

The reliability-gated modality allocation is obtained by a row-wise softmax [88]:

$$\boldsymbol{\beta}_e = \mathrm{softmax}\left(\mathbf{s}_e^M\right), \tag{28}$$

where $\boldsymbol{\beta}_e = \left[\beta_e^G, \beta_e^F, \beta_e^T, \beta_e^C\right]^{\mathrm{T}}$ contains the modality weights of segment e.

3.4. Mini-batch learning and pressure inference

The fused graph representation must be connected to a scalable learning procedure that can handle large road networks and many segment-time samples. This subsection describes the mini-batch stochastic pressure inference scheme. For each segment-time sample, graph embeddings are combined with dynamic traffic evidence, reliability indicators, camera state, and static road descriptors. A multi-hidden-layer inference head maps these fused features to a bounded pressure estimate. During training, stochastic perturbations are applied to the feature space to improve robustness against noisy measurements, source incompleteness, and camera-related degradation. The learning process therefore serves both prediction and regularization purposes, making the inferred pressure field more stable under realistic information-fusion conditions.

For each segment-time sample (e,τ), the final inference feature vector combines graph embedding, dynamic traffic evidence, reliability evidence, camera evidence, and static road evidence. It is defined as:

$$\mathbf{x}_{e,\tau} = \left[\left(\mathbf{h}_e^G\right)^{\mathrm{T}}, \widetilde{F}_{e,\tau}, \widetilde{T}_{e,\tau}, \widetilde{U}_{e,\tau}^{-1}, \widetilde{\Omega}_{e,\tau}, \widetilde{R}_{e,\tau}, \rho_{e,\tau}, C_e, A_e^C, \widetilde{\ell}_e, \widetilde{\kappa}_e, \widetilde{\chi}_e\right]^{\mathrm{T}}. \tag{29}$$

Let $\mathbf{X} \in \mathbb{R}^{N_S \times D_X}$ be the sample-feature matrix, where N_S is the number of segment-time samples and D_X is the feature dimension. The corresponding weak target vector is $\in \mathbb{R}^{N_S}$, where each entry is $Y_{e,\tau}$. Features are standardized using the training-set mean $\boldsymbol{\mu}_X$ and standard deviation $\boldsymbol{\sigma}_X$:

$$\widehat{\mathbf{x}}_{e,\tau} = \frac{\mathbf{x}_{e,\tau} - \boldsymbol{\mu}_X}{\boldsymbol{\sigma}_X + \varepsilon}. \tag{30}$$

The pressure inference head is a multi-hidden-layer neural mapping. The first hidden layer is:

$$\mathbf{a}_{e,\tau}^{(1)} = \mathrm{ReLU}\left(\mathbf{W}^{(1)}\widehat{\mathbf{x}}_{e,\tau} + \mathbf{b}^{(1)}\right), \tag{31}$$

and the second hidden layer is:

$$\mathbf{a}_{e,\tau}^{(2)} = \mathrm{ReLU}\left(\mathbf{W}^{(2)}\mathbf{a}_{e,\tau}^{(1)} + \mathbf{b}^{(2)}\right). \tag{32}$$

The predicted pressure is bounded by a sigmoid output:

$$\widehat{Y}_{e,\tau} = \sigma\left(\left(\mathbf{w}^{(3)}\right)^{\mathrm{T}}\mathbf{a}_{e,\tau}^{(2)} + b^{(3)}\right). \tag{33}$$

The supervised objective is the mean squared error between predicted pressure and weak pressure target:

$$\mathscr{L}_{\mathrm{MSE}} = \frac{1}{|\mathscr{B}|}\sum_{(e,\tau)\in\mathscr{B}}\left(\widehat{Y}_{e,\tau} - Y_{e,\tau}\right)^2, \tag{34}$$

where $\mathscr{B}$ denotes a mini-batch of segment-time samples. The mini-batch is sampled from both active segment-time states and randomly selected background states. This mixture prevents the model from learning only frequently traversed segments and improves coverage over camera-sparse or low-activity road elements.

Stochastic feature perturbation is applied during training. Let $\widehat{\mathbf{x}}_n$ denote the standardized feature vector of sample n. Its perturbed version is:

$$\widehat{\mathbf{x}}_n^P = \widehat{\mathbf{x}}_n \odot \mathbf{d}_n + \epsilon_n, \tag{35}$$

where $\odot$ is element-wise multiplication, $\mathbf{d}_n$ is a Bernoulli feature-retention vector, and ϵ_n is Gaussian feature noise. Their distributions are:

$$\begin{aligned} d_{nj} &\sim \mathrm{Bernoulli}(p_D), \\ \epsilon_{nj} &\sim \mathcal{N}\left(0, \sigma_P^2 s_j^2\right), \end{aligned} \tag{36}$$

where p_D is the feature-retention probability, σ_P controls perturbation magnitude, and s_j is the empirical standard deviation of feature j. This stochastic training strategy increases robustness to missing sources, noisy traffic evidence, and camera-observation instability.

Let Θ collect all trainable parameters of the inference head. The stochastic gradient update is based on Adam-type moment estimation [89]:

$$\begin{aligned} \mathbf{g}_r &= \nabla_\Theta \mathscr{L}_{MSE}(\Theta_r), \\ \mathbf{m}_r &= \beta_1 \mathbf{m}_{r^-} + (1-\beta_1)\mathbf{g}_r, \\ \mathbf{v}_r &= \beta_2 \mathbf{v}_{r^-} + (1-\beta_2)(\mathbf{g}_r \odot \mathbf{g}_r), \\ \widehat{\mathbf{m}}_r &= \frac{\mathbf{m}_r}{1-\beta_1^r} \\ \widehat{\mathbf{v}}_r &= \frac{\mathbf{v}_r}{1-\beta_2^r} \\ \Theta_{r^+} &= \Theta_r - \eta\frac{\widehat{\mathbf{m}}_r}{\sqrt{\widehat{\mathbf{v}}_r} + \varepsilon}, \end{aligned} \tag{37}$$

where r is the iteration index, η is the learning rate, β_1 and β_2 are exponential decay factors, and r^- and r^+ indicate the previous and updated iterations.

The fused priority score combines predicted pressure, unreliability, flow exposure, and structural exposure:

$$\Pi_e = \mathcal{N}\left(\widehat{P}_e\left(1-\rho_{e,\tau^*}\right)\left[1+\mathcal{N}(\overline{F}_e)\right]\left[1+\alpha_D\mathcal{N}\left(d_e^E\right)\right]\right), \tag{38}$$

where α_D controls the contribution of structural exposure. This score prioritizes road elements that are predicted to be pressure-intensive, uncertain, demand-loaded, and structurally influential.

3.5. Robustness, counterfactual inference, and decision-oriented diagnostics

Pressure inference is useful for road operation and maintenance only when its outputs are stable, interpretable, and actionable. This subsection therefore extends the model from prediction to diagnostic decision support. Robustness analysis examines how pressure estimates change under camera dropout, trajectory degradation, feature perturbation, flow noise, and geometry noise. Counterfactual inference evaluates how pressure fields respond to hypothetical changes in camera retention, demand scaling, and capacity-related conditions. Decision-oriented diagnostics further quantify source-channel sensitivity, uncertainty, propagation gaps, residual structure, ordinal pressure classification, and

top-risk rank stability. These analyses help convert the inferred pressure field into maintenance priorities, intervention values, and reliability-aware operational insights.

The framework evaluates robustness by simulating information degradation in different source channels. Camera dropout is represented by a Bernoulli retention variable b_n^C for sample n. The perturbed camera feature is:

$$\begin{aligned} C_n^P &= C_n b_n^C, \\ b_n^C &\sim \text{Bernoulli}(p_C), \end{aligned} \tag{39}$$

where p_C is the camera-retention probability. The perturbed reliability is adjusted as:

$$\rho_n^P = \text{clip}_{[0,1]}\left(\rho_n g_C(C_n^P)\right), \tag{40}$$

where $g_C(\cdot)$ is a monotone camera-reliability adjustment function. Trajectory dropout is represented by scaling flow and recurrence features and perturbing irregularity:

$$\begin{aligned} \widetilde{F}_n^P &= clip_{[0,1]}\left(r_T \widetilde{F}_n + \epsilon_n^F\right), \\ \widetilde{R}_n^P &= clip_{[0,1]}\left(r_T \widetilde{R}_n + \epsilon_n^R\right), \\ \widetilde{\Omega}_n^P &= clip_{[0,1]}\left(\widetilde{\Omega}_n + \epsilon_n^\Omega\right), \end{aligned} \tag{41}$$

where r_T is a trajectory-retention factor, and ϵ_n^F, ϵ_n^R, and ϵ_n^Ω are perturbation terms. Flow-channel noise is expressed as:

$$\mathbf{x}_{n,\mathscr{F}}^P = \text{clip}_{[0,1]}\left(\mathbf{x}_{n,\mathscr{F}} + \epsilon_n^F\right), \tag{42}$$

where $\mathbf{x}_{n,\mathscr{F}}$ contains flow, recurrence, and irregularity-related components. Geometry-channel noise is:

$$\mathbf{x}_{n,\mathscr{G}}^P = \text{clip}_{[0,1]}\left(\mathbf{x}_{n,\mathscr{G}} + \epsilon_n^G\right), \tag{43}$$

where $\mathbf{x}_{n,\mathscr{G}}$ contains length and curvature descriptors.

Stochastic uncertainty is estimated by repeated perturbations. Let $\widehat{y}_n^{(m)}$ be the prediction of sample n under the m-th stochastic perturbation. The uncertainty score is:

$$\mathscr{U}_n = \text{Std}_m\left(\widehat{y}_n^{(m)}\right) + \lambda_\rho(1-\rho_n), \tag{44}$$

where λ_ρ controls the contribution of observation unreliability. The first term captures model response variability under perturbation, while the second term injects source-confidence awareness.

Counterfactual inference is used to evaluate how pressure would change under hypothetical camera retention, demand scaling, and capacity retention. Let r_C, r_D, and r_Q be camera-retention, demand-scale, and capacity-retention factors. Under a camera-demand counterfactual intervention, the modified feature vector is:

$$\mathbf{x}_{e,\tau}^{CD}(r_C, r_D) = \Gamma_{CD}\left(\mathbf{x}_{e,\tau}; r_C, r_D\right), \tag{45}$$

where Γ_{CD} scales camera-related reliability and demand-related flow and recurrence channels. The mean counterfactual pressure surface is:

$$\mathscr{P}^{CD}(r_C, r_D) = \frac{1}{N_E}\sum_{e\in\mathscr{E}} f_\Theta\left(\mathbf{x}_{e,\tau^*}^{CD}(r_C, r_D)\right), \tag{46}$$

where $f_\Theta(\cdot)$ denotes the trained pressure inference function. Under a capacity-retention intervention, traversal-time burden and inverse-speed pressure are rescaled:

$$\mathbf{x}_{e,\tau}^{QD}(r_Q, r_D) = \Gamma_{QD}\left(\mathbf{x}_{e,\tau}; r_Q, r_D\right). \tag{47}$$

The top-risk counterfactual pressure is:

$$\mathscr{P}_q^{QD}(r_Q, r_D) = \frac{1}{\left|\text{Top}_q\left(\widehat{\mathbf{P}}^{QD}\right)\right|}\sum_{e\in\text{Top}_q\left(\widehat{\mathbf{P}}^{QD}\right)} \widehat{P}_e^{QD}(r_Q, r_D), \tag{48}$$

where $\widehat{\mathbf{P}}^{QD}$ is the vector of counterfactual predictions under the capacity-demand intervention. The pressure relief from a counterfactual intervention is:

$$\mathscr{R}^I = \overline{P}^{\text{base}} - \overline{P}^{\text{cf}}, \tag{49}$$

where $\overline{P}^{\text{base}}$ and $\overline{P}^{\text{cf}}$ are the mean baseline and counterfactual pressure values.

Local source sensitivity is estimated by finite differences. For feature channel j, the sensitivity of segment e is:

$$S_{e,j} = \frac{\left|f_\Theta\left(\mathbf{x}_{e,\tau^*} + \delta\mathbf{e}_j\right) - f_\Theta\left(\mathbf{x}_{e,\tau^*}\right)\right|}{\delta}, \tag{50}$$

where δ is a small perturbation magnitude and $\mathbf{e}_j$ is the unit vector for channel j. This yields flow-channel sensitivity and camera-channel sensitivity by selecting the corresponding feature indices.

Graph propagation diagnostics compare local pressure with neighboring pressure. The neighbor-averaged prediction is:

$$\widehat{P}_e^N = \frac{\sum_{e'\in\mathscr{E}} A_{e,e'}^H \widehat{P}_{e'}}{D_e^H + \varepsilon}, \tag{51}$$

and the propagation gap is:

$$G_e^P = \widehat{P}_e - \widehat{P}_e^N. \tag{52}$$

A positive gap indicates a local pressure elevation relative to adjacent road elements, while a negative gap indicates that the segment is embedded in a higher-pressure neighborhood.

The local intervention value combines flow sensitivity, camera sensitivity, uncertainty, and unreliability:

$$V_e^I = \mathscr{N}\left(\theta_{I,F} S_{e,F}\left(1-\rho_{e,\tau^*}\right) + \theta_{I,C} S_{e,C}\widehat{P}_e + \theta_{I,U}\mathscr{U}_e\right), \tag{53}$$

where $S_{e,F}$ and $S_{e,C}$ are flow-channel and camera-channel sensitivities, and $\theta_{I,F}$, $\theta_{I,C}$, and $\theta_{I,U}$ are intervention-fusion coefficients.

Together, the robustness, counterfactual, and diagnostic formulations extend the pressure-inference model from a predictive tool to an information-fusion decision system. The framework estimates pressure, measures source reliability, tests degradation stability, identifies sensitivity channels, evaluates counterfactual intervention effects, and provides graph-aware uncertainty diagnostics for road operation and maintenance planning.

3.6. Algorithmic implementation of the proposed SHGAT framework

To provide a concise and reproducible description of the overall computational procedure, the principal implementation steps of the proposed SHGAT framework are summarized in Algorithm 1. The algorithm integrates road-network preprocessing, trajectory aggregation, operational evidence construction, observation-reliability estimation, weak pressure generation, reliability-aware graph attention, modality fusion, stochastic mini-batch learning, and pressure inference within a unified workflow. It further incorporates uncertainty quantification, robustness assessment, and maintenance-priority evaluation, thereby linking the graph-learning procedure with the reliability-oriented diagnostic and decision-support objectives of the framework.

4. Results

Based on the parameter settings in Tables 2 and 3, a series of numerical experiments was conducted to validate the effectiveness of the analytical framework in this study. The empirical analysis uses a public large-scale vehicle trajectory dataset constructed from traffic-camera videos and road-network information. In this study, the selected Shenzhen trajectory record dated November 4, 2020 [90] is adopted as

Algorithm 1
Reliability-informed SHGAT training, inference, and independent cross-city validation.

Input: Development-city road-network, trajectory, and camera data; external-city data; SHGAT hyperparameters; directed road graph $\mathcal{G}$; hybrid segment adjacency $\mathbf{A}^{\mathrm{H}}$; number of temporal windows.

Output: Trained pressure-inference function; segment-time pressure predictions $\widehat{Y}_{e,\tau}$, stochastic uncertainty, maintenance-priority score Π_{e}, robustness/diagnostic quantities, and independent external-validation metrics.

1: Read and clean node/edge records; construct $\mathcal{G}$ and $\mathbf{A}^{\mathrm{H}}$ from shared segment endpoints. Derive road class, length, curvature, structural degree, camera coverage, and camera asymmetry.
2: Map trajectory point sequences to directed road segments and aggregate them into multiple segment-time windows. Compute flow $F_{e,\tau}$, traversal-time burden, inverse-speed evidence, temporal irregularity $\Omega_{e,\tau}$, recurrence $R_{e,\tau}$, and bottleneck evidence $B_{e,\tau}$.
3: Normalize heterogeneous descriptors and calculate observation reliability $\rho_{e,\tau}$ and observation uncertainty $U^{O}_{e,\tau}$. Fuse the operational evidence into the weak pressure target $Y_{e,\tau}$.
4: Assemble the static hybrid feature matrix $\mathbf{X}^{\mathrm{H}}$ and project it to the SHGAT hidden space.
5: For attention layer $l = 1, \ldots, L$:
6: Compute query/key/value representations and reliability-aware attention coefficients using feature similarity, neighboring reliability, length contrast, and road-class compatibility.
7: Aggregate hybrid-neighbor messages, apply residual self-retention and layer normalization, and update the hidden representation.
8: End the attention loop and obtain the final graph embedding matrix $\mathbf{H}^{\mathrm{G}}$ and segment embedding $\mathbf{h}^{\mathrm{G}}_{e}$.
9: Compute geometry-, flow-, temporal-, and camera-modality scores; convert them to reliability-gated softmax weights and assemble the final segment-time feature vector $\mathbf{x}_{e,\tau}$.
10: For the development city, sample active and background segment-time states, split training/validation data, and standardize features using the training statistics $\boldsymbol{\mu}_X$ and $\boldsymbol{\sigma}_X$.
11: Train the multilayer perceptron pressure head by stochastic mini-batches: apply Bernoulli feature dropout and Gaussian feature noise, predict $\widehat{Y}_{e,\tau}$, minimize MSE against $Y_{e,\tau}$, and update the trainable parameters with Adam.
12: Freeze the trained model and training normalization statistics. Perform network-wide pressure inference; estimate stochastic uncertainty by repeated perturbations and derive Π_{e} together with robustness, sensitivity, propagation, and counterfactual diagnostics.
13: For independent validation, repeat Steps 1–9 with the same feature definitions, SHGAT depth, coefficients, and perturbation settings; map road-class scalars to development city training scale.
14: Apply the frozen trained model with the frozen $\boldsymbol{\mu}_X$ and $\boldsymbol{\sigma}_X$ directly to the independent validation data; do not retrain, fine-tune, or re-estimate normalization statistics. Compute external MAE, RMSE, correlation, pressure, uncertainty, and priority fields.
15: **Return** the trained SHGAT inference system and the in-domain/external validation outputs for reliability-aware maintenance screening.

Table 2
Parameter settings used in the study.

Symbol	Meaning	Value
N_{epo}	Number of mini-batch training epochs	180
$\mathscr{B}$	Mini-batch size	1024
γ	Learning rate of the Adam optimizer	2.5×10^{-3}
α_R	Neighbor-flow recurrence coefficient in $R_{e,\tau}$	0.5
$\theta_{\rho,0}$	Intercept in observation-reliability construction	-1.25
$\theta_{\rho,C}$	Camera-support weight in observation reliability	1.85
$\theta_{\rho,F}$	Flow-evidence weight in observation reliability	1.15
$\theta_{\rho,\Omega}$	Temporal-irregularity penalty in observation reliability	1.10
η_O	Observation-uncertainty scaling coefficient	0.5
$\theta_{Y,0}$	Intercept in the weak pressure proxy	-2.10
$\theta_{Y,F}$	Flow-exposure weight in the weak pressure proxy	2.10
$\theta_{Y,T}$	Traversal-time burden weight in the weak pressure proxy	1.05
$\theta_{Y,U}$	Inverse-speed burden weight in the weak pressure proxy	1.35
$\theta_{Y,\Omega}$	Temporal-irregularity weight in the weak pressure proxy	1.55
$\theta_{Y,R}$	Recurrence-flow weight in the weak pressure proxy	1.50
$\theta_{Y,B}$	Bottleneck-contrast weight in the weak pressure proxy	1.05
$\theta_{Y,O}$	Observation-uncertainty weight in the weak pressure proxy	1.20
$N^{\mathrm{val}}_{\mathrm{MC}}$	Monte Carlo repetitions for test uncertainty	12
$N^{\mathrm{peak}}_{\mathrm{MC}}$	Monte Carlo repetitions for peak-window uncertainty	10
δ_{FD}	Finite-difference step for source-channel sensitivity	0.05

the study-day trajectory source, and the full-day observation period is divided into 24 temporal windows for segment-level operational evidence construction. Fig. 4 shows the spatial coverage of the selected Shenzhen trajectory study area, where the dashed boundary indicates the vehicle-trajectory coverage region used in the numerical experiments. This spatial reference is important because all subsequent flow exposure, traversal burden, camera-sparse pressure, graph propagation, and intervention diagnostics are interpreted within this selected urban road-network domain. To further evaluate cross-city transferability and external validity, an independent Jinan dataset collected on October 17, 2022 is additionally introduced as a second case. The SHGAT model trained on the Shenzhen data is directly transferred to the Jinan road network without retraining, fine-tuning, or parameter adjustment, thereby providing an independent assessment under a distinct urban road-network and traffic-observation environment. The principal data characteristics of the Shenzhen development case and the Jinan external-validation case are summarized in Table 4, including the numbers of road-network nodes, valid directed segments, processed trajectories, camera density, missing-observation ratio, and road-class distribution. The comparison highlights substantial differences in network scale, sensing coverage, trajectory availability, and structural composition between the two cities, providing a heterogeneous basis for evaluating both in-domain performance and cross-city generalization.

The implemented model follows the SHGAT configuration summarized in Table 3. In this setting, directed road segments are treated as graph entities, and the hybrid neighborhood is constructed from shared intersection endpoints. The model first extracts a static hybrid segment representation from length, road class, camera support, camera asymmetry, edge degree, curvature, mean flow, mean temporal irregularity, mean reliability, and mean weak pressure. This representation is then projected into a hidden space and propagated through a three-layer stochastic hybrid graph attention encoder. The attention mechanism is not a purely topological aggregation rule; it is reliability-gated and structure-aware, since neighboring segments are weighted according to similarity, observation reliability, length contrast, and road-class compatibility. Fig. 5 summarizes this architecture, including the edge-level hybrid graph input, the initial projection, the three-layer SHGAT encoder, the modality fusion module, the segment-time feature assembly, and the stochastic mini-batch training process. After the SHGAT embedding is obtained, each segment-time sample is represented by the graph embedding together with dynamic and contextual variables, including flow exposure, traversal time, inverse-speed burden, temporal irregularity, recurrence, reliability, camera state, camera asymmetry, length, class, and curvature. These features are passed to a multilayer perceptron pressure-inference head and optimized against the weak pressure target derived from trajectory evidence. During training, mini-batch sampling, feature noise, and feature dropout are introduced to improve stability under incomplete and noisy multi-source observations.

Furthermore, the computational efficiency of the proposed framework was evaluated in terms of theoretical complexity and actual runtime. For N road segments, M retained hybrid-neighbor relations, hidden dimension d, and L attention layers, the graph-attention stage has a computational complexity of approximately $O\left[Nd^2 + Md\right]$, while its principal memory requirement scales as $O[Nd + M]$; the subsequent mini-batch inference head scales linearly with the number of processed samples for a fixed network architecture. Since the hybrid neighborhood

Table 3
SHGAT architecture, fusion, and learning settings.

Setting group	Configuration	Implementation value	Role in the framework
Segment entity	Graph node type	Directed road segment	Treats road segments as edge-level graph entities
Hybrid adjacency	Neighborhood rule	Shared start or end node	Constructs segment-to-segment propagation relations
Neighbor sparsification	Maximum incident segment sampling	18 per road node	Controls graph density and computational cost
Static SHGAT input	$\mathbf{x}_e^H$ dimension	10	Contains length, road class, camera state, asymmetry, edge degree, curvature, mean flow, mean irregularity, mean reliability, and mean weak pressure
Initial projection	$\mathbf{W}^{(0)}$ mapping	$10 \rightarrow 10$	Maps static segment features to hidden graph representation
Hidden dimension	D_Z	10	Embedding size of each segment-node
Number of SHGAT layers	L	3	Stacked hybrid graph attention propagation
Query-key-value projection	**Q**, **K**, **V** mapping	$10 \rightarrow 10$ per layer	Computes attention similarity and neighbor messages
Attention reliability bonus	$\theta_{a,\rho}$	0.35	Increases attention to reliable neighboring segments
Length-contrast penalty	$\theta_{a,l}$	0.20	Suppresses structurally dissimilar neighbors
Road-class mismatch penalty	$\theta_{a,\kappa}$	0.10	Penalizes incompatible road-class neighbors
Attention-score clipping	Score range	[-20, 20]	Avoids numerical overflow in attention normalization
Residual self-retention	Self coefficient	0.45	Preserves the current segment representation
Neighbor message mixing	Neighbor coefficient	0.55	Injects attention-weighted neighboring information
Nonlinear activation	Graph layer activation	tanh	Bounds the graph-hidden representation
Normalization	Graph layer normalization	Row-wise layer normalization	Stabilizes multi-layer propagation
Final SHGAT embedding	H_{edge}	10	Segment-level graph embedding used in sample features
Modality channels	Source groups	Geometry, flow, temporal, camera	Supports reliability-gated source interpretation
Modality allocation	Normalization rule	Row-wise softmax	Produces segment-level modality weights
Training sample construction	Feature dimension	21	Combines SHGAT embedding with segment-time features
Segment-time appended features	Dynamic and context features	Flow, mean traversal time, inverse speed, irregularity, recurrence, reliability, camera indicator, camera asymmetry, length, class, curvature	Forms the multilayer perceptron (MLP) input after SHGAT embedding
Active-background sampling	Sample composition	About 70% active samples and 30% random background samples	Balances observed traversal events and background segment-time states
Train-validation-test split	Nominal split	70% training, 15% validation, and 15% test	Evaluates generalization on a holdout set
Pressure head	MLP architecture	$21 \rightarrow 48 \rightarrow 24 \rightarrow 1$	Maps fused features to predicted pressure
Hidden activations	MLP hidden layers	ReLU	Introduces nonlinear pressure inference
Output activation	Pressure output	Sigmoid	Bounds predicted pressure in [0, 1]
Loss function	Training objective	MSE	Fits predicted pressure to weak pressure target
Stochastic feature noise	Input perturbation	0.030	Improves robustness to noisy evidence
Feature dropout	Input dropout rate	0.050	Improves robustness to missing or degraded features
Optimizer	Parameter update rule	Adam	Updates pressure-head parameters

is sparsified by limiting local segment connections, the number of retained relations grows approximately with network size, supporting application to larger urban road networks without constructing dense global attention matrices. The experiments were conducted on a workstation equipped with an NVIDIA GeForce RTX 4070 GPU, an Intel Core i7-13700F CPU, and 48 GB of system memory. Under this configuration, model training required 18 min 32 s and inference required 6 min 12 s, indicating that the complete workflow can be executed within a practical computational budget for network-scale road assessment.

4.1. Road-network observability and operational risk evidence

This subsection examines the empirical reliability structure of the road-network evidence used for pressure inference. The objective is not only to describe the spatial and temporal coverage of the data, but also to determine whether the selected network contains the type of heterogeneity required for reliability-oriented maintenance assessment. Therefore, the analysis begins with road-class composition, segment length, directional loading, and camera observability, and then proceeds to flow exposure, traversal-time burden, temporal irregularity, and pressure-threshold exceedance. These results clarify how raw trajectory records are converted into interpretable segment-level risk evidence under uneven observation support.

Fig. 6 establishes the empirical basis for the subsequent pressure-inference analysis by showing that the selected road network is structurally heterogeneous, camera-sparse, and temporally uneven. Panel (a) shows dense local connectivity in the directed road graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, but camera-covered nodes are unevenly distributed and mainly

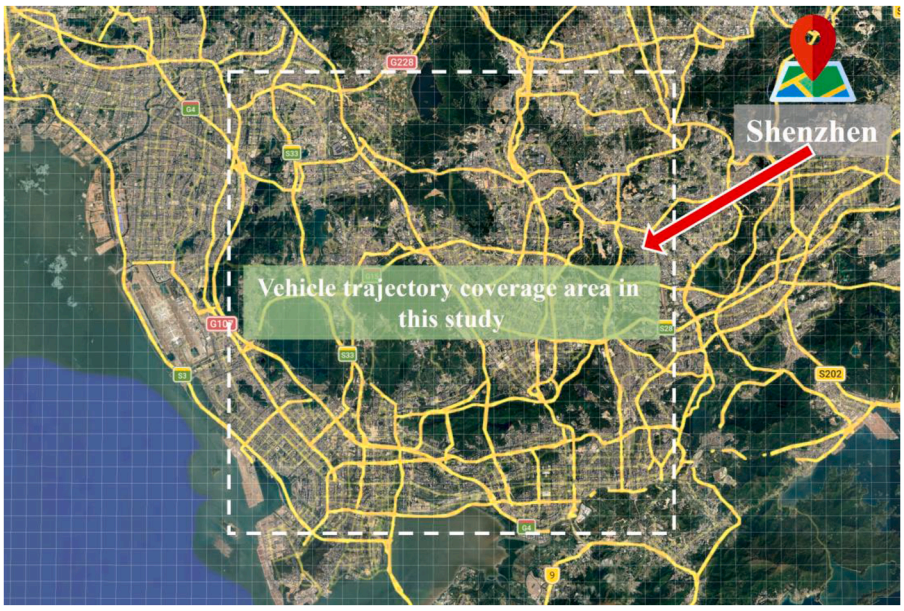


Fig. 4. Spatial coverage of the selected Shenzhen trajectory study area.

Table 4
Dataset characteristics of the Shenzhen development case and Jinan independent external-validation case.

Characteristic	Shenzhen	Jinan
Observation date	November 4, 2020	October 17, 2022
Number of nodes	11,933	8,908
Number of valid directed road segments	27,410	23,312
Number of trajectories used	3,823	43,461
Matched segment-traversal events	114,860	450,000
Camera nodes	686	1,651
Camera density	5.75%	18.53%
Missing-data ratio	96.15%	91.83%
Number of road classes	13	13
Residential	7,597 (27.72%)	7,551 (32.39%)
Tertiary	6,166 (22.50%)	5,767 (24.74%)
Secondary	3,779 (13.79%)	4,151 (17.81%)
Primary	3,535 (12.90%)	2,813 (12.07%)
Unclassified	3,348 (12.21%)	452 (1.94%)
Trunk	471 (1.72%)	1,283 (5.50%)
Primary link	924 (3.37%)	76 (0.33%)
Trunk link	379 (1.38%)	413 (1.77%)
Motorway link	414 (1.51%)	186 (0.80%)
Secondary link	279 (1.02%)	34 (0.15%)
Motorway	224 (0.82%)	128 (0.55%)
Living street	209 (0.76%)	449 (1.93%)
Tertiary link	85 (0.31%)	9 (0.04%)

concentrated in several central and corridor-like areas. This indicates biased observation reliability and supports treating camera availability as a reliability signal rather than a simple binary attribute. Panel (b) shows that residential, tertiary, secondary, and primary roads dominate the edge population, whereas trunk-related and motorway-link classes are fewer but generally have larger mean segment lengths ℓ_e. Quantitatively, residential and tertiary roads contain approximately 7,600 and 6,200 edges, respectively, whereas several trunk- and link-related classes contain fewer than 1,000 edges despite mean segment lengths approaching 500 m. Thus, the network combines many short urban access links with fewer longer high-capacity links, leading to class-dependent maintenance exposure.

Panels (c)-(e) further indicate that both pressure and observability depend on road class and direction. The mean weak pressure $\overline{Y}_e$ varies with bearing direction, suggesting that pressure is affected by corridor orientation and network layout, not only by total traffic volume. Across the directional groups, the mean weak pressure varies from approximately 0.44 to 0.49, with the highest level occurring around the 130° bearing range. Panel (d) confirms that high camera observability is limited to specific spatial zones, while panel (e) reveals a mismatch between class-level camera-weighted coverage and peak-window flow share. Some classes are well observed but carry limited peak demand, whereas others contribute more strongly to peak flow and require greater attention in pressure inference. Panel (f) shows that the active-edge ratio remains low and relatively stable, while total traversals vary more sharply over time, indicating that daily demand is concentrated on a limited subset of repeatedly used segments. Overall, Fig. 6 shows that operation-maintenance pressure must be inferred from the joint effects of road class, geometry, camera observability, directional loading, and temporally concentrated trajectory demand.

Fig. 7 shows how raw trajectories are transformed into operational evidence with clear temporal and spatial heterogeneity. Panel (a) indicates that $\log(1+F_{e,\tau})$ is highly concentrated on a limited subset of sampled edges, while many segments remain weakly active or inactive across most time windows. This confirms that the trajectory-derived flow matrix $\mathbf{F}$ is sparse and uneven rather than uniformly distributed across the network. Panel (b) further shows that flow-weighted traversal-time burden and flow-weighted reliability are only partially synchronized. Traversal burden remains relatively stable during the main activity period but rises near the end of the day, while reliability also increases at boundary windows. During most central time windows, the flow-weighted traversal-time burden remains approximately 47–54, before increasing sharply to about 82 near the end of the observation period. This suggests that high observation confidence does not necessarily imply low operational burden.

Panel (c) highlights the temporal instability of active road usage. Active-edge turnover captures the replacement of active segments between adjacent time windows, while novelty pressure-load share measures the contribution of newly active segments to the current pressure-weighted load. Their peaks in the early transition window and near the end of the day indicate that pressure formation is influenced not only by persistent flow corridors, but also by abrupt changes in the active segment set. Panel (d) provides a micro-level explanation: most observed states are concentrated in the low-speed and low-to-moderate irregularity region, but a visible spread toward higher irregularity shows that slow movement and unstable traversal behavior jointly contribute to pressure evidence.

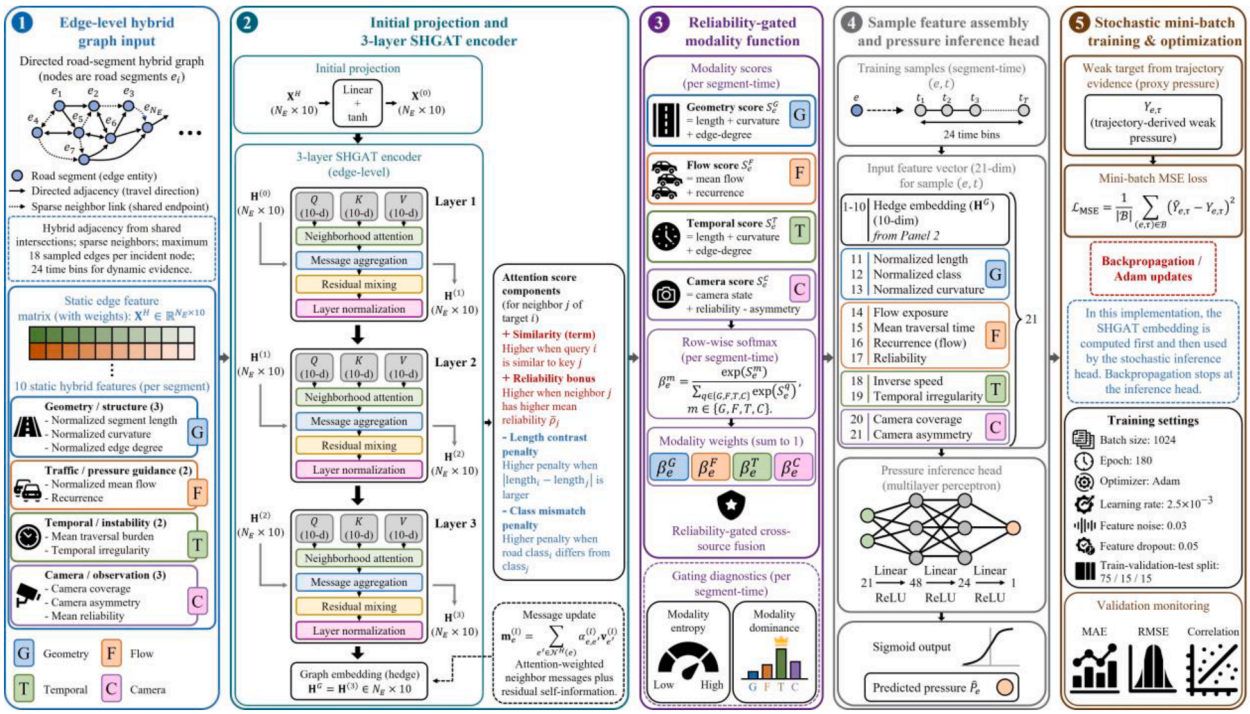


Fig. 5. Structure of the reliability-gated SHGAT for road operation-maintenance pressure inference.

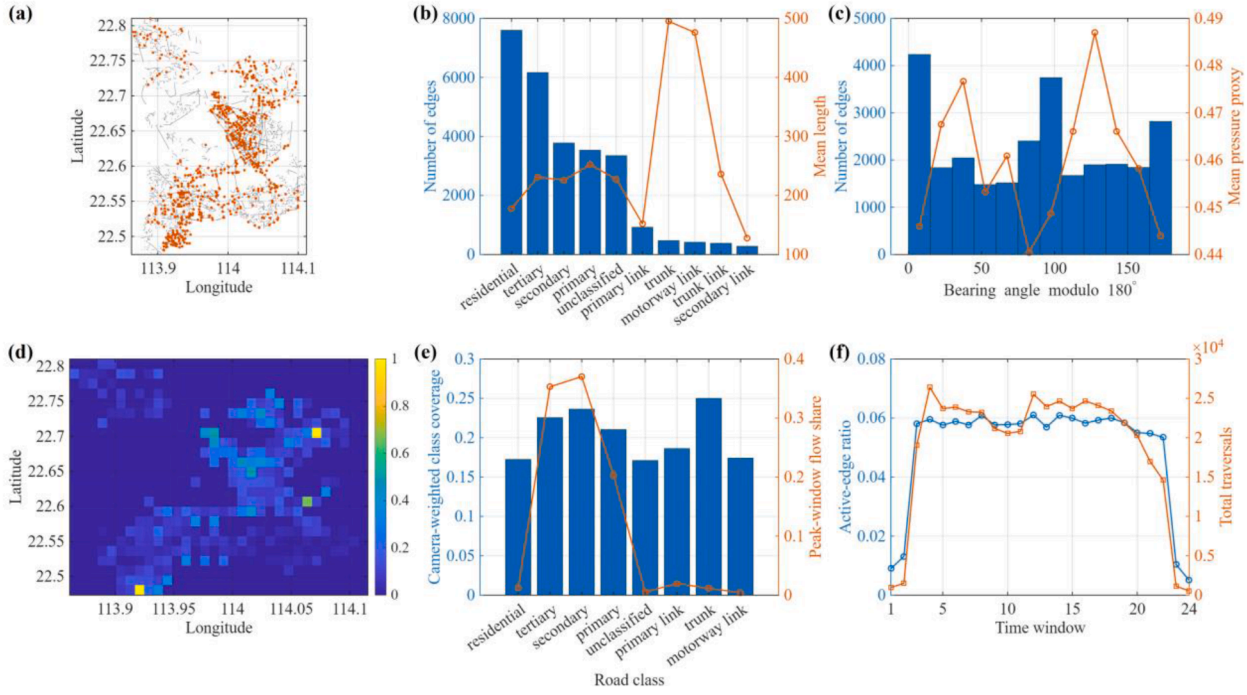


Fig. 6. Road-network structure, camera observability, and temporal demand coverage: (a) directed road-network graph $\mathscr{G} = (\mathscr{V}, \mathscr{E})$, where intersections are represented by $v_i \in \mathscr{V}$ and directed road segments by $e = (o_e, d_e) \in \mathscr{E}$, with camera-covered nodes highlighted, (b) road-class distribution and mean segment length ℓ_e, (c) directional anisotropy of road loading and mean weak pressure $\overline{Y}_e$, (d) spatial camera-observability field derived from node-level camera availability, (e) class-level camera-weighted coverage and peak-window flow share, (f) active-edge ratio and total traversals over the temporal window set $\mathscr{T}$.

Panels (e) and (f) connect these temporal patterns to spatial and observational regimes. The mean temporal irregularity $\overline{\Omega}_e$ is spatially localized, with high-irregularity segments clustering in central and corridor-like areas, which supports propagating temporal instability through the hybrid segment graph rather than treating it as independent noise. More importantly, the exceedance curves in panel (f) show that camera-sparse segments retain higher exceedance probability over a broad intermediate pressure-threshold range. At a pressure threshold of approximately 0.40, the exceedance probability remains above 0.9 for camera-sparse segments, compared with only about 0.27 for camera-covered segments. This indicates that camera sparsity is not only a data limitation, but is also related to hidden pressure exposure requiring

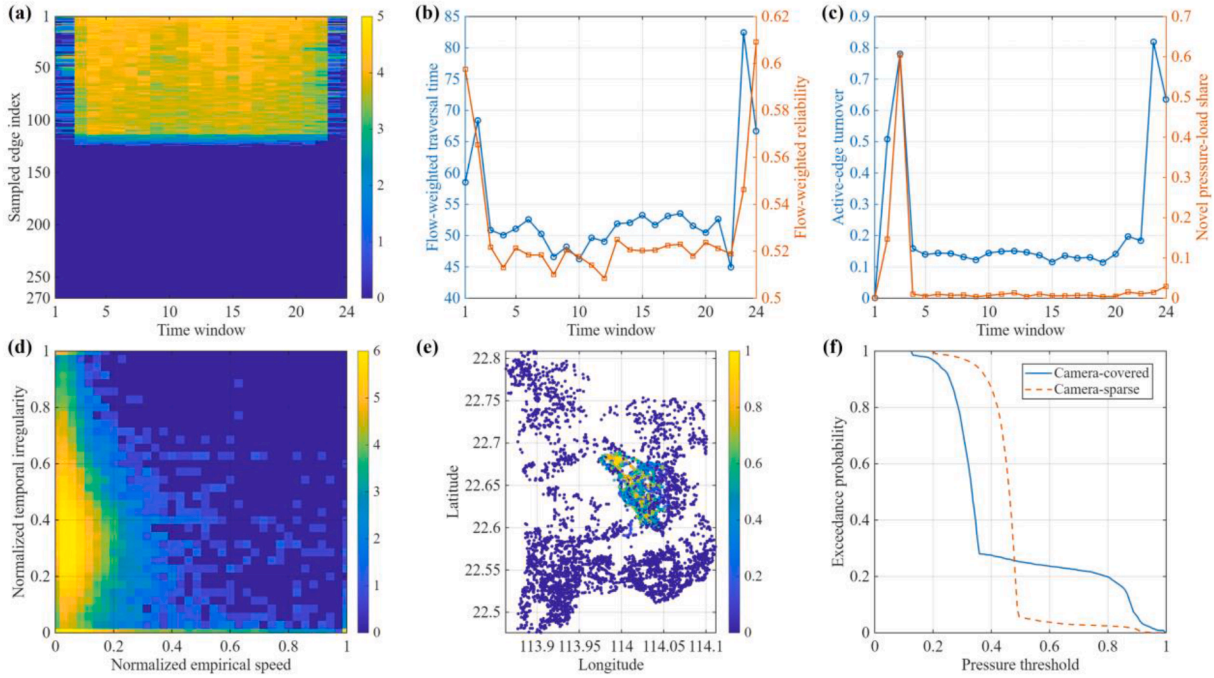


Fig. 7. Trajectory-derived operational evidence and pressure heterogeneity: (a) logarithmic segment-time flow evidence $\log(1+F_{e,\tau})$ sampled from the trajectory-derived flow matrix F, (b) flow-weighted traversal-time burden $\overline{T}_{e,\tau}$ and observation reliability $\rho_{e,\tau}$, (c) active-edge turnover and novelty pressure-load share over the temporal window set $\mathscr{T}$, (d) joint density of normalized empirical speed $\widetilde{U}_{e,\tau}$ and normalized temporal irregularity $\widetilde{\Omega}_{e,\tau}$, (e) spatial distribution of mean temporal irregularity $\overline{\Omega}_e$, (f) exceedance probability of mean weak pressure $\overline{Y}_e$ for camera-covered and camera-sparse road segments.

reliability-aware fusion. Overall, Fig. 7 demonstrates that the weak pressure field is supported by complementary trajectory-derived signals, including concentrated flow, traversal burden, active-set turnover, temporal irregularity, and camera-dependent pressure heterogeneity.

4.2. *Predictive reliability, source contribution, and diagnostic validity*

This subsection evaluates the internal behavior and predictive credibility of the proposed SHGAT framework. The analysis focuses on three connected aspects: how hybrid graph attention captures

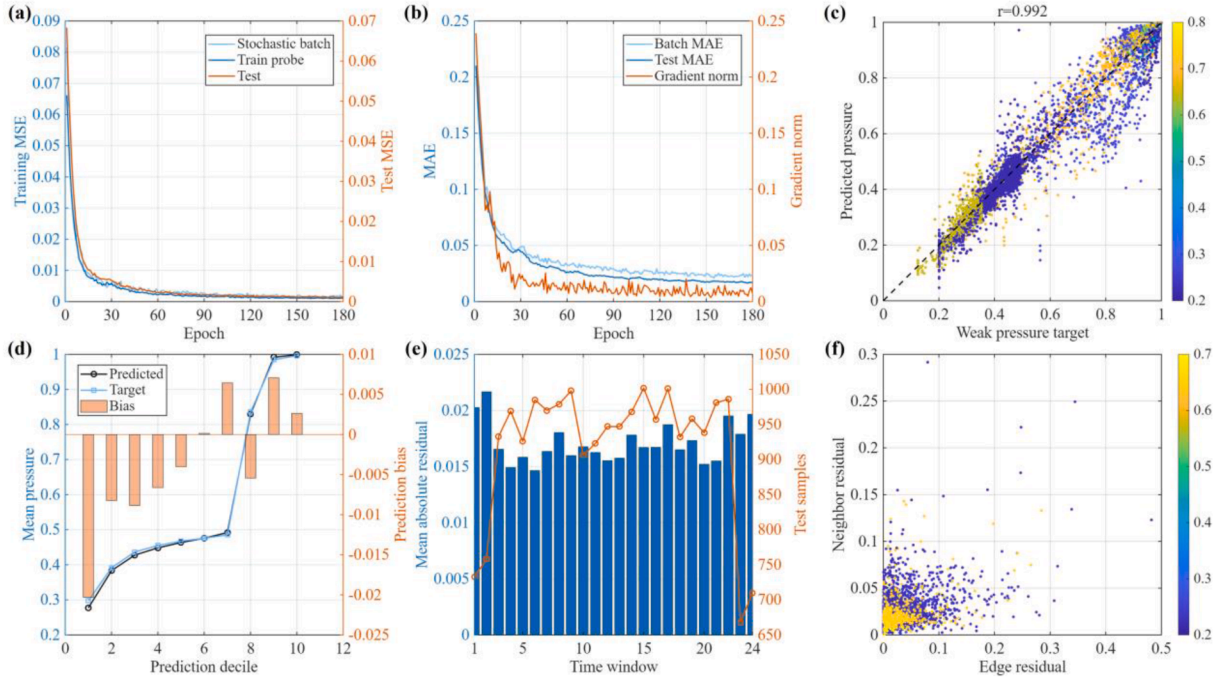


Fig. 8. Mini-batch stochastic learning and pressure-inference performance: (a) stochastic mini-batch, train-probe, and test MSE across training epochs, (b) batch MAE, test MAE, and gradient norm, (c) agreement between predicted pressure $\widehat{Y}_{e,\tau}$ and weak pressure target $Y_{e,\tau}$, with samples colored by observation reliability $\rho_{e,\tau}$, (d) decile-wise pressure calibration and prediction bias, (e) temporal distribution of test residual exposure over the temporal window set $\mathscr{T}$, (f) residual spillover over the hybrid adjacency A^H, comparing edge-level residuals with neighboring residuals.

neighborhood propagation, how reliability-gated modality fusion allocates information across geometry, flow, temporal, and camera channels, and how the mini-batch stochastic pressure-inference head learns the weak pressure target. In addition to reporting training and test performance, this subsection further examines calibration, residual spillover, ordinal pressure classification, source ablation, and uncertainty-residual alignment. These diagnostics provide further insight into how the model combines information across the graph and different sources, while also revealing its error characteristics.

Fig. 8 evaluates whether the proposed SHGAT-based pressure inference head can stably learn the weak pressure field from stochastic mini-batches. Panels (a) and (b) show that both MSE and MAE decrease rapidly in the early epochs and then enter a stable convergence phase. The test MSE decreases from approximately 0.07 initially to below 0.01 within about 30 epochs, while the test MAE subsequently stabilizes at approximately 0.02. Here, MSE and MAE measure the squared and absolute differences between the predicted pressure $\widehat{Y}_{e,\tau}$ and the weak target $Y_{e,\tau}$, respectively. The test curves remain close to the train-probe curves, indicating that stochastic mini-batch perturbations support a stable mapping from fused features to pressure estimates rather than overfitting the sampled training points. The declining gradient norm further suggests that the optimizer reaches a numerically stable region after the initial update stage.

Panel (c) provides direct evidence of predictive agreement. The predicted pressure $\widehat{Y}_{e,\tau}$ aligns closely with the weak pressure target $Y_{e,\tau}$, with a high correlation of $r = 0.992$. The scatter remains near the diagonal across most pressure levels, showing that the model captures both moderate- and high-pressure states. The color distribution also indicates that samples with different observation reliabilities $\rho_{e,\tau}$ are not collapsed into a single regime, suggesting that reliability contributes to the learned pressure structure. Panel (d) confirms this from a calibration perspective, as the predicted and target decile means nearly overlap, with only small bias in the lowest and transition deciles. Across the deciles, the prediction bias remains within approximately ± 0.02, with most groups exhibiting substantially smaller deviations. This indicates that the model preserves the ordinal pressure structure needed for maintenance prioritization.

Panels (e) and (f) examine residual exposure in time and graph space. Temporal residuals remain low and relatively stable across most windows, despite variations in the number of test samples, suggesting that the learned pressure function is not dominated by a specific temporal period. In panel (f), edge residuals are compared with hybrid-neighborhood residuals defined through $\mathbf{A}^H$. Most points are concentrated near the lower-left region, indicating that large edge errors and large neighboring errors rarely occur simultaneously. This weak residual spillover suggests that the hybrid attention structure does not introduce strong spatially coherent error amplification. Overall, Fig. 8 shows that the mini-batch stochastic learning procedure produces an accurate, well-calibrated, and graph-stable pressure field.

Fig. 9 examines the internal behavior of the proposed SHGAT fusion mechanism rather than only its final prediction accuracy. Panel (a) shows that the hybrid-neighbor count D_e^H is not simply determined by flow exposure. Many segments with moderate or high mean flow have only limited hybrid neighbors, while some low-flow segments appear in more connected local structures. Across most of the observed flow range, the hybrid-neighbor count is concentrated at approximately 4–13, although a small number of structurally exceptional segments reach values above 20. This confirms that the hybrid segment graph $\mathscr{G}^H$ provides structural information that cannot be replaced by traffic volume alone. In panel (b), the incoming attention mass M_e^A shows a broad relationship with the mean weak pressure $\overline{Y}_e$. High-pressure segments do not always receive the largest incoming attention, but larger attention accumulation tends to occur in structurally richer neighborhoods. This suggests that the attention mechanism captures both local propagation and structural exposure instead of merely amplifying high-flow links.

Panels (c) and (d) clarify the reliability-gated multi-source fusion behavior. In panel (c), modality entropy E_e^M decreases as modality

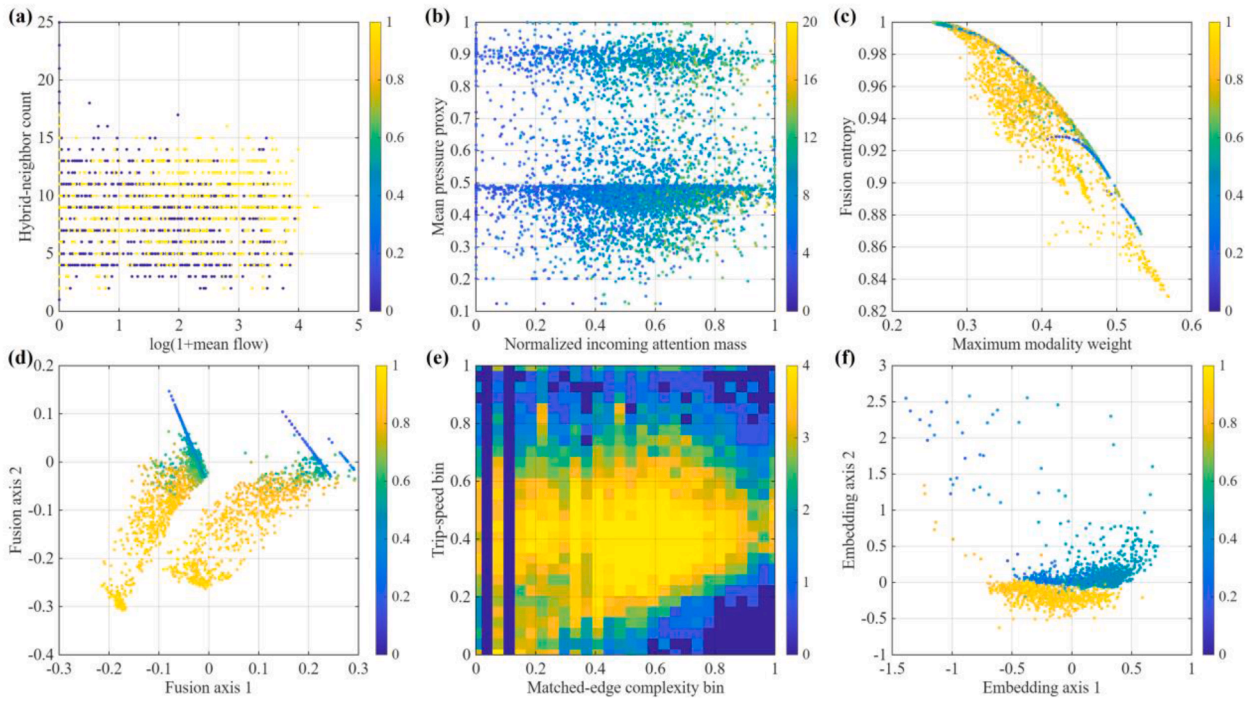


Fig. 9. Hybrid graph attention and reliability-gated multi-source fusion behavior: (a) relationship between flow exposure and hybrid-neighbor count D_e^H in the hybrid segment graph $\mathscr{G}^H$, (b) incoming attention mass M_e^A and mean weak pressure $\overline{Y}_e$, (c) modality dominance D_e^M, modality entropy E_e^M, and pressure tendency, (d) low-dimensional manifold of modality-weight deviations derived from the modality-weight vector β_e, (e) trajectory efficiency-complexity density based on matched-edge complexity and trip-speed information, (f) hybrid graph embedding H^G projected onto two embedding axes and colored by $\overline{Y}_e$.

dominance D_e^M increases, which is consistent with the stronger control of one source channel when the maximum modality weight becomes larger. Most points remain in a relatively high-entropy range, indicating that the model generally integrates geometry, flow, temporal, and camera evidence in a balanced manner. For maximum modality weights below approximately 0.50, the fusion entropy is predominantly above 0.90, whereas it decreases toward about 0.83 as the dominant modality weight approaches 0.57. However, higher-pressure segments are more often associated with stronger modality dominance, suggesting that local pressure may become source-specific when one evidence channel is particularly informative. Panel (d) further shows that the centered modality-weight vector $\boldsymbol{\beta}_e$ forms several branch-like source-allocation regimes rather than a random cloud, providing interpretability for different pressure levels.

Panels (e) and (f) connect fusion behavior with trajectory and embedding structure. Panel (e) shows that most trips are concentrated around intermediate matched-edge complexity and trip-speed states, supporting the stability of trajectory-derived operational evidence. Panel (f) shows that the hybrid graph embedding $\mathbf{H}^G$ forms a compact dominant manifold with several sparse outlying regions. The pressure-related color gradient indicates that SHGAT learns an embedding space in which segments with different pressure levels occupy distinguishable regions. Overall, Fig. 9 demonstrates that the proposed model separates topology from flow, balances multi-source evidence through reliability-gated modality weights, and produces graph embeddings with interpretable pressure information.

Fig. 10 evaluates whether the continuous pressure prediction $\widehat{Y}_{e,\tau}$ can support reliable ordinal and high-risk decisions. Panel (a) shows clear diagonal dominance in the normalized confusion matrix. The low- and high-pressure classes are identified with high consistency, with diagonal ratios around 0.93 and 0.90, respectively, while the medium-pressure class has a lower ratio near 0.71 and more leakage to adjacent classes. This is expected because medium-pressure samples lie near the transition region of the weak pressure field, where trajectory burden, reliability, and temporal irregularity do not form a sharply separated boundary. Thus, the main classification uncertainty is concentrated around ordinal transitions rather than occurring randomly.

Panel (b) further assesses high-risk detection under different pressure thresholds. MCC reflects the balance among true positives, false positives, false negatives, and true negatives, while balanced accuracy averages the true-positive and true-negative rates. Both indicators increase rapidly when the high-risk threshold enters the main separation range, indicating that the model can distinguish high-pressure states once the threshold is not overly permissive. Around a high-risk threshold of 0.50, MCC reaches approximately 0.75 and balanced accuracy approximately 0.94, after which both metrics remain at relatively high levels over most of the upper-threshold range. The high balanced-accuracy plateau suggests that the predicted pressure field preserves the ranking structure of high-risk segments, although extremely high thresholds become less informative because only a few samples remain classified as high risk.

Panels (c)-(f) reveal where the remaining errors occur. The directed transition graph in panel (c) shows that misclassifications mainly occur between adjacent ordinal classes, especially low-to-medium and medium-to-high transitions, rather than through severe low-to-high or high-to-low jumps. Among the directed misclassification transitions, medium-to-low and medium-to-high account for 38.1% and 25.2%, respectively, whereas the low-to-high transition accounts for only 1.8%. Panel (d) shows that the error rate decreases as the prediction margin increases, meaning that decisions farther from the ambiguous central region are more reliable. Panel (e) indicates small but temporally nonuniform prediction bias, with slightly negative bias more often appearing in low- or medium-pressure classes. Panel (f) maps the spatial contrast between false-positive (FP) and false-negative (FN) high-pressure decisions, showing that residual errors are localized rather than uniformly scattered. Overall, Fig. 10 supports the use of the inferred pressure field for ordinal maintenance prioritization, while indicating that medium-pressure and spatially localized transition areas require more cautious interpretation.

Fig. 11 evaluates the contribution of each information source to

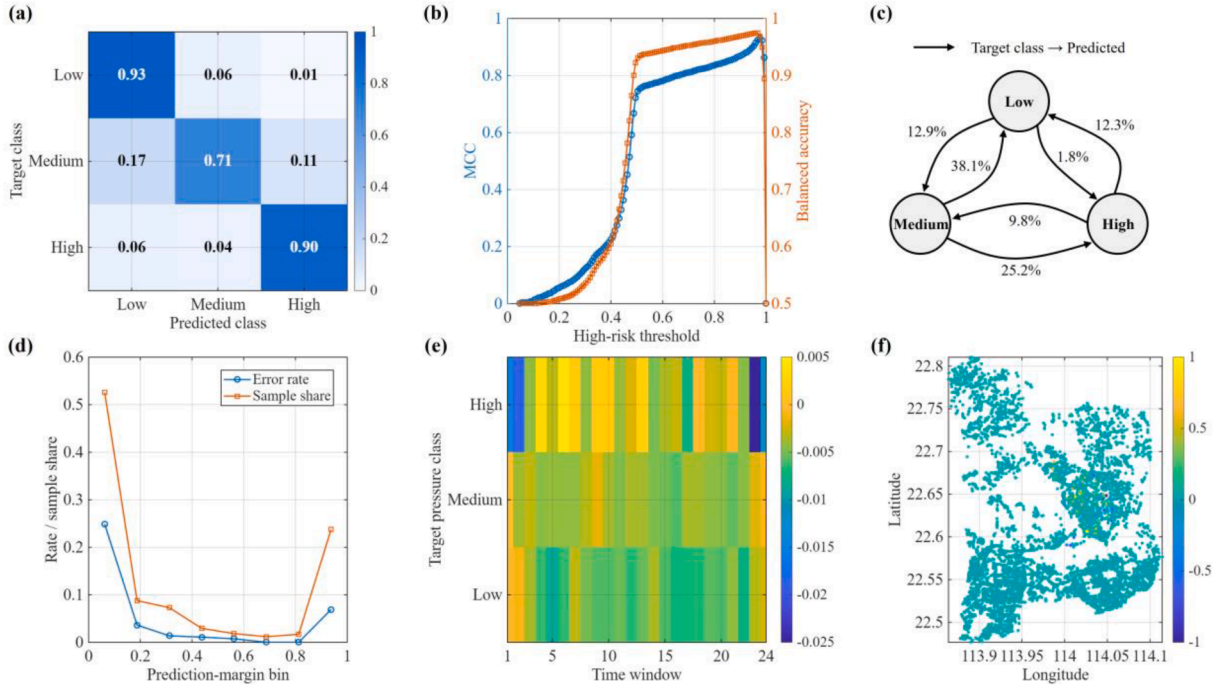


Fig. 10. Ordinal pressure-class diagnostics and spatial error contrast: (a) normalized ordinal confusion matrix for low-, medium-, and high-pressure classes derived from $Y_{e,\tau}$ and $\widehat{Y}_{e,\tau}$, (b) threshold-dependent MCC and balanced accuracy for high-risk pressure identification, (c) directed ordinal misclassification transition graph from target class to predicted class, (d) prediction-margin reliability of discrete pressure decisions, (e) temporal prediction-bias field by target pressure class, (f) spatial contrast between false-positive (FP) and false-negative (FN) high-pressure decisions.

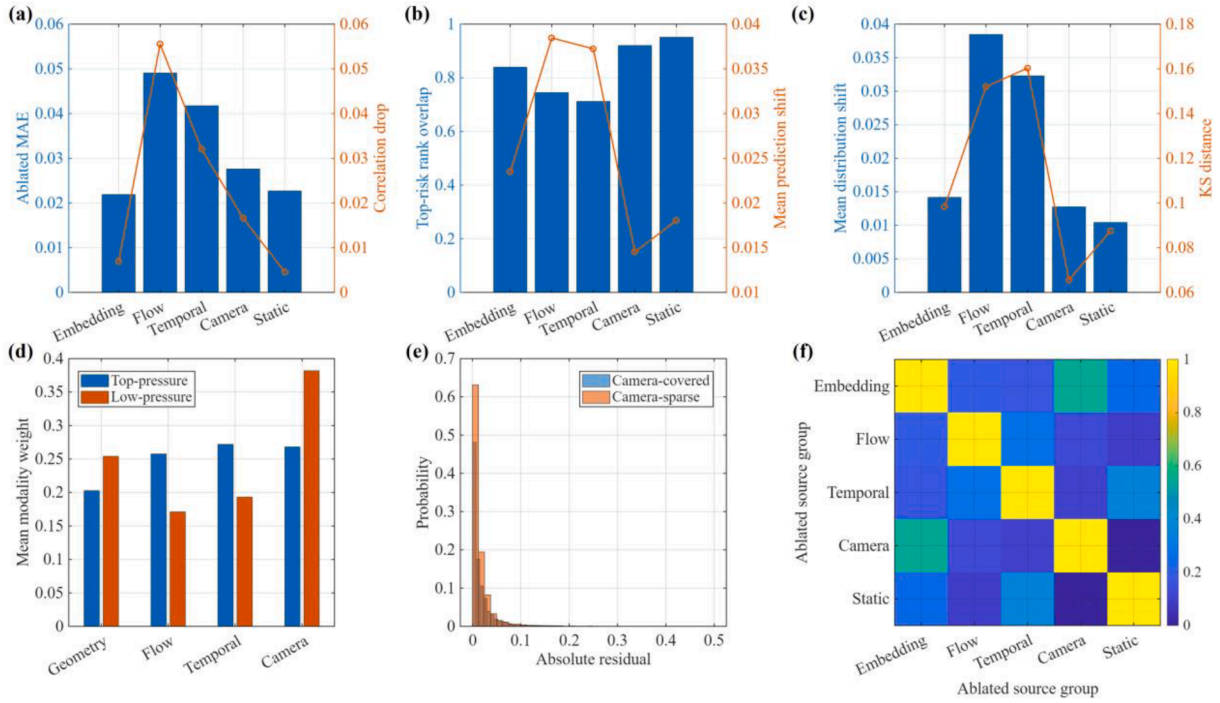


Fig. 11. Source ablation and cross-source contribution diagnostics: (a) MAE and correlation degradation after removing embedding, flow, temporal, camera, or static feature groups, (b) top-risk rank overlap and mean prediction shift after source ablation, (c) distributional distortion and KS distance after source removal, (d) mean modality weights β_e for high-pressure and low-pressure cohorts over geometry, flow, temporal, and camera channels, (e) residual distributions under camera-covered and camera-sparse observation regimes, (f) cross-source ablation-shift correlation matrix.

pressure inference by systematically removing feature groups from the trained model. In the ablation test, the selected source group is replaced by its empirical median value, and the resulting predictions are compared with the baseline predictions. Panels (a)-(c) show that the flow and temporal groups cause the strongest degradation. Removing the flow group produces the largest MAE increase and correlation drop, while removing the temporal group also leads to clear performance loss. Specifically, flow ablation increases MAE by approximately 0.049 and reduces correlation by about 0.055, compared with corresponding changes of approximately 0.042 and 0.032 under temporal ablation. This indicates that the inferred pressure field is mainly governed by trajectory-derived operational evidence, including flow exposure, recurrence, traversal burden, and irregularity. By contrast, the camera and static groups cause smaller pointwise degradation, suggesting that they mainly provide reliability conditioning and structural regularization rather than dominant direct prediction signals.

Panel (b) adds a ranking perspective. Although ablation causes visible prediction shifts, the top-risk rank overlap remains relatively high for most source removals, especially for camera and static features. The top-risk overlap remains approximately 0.92 after camera ablation and 0.95 after static-feature ablation, whereas it decreases to about 0.74 and 0.71 after removing flow and temporal information, respectively. This means that the highest-pressure ordering is not fully destabilized by removing a single source group. However, the larger shifts under flow and temporal ablation confirm their importance for pressure magnitude calibration. Panel (c) further shows that flow and temporal removal produce stronger Kolmogorov-Smirnov distances, indicating that these two channels reshape the overall pressure distribution rather than only affecting isolated samples.

Panels (d)-(f) provide complementary evidence from fusion weights and residual behavior. The modality-weight vector $\boldsymbol{\beta}_e$ shows that high-pressure segments assign more weight to flow and temporal channels, whereas low-pressure segments rely more on camera-related information. This is consistent with the interpretation that high-pressure inference depends on dynamic service burden, while low-pressure states are more closely related to observation and coverage conditions. Panel (e) shows that both camera-covered and camera-sparse residuals are concentrated near low values, although camera-sparse samples exhibit slightly heavier residual exposure. Panel (f) further indicates that ablation-induced shifts are not independent across sources: embedding and camera shifts are moderately related, while flow, temporal, and static groups show distinct but interacting effects. Overall, Fig. 11 supports the multi-source SHGAT design, where flow and temporal evidence drive pressure magnitude, camera information modulates reliability, and graph embedding with static structure stabilizes cross-source inference.

Fig. 12 further examines whether the predictive performance of SHGAT originates from the complete reliability-informed architecture rather than from an isolated component or a particular network depth. As shown in Fig. 12 (a), the full SHGAT achieves the lowest errors, with an MAE of approximately 0.017 and an RMSE of approximately 0.033, while maintaining a correlation close to 0.992. Removing the reliability gate, modality fusion, or stochastic learning produces a pronounced deterioration in both error and correlation, whereas eliminating the hybrid graph causes a smaller but still detectable degradation. This hierarchy indicates that reliability-conditioned information allocation and stochastic regularization provide substantial contributions to the final inference quality, while hybrid neighborhood propagation supplies complementary structural information. The depth analysis in Fig. 12 (b) identifies three attention layers as the most favorable configuration, simultaneously yielding the minimum MAE and the highest correlation and composite score. Two layers provide insufficient propagation capacity, whereas increasing the depth to four or five layers leads to progressively weaker performance, suggesting that excessive propagation begins to dilute locally discriminative pressure information rather than adding useful network context.

For a broader architectural benchmark, three representative graph-learning approaches were considered under the same data and training protocol. GraphSAGE [91] represents neighborhood aggregation by combining each road segment with information summarized

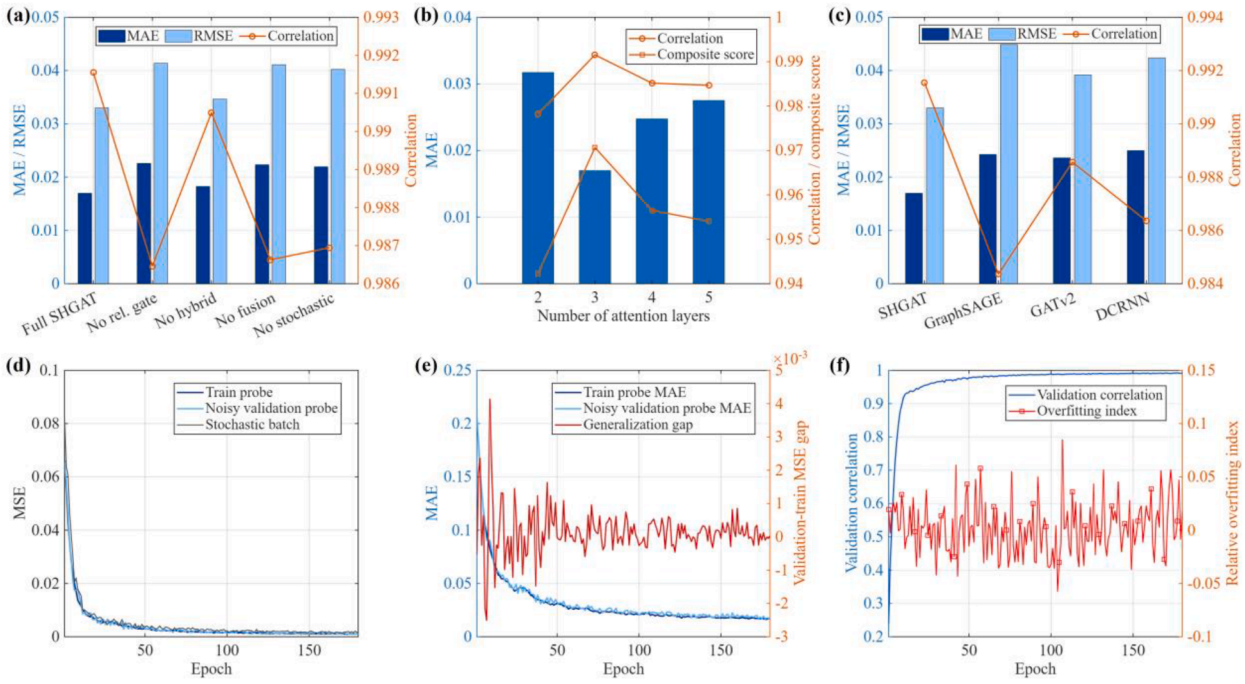


Fig. 12. Architectural contribution, attention-depth sensitivity, benchmark comparison, and generalization behavior of the proposed SHGAT: (a) component-wise architectural ablation in terms of MAE, RMSE, and correlation, (b) sensitivity of predictive performance to the number of attention layers, (c) standardized comparison with GraphSAGE-, GATv2-, and DCRNN-based graph encoders, (d) raw MSE trajectories for the training probe, perturbed validation probe, and stochastic mini-batches, (e) raw MAE evolution and validation–training generalization gap, (f) evolution of validation correlation and the relative overfitting index.

from its adjacent segments, whereas GATv2 [92] introduces dynamic attention to allow neighbor importance to depend more flexibly on the interacting graph entities. DCRNN-based [93] encoding provides a diffusion-oriented alternative that propagates information through directed graph transitions and is particularly relevant to transportation-network modeling. Against these representative aggregation, attention, and diffusion mechanisms, Fig. 12 (c) shows that SHGAT provides the best overall predictive accuracy, with consistently lower MAE and RMSE and higher correlation. GATv2 is the closest competing architecture, whereas GraphSAGE and DCRNN exhibit more

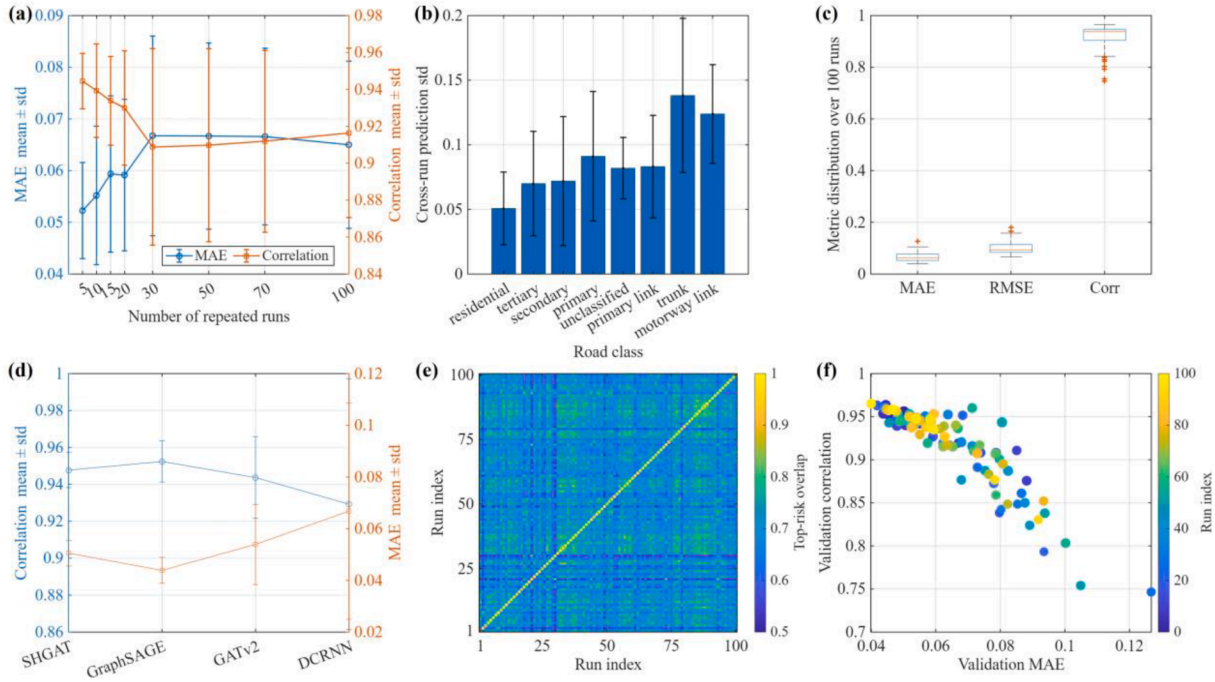


Fig. 13. Cross-run reliability and stochastic stability of SHGAT inference: (a) cumulative MAE and correlation statistics with increasing numbers of repeated runs, (b) road-class-specific dispersion of peak-window predictions across runs, (c) distributions of MAE, RMSE, and correlation over 100 independent runs, (d) repeated-run stability of SHGAT and alternative graph encoders, (e) pairwise top-risk overlap among 100 SHGAT runs, (f) run-wise relationship between validation MAE, correlation, and top-risk consistency.

evident performance losses. This comparison indicates that the observed improvement cannot be attributed solely to generic neighborhood aggregation, dynamic graph attention, or diffusion-based propagation. Instead, explicitly incorporating observation reliability, structural compatibility, and reliability-gated information fusion provides additional benefit for representing heterogeneous and unevenly observed road-network evidence.

The optimization behavior in Fig. 12 (d)–(f) further supports the stability of this architectural advantage. Both MSE and MAE decrease rapidly during the early training stage and subsequently enter a gradual convergence regime, while the perturbed validation probe follows the training trajectory without developing persistent divergence. The validation–training MSE gap remains close to zero throughout most epochs despite local stochastic fluctuations. Meanwhile, the validation correlation rises rapidly above 0.9 and approaches approximately 0.99 as training proceeds, whereas the relative overfitting index fluctuates around zero without a systematic upward trend. These results indicate that the selected architecture reaches a stable high-correlation regime without progressive train–validation separation, supporting the adopted three-layer configuration and stochastic training strategy.

Fig. 13 evaluates whether the inferred pressure field remains statistically reliable under repeated stochastic training. Panel (a) shows that the cumulative performance estimates vary noticeably when only a small number of runs are considered, but become substantially more stable as the repetition count increases. The mean MAE gradually settles around 0.06–0.07 after approximately 30–50 runs, while the mean correlation stabilizes slightly above 0.90. The persistent error bars indicate that stochastic initialization, mini-batch sampling, dropout, and feature perturbation produce non-negligible run-to-run variability, making a single realization insufficient for characterizing model reliability. Panel (b) further reveals that this variability is structurally heterogeneous across the road network. Residential and tertiary roads exhibit comparatively small prediction dispersion, whereas trunk and motorway-link segments show markedly larger cross-run standard deviations. This suggests that stochastic uncertainty is amplified on road classes with stronger network influence or more heterogeneous operational states, rather than being uniformly distributed across segment types.

The distributional results in Fig. 13 (c) reinforce this interpretation. Most of the 100 runs are concentrated within relatively narrow MAE and RMSE ranges, while correlation remains predominantly above 0.90; only a limited number of runs produce substantially weaker agreement. Panel (d) shows that the repeated-run behavior also depends on the graph encoder. SHGAT remains competitive with GraphSAGE and exhibits a favorable balance between predictive correlation and error, whereas GATv2 and particularly DCRNN show larger mean errors under the repeated-run protocol. GraphSAGE displays a slightly higher mean correlation and lower mean MAE in this specific repeated-run experiment, indicating that stochastic stability and best single-configuration predictive accuracy are not necessarily identical criteria. This complements the standardized comparison in Fig. 12 (c), where SHGAT achieves the strongest overall predictive performance under the main evaluation protocol.

Panels (e) and (f) examine whether stochastic variation materially changes the segments identified as high risk. The pairwise overlap matrix in Fig. 13 (e) is dominated by moderate-to-high overlap values, with most off-diagonal entries remaining well above 0.6, indicating that the top-risk set is considerably more stable than the raw prediction errors alone would suggest. Several block-like patterns nevertheless reveal groups of runs with stronger mutual consistency, reflecting sensitivity to particular stochastic realizations. Panel (f) shows a clear inverse relationship between validation MAE and correlation: most runs cluster in the high-correlation, low-error region, whereas a small number of unfavorable runs exhibit both increased MAE and reduced correlation. The associated marker sizes further indicate that runs with stronger predictive agreement generally preserve a larger fraction of the reference top-risk set. Collectively, these results show that the SHGAT framework maintains broadly consistent maintenance-priority identification across stochastic realizations, while also quantifying the residual variability that should be retained in reliability-oriented decision making.

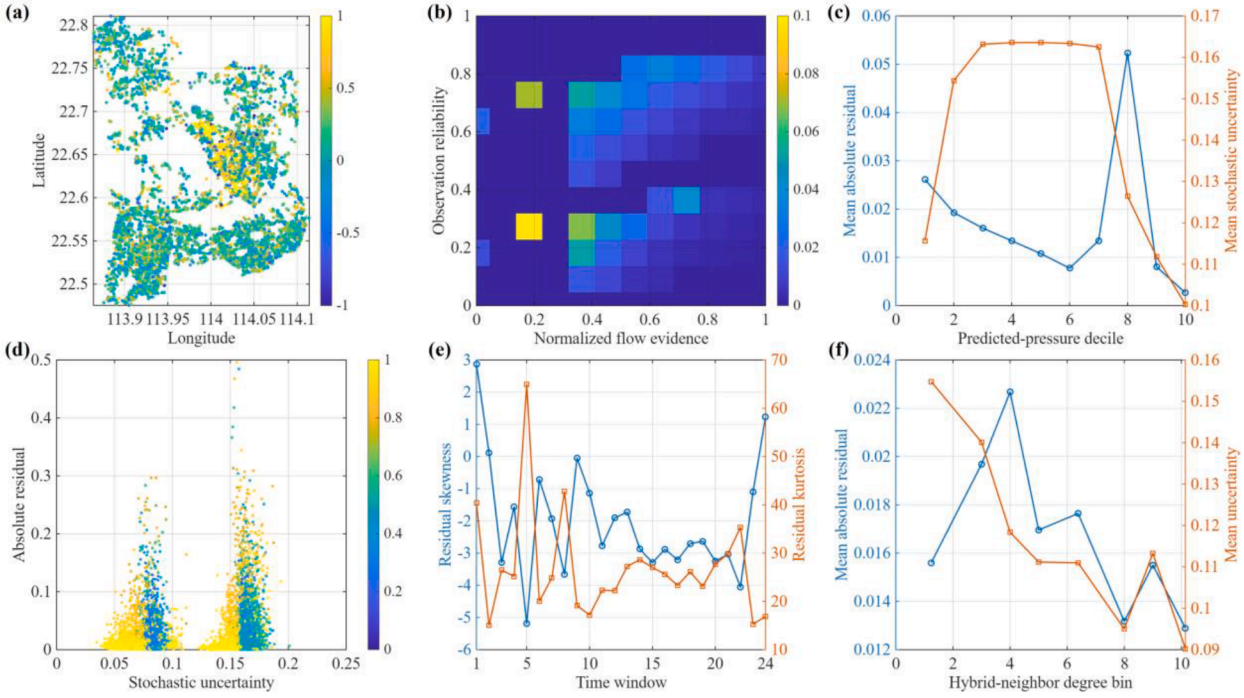


Fig. 14. Residual topology, uncertainty alignment, and error mechanisms: (a) spatial residual sign-coherence field over the hybrid segment graph, (b) mean residual over normalized flow evidence $\widetilde{F}_{e,\tau}$ and observation reliability $\rho_{e,\tau}$, (c) mean absolute residual and stochastic uncertainty by predicted-pressure decile, (d) sample-level alignment between stochastic uncertainty and absolute residual, (e) skewness and kurtosis of temporal residuals over the temporal window set $\mathcal{T}$, (f) error and uncertainty variation over hybrid-neighbor degree D_e^H.

Fig. 14 provides a diagnostic view of where SHGAT inference errors arise and whether they are spatially, probabilistically, or structurally organized. Panel (a) maps the residual sign-coherence field over the hybrid segment graph, measuring the agreement between the residual sign of a segment and that of its hybrid neighbors. The map shows mixed positive and negative coherence rather than one large coherent region, suggesting that the model does not produce a systematic spatial bias over the whole study area. However, several local clusters remain visible around dense central corridors, indicating that local topology and traffic concentration still influence residual organization.

Panel (b) links residual magnitude to normalized flow evidence $\widetilde{F}_{e,\tau}$ and observation reliability $\rho_{e,\tau}$. Large residuals are not uniformly distributed across the flow-reliability plane, but mainly appear in intermediate states where flow evidence and reliability are partially mismatched. This indicates that high flow alone does not guarantee low error, and high reliability does not remove ambiguity when the operational state is dynamically complex. Panel (c) further shows that stochastic uncertainty and residuals vary differently across predicted-pressure deciles. Uncertainty remains relatively high in the middle deciles, while the mean absolute residual is generally small except for a local increase in the upper pressure range. Across approximately the third to seventh deciles, mean stochastic uncertainty remains near 0.16 while the mean absolute residual is mostly below 0.02, before a localized residual peak of about 0.052 appears at the eighth decile. Thus, the uncertainty estimate is not merely a copy of the residual, but reflects sensitivity to stochastic feature perturbations.

Panels (d)-(f) further examine uncertainty, residual shape, and graph degree. In panel (d), most samples are concentrated in low-residual regions, while larger residuals tend to occur under moderate stochastic uncertainty, indicating that some errors arise from structural or regime-transition effects rather than input noise alone. Panel (e) shows that temporal residual distributions are non-Gaussian and vary over the temporal window set $\mathcal{T}$, with changing skewness and kurtosis. Panel (f) indicates that hybrid-neighbor degree D_e^H also affects error and uncertainty: higher connectivity does not always reduce residuals, but uncertainty generally decreases from sparse to intermediate degree regimes before slightly fluctuating. Mean uncertainty decreases from approximately 0.15 in the sparsest degree bin to about 0.09 in the highest-connectivity bin, while mean absolute residual remains within a comparatively narrow range of approximately 0.013–0.023. Overall, Fig. 14 shows that the remaining errors are governed by local graph topology, flow-reliability mismatch, stochastic feature sensitivity, and temporal regime variation, supporting the need for reliability-aware diagnostics beyond standard accuracy metrics.

4.3. Robustness, counterfactual resilience, and maintenance decision support

This subsection shifts the analysis from model fitting to operational usefulness. Since road operation-maintenance applications require stable decisions under incomplete, noisy, and unevenly reliable observations, the proposed framework is examined under source degradation, stochastic perturbation, and counterfactual intervention scenarios. The results evaluate how pressure predictions respond to camera dropout, trajectory degradation, flow noise, geometry noise, demand scaling, capacity variation, and sensor-related changes. The subsection further connects uncertainty, source-channel sensitivity, propagation gap, and intervention value to maintenance-priority mapping. In this way, the inferred pressure field is translated into decision-oriented evidence for identifying high-risk segments, evaluating intervention benefits, and designing reliability-aware maintenance portfolios.

Fig. 15 evaluates whether the learned pressure field responds reasonably under controlled counterfactual changes in sensing, demand, and capacity conditions. Panel (a) shows that the network-level mean pressure mainly increases with the demand-scaling factor, while camera-retention variation produces a weaker but visible effect. This indicates that the predicted operation-maintenance pressure is primarily demand-sensitive, but its magnitude is also conditioned by retained observation reliability. In practical terms, higher demand under incomplete camera retention leads to a higher inferred burden, which is

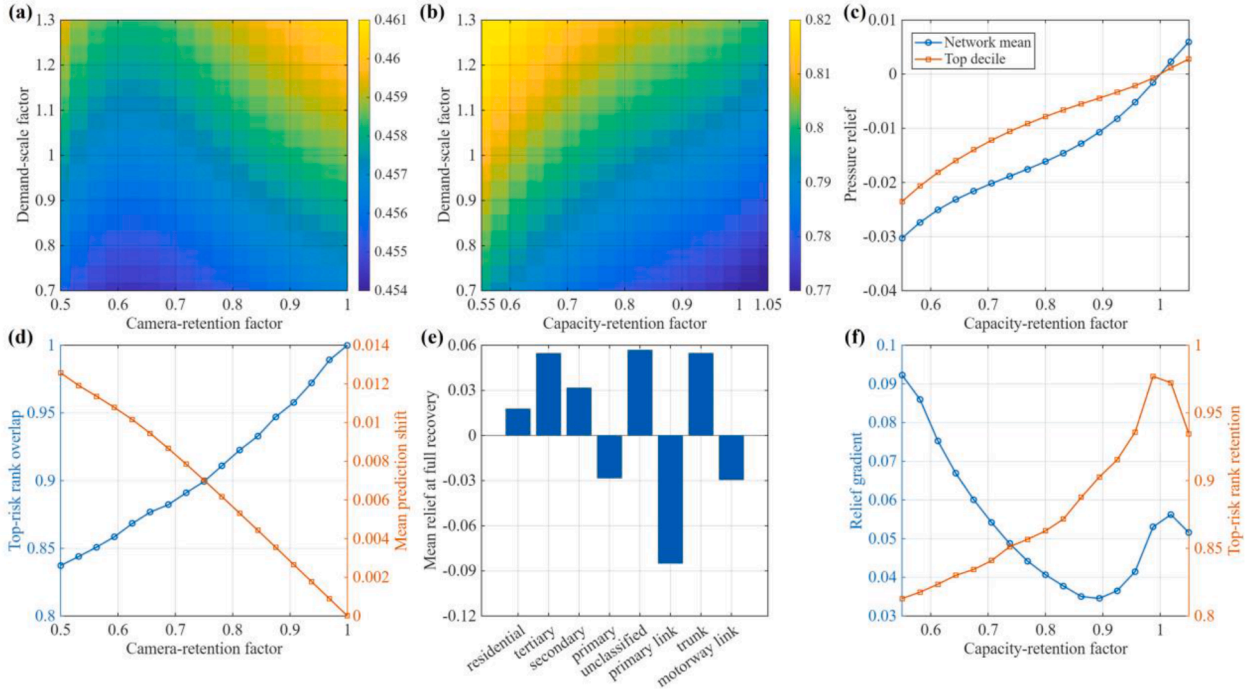


Fig. 15. Counterfactual pressure response under camera, demand, and capacity variation: (a) counterfactual mean pressure under camera-retention and demand-scaling scenarios, (b) top-decile predicted pressure under capacity-retention and demand-scaling scenarios, (c) marginal pressure relief under capacity recovery, (d) top-risk rank overlap and mean prediction shift under camera-retention variation, (e) road-class-specific recovery benefit, (f) capacity-retention relief gradient and top-risk rank stability.

consistent with the reliability-gated design of the framework.

Panel (b) focuses on the top-decile predicted pressure under capacity-retention and demand-scaling scenarios. Compared with the network mean, the high-pressure subset is more sensitive to capacity and demand variation. Lower capacity retention combined with higher demand scaling clearly elevates the top-decile pressure, which is important because maintenance planning is usually driven by the most stressed road segments rather than by the network average. Panel (c) further confirms this distinction. Pressure relief, computed as the difference between baseline and counterfactual pressure after capacity recovery, improves with increasing capacity retention for both network mean and top-decile pressure. As the capacity-retention factor increases from approximately 0.55 to 1.05, network-mean pressure relief rises from about −0.031 to 0.006, while the top-decile relief increases from roughly −0.024 to 0.003. The smoother top-decile curve suggests that high-risk segments preserve a more stable response pattern under capacity recovery.

Panels (d)-(f) provide decision-oriented interpretation. Panel (d) shows that reducing camera retention causes only moderate changes in the top-risk set: rank overlap remains high, while the mean prediction shift decreases as camera retention approaches the baseline. Even at a camera-retention factor of 0.50, the top-risk rank overlap remains approximately 0.84 and increases progressively to 1.00 as full camera retention is restored. This suggests that the high-risk ordering is not overly fragile to camera degradation, although camera information still affects pressure magnitude. Panel (e) shows that recovery benefits are road-class-dependent, with some classes showing positive mean relief and others showing weaker or negative relief, likely reflecting differences in flow concentration, geometric role, and camera-supported reliability. Panel (f) shows that marginal relief is strongest at lower capacity-retention levels and decreases as capacity approaches the baseline, while top-risk rank retention remains high over most of the range. Overall, Fig. 15 demonstrates that SHGAT supports counterfactual reasoning by distinguishing average and top-risk pressure, identifying nonlinear capacity benefits, and preserving maintenance-priority structure under plausible sensing and demand perturbations.

Fig. 16 translates the inferred pressure field into maintenance-oriented decision evidence. Panel (a) maps the local intervention-value field, which combines source-channel sensitivity, stochastic uncertainty, and reliability-conditioned pressure information. High values are localized in specific corridors and dense central areas rather than uniformly distributed, indicating that intervention value is jointly shaped by pressure sensitivity and observational weakness. Panel (b) focuses on camera-sparse pressure exposure at the peak-demand window τ^*, where τ^* denotes the temporal window with the largest total trajectory demand. The high-pressure region is concentrated in a camera-sparse central corridor, suggesting that the framework can reveal hidden operational burden where direct camera support is limited. Within this central camera-sparse corridor, the predicted peak-window pressure frequently exceeds approximately 0.8 and locally approaches 1.0.

Panel (c) evaluates intervention efficiency through length-normalized capture curves, which measure the fraction of total pressure, stochastic uncertainty, or sensor-relief objective captured when segments are selected by priority per unit length. The pressure curve rises fastest, indicating that a relatively small selected length fraction can cover a large share of the predicted pressure burden. Selecting approximately 20% of the total road length captures about 50% of the pressure objective, compared with roughly 44% of stochastic uncertainty and 34% of the sensor-relief objective. In contrast, the uncertainty and sensor-relief curves increase more gradually, suggesting that uncertainty reduction and sensing improvement require broader spatial coverage than pressure capture alone. This distinction shows that pressure-oriented and sensing-oriented interventions may prioritize different road segments.

Panels (d)-(f) further explain how the framework supports practical

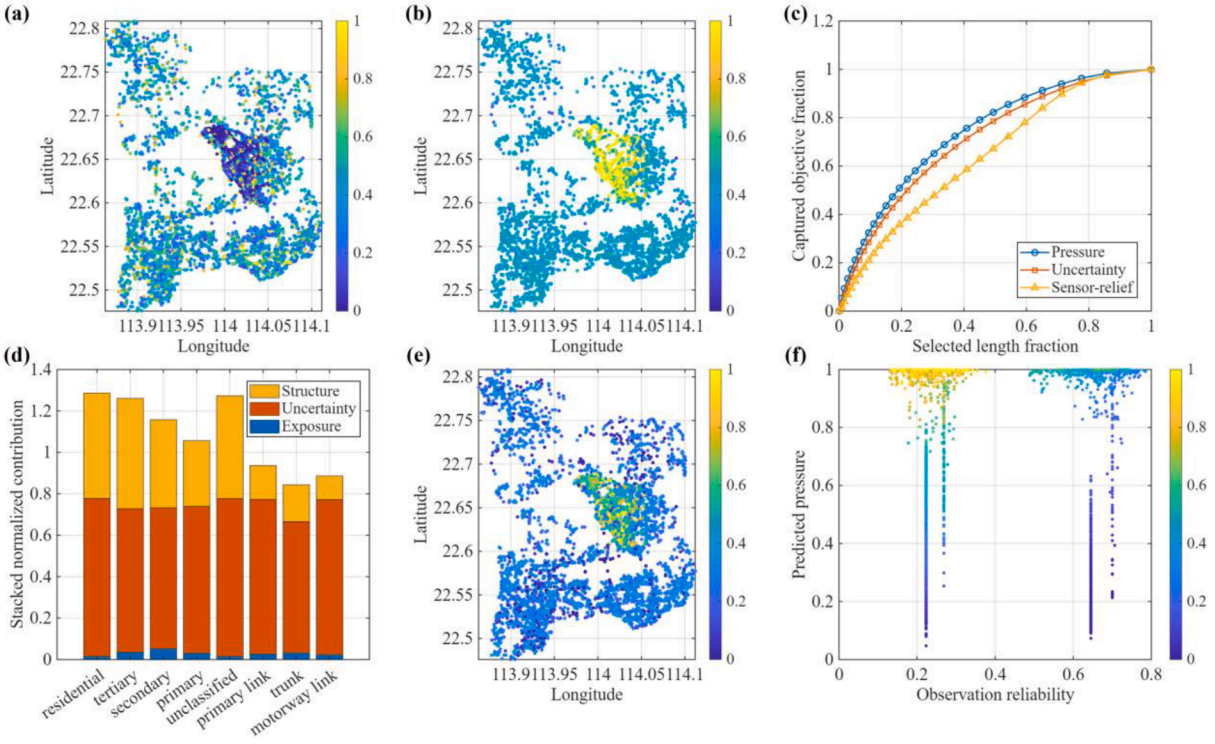


Fig. 16. Maintenance-oriented pressure priority and intervention value: (a) local intervention-value field derived from sensitivity, uncertainty, and reliability-aware pressure information, (b) camera-sparse pressure exposure represented by predicted pressure $\hat{Y}_{e,\tau^*}$ at the peak-demand window τ^*, (c) length-normalized capture curves for pressure, stochastic uncertainty, and sensor-relief objectives, (d) road-class hidden-risk decomposition by exposure, uncertainty, and structural contribution, (e) fused maintenance priority field, (f) decision quadrant between observation reliability ρ_{e,τ^*} and predicted pressure $\hat{Y}_{e,\tau^*}$.

prioritization. Panel (d) decomposes hidden risk by road class into exposure, uncertainty, and structural contribution. Uncertainty dominates for most classes, while structural contribution varies by class, indicating that maintenance priority cannot be inferred from traffic exposure alone. Panel (e) presents the fused maintenance priority field, where high-priority segments remain spatially concentrated rather than randomly scattered. Panel (f) compares predicted pressure $\widehat{Y}_{e,\tau^*}$ with observation reliability ρ_{e,τ^*}. The most critical region corresponds to high pressure and relatively low reliability, where strong operational burden coincides with insufficient observability. Overall, Fig. 16 shows that SHGAT converts multi-source fusion outputs into actionable evidence for selecting maintenance targets, identifying camera-sparse risks, and balancing pressure reduction with uncertainty-aware sensing improvement.

Fig. 17 examines the robustness of the proposed pressure-inference framework under systematic source degradation and perturbation. Panel (a) shows the prediction shift under joint camera dropout and trajectory dropout, where prediction shift denotes the mean absolute change between perturbed and baseline inferred pressure. The surface increases smoothly as both dropout probabilities grow, indicating gradual rather than abrupt stability loss. Even when both camera and trajectory dropout probabilities increase to 0.50, the mean prediction shift remains approximately 0.025, compared with about 0.010 under the mildest degradation condition. This suggests that SHGAT retains considerable resilience when observation support is partially removed, although trajectory degradation has a slightly stronger effect than camera degradation under higher perturbation levels. Panel (b) presents the MAE surface under joint flow-noise and geometry-noise perturbation. The dominant gradient follows the flow-noise direction more strongly than the geometry-noise direction, indicating that pressure inference is more sensitive to distortion in dynamic traffic evidence than to moderate corruption of static geometric descriptors.

Panels (c) and (d) further clarify the structural robustness of high-risk identification. In panel (c), high-pressure rank overlap remains relatively high over a broad range of camera and trajectory degradation, although it gradually declines as perturbations intensify. This is important because maintenance prioritization depends more on stable relative ordering than on exact pointwise equality. Panel (d) evaluates the non-additive dropout interaction, namely the extra degradation beyond the sum of the two marginal effects. The interaction field is mostly moderate, with only a few localized high-response regions, suggesting that the joint impact of camera and trajectory loss is largely predictable and does not usually cause severe compounding instability.

Panels (e) and (f) summarize the budgeted and diagonal robustness behavior. In panel (e), both worst-case prediction shift and worst-case rank loss increase steadily with perturbation budget, but the smooth increase indicates graceful degradation under bounded adverse conditions. As the perturbation budget increases from 0 to 1, the worst-case prediction shift rises from approximately 0.009 to 0.026, while the corresponding rank loss increases from about 0.03 to 0.13. Panel (f) shows the diagonal degradation profile under equalized flow-geometry noise. MAE rises monotonically, while the local degradation curvature remains small at low-to-moderate noise and becomes more pronounced only at larger perturbation levels. This means that the framework is stable under mild noise but deteriorates more rapidly after perturbations exceed a higher threshold. Overall, Fig. 17 shows that the proposed method is robust to realistic source degradation, while dynamic traffic information remains the most critical component for accurate and stable pressure inference.

Fig. 18 provides a propagation-aware view of uncertainty and deployment priority at the peak-demand window τ^*. In panel (a), the stochastic uncertainty field is obtained from Monte Carlo perturbations of the dynamic inputs, with low observation reliability further increasing the final uncertainty score. The map shows that uncertainty is not uniformly distributed, but is mainly concentrated in the dense central corridor, while many peripheral segments remain at much lower levels. Quantitatively, the uncertainty in the central corridor commonly reaches about 0.14–0.18, whereas many peripheral segments remain below approximately 0.08. This indicates that the most operationally important region is also where inference confidence is relatively weaker,

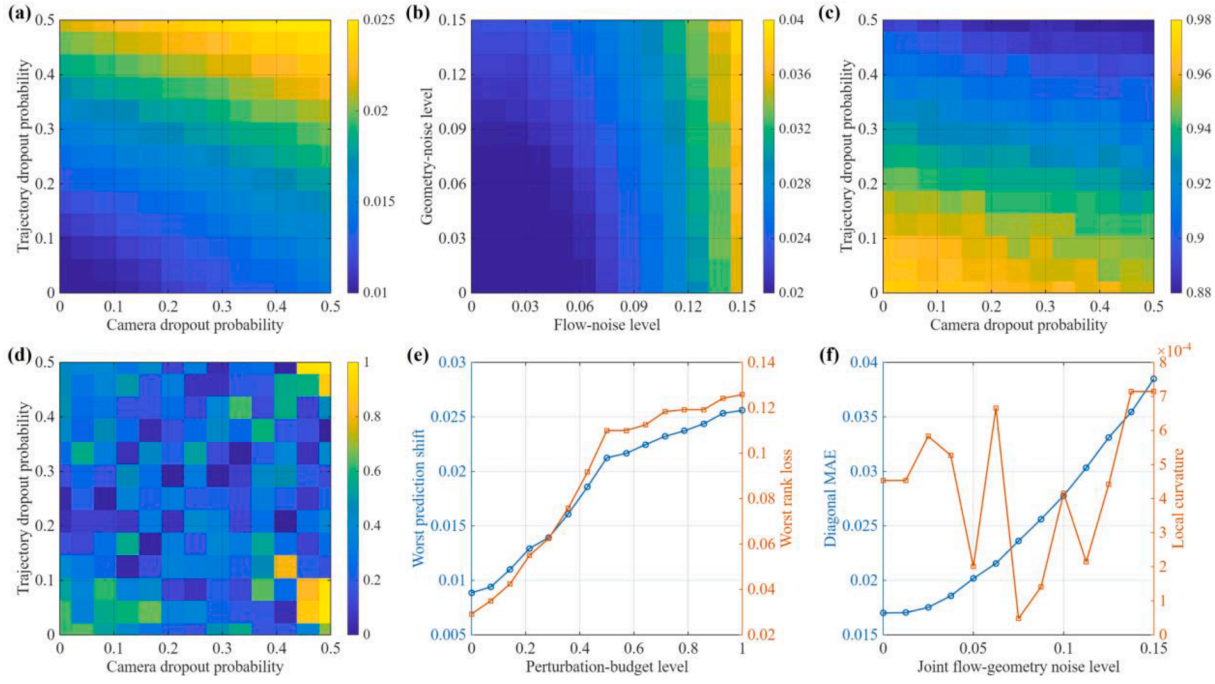


Fig. 17. Robustness of pressure inference under source degradation and perturbation: (a) prediction shift under joint camera dropout and trajectory dropout, (b) MAE surface under joint flow-noise and geometry-noise perturbation, (c) high-pressure rank overlap under camera and trajectory degradation, (d) non-additive dropout interaction beyond marginal degradation, (e) worst-case prediction shift and rank loss under bounded perturbation budget, (f) diagonal degradation profile under joint flow-geometry noise, including MAE and local curvature of degradation.

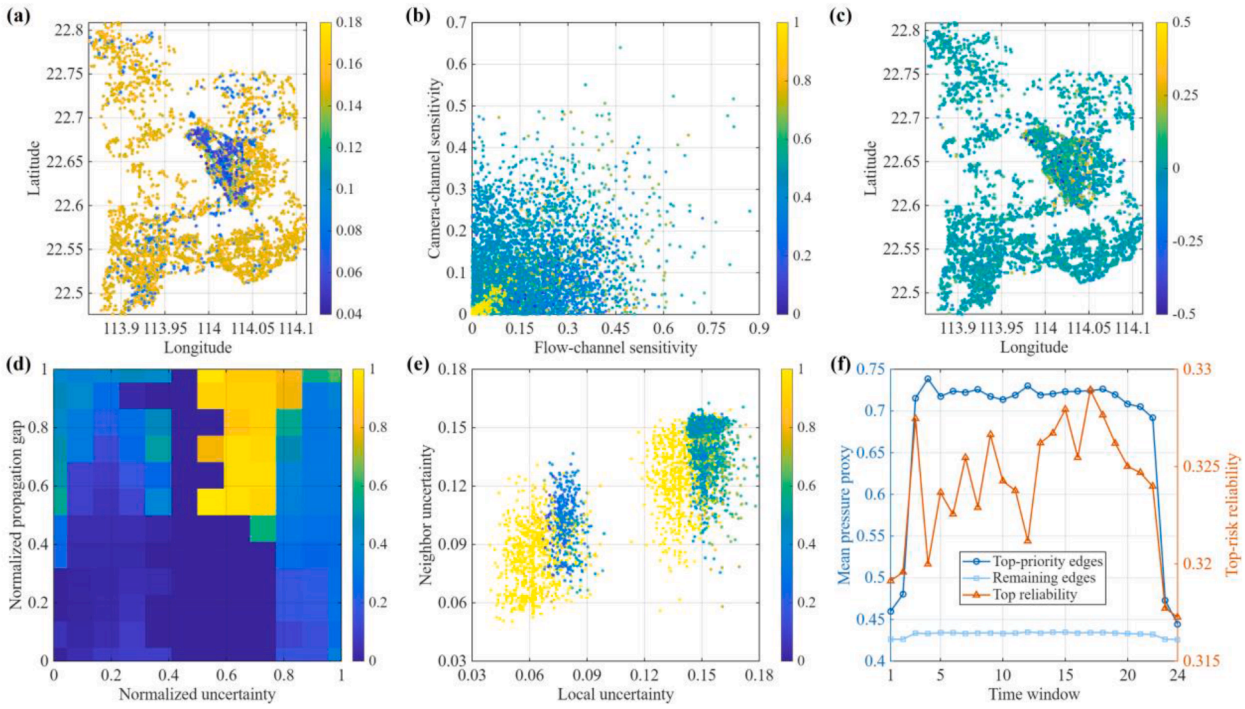


Fig. 18. Stochastic uncertainty, source sensitivity, and propagation-aware deployment: (a) peak-window stochastic uncertainty field obtained from Monte Carlo perturbations of dynamic features, (b) local source-channel sensitivity with respect to flow evidence and camera-related reliability, (c) local pressure propagation gap between $\widehat{Y}_{e,\tau^*}$ and its hybrid-neighborhood average, (d) deployment priority over normalized uncertainty and propagation-gap states, (e) uncertainty clustering over the hybrid adjacency A^H, (f) temporal signature of top-priority road segments, including mean pressure proxy $Y_{e,\tau}$, remaining-edge pressure, and top-risk reliability $\rho_{e,\tau}$.

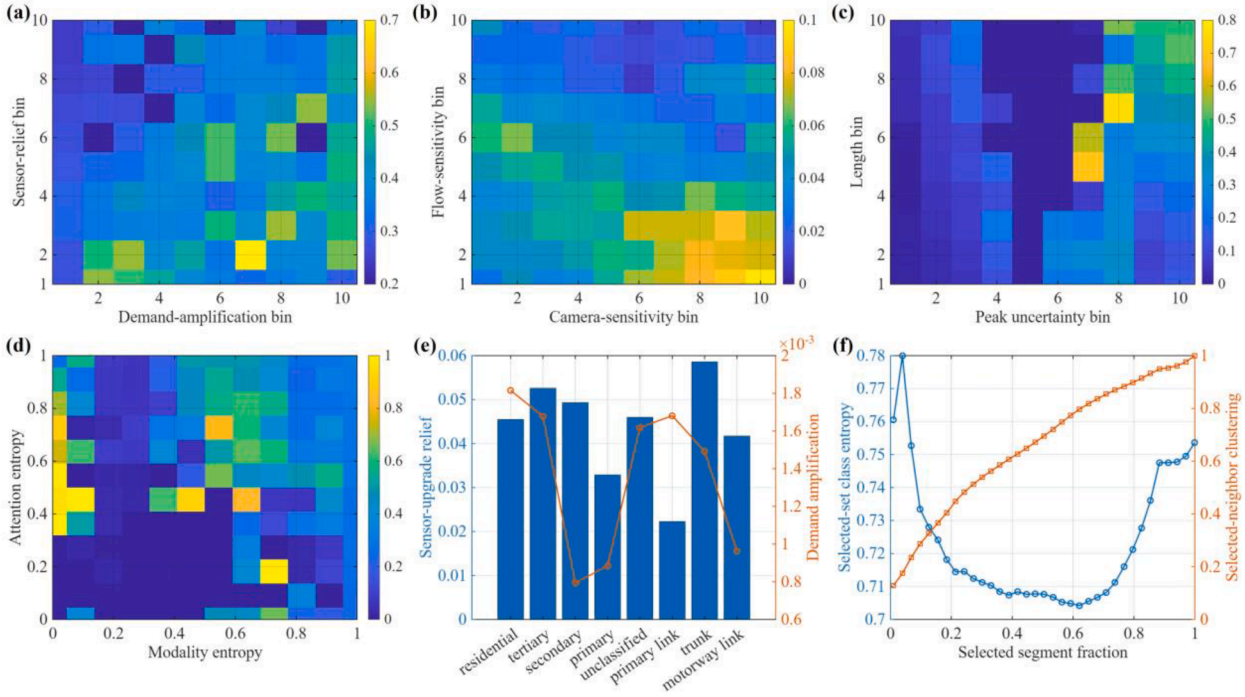


Fig. 19. Decision-oriented intervention diagnostics and portfolio structure: (a) priority matrix over sensor-relief and demand-amplification states, (b) sensor-relief matrix over camera-channel and flow-channel sensitivities, (c) smoothing-need matrix over peak uncertainty and segment length ℓ_e, (d) priority surface over modality entropy E_e^M and attention entropy E_e^A, (e) class-level sensor-upgrade relief and demand amplification, (f) intervention-portfolio diversity and graph clustering as the selected segment fraction increases.

making maintenance-oriented decision support particularly necessary. Panel (b) shows that local source-channel sensitivity is more widely spread along the flow channel than along the camera-related channel. Since sensitivities are computed by perturbing the corresponding input channel and measuring the induced pressure change, this result suggests that the inferred pressure field is more strongly driven by flow evidence, although some segments still show non-negligible camera sensitivity.

Panels (c)-(e) clarify how uncertainty interacts with hybrid-graph propagation. The local pressure propagation gap in panel (c), defined as the difference between the local prediction $\widehat{Y}_{e,\tau^*}$ and the average prediction over hybrid neighbors, is spatially concentrated rather than random. Most segments remain close to zero, but several zones in the main corridor show positive or negative departures, indicating locally amplified or suppressed pressure relative to surrounding segments. Panel (d) shows that deployment priority becomes highest when both uncertainty and the normalized propagation gap are high, meaning that segments are most intervention-worthy when they are uncertain and structurally inconsistent with their neighborhood. Panel (e) further confirms uncertainty clustering over the hybrid adjacency $\mathbf{A}^H$, where high-uncertainty segments tend to connect with neighbors of similarly high uncertainty. This supports area-based maintenance planning rather than purely isolated segment-wise intervention.

Panel (f) summarizes the temporal signature of top-priority segments. The mean pressure proxy $Y_{e,\tau}$ for top-priority edges remains consistently higher than that of the remaining edges across most temporal windows $\mathcal{T}$, with a stable separation during the main operating period. Across the main operating windows, the mean pressure proxy of top-priority edges stays around 0.71–0.73, compared with roughly 0.43 for the remaining edges, while the corresponding top-risk reliability fluctuates only slightly around 0.32. In contrast, their reliability curve varies only mildly and stays at a moderate level. This shows that the selected priority segments are not simply the least reliable ones, but those where sustained pressure burden, uncertainty, and propagation-related inconsistency coincide. Overall, Fig. 18 demonstrates that the proposed framework identifies deployment targets through the joint assessment of predictive burden, stochastic instability, and graph-mediated spatial context, making the resulting intervention priorities more informative than pressure-only ranking.

Fig. 19 translates the previous inference and robustness results into a decision-oriented intervention perspective. Panel (a) shows the priority matrix over sensor-relief and demand-amplification states. Sensor-relief measures the pressure reduction produced by improving camera support, whereas demand amplification measures the pressure increase under intensified demand. The highest priority appears where demand sensitivity and observation-improvement benefit coexist, indicating that effective intervention targets are jointly shaped by operational stress and observational recoverability. Panel (b) further maps sensor-relief against camera-channel and flow-channel sensitivities. The strongest relief occurs where camera sensitivity is high but flow sensitivity is low to moderate, suggesting that sensor upgrades are most valuable when the main limitation is observational rather than purely traffic-driven. Panel (c) shows that larger smoothing-need values are concentrated in high-uncertainty bins, especially for medium-to-long segments, indicating that uncertainty-aware smoothing or local coordination should focus on structurally important segments instead of being applied uniformly.

Panels (d) and (e) relate decision priority to information fusion and road-class heterogeneity. In panel (d), the priority surface is shown over modality entropy E_e^M and attention entropy E_e^A, representing source-allocation balance and attention dispersion over hybrid neighbors, respectively. The non-monotonic surface indicates that high-priority segments are not simply those with maximum fusion uncertainty, but those located in entropy regimes where cross-source evidence remains mixed and structurally informative. Panel (e) shows that intervention leverage differs by road class. Sensor-upgrade relief is stronger for classes such as trunk and tertiary roads, whereas demand amplification is more pronounced for several urban service-oriented classes. The sensor-upgrade relief reaches approximately 0.059 for trunk roads and 0.053 for tertiary roads, compared with only about 0.023 for unclassified roads. This suggests that maintenance strategies should be class-sensitive because sensing improvement and demand vulnerability are unevenly distributed across the road hierarchy.

Panel (f) summarizes the intervention portfolio as the selected segment fraction increases. Selected-neighbor clustering increases steadily and approaches 1, showing that larger portfolios become more spatially and topologically contiguous. Meanwhile, selected-set class entropy decreases from about 0.76 to a minimum near 0.70 at a selected fraction of approximately 0.6, before recovering to about 0.75 as the portfolio approaches full coverage, indicating a shift from an initially diverse candidate pool to a concentrated core set, followed by renewed diversification as lower-priority segments are added. This suggests that the highest-value intervention portfolio is not merely a list of isolated critical edges, but a structured subset with thematic selectivity and increasing spatial coherence, which is useful for coordinated maintenance deployment and staged operational planning.

4.4. Cross-city external validation and generalization assessment

To further assess the generalization capability of the proposed framework beyond the original Shenzhen case, an independent cross-city validation was conducted using a separate dataset from Jinan, China [90]. The SHGAT model trained using the Shenzhen dataset was directly transferred to the Jinan road network without retraining, fine-tuning, or city-specific parameter adjustment. Accordingly, the Jinan data were used exclusively for external inference and validation, providing a more stringent evaluation of whether the learned reliability-aware graph representation can remain effective under changes in road-network topology, traffic conditions, and observation patterns.

The independent validation dataset was collected in Jinan on October 17, 2022. As shown in Fig. 20, the trajectory observations cover a broad urban road-network area containing multiple road classes and spatially heterogeneous traffic corridors. The dashed boundary delineates the principal vehicle-trajectory coverage area used in the present external-validation case. Compared with the Shenzhen dataset used for model development, the Jinan case represents an independent spatial and temporal setting, thereby enabling the transferability of the trained framework to be examined under a different urban road-network environment.

Fig. 21 evaluates the cross-city transferability of the Shenzhen-trained SHGAT model under a fully frozen setting. As shown in Fig. 21 (a), the Jinan validation case contains a spatially extensive urban road network with unevenly distributed camera observations, providing a structurally distinct sensing environment from the Shenzhen training domain. Despite this change in network topology and observation configuration, panel (b) shows that the transferred model preserves strong predictive agreement on the independent Jinan dataset, with a correlation of $r = 0.978$. The predicted values remain closely aligned with the one-to-one reference across most of the target range, indicating that the reliability-informed representation learned in Shenzhen captures pressure-related relationships that remain valid under a different urban network and traffic environment. The color distribution further suggests that observation reliability affects the spread of the transferred predictions without disrupting the overall consistency between the target and inferred pressure.

The calibration results in Fig. 21 (c) provide a more detailed view of this transfer behavior. The predicted and target decile means follow a similar monotonic pattern, particularly in the upper deciles where the inferred pressure rises sharply and both curves approach unity. The bias remains predominantly negative across most deciles, indicating a mild conservative tendency of the frozen model on the Jinan case, with the

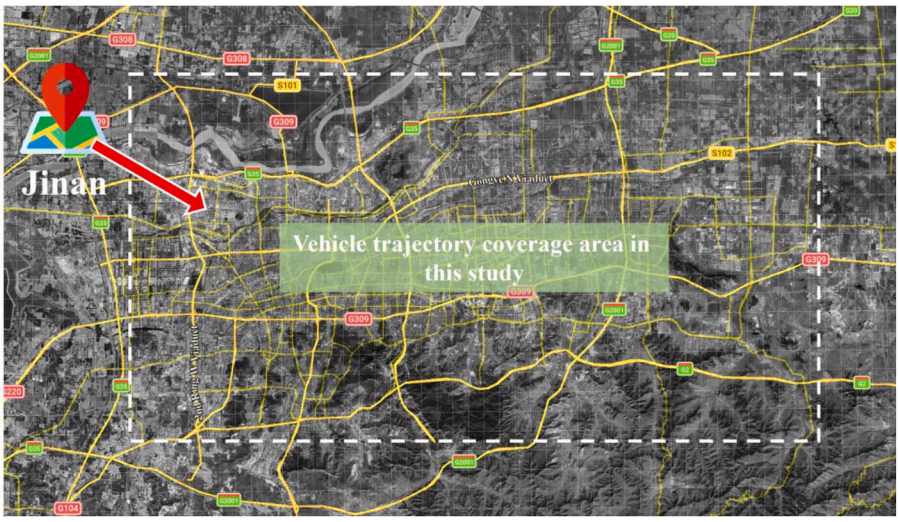


Fig. 20. Study area and vehicle trajectory coverage of the Jinan external-validation dataset on October 17, 2022.

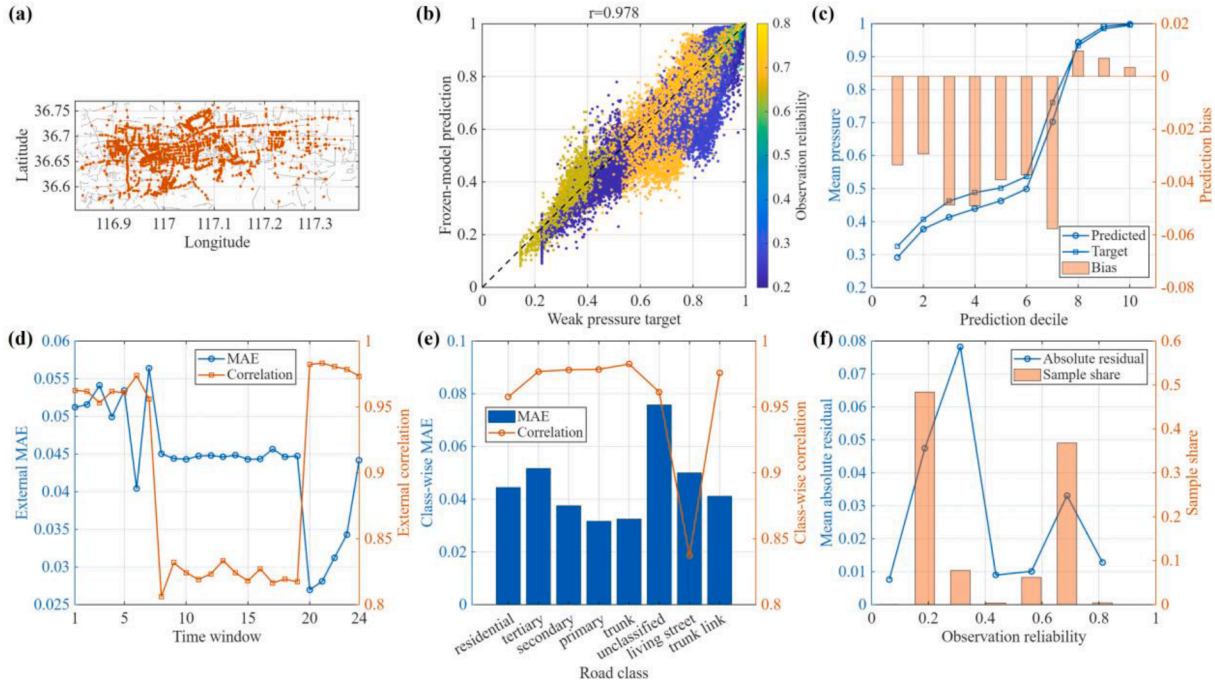


Fig. 21. Independent external validation of the frozen SHGAT model on the Jinan dataset: (a) Jinan road network and camera observations, (b) agreement between the weak pressure target and frozen-model prediction, colored by observation reliability, (c) external decile calibration and prediction bias, (d) time-resolved external validation performance, (e) road-class-specific external validation, (f) reliability-conditioned external residual behavior.

largest underestimation occurring in the intermediate-to-high pressure range rather than at the extreme high-pressure end. This behavior shows that the transferred model largely preserves ordinal discrimination under cross-city domain shift, while exhibiting moderate compression of some intermediate pressure levels.

Panels (d) and (e) reveal that the external performance is not temporally or structurally uniform. Across the 24 time windows, the MAE remains within a relatively narrow range for most periods, whereas the correlation exhibits stronger temporal variation and recovers to a high level in the later windows. This indicates that the transferred model is more sensitive to temporal operating regimes than to global error inflation. The road-class analysis in Fig. 21 (e) shows a similar heterogeneity: primary and trunk-related categories generally exhibit lower MAE and stronger correlation, whereas unclassified roads and living-street-type segments present weaker transfer performance. This suggests that cross-city generalization is stronger for road classes with more regular operational patterns and clearer structural roles.

Fig. 21 (f) further links the residual structure to observation reliability. The residual magnitude does not decrease monotonically with reliability; instead, larger errors occur within several intermediate reliability intervals that also contain substantial sample shares. This indicates that external prediction difficulty is governed not only by nominal observation confidence, but also by the concentration of challenging traffic states within particular reliability ranges. Overall, the results demonstrate that the frozen SHGAT model retains strong cross-city predictive consistency while preserving interpretable temporal, structural, and reliability-dependent error patterns under the independent Jinan case.

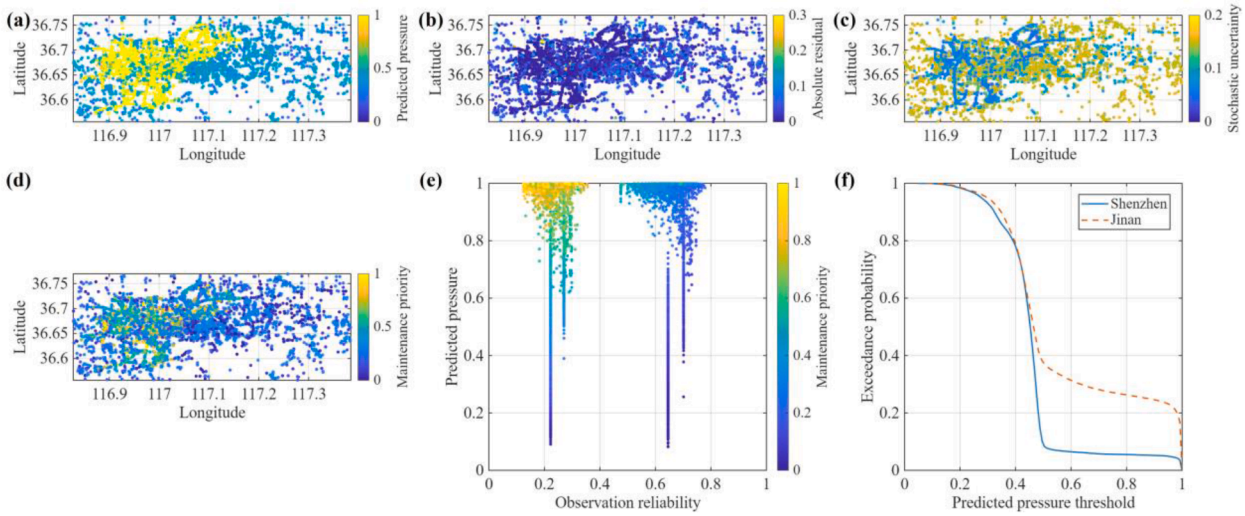


Fig. 22. Spatial transferability and externally transferred maintenance inference of the frozen SHGAT model in Jinan: (a) peak-window predicted pressure field, (b) spatial distribution of absolute residuals, (c) stochastic uncertainty field under frozen-model perturbations, (d) transferred maintenance-priority field, (e) pressure–reliability decision space colored by maintenance priority, (f) cross-city comparison of peak-pressure exceedance curves between Shenzhen and Jinan.

Fig. 22 extends the external validation from aggregate prediction accuracy to the spatial and decision-oriented behavior of the transferred model. The predicted peak-window pressure field in Fig. 22 (a) is clearly heterogeneous rather than spatially diffuse, with several concentrated high-pressure corridors embedded within a broader medium- to low-pressure background. This indicates that the frozen model retains the ability to differentiate localized operational burden under a new city setting. The corresponding residual map in Fig. 22 (b) shows that most segments remain within a relatively low-error range, while larger absolute residuals appear only in limited spatial pockets. The transferred pressure field is therefore not only spatially structured, but also spatially credible over most of the observed network.

The uncertainty field in Fig. 22 (c) provides an important complement to the residual analysis. The stochastic uncertainty is not uniformly distributed and tends to increase in portions of the network where prediction variability is more sensitive to perturbation. However, these zones do not coincide perfectly with the largest residual clusters in Fig. 22 (b), implying that instantaneous prediction error and stochastic sensitivity capture different aspects of transfer behavior. In practical terms, this distinction is valuable: some segments are difficult because they are systematically harder to fit, whereas others are mainly sensitive to stochastic variation in the transferred inference process.

The maintenance-priority field in Fig. 22 (d) shows that the externally inferred priority pattern remains selective and spatially organized, rather than merely reproducing the raw pressure map. This is further clarified by the pressure–reliability decision space in Fig. 22 (e). High maintenance priority concentrates primarily in the region combining elevated predicted pressure with lower or moderate observation reliability, whereas segments with similarly high pressure but stronger observational support are assigned comparatively lower priority colors. This confirms that the decision score is not a simple pressure ranking; it incorporates the interaction between latent operational burden and observation confidence, thereby preserving the reliability-aware design principle of the framework under cross-city transfer.

The cross-city exceedance comparison in Fig. 22 (f) further highlights the distinct pressure distribution learned for the Jinan case. The Shenzhen curve decreases sharply around the mid-pressure range, whereas the Jinan curve decays more gradually and stays higher over a broader range of thresholds. This indicates that high-pressure segments are more persistently represented in the transferred Jinan distribution, suggesting a heavier upper tail in the inferred maintenance-pressure profile. In other words, the Jinan case is not only predictable by the frozen model, but also exhibits a different pressure-risk structure from the Shenzhen training domain. Taken together, Fig. 22 shows that the proposed framework preserves spatial discrimination, interpretable uncertainty, and reliability-informed prioritization when transferred to an independent city, thereby supporting both its external validity and its practical relevance for cross-network maintenance screening.

5. Discussion

The weak pressure field provides a structured supervisory signal derived from trajectory-based operational evidence, and several of its constituent variables also enter the subsequent inference process. This design places the model within a weakly supervised learning setting, where the principal objective is to learn a stable network-level representation of operational maintenance pressure from incomplete and heterogeneous observations. The SHGAT architecture expands the information available to each segment through hybrid-neighborhood propagation, road geometry, road class, observation reliability, modality allocation, and stochastic perturbation, allowing the inferred pressure field to reflect broader spatial and contextual relationships than the local weak-label formulation alone. Holdout evaluation, component ablation, robustness analysis, and independent cross-city transfer further assess the stability of this learned representation under changes in samples, model components, sensing conditions, and network environments. The resulting pressure field is therefore most appropriately interpreted as an operational maintenance-risk proxy, with future incorporation of pavement-condition surveys, inspection records, and maintenance histories providing an additional basis for physical calibration.

For road maintenance management, the inferred pressure, stochastic uncertainty, observation reliability, intervention value, and maintenance-priority score can be combined into a hierarchical screening process. Network segments with persistently elevated pressure and priority can be identified for earlier inspection, while areas characterized by high operational burden and weak sensing support can receive additional monitoring before intervention decisions are finalized. Spatial clustering and propagation diagnostics can further reveal connected critical corridors, which is useful when maintenance activities are organized at corridor or network level rather than for isolated segments. The capture-curve and intervention analyses also provide a basis for allocating limited maintenance and sensing resources toward subsets of the network that account for a comparatively large share of

the inferred burden. Counterfactual responses under different sensing, demand, and capacity conditions add another layer of decision information by showing how candidate interventions may alter the spatial distribution of pressure and priority.

Agency-level implementation can be organized around routinely available road-network, trajectory, traffic-observation, and camera-deployment data. Directed road connectivity, road class, geometry, and sensing availability provide the static network description, while time-referenced vehicle movements or equivalent segment-level traffic observations supply the dynamic evidence required for periodic pressure updating. The current framework uses a compact graph-learning and multilayer inference structure with sparse neighborhood construction, making network-wide inference suitable for repeated operational assessment once data preprocessing has been completed. In practice, major implementation considerations are likely to arise from trajectory access, map matching, synchronization among heterogeneous data sources, evolving road-network topology, uneven sensing quality, and transfer across operating environments. Periodic updating with newly available inspection and maintenance information can progressively strengthen local calibration and maintain the relevance of the inferred priorities as network conditions evolve.

6. Conclusion

This study developed a reliability-informed stochastic hybrid graph attention network (SHGAT) framework for uncertainty-aware road operation-maintenance risk inference under heterogeneous and camera-sparse traffic-network observations. The framework transformed full-day vehicle trajectories into segment-time operational evidence, including flow exposure, traversal-time burden, inverse-speed pressure, temporal irregularity, neighbor-induced recurrence, bottleneck contrast, observation reliability, and uncertainty. These evidence components were integrated with road geometry, road class, topology, and camera availability to construct a weak operation-maintenance pressure field and a hybrid segment graph. In this formulation, the inferred pressure is treated as a reliability-risk proxy rather than a direct pavement-damage label, which makes the framework suitable for early screening and decision support when direct maintenance records or condition measurements are unavailable.

The empirical results showed that road operation-maintenance pressure is highly heterogeneous in both space and time. Pressure-related evidence was concentrated in limited corridors and segment groups, while camera observability was unevenly distributed across the network. The SHGAT architecture captured reliability-aware neighborhood propagation and source-specific fusion behavior, and the stochastic mini-batch learning process achieved stable convergence and strong agreement between predicted pressure and weak pressure targets. Source ablation and modality analysis further indicated that flow and temporal evidence dominate the inferred pressure magnitude, whereas camera and static information mainly contribute to reliability conditioning, structural stabilization, and interpretability. These findings suggest that maintenance-oriented network risk cannot be reliably inferred from traffic exposure alone, but requires the joint interpretation of demand, temporal instability, observability, topology, and source reliability.

The robustness and decision-oriented analyses further demonstrated the value of the proposed framework for reliability-based maintenance planning. The inferred pressure field remained relatively stable under camera dropout, trajectory degradation, flow noise, and geometry perturbation, while counterfactual experiments revealed how pressure responds to changes in sensing, demand, and capacity conditions. The resulting uncertainty fields, source-channel sensitivities, propagation gaps, intervention values, and priority portfolios provide interpretable evidence for identifying camera-sparse high-pressure segments and supporting risk-informed maintenance decisions. A limitation of the present study is that the weak pressure target is derived from operational observations rather than direct pavement-condition records or maintenance-failure records. Future work should therefore integrate multi-day trajectories, inspection data, pavement distress measurements, accident or service-disruption records, and field maintenance logs to further calibrate and validate the inferred reliability-risk indicators. Overall, the proposed SHGAT framework offers a data-driven and uncertainty-aware approach for transforming heterogeneous traffic-network observations into actionable road infrastructure reliability intelligence.

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.

Acknowledgments

This research is supported by the China Scholarship Council (CSC., No. 202308080042).

Appendix A. Auxiliary road-network and trajectory evidence construction

The road-network representation and trajectory records are transformed into a set of geometric, topological, and operational descriptors before being incorporated into the pressure-inference framework. This appendix presents the detailed formulations used to construct these descriptors, including road geometry, traversal characteristics, temporal variability, neighboring flow recurrence, and bottleneck-related evidence.

Let $\mathscr{T}_m$ denote the m-th vehicle trajectory. Each trajectory is represented by an ordered sequence of node-time observations:

$$\mathscr{T}_m = \left\{ \left(v_{m,q}, t_{m,q} \right) \right\}_{q \in \mathscr{C}_m}, \tag{A1}$$

where $v_{m,q} \in \mathscr{V}$ is the visited node, $t_{m,q}$ is the associated timestamp, and $\mathscr{C}_m$ is the ordered point index set of trajectory m. Consecutive node pairs are matched to directed road segments whenever a direction-compatible segment exists. The traversal event indicator for segment e and trajectory step q is:

$$I_e^{m,q} = \mathbb{I}\left\{ \left(v_{m,q}, v_{m,q^+} \right) \sim e \right\}, \tag{A2}$$

where q^+ is the successor of q, $\sim$ denotes direction-compatible segment matching, and $\mathbb{I}\{\cdot\}$ is the indicator function. The traversal duration associated with this event is:

$$\Delta t_{m,q} = t_{m,q^+} - t_{m,q}. \tag{A3}$$

The observation horizon is partitioned into a sequence of time windows $\{\mathscr{S}_\tau\}_{\tau \in \mathscr{T}}$, where τ indexes a temporal bin. The segment-time indicator is:

$$I_{e,\tau}^{m,q} = I_e^{m,q} II\left[t_{m,q} \in \mathscr{J}_\tau\right] \tag{A4}$$

The trajectory-derived flow matrix $\mathbf{F} \in \mathbb{R}^{N_E \times |\mathscr{T}|}$ is defined by:

$$F_{e,\tau} = \sum_m \sum_{q \in \mathscr{Q}_m} I_{e,\tau}^{m,q}. \tag{A5}$$

The cumulative traversal-time matrix $\mathbf{S}^T$ and squared traversal-time matrix $\mathbf{S}^{T2}$ are:

$$S_{e,\tau}^T = \sum_m \sum_{q \in \mathscr{Q}_m} \Delta t_{m,q} I_{e,\tau}^{m,q}, \tag{A6}$$

$$S_{e,\tau}^{T2} = \sum_m \sum_{q \in \mathscr{Q}_m} \left(\Delta t_{m,q}\right)^2 I_{e,\tau}^{m,q}. \tag{A7}$$

The mean traversal-time burden is obtained as:

$$\overline{T}_{e,\tau} = \frac{S_{e,\tau}^T}{F_{e,\tau} + \varepsilon}, \tag{A8}$$

and the traversal-time variance is:

$$V_{e,\tau}^T = \frac{S_{e,\tau}^{T2}}{F_{e,\tau} + \varepsilon} - \left(\overline{T}_{e,\tau}\right)^2. \tag{A9}$$

The temporal irregularity index is then:

$$\Omega_{e,\tau} = \frac{\sqrt{\max\left(V_{e,\tau}^T, 0\right)}}{\overline{T}_{e,\tau} + \varepsilon}. \tag{A10}$$

The empirical segment speed is:

$$U_{e,\tau} = \frac{\ell_e}{\overline{T}_{e,\tau} + \varepsilon}. \tag{A11}$$

Since high speed usually indicates lower operational burden, the inverse-speed pressure evidence is represented by:

$$\widetilde{U}_{e,\tau}^{-1} = 1 - \mathscr{N}\left(U_{e,\tau}\right). \tag{A12}$$

The normalized dynamic flow, traversal-time burden, and time-irregularity features are:

$$\widetilde{F}_{e,\tau} = \mathscr{N}\left(\log\left(1 + F_{e,\tau}\right)\right), \tag{A13}$$

$$\widetilde{T}_{e,\tau} = \mathscr{N}\left(\overline{T}_{e,\tau}\right),\ \widetilde{\Omega}_{e,\tau} = \mathscr{N}\left(\Omega_{e,\tau}\right). \tag{A14}$$

The hybrid neighbor count of segment e is:

$$D_e^H = \sum_{e' \in \mathscr{E}} A_{e,e'}^H. \tag{A15}$$

The neighbor flow exposure is:

$$F_{e,\tau}^N = \sum_{e' \in \mathscr{E}} A_{e,e'}^H F_{e',\tau}, \tag{A16}$$

and the corresponding average neighboring flow is:

$$\overline{F}_{e,\tau}^N = \frac{F_{e,\tau}^N}{D_e^H + \varepsilon}. \tag{A17}$$

A recurrence-enhanced flow descriptor is introduced as:

$$R_{e,\tau} = F_{e,\tau} + \alpha_R \overline{F}_{e,\tau}^N, \tag{A18}$$

where α_R is a non-negative recurrence coefficient controlling the contribution of neighboring demand. Its normalized form is:

$$\widetilde{R}_{e,\tau} = \mathscr{N}\left(\log\left(1 + R_{e,\tau}\right)\right). \tag{A19}$$

The local bottleneck contrast is defined by comparing the segment flow with the total local flow around it:

$$B_{e,\tau} = \frac{F_{e,\tau}}{F_{e,\tau} + F_{e,\tau}^N + \varepsilon}. \tag{A20}$$

Its normalized form is denoted by $\widetilde{B}_{e,\tau} = \mathscr{N}\left(B_{e,\tau}\right)$. This quantity is high when a segment absorbs a large proportion of its local neighborhood demand, which is consistent with a bottleneck-like or concentration-like operational role.

Time-aggregated descriptors are also used to support static embedding and visualization. The mean pressure, mean reliability, mean flow, and mean irregularity of segment e are:

$$\begin{aligned}\overline{Y}_e &= \frac{1}{|\mathscr{T}|}\sum_{\tau\in\mathscr{T}} Y_{e,\tau},\\ \overline{\rho}_e &= \frac{1}{|\mathscr{T}|}\sum_{\tau\in\mathscr{T}} \rho_{e,\tau},\\ \overline{F}_e &= \frac{1}{|\mathscr{T}|}\sum_{\tau\in\mathscr{T}} F_{e,\tau},\\ \overline{\Omega}_e &= \frac{1}{|\mathscr{T}|}\sum_{\tau\in\mathscr{T}} \widetilde{\Omega}_{e,\tau}.\end{aligned} \tag{A21}$$

Appendix B. Auxiliary interpretability, evaluation, and diagnostic formulations

The inferred pressure field is further examined from the perspectives of model interpretation, predictive performance, robustness, temporal behavior, and classification consistency. This appendix presents the corresponding quantitative formulations used to characterize attention and modality behavior, prediction accuracy, perturbation response, residual structure, and distributional changes under different evaluation conditions.

The modality entropy is:

$$E_e^M = -\frac{\sum_{m\in\mathscr{M}} \beta_e^m \log\left(\beta_e^m + \varepsilon\right)}{\log|\mathscr{M}|}, \tag{B1}$$

where $\mathscr{M} = \{G, F, T, C\}$. The modality dominance is:

$$D_e^M = \max_{m\in\mathscr{M}} \beta_e^m. \tag{B2}$$

High modality entropy indicates balanced cross-source fusion, whereas high modality dominance indicates that one evidence channel controls the pressure interpretation.

Prediction quality is evaluated through continuous agreement measures. Given validation predictions $\widehat{\mathbf{y}}$ and targets $\mathbf{y}$, the mean absolute error (MAE) is:

$$\mathrm{MAE} = \frac{1}{N}\sum_{n\in\mathscr{V}_a} |\widehat{y}_n - y_n|, \tag{B3}$$

where $\mathscr{V}_a$ is the validation sample set and $N = |\mathscr{V}_a|$. The root mean squared error (RMSE) is:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{n\in\mathscr{V}_a} (\widehat{y}_n - y_n)^2}. \tag{B4}$$

The Pearson correlation coefficient is defined as [94]:

$$\mathrm{Corr} = \frac{\sum_{n\in\mathscr{V}_a} (\widehat{y}_n - \overline{\widehat{y}})(y_n - \overline{y})}{\sqrt{\sum_{n\in\mathscr{V}_a} (\widehat{y}_n - \overline{\widehat{y}})^2}\sqrt{\sum_{n\in\mathscr{V}_a} (y_n - \overline{y})^2} + \varepsilon}, \tag{B5}$$

where $\overline{\widehat{y}}$ and $\overline{y}$ are validation means.

For network-level mapping, a representative time window τ^* is selected from the temporal demand profile. The peak-window prediction of segment e is:

$$\widehat{P}_e = \widehat{Y}_{e,\tau^*}. \tag{B6}$$

A structural baseline is introduced to distinguish learned operational pressure from static structural predisposition:

$$P_e^S = \sigma\left(\boldsymbol{\theta}_S^{\mathrm{T}}\left[\widetilde{\ell}_e, \widetilde{\chi}_e, \widetilde{d}_e^E, \widetilde{\kappa}_e\right]^{\mathrm{T}} + b_S\right), \tag{B7}$$

where $\boldsymbol{\theta}_S$ and b_S define the static baseline mapping. The excess pressure is:

$$P_e^X = \widehat{P}_e - P_e^S. \tag{B8}$$

Given a base prediction $\widehat{y}_n$ and perturbed prediction $\widehat{y}_n^P$, the prediction-shift consistency measure is:

$$\Delta_{\mathrm{cons}} = \frac{1}{N_R}\sum_{n\in\mathscr{R}} |\widehat{y}_n^P - \widehat{y}_n|, \tag{B9}$$

where $\mathscr{R}$ is the robustness-evaluation sample set and $N_R = |\mathscr{R}|$. The perturbed MAE is:

$$\mathrm{MAE}^P = \frac{1}{N_R}\sum_{n\in\mathscr{R}}\left|\widehat{y}_n^P - y_n\right|. \tag{B10}$$

High-pressure rank stability is measured by top-set overlap. Let $\mathrm{Top}_q(\mathbf{a})$ be the index set of the largest q-fraction of vector $\mathbf{a}$. The rank-overlap score is:

$$\mathscr{O}_q = \frac{\left|\mathrm{Top}_q(\widehat{\mathbf{y}}) \cap \mathrm{Top}_q(\widehat{\mathbf{y}}^P)\right|}{\left|\mathrm{Top}_q(\widehat{\mathbf{y}})\right|}. \tag{B11}$$

Joint perturbation surfaces are computed by evaluating prediction shift or error over two degradation axes. For camera dropout level a and trajectory dropout level b, the shift surface is:

$$\mathscr{S}^{CT}(a,b) = \frac{1}{N_R}\sum_{n\in\mathscr{R}}\left|\widehat{y}_n^P(a,b) - \widehat{y}_n\right|. \tag{B12}$$

The non-additive interaction between camera and trajectory degradation is:

$$\mathscr{I}^{CT}(a,b) = \left|\mathscr{S}^{CT}(a,b) - \overline{\mathscr{S}}^{CT}_{a,\cdot} - \overline{\mathscr{S}}^{CT}_{\cdot,b} + \overline{\mathscr{S}}^{CT}_{\cdot,\cdot}\right|, \tag{B13}$$

where the barred terms denote row, column, and global averages over the perturbation surface. A high interaction value indicates that simultaneous camera and trajectory degradation causes more complex behavior than the sum of separate degradations.

Temporal diagnostics summarize demand concentration. The active-edge ratio at time τ is:

$$A_\tau^R = \frac{\sum_{e\in\mathscr{E}} II\left[F_{e,\tau} > 0\right]}{N_E}. \tag{B14}$$

The total temporal demand is:

$$F_\tau^\Sigma = \sum_{e\in\mathscr{E}} F_{e,\tau}. \tag{B15}$$

The temporal flow concentration index is:

$$H_\tau = \sum_{e\in\mathscr{E}}\left(\frac{F_{e,\tau}}{F_\tau^\Sigma + \varepsilon}\right)^2. \tag{B16}$$

The flow-weighted reliability is:

$$\overline{\rho}_\tau^F = \frac{\sum_{e\in\mathscr{E}}\rho_{e,\tau}F_{e,\tau}}{F_\tau^\Sigma + \varepsilon}. \tag{B17}$$

Residual diagnostics are introduced to examine spatial error coherence. Let r_e be the mean signed residual of segment e over validation samples assigned to that segment. The residual sign-coherence score is:

$$C_e^R = \mathrm{sgn}(r_e)\frac{\sum_{e'\in\mathscr{E}} A_{e,e'}^H \mathrm{sgn}(r_{e'})}{D_e^H + \varepsilon}. \tag{B18}$$

A positive value indicates that the segment and its neighbors tend to have residuals with the same sign, suggesting spatially coherent underestimation or overestimation. A negative value indicates local disagreement in residual direction.

For ordinal diagnostic analysis, continuous pressure targets and predictions can be discretized by empirical quantile thresholds. Let $\mathscr{C}(y)$ map a continuous value y to an ordinal pressure class. The confusion matrix is:

$$\mathbf{M}_{a,b} = \sum_{n\in\mathscr{T}_a} II[\,\mathscr{C}(y_n) = a,\ \mathscr{C}(\widehat{y}_n) = b], \tag{B19}$$

where a and b denote target and predicted classes. For high-pressure detection, the Matthews correlation coefficient (MCC) is [95]:

$$\mathrm{MCC} = \frac{TP\cdot TN - FP\cdot FN}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)} + \varepsilon}, \tag{B20}$$

where TP, FP, FN, and TN are the standard binary decision counts. The balanced accuracy (BA) is:

$$\mathrm{BA} = \frac{1}{2}\left(\frac{TP}{TP+FN+\varepsilon} + \frac{TN}{TN+FP+\varepsilon}\right). \tag{B21}$$

Finally, higher-order residual statistics are computed for each time window. If $r_{n,\tau}$ is the signed residual of validation sample n in time window τ, its skewness is:

$$\mathrm{Skew}_\tau = \frac{1}{N_\tau}\sum_{n\in\mathscr{T}_\tau}\left(\frac{r_{n,\tau} - \bar{r}_\tau}{s_\tau + \varepsilon}\right)^3, \tag{B22}$$

and its kurtosis is:

$$\mathrm{Kurt}_\tau = \frac{1}{N_\tau}\sum_{n\in\mathscr{V}_\tau}\left(\frac{r_{n,\tau}-\bar{r}_\tau}{s_\tau+\varepsilon}\right)^4, \tag{B23}$$

where $\mathscr{V}_\tau$ is the validation subset in time window τ, $N_\tau = |\mathscr{V}_\tau|$, $\bar{r}_\tau$ is the mean residual, and s_τ is the residual standard deviation. These diagnostics reveal whether model errors are temporally symmetric or dominated by rare extreme deviations.

In addition to pointwise prediction degradation, the framework evaluates whether removing one source group changes the global distribution of pressure predictions. Let $\mathscr{G}_{\mathrm{ab}}$ denote the set of ablated source groups, including the graph embedding, flow, temporal, camera, and static feature groups. For a source group $g \in \mathscr{G}_{\mathrm{ab}}$, let $\widehat{Y}_i$ be the baseline prediction of the i-th test sample and let $\widehat{Y}_i^{(-g)}$ be the corresponding prediction after replacing the features of group g by their empirical median values. The empirical cumulative distribution functions of the baseline and ablated predictions are defined as:

$$\begin{aligned}\widehat{\mathscr{F}}(z) &= \frac{1}{N_{te}}\sum_{N_{te}}^{i=1} II\{\widehat{Y}_i \le z\},\\ \widehat{\mathscr{F}}^{(-g)}(z) &= \frac{1}{N_{te}}\sum_{N_{te}}^{i=1} II\{\widehat{Y}_i^{(-g)} \le z\},\end{aligned} \tag{B24}$$

where N_{te} is the number of test samples, z is a pressure-threshold variable, and $\mathbb{I}\{\cdot\}$ is the indicator function. The Kolmogorov-Smirnov (KS) distance [96] for source group g is then computed as:

$$D_{\mathrm{KS}}^{(g)} = \sup_{z\in[0,1]} |\widehat{\mathscr{F}}(z) - \widehat{\mathscr{F}}^{(-g)}(z)|. \tag{B25}$$

A larger $D_{\mathrm{KS}}^{(g)}$ indicates that removing source group g causes a stronger distributional distortion in the predicted pressure field. Therefore, this metric complements the mean absolute error and correlation degradation by measuring whether a source group changes the overall pressure-ranking distribution rather than only the samplewise prediction errors.

Data availability

Data will be made available on request.

References

[1] Gyergyay B, Gomari S, Friedrich M, Sonnleitner J, Olstam J, Johansson F. Automation-ready framework for urban transport and road infrastructure planning. Transp Res Procedia 2019;41:88–97. https://doi.org/10.1016/j.trpro.2019.09.018.

[2] Scolamiero V, Boccardo P. A BIM-based digital twin framework for urban roads: integrating MMS and municipal geospatial data for ai-ready urban infrastructure management. Sens 2026;26(3). https://doi.org/10.3390/s26030947.

[3] Hasibuan GCR, Anas MR, Dalimunthe NIP, Fath MTA, Dewi RA, Syafridon GGA, et al. A global perspective on road condition assessment and maintenance: trends in research and technology integration. Transp Res Interdiscip Perspect 2026;36:101853. https://doi.org/10.1016/j.trip.2026.101853.

[4] Chiou S-W. Post-congestion recovery optimization for mixed traffic road networks with uncertain demand. Reliab Eng Syst Saf 2026;275:112787. https://doi.org/10.1016/j.ress.2026.112787.

[5] Yang S, Jia B, Yang K, Tao X, Gao Z, Ji H. Enhancing lifecycle resilience of urban road networks under uncertainty: a two-stage stochastic programming approach for integrated reinforcement and restoration. Reliab Eng Syst Saf 2026;275:112847. https://doi.org/10.1016/j.ress.2026.112847.

[6] Mbiyana K, Kans M, Campos J, Håkansson L. Literature review on gravel road maintenance: current state and directions for future research. Transp Res Rec: J Transp Res Board 2023;2677(5):506–22. https://doi.org/10.1177/03611981221133102.

[7] Wubuli A, Li F, Cao S, Zhang L. Timing of preventive highway maintenance: a study from the whole life cycle perspective. Sustainability 2025;17(3):1009. https://doi.org/10.3390/su17031009.

[8] Feng D. A DEM-based particle–force chain informatics framework for data-driven evaluation of pavement pre-compaction. Adv Eng Inform 2026;71:104402. https://doi.org/10.1016/j.aei.2026.104402.

[9] Ai Q, Xu C, Tong H, Peng C, Wen H, Levin M. Daily resilience of urban road transportation systems: Assessment, enhancement, and recovery. Reliab Eng Syst Saf 2026;275:112788. https://doi.org/10.1016/j.ress.2026.112788.

[10] Wang B, Nguyen H, Lu Q, Zhang Q, Li M. A reinforcement learning approach to network-level maintenance planning of co-located road and water infrastructures with interdependencies. Reliab Eng Syst Saf 2026;267:111889. https://doi.org/10.1016/j.ress.2025.111889.

[11] Intini P, Blasi G, Fracella F, Francone A, Vergallo R, Perrone D. Predicting traffic volumes on road infrastructures in the context of multi-risk assessment frameworks. Int J Disaster Risk Reduct 2025;117:105139. https://doi.org/10.1016/j.ijdrr.2024.105139.

[12] Park JE, Byun W, Kim Y, Ahn H, Shin DK. The impact of automated vehicles on traffic flow and road capacity on urban road networks. J Adv Transp 2021;2021:1–10. https://doi.org/10.1155/2021/8404951.

[13] Azargoshasbi R, Esfeh MA, Kattan L. Urban road network resilience assessment framework: Integrating spatiotemporal analysis with the resilience triangle and temporal performance indicators. Reliab Eng Syst Saf 2026;266:111742. https://doi.org/10.1016/j.ress.2025.111742.

[14] Chen Z, Zheng C, Tao T, Wang Y. Reliability analysis of urban road traffic network under targeted attack strategies considering traffic congestion diffusion. Reliab Eng Syst Saf 2024;248:110171. https://doi.org/10.1016/j.ress.2024.110171.

[15] Cong P, Deng M, Xiao Y, Zhu Y, Zhang X. Trajectory prediction based on the dynamic characteristics and coupling relationships among vehicles in highway scenarios. Eng Appl Artif Intell 2025;140:109718. https://doi.org/10.1016/j.engappai.2024.109718.

[16] Li Y, Cheng Z, Yao X, Kong Z, Wang Z, Liu M. Multi-objective optimal deployment of road traffic monitoring cameras: a case study in wujiang, china. Sustainability 2023;15(15):12011. https://doi.org/10.3390/su151512011.

[17] Monfared M, Rezazadeh M, Alipour Z. Road networks reliability estimations and optimizations: A Bi-directional bottom-up, top-down approach. Reliab Eng Syst Saf 2022;222:108427. https://doi.org/10.1016/j.ress.2022.108427.

[18] Zhang D, Zhang L. Reproducing traffic loads on long-span bridges via fusion of sparse cameras and weigh-in-motion information. Comput-Aided Civ Infrastruct Eng 2026;41:100022. https://doi.org/10.1016/j.cacaie.2026.100022.

[19] Duan J, Li D, Huang H-J. Reliability of the traffic network against cascading failures with individuals acting independently or collectively. Transp Res C: Emerg Technol 2023;147:104017. https://doi.org/10.1016/j.trc.2023.104017.

[20] Gao P, Shuai B, Zhang R, Wang B. STCAGNN-RNKDE: a traffic accident prediction model for spatiotemporal combinatorial attention graph neural networks using Ripley's K and network kernel density estimation. Reliab Eng Syst Saf 2026;265:111593. https://doi.org/10.1016/j.ress.2025.111593.

[21] Li B, Yuan R, Sha X, Zhao Z, Yang Y, Su M. Dynamic graph neural networks for improving the reliability of traffic flow prediction: progress and prospects. Reliab Eng Syst Saf 2026;269:112142. https://doi.org/10.1016/j.ress.2025.112142.

[22] Fernández-Caballero A, Gómez FJ, López-López J. Road-traffic monitoring by knowledge-driven static and dynamic image analysis. Expert Syst Appl 2008;35(3):701–19. https://doi.org/10.1016/j.eswa.2007.07.017.

[23] Li Y, Huang M. Identification of critical road links based on static and dynamic features fusion. Appl Sci 2023;13(10):5994. https://doi.org/10.3390/app13105994.

[24] Lin X, Lu Q, Zhao P, Chen L, Tang J, Guan D, et al. Field-theory inspired physics-informed graph neural network for reliable traffic flow prediction under urban flooding. Reliab Eng Syst Saf 2026;265:111487. https://doi.org/10.1016/j.ress.2025.111487.

[25] Kumar S, Agrawal S, Asaithambi G, Kanagaraj V, Treiber M. Stop-and-go trajectory data of disordered traffic using a swarm of unmanned aerial vehicles. Transp Res Procedia 2025;86:167–74. https://doi.org/10.1016/j.trpro.2025.04.022.
[26] Yang G, Jiang Y, Wang S, Chen K. VinsFusion-line: binocular vision inertial navigation real-time slam system based on line features. Cybern Syst 2026;57(2):350–75. https://doi.org/10.1080/01969722.2025.2606002.
[27] Rastogi T, Simoni MD, Karlström A. Model-based traffic state estimation using camera-equipped probe vehicles. Eur Transp Res Rev 2025;17(1). https://doi.org/10.1186/s12544-025-00761-6.
[28] Qiu J, Fu S, Ou J, Tang K, Qu X, Liang S, et al. Estimating link flow through link speed with sparse flow data sampling. Comput-Aided Civ Infrastruct Eng 2025;40(2):181–97. https://doi.org/10.1111/mice.13323.
[29] Liu Y, Jia C, Rasouli S, Gong J, Feng T, Wong M, et al. SAFER-predictor: sparse adversarial training framework for robust traffic prediction under missing and noisy data. Commun Transp Res 2025;5:100192. https://doi.org/10.1016/j.commtr.2025.100192.
[30] Gu X, Chen J. Progressive degradation-aware distillation for robust indoor object detection. Sci Rep 2026. https://doi.org/10.1038/s41598-026-63984-0.
[31] Xu C, Chen D, Zhang X, Wang Q, Yu J, Wang S, et al. High temperature Nb-Si alloys using data science: optimization of fracture toughness and high-temperature strength. Nat Commun 2026. https://doi.org/10.1038/s41467-026-75353-6.
[32] Feng D. A boundary element method (BEM)-enhanced physics-informed neural network (PINN) framework for thin-walled geocell-reinforced geotechnical foundations. Comp Geotech 2027;201:108588. https://doi.org/10.1016/j.compgeo.2026.108588.
[33] Wang X, Zhang R, Wu L-Y. Performance evaluation of PBFT-based blockchain via a stage-Informed queueing model. J King Saud Univ Comp Info Sci 2026;38:507. https://doi.org/10.1007/s44443-026-00913-1.
[34] Wu J, Wei Z, Liu C, Ma Q, Wei Z, Wu S, et al. ADH-CLSNet: cross-latitude spatiotemporal network for atmospheric duct heights diagnosis. IEEE Transactions on Antennas and Propagation 2026; Early Access. 2026. https://doi.org/10.1109/TAP.2026.3700272.
[35] Feng L, Guo M, Wang W, Chen Y, Shi Q, Guo W, et al. Comparative analysis of machine learning methods and a physical model for shallow landslide risk modeling. Sustainability 2023;15(1):6. https://doi.org/10.3390/su15010006.
[36] Feng D, Guan J, Fu C, Hu Y, Otto F, Hernandez AG, et al. Augmentation of 3D virtual aggregate database using deep convolutional Wasserstein generative adversarial networks. Adv Eng Inform 2026;69:103994. https://doi.org/10.1016/j.aei.2025.103994.
[37] Rajagopal M, Sivasakthivel R, Anitha G, Arunachalam KP, Loganathan K, Abbas M, et al. An efficient intelligent transportation system for traffic flow prediction using meta-temporal hyperbolic quantum graph neural networks. Sci Rep 2025;15(1):27476. https://doi.org/10.1038/s41598-025-10794-5.
[38] Yao S, Li Z, Guan R, Xia M, Hu F, Cao K, et al. GIF-Calib: geometry-intensity fusion for 4D radar-camera calibration in autonomous driving vehicles. Inf Fusion 2026;136:104510. https://doi.org/10.1016/j.inffus.2026.104510.
[39] Yan J, Chen B, Zhang F, Cheng Y, Wang H, Wang H, et al. Meta-learning-based graph convolutional wavelet network for intelligent dynamic modeling of high-speed rail subsystems. IEEE Trans Veh Technol 2026:1–16. https://doi.org/10.1109/TVT.2026.3714895.
[40] Scarselli F, Gori M, Tsoi AC, Hagenbuchner M, Monfardini G. The graph neural network model. IEEE Trans Neural Netw 2009;20(1):61–80. https://doi.org/10.1109/TNN.2008.2005605.
[41] Vrahatis AG, Lazaros K, Kotsiantis S. Graph attention networks: a comprehensive review of methods and applications. Future Internet 2024;16(9):318. https://doi.org/10.3390/fi16090318.
[42] Huang C, Guo T, Yang N, Antoniou C. Augmenting spatio-temporal dependencies for gnn-based short-term traffic flow prediction. Transp Res Procedia 2026;95:592–9. https://doi.org/10.1016/j.trpro.2026.02.075.
[43] Huang H, Lee J-H, Ge Y, Roh S-B, Zhao X. Adaptive spatial–temporal graph attention network for real-time traffic forecasting. Eng Appl Artif Intell 2026;163:112883. https://doi.org/10.1016/j.engappai.2025.112883.
[44] Khaled A, Elsir AMT, Wang P, Shen Y, Zhang Q. A graph-based approach for traffic prediction using similarity and causal relations between nodes. Knowl-Based Syst 2024;296:111913. https://doi.org/10.1016/j.knosys.2024.111913.
[45] Gao J, Luo S, Tian J, Zhou C, An L. Traffic flow prediction for highways based on a multi-task spatiotemporal graph network. Transp Saf Environ 2025;7(1). https://doi.org/10.1093/tse/tdaf003.
[46] Li M, Shi Z, Ma W. Spatiotemporal traffic prediction with missing data: a self-imputation assisted prediction model. Transp Res Interdiscip Perspect 2026;37:101937. https://doi.org/10.1016/j.trip.2026.101937.
[47] Sun T, Zhu S, Hao R, Sun B, Xie J. Traffic missing data imputation: a selective overview of temporal theories and algorithms. Mathematics 2022;10(14):2544. https://doi.org/10.3390/math10142544.
[48] Wu T, Zhu Y, Xiang L, Qin J, Wan Y. A topological model of trajectories with road network space. Trans GIS 2022;26(4):1847–78. https://doi.org/10.1111/tgis.12930.
[49] Bian Y, Liu J, Su X, Tang Y. Construction of a driving route inference model integrating road network topology and traffic dynamics. IJGI 2026;15(2):84. https://doi.org/10.3390/ijgi15020084.
[50] Liu R, Shin S-Y. A review of traffic flow prediction methods in intelligent transportation system construction. Appl Sci 2025;15(7):3866. https://doi.org/10.3390/app15073866.
[51] Emenike CC, Bhero E, Twala B. Optimized predictive modelling for highway traffic flow using real world data. Sci Afr 2026:e03330. https://doi.org/10.1016/j.sciaf.2026.e03330.
[52] Chen J, Shu Q, Lu Y, Zhang Y, Wang Y. QCTF: a quantized communication and transferable fusion framework for multi-agent collaborative perception. IEEE Trans Intell Transp Syst 2025;26(10):15013–27. https://doi.org/10.1109/TITS.2025.3574725.
[53] Rowan D, He H, Hui F, Yasir A, Mohammed Q. A systematic review of machine learning-based microscopic traffic flow models and simulations. Commun Transp Res 2025;5:100164. https://doi.org/10.1016/j.commtr.2025.100164.
[54] Zhang H, Dong H, TSGDiff Yang Z. Traffic state generative diffusion model using multi-source information fusion. Transp Res C: Emerg Technol 2025;174:105081. https://doi.org/10.1016/j.trc.2025.105081.
[55] Gu J, Jia Z, Cai T, Song X, Mahmood A. Dynamic correlation adjacency-matrix-based graph neural networks for traffic flow prediction. Sens 2023;23(6). https://doi.org/10.3390/s23062897.
[56] Khairy M, Mokhtar HM, Abdalla M. Adaptive traffic prediction model using graph neural networks optimized by reinforcement learning. Int J Cogn Comput Eng 2025;6:431–40. https://doi.org/10.1016/j.ijcce.2025.02.001.
[57] Wang C, Tian R, Hu J, Ma Z. A trend graph attention network for traffic prediction. Inf Sci 2023;623:275–92. https://doi.org/10.1016/j.ins.2022.12.048.
[58] Komarovsky S, Haddad J. Spatio-temporal graph convolutional neural network for traffic signal prediction in large-scale urban networks. Transp Res Interdiscip Perspect 2025;32:101482. https://doi.org/10.1016/j.trip.2025.101482.
[59] Zhu J, Yuan J, Ye F, Nguyen T-T, Wang R, Zeng W, et al. Traffic prediction using an active causality recurrent graph convolutional network. Expert Syst Appl 2026;298:129506. https://doi.org/10.1016/j.eswa.2025.129506.
[60] Ma W, Chu Z, Chen H, Li M. Spatio-temporal envolutional graph neural network for traffic flow prediction in UAV-based urban traffic monitoring system. Sci Rep 2024;14(1):26800. https://doi.org/10.1038/s41598-024-78335-0.
[61] Tong Q, Yang Z, Chai G, Wang Y, Qi Z, Wang F, et al. Driving state evaluation of intelligent connected vehicles based on centralized multi-source vehicle road collaborative information fusion. Env Dev Sustain 2025;27(10):24107–26. https://doi.org/10.1007/s10668-021-01957-1.
[62] Yang H, Hu K, Wang H, Hong W, Wang X, Wang H, et al. BCLIP-ADer: a bayesian prompt contrastive language-image pretraining method for catenary component anomaly detection in electrified railways. IEEE Trans Transp Electrif 2026;1. https://doi.org/10.1109/TTE.2026.3689522.
[63] Duan F, Wang H, Yang H, Wei C, Zhang C, Song Y, et al. MRFM-IFCOS: an anchor-free interactive detector based on multireceptive field mamba for detecting catenary support components. IEEE Trans Instrum Meas 2025;74:1–15. https://doi.org/10.1109/TIM.2025.3637977.
[64] Feng Y, Zhao Y, Zhang X, Batista SFA, Demiris Y, Angeloudis P. Predicting spatio-temporal traffic flow: a comprehensive end-to-end approach from surveillance cameras. Transp B: Transp Dyn 2024;12(1). https://doi.org/10.1080/21680566.2024.2380915.
[65] Yang B, Liu Z, Cai Z, Li D, Su X, Guo L, et al. A novel traffic flow reduction method based on incomplete vehicle history spatio-temporal trajectory data. IJGI 2022;11(3):209. https://doi.org/10.3390/ijgi11030209.
[66] Shafiei S, Zhu Q, Qin AK, Mihaita SA, Grzybowska H, Dia H. Road traffic attributes prediction using deep learning hybridization by the traffic fundamental diagram. J Intell Transp Syst 2025:1–16. https://doi.org/10.1080/15472450.2025.2577735.
[67] Chen J, Ren C, Hu Y, Zhang Y, Lu Y, Li Q, et al. Dual-centralized Q -network-based reinforcement learning for cooperative path planning of multiple UAVs. IEEE Trans Intell Transp Syst 2025;26(9):13232–46. https://doi.org/10.1109/TITS.2025.3587392.
[68] Wang B, Wang J. ST-MGAT:spatio-temporal multi-head graph attention network for Traffic prediction. Phys A: Stat Mech Appl 2022;603:127762. https://doi.org/10.1016/j.physa.2022.127762.
[69] Remmouche B, Boukraa D, Zakharova A, Bouwmans T, Taffar M. Long-term spatio-temporal graph attention network for traffic forecasting. Expert Syst Appl 2025;288:128244. https://doi.org/10.1016/j.eswa.2025.128244.
[70] Xu X, Fei L, Zhao Y, Lü X. Optimal graph information fused graph attention network for traffic flow forecasting. J Adv Transp 2025;2025(1). https://doi.org/10.1155/atr/5195875.
[71] Chen Y, Huang J, Xu H, Guo J, Su L. Road traffic flow prediction based on dynamic spatiotemporal graph attention network. Sci Rep 2023;13(1):14729. https://doi.org/10.1038/s41598-023-41932-6.
[72] Wu X, Dong J, Bao W, Zou B, Wang L, Wang H. Augmented intelligence of things for emergency vehicle secure trajectory prediction and task offloading. IEEE Internet Things J 2024;11(22):36030–43. https://doi.org/10.1109/JIOT.2024.3417285.
[73] Yasuda S, Katayama H, Nakanishi W, Iryo T. Trajectory data-driven network representation for traffic state prediction using deep learning. Int J ITS Res 2024;22(1):136–45. https://doi.org/10.1007/s13177-023-00383-z.
[74] Liu Y, Ge Q, Luo W, Huang Q, Zou L, Wang H, et al. GraphMM: graph-based vehicular map matching by leveraging trajectory and road correlations. IEEE Trans Knowl Data Eng 2024;36(1):184–98. https://doi.org/10.1109/TKDE.2023.3287739.
[75] Tang W, Chen J, Sun C, Wang H, Li G. Traffic demand estimations considering route trajectory reconstruction in congested networks. Algorithms 2022;15(9):307. https://doi.org/10.3390/a15090307.
[76] Jiang Y, Li X. Travel time prediction based on historical trajectory data. Ann GIS 2013;19(1):27–35. https://doi.org/10.1080/19475683.2012.758173.
[77] Liu J, Yuan Y. Exploring dynamic urban mobility patterns from traffic flow data using community detection. Ann GIS 2024;30(4):435–54. https://doi.org/10.1080/19475683.2024.2324393.

[78] Liljenström C, Björklund A, Toller S. Including maintenance in life cycle assessment of road and rail infrastructure—a literature review. Int J Life Cycle Assess 2022;27(2):316–41. https://doi.org/10.1007/s11367-021-02012-x.
[79] Daulat S, Rokstad MM, Klein-Paste A, Langeveld J, Tscheikner-Gratl F. Challenges of integrated multi-infrastructure asset management: a review of pavement, sewer, and water distribution networks. Struct Infrastruct Eng 2024;20(4):546–65. https://doi.org/10.1080/15732479.2022.2119480.
[80] Zhang Y, Ma H, Kwan M-P. How drivers' depth perception of environmental features influences traffic speed. Geo-spat Inf Sci 2025:1–17. https://doi.org/10.1080/10095020.2025.2584862.
[81] Da G, Wang C, Peng R. Reliability of wireless sensor networks with probabilistic functional dependence: a lifetime-based approach. Reliab Eng Syst Saf 2027;277:113122. https://doi.org/10.1016/j.ress.2026.113122.
[82] Törő O, Esztergár-Kiss D, Schmöcker J-D, Tettamanti T. Estimating traffic from crowdsourced location activity: a new data source perspective. Transp Res Interdiscip Perspect 2026;37:101973. https://doi.org/10.1016/j.trip.2026.101973.
[83] Xiao DX, Wu Z. Using systematic indices to relate traffic load spectra to pavement performance. Int J Pavement Res Technol 2016;9(4):302–12. https://doi.org/10.1016/j.ijprt.2016.07.005.
[84] Woodard D, Nogin G, Koch P, Racz D, Goldszmidt M, Horvitz E. Predicting travel time reliability using mobile phone GPS data. Transp Res C: Emerg Technol 2017;75:30–44. https://doi.org/10.1016/j.trc.2016.10.011.
[85] Wei L, Chen P, Mei Y, Wang Y. Turn-level network traffic bottleneck identification using vehicle trajectory data. Transp Res C: Emerg Technol 2022;140:103707. https://doi.org/10.1016/j.trc.2022.103707.
[86] Veličković Petar, Cucurull Guillem, Casanova Arantxa, Romero Adriana, Liò Pietro, Bengio Yoshua. Graph attention networks. Int Conf Learn Represent 2018.
[87] Ying C, Cai T, Luo S, Zheng S, Ke G, He Di, et al. Do transformers really perform badly for graph representation? Adv Neural Inf Process Syst 2021;34:28877–88.
[88] Bridle JS. Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition. Neurocomputing. Berlin, Heidelberg: Springer; 1990. p. 227–36.
[89] Kingma Diederik P, Ba Jimmy. Adam: a method for stochastic optimization. In: International Conference on Learning Representations; 2014.
[90] Yu F, Yan H, Chen R, Zhang G, Liu Y, Chen M, et al. City-scale vehicle trajectory data from traffic camera videos. Sci Data 2023;10(1):711. https://doi.org/10.1038/s41597-023-02589-y.
[91] Hamilton WL, Ying R, Leskovec J. Inductive representation learning on large graphs. In: NIPS'17: Proceedings of the 31st International Conference on Neural Information Processing System; 2017. p. 1025–35. https://doi.org/10.48550/arXiv.1706.02216.
[92] Brody Shaked, Alon Uri, Yahav Eran. How attentive are graph attention networks?. In: International Conference on Learning Representations; 2021.
[93] Li Yaguang, Yu Rose, Shahabi C, Liu Yan. Diffusion convolutional recurrent neural network: data-driven traffic forecasting. Int Conf Learn Represent 2017.
[94] Pearson KVII. Mathematical contributions to the theory of evolution.—III. regression, heredity, and panmixia. Philos Trans Math Phys Eng Sci 1896;(187):253–318. https://doi.org/10.1098/rsta.1896.0007.
[95] GU Yule. On the methods of measuring association between two attributes. J R Stat Soc 1912;75(6):579. https://doi.org/10.2307/2340126.
[96] Smirnov N. On the estimation of discrepancy between empirical curves of distribution for two independent samples. Bull Math L'Université Mosc 1939;2(2):3–11.