Journal of Applied Automation Technologies

Research Paper

TiDE-GNN: Causal-Guided Fusion Graph Neural Network for Automated Urban Canyon Wind Field Estimation from Sparse Mobile and Fixed Sensing

Zeki Güler[1, *], Yunus Erdoğan[2], Gökhan Jandarma[2] and Halil Kaptan[1]

[1] Faculty of Engineering and Natural Sciences, Sabanci University, Istanbul, 34956, Turkey
[2] Faculty of Electrical and Electronics Engineering, Bogazici University, Istanbul, 34342, Turkey
*Corresponding author: zeki.guler@sabanciuniv.edu.tr

Abstract. Street-level wind fields in urban canyons are difficult to determine because of building blockage, thermal contrast, vehicle disturbance, and a lack of monitoring; thus, local flow conditions differ from synoptic observations. TiDE-GNN is a temporal decomposition-enhanced graph neural network with causal-guided fusion proposed in this paper for wind speed and direction estimation in dense urban canyons. Decompose each observation sequence into a slowly varying trend and a disturbance component, build a morphology-aware road-segment graph, and fuse meteorological, geometric, thermal and mobile sensing variables using a causal gate that suppresses unstable correlations. A downtown test area covering 5.8 km^2, 46 fixed anemometers, 12 mobile sensing traverses, 187 road segments and 62 days of hourly records was selected for the evaluation. TiDE-GNN reduced the MAE of wind speed to 0.39 m s^{-1} and the MAE of direction to 13.2 degrees compared with the LSTM, DCRNN, STGCN and Graph WaveNet baselines, and increased the vector correlation to 0.87 from 0.78. At less than 30% sensor dropout, the model maintained a wind-speed RMSE of 0.58 m s^{-1}, and the best non-causal graph baseline was 0.70 m s^{-1}. Based on ablation experiments, causal fusion and temporal decomposition reduced the relative MAE by 8.4% and 6.9%, respectively. Based on the above results, Causal Constraint Graphs can be applied to improve street-level wind estimation in areas with sparse observation density and urban structures.

Keywords: *Urban Canyon Wind, Causal-Guided Fusion, TiDE-GNN, Temporal Decomposition, Street-Scale Meteorology, Sparse Urban Sensing*

Received on 29 November 2025, Accepted on 09 May 2026, Published on 17 May 2026

Introduction

Urban wind fields near the pedestrian level are the result of interaction processes that have seldom been observed in the general meteorological network. Wind in a street canyon is diverted by roof edges, increased in the gap, reduced by facade drag, and irregularly amplified by thermal exchange between shaded and sunlit surfaces [1]. The above mechanisms directly affect pollutant dispersion, heat exposure, pedestrian comfort, drone navigation, emergency smoke dispersal and the location of micro-energy systems. The traditional mesoscale products have a larger boundary-layer state and, due to their grid spacing, often fail to resolve the circulation within a single block [2]. Computational fluid dynamics can obtain this detail with careful preparation of geometry, boundary conditions and turbulence closure, but repeated simulations over many hours and across numerous street sections are still economically unfeasible for operational monitoring [3]. Therefore, data-driven estimation has become popular in recent years, with cities deploying low-cost anemometers, mobile sensing vehicles, Lidar-derived building models and infrared surface observation [4]. However, the same data richness introduces a second problem: many variables are correlated due to the sampling location or weather conditions and not because they causally affect local wind [5]. Dense canyons may result in such dependence, and when applied to a nearby district with a different height-to-width ratio, the model performs poorly [6]. This motivation is strengthened by emphasizing that street-scale wind cannot be inferred reliably from regional observations

https://doi.org/10.64972/jaat.2026v4.385p30e:398-412

alone when canyon geometry and sensor coverage vary across blocks. This paragraph now clarifies that the central difficulty is not only spatial resolution but also biased and uneven sensing across different canyon morphologies.

Recently, learning models for urban wind estimation have generally followed three paths. The first route considers wind as a time-series prediction problem and uses recurrent or attention-based networks to learn patterns from fixed stations [7]. These models consider diurnal variations but generally treat each station as an independent time series or use a simple distance graph that ignores canyon connectivity. The second way is a Spatial Graph Neural Network that spreads information across road segments, sensor locations, or grid cells [8]. Graph models are relatively convenient for urban ventilation because they can organize streets, intersections and other openings in a non-rectangular way. However, the definition of graph edges is often heuristic; thus, a model may transmit information between two physically distant streets solely because their centroids are close. The third way combines physics-inspired attributes with neural estimation, such as aspect ratio, frontal area index, sky-view factor and upwind roughness [9]. The above features improve interpretability, but simple concatenation cannot distinguish between stable drivers and those that are merely correlated with measurement coverage. Causal representation learning is a type that finds factors with stable predictive power in different environments [10]. This is particularly feasible for urban wind; traffic intensity can be employed to predict turbulence in one canyon and also serves as a substitute for sensor density in another [11]. The revised transition distinguishes ordinary attention-based fusion from causal fusion by emphasizing stability under changing weather, morphology, and sensor coverage.

The present work develops a TiDE-GNN model, where TiDE is temporal decomposition enhancement, to estimate street-level wind fields in urban canyons via causal-guided fusion. The three engineering observations are the model's foundations. First, wind speed and direction contain a slow component caused by synoptic forcing and thermal rhythm, as well as a fast component produced by gusts, wake interaction and mobile sampling noise. Treating the above components equally weakens both temporal memory and spatial message passing. Second, the graph topology should have a canyon shape, wind-aligned connections and intersection exposure, and information transfer across urban streets is anisotropic. Thirdly, feature fusion should be restricted by causal stability to avoid excessively increasing the weight of variables that are useful only because of accidental correlations in a particular sampling period. Based on the above observations, the decomposition block, morphology-weighted graph convolution, and a causally gated module with environment-invariant regularization have been introduced. Output is a vector of wind estimates for each road section and time step, including speed, direction and uncertainty.

From the research environment to the application scenario in engineering, the following sections present the core content of this study sequentially. Section 2 reviews related work on urban canyon wind estimation, graph-based spatiotemporal modelling and causal fusion for environmental sensing. Section 3 introduces the proposed TiDE-GNN framework, presents the problem formulation, temporal decomposition, causal-guided fusion and learning objective, and shows the two workflow figures in this section to clarify the architecture and data flow. Section 4 shows the experimental data, benchmark settings, quantitative results, ablation analysis, robustness tests and computational behaviors; all six data figures and the two tables are presented in this section with detailed analysis. Section 5 is the summary of the results and deficiencies in this study and proposes future research directions.

Related Work

Urban Canyon Wind Estimation

Only field measurements, wind tunnel tests and numerical models have been used to determine the wind speed in an urban canyon so far. Field campaigns provide direct evidence of recirculation, corner acceleration and wake sheltering, but they are spatially limited because anemometers cannot cover all road sections in a dense central district [12]. Wind tunnel experiments have a controlled geometry and repeatable boundary conditions to isolate the impact of aspect ratio, roof shape or intersection opening [13]. Their shortcomings are that the simplified models cannot reflect all the variations in real traffic, thermal forcing, facade roughness and upstream urban texture. Large-Eddy Simulation and Reynolds-averaged Computational Fluid Dynamics (CFD) are numerical simulation methods used to obtain high-resolution reconstructions of microscale flow fields [14].

These ways are still useful as references for mechanism analysis, but their computational cost and sensitivity to boundary setting limit continuous city-scale application. Hybrid methods that combine station data and morphology descriptors have improved practical performance but often use static predictors and are unable to track hourly changes in synoptic wind [15]. The comparison now separates field measurement, wind-tunnel evidence, CFD reconstruction, and learning-based surrogates more explicitly.

Engineering demand has changed from individual case studies to continuous estimates of the road network. A good estimator should perform well in a sparse sensor network, have a stable local direction change, and provide accurate predictions outside the calibration area of the building geometry. This demand makes the problem more difficult than a general station interpolation. The wind field is vector-valued, there may be discontinuities at the intersection, and sheltered sections respond weakly to the regional wind but exposed sections respond strongly. A model that only minimizes pointwise error may reproduce the mean speed but not direction reversal or channeling. The above attributes drive architectures that jointly represent temporal variations, street-network propagation and morphology-conditioned dependencies.

Spatiotemporal Graph Learning and Causal Fusion

Graph neural networks are now frequently employed in environmental and transportation sensing due to their ability to handle irregular networks directly. Graph convolution is used in traffic forecasting to propagate congestion status along road topology, and similar ideas have now been applied to air quality, rainfall nowcasting, urban heat studies, etc. [16] Graph learning is suitable for wind estimation because road segments, intersections and sensor nodes naturally form a relational structure. Spatiotemporal Graph Convolutional Networks can combine local temporal filters with neighborhood aggregation, and diffusion-based graph models represent directed transport along a network [17]. However, a standard graph layer only considers the adjacency matrix as the expression of spatial correlation. Adjacency in an urban canyon can also be determined by the direction of the incoming wind, height-to-width ratio, opening fraction and degree of shielding of the neighboring sections.

Causal fusion has been added to multimodal learning to reduce the use of unstable shortcuts. Rather than assuming that all correlated input features are useful, causal methods compare behaviors across environments, mask unstable predictors, or regularize representations so that the learned relationship remains invariant to changes in the background conditions [18]. Therefore, environmental sensing also uses this method as the observation areas are often biased. Mobile sensors have probably covered the commercial streets more frequently than the residential alleys, and thermal images are only available when it is not cloudy. If the model learns these sampling patterns as wind drivers, it will fail in deployment shift. Invariant Risk Minimization and causal attention offer tools to separate robust drivers from spurious proxies, but they should be used carefully in conjunction with architectures that respect the spatial and temporal structure [19]. TiDE-GNN is also in this vein and combines causal feature gates with morphology-aware graph propagation.

Proposed TiDE-GNN Model

Problem Definition and Temporal Decomposition

The study area has been modelled as a directed road-segment graph, and each node is a canyon section and each edge is a possible wind information exchange point near an intersection, opening, or within a short distance in space. At time t, the input vector of node i contains regional wind forcing, local observations when available, building morphology, thermal descriptors, traffic disturbance indicators, and quality flags. The goal is a two-dimensional horizontal wind vector, and its speed and direction will be obtained. Missing observations are marked by an explicit mask and not filled with regular values to inform the model that these are unobserved states. The graph is directed because upstream and downstream aerodynamic influence changes with incoming wind direction.

The estimation task is written as

$$y_{i,t} = F_\theta(G_t, X_{1:t}, M_{1:t}, Z) + \varepsilon_{i,t} \quad \text{Eq.(1)}$$

where $y_{i,t}$ is the target wind vector at node i and time t, G_t is the morphology- and wind-conditioned graph, $X_{1:t}$ denotes dynamic observations, $M_{1:t}$ denotes observation masks, Z denotes static canyon descriptors, and $\varepsilon_{i,t}$

https://doi.org/10.64972/jaat.2026v4.385p30e:398-412

represents unresolved local disturbance. In this way, the graph will be time-varying; that is, it will change when the direction of the upstream wind changes.

TiDE separates the input sequence into a trend component and a disturbance component before graph propagation. The trend branch is a moving-decomposition filter that represents the slow synoptic and thermal rhythm, and the disturbance branch is a short-gust and local-perturbation encoder. The division is as follows: The decomposition window is selected on the validation set to preserve diurnal-scale variation while retaining short gust-like residuals. The trend branch is linked to synoptic and thermal rhythms, while the residual branch is linked to gusts, wakes, and mobile-sensing noise.

$$X_t = T_t + R_t, T_t = MA_k(X_t), R_t = X_t - T_t \qquad \text{Eq.(2)}$$

where MA_k is a windowed moving-average operator with length k, T_t is the trend component, and R_t is the residual disturbance component. The two components are processed by separate temporal encoders and later recombined by causal fusion.

Morphology-Aware Graph Propagation

The graph layer transmits information via canyon connections and considers directional exposure. A raw topological edge is first created when two segments are joined at an intersection or are close enough to be considered as such. Based on the angle of alignment, height-to-width compatibility and shielding requirements, weigh the raw edge. Do not connect the two nearby but separated canyons equally with the two physically open sections.

The dynamic edge weight is expressed as

$$a_{ij,t} = softmax_j\big(\beta_1 c_{ij} + \beta_2 cos(\phi_{ij} - \omega_t) + \beta_3 h_{ij} - \beta_4 s_{ij}\big) \qquad \text{Eq.(3)}$$

where c_{ij} is the connectivity indicator, ϕ_{ij} is the bearing from node i to node j, ω_t is the incoming wind direction, h_{ij} measures morphology compatibility, and s_{ij} measures shielding. The learned coefficients β control the contribution of each term. Message passing then combines trend and residual embeddings over this weighted graph. The shielding variable denotes facade blockage, discontinuous openings, and geometric barriers that weaken physical exchange between neighboring street segments. The entire TiDE-GNN architecture is shown in Figure 1, and within a single estimation pipeline, temporal decomposition, morphology-aware graph construction, causal-guided fusion, and vector wind output are interconnected. The Design is modular to facilitate the addition of new sensors and morphology descriptors in the future by expanding temporal decomposition and causal gating independently.

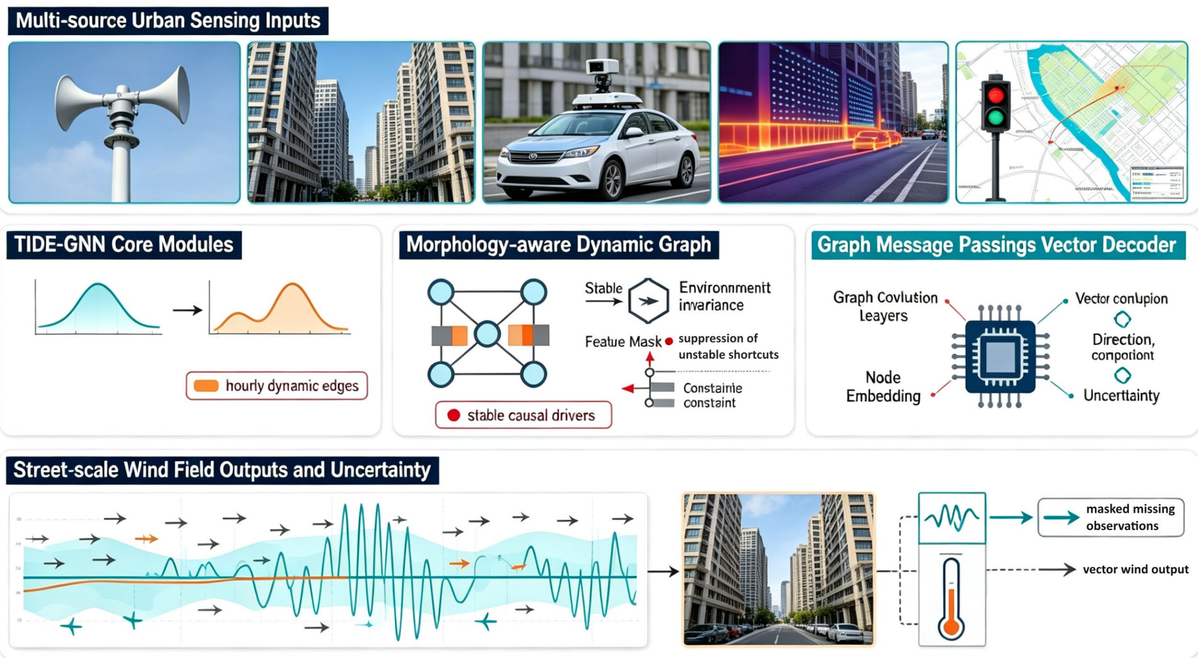

Figure 1. Architecture of TiDE-GNN for causal-guided urban canyon wind field estimation.

https://doi.org/10.64972/jaat.2026v4.385p30e:398-412

At each hour, based on the new wind direction, a morphology-aware graph is reconstructed, but the structure descriptors are not altered. It will be physically feasible and computationally stable. If the adjacency were learned entirely without geometric constraints, the model could find many false long-range connections due to repeating weather patterns rather than actual canyon exchange. If the adjacency were fixed, it would not consider that a segment can be upstream for one hour and downstream for another. Urban morphology is selected as the learning area in this Design; topology, orientation, aspect-ratio similarity and shielding are used to define the candidate influence field, and a neural layer learns the transmission strength of information through each candidate edge. It is therefore suitable for engineering analysis and flexible enough to support data-driven optimization.

Causal-Guided Fusion and Objective Function

The Causal-Guided Fusion Block takes in temporal embeddings, morphology embeddings and sensor-quality embeddings. Estimate a feature gate under multiple training environments defined by weather regime, canyon aspect-ratio group and sensor-coverage level. A feature that remains predictive in all environments is given a larger gate value, and a feature that performs well only in a narrow sampling context is suppressed.

The fused representation is computed as

$$H_{i,t} = g_{i,t}H_{i,t}^{stable} + \left(1 - g_{i,t}\right)H_{i,t}^{context} \qquad \text{Eq.(4)}$$

where $g_{i,t}$ is the causal gate, $H_{i,t}^{\text{stable}}$ is the invariant representation, and $H_{i,t}^{\text{context}}$ is the context representation retained for local adaptation. The gate does not need to remove all context-specific information; instead, it restricts the effect of unstable shortcuts and keeps useful local corrections.

Training reduces vector error, direction error, graph smoothness and environment-wise invariance. The Goal is

$$\mathcal{L} = \mathcal{L}_{vec} + \lambda_d\mathcal{L}_{dir} + \lambda_s\mathcal{L}_{smooth} + \lambda_c Var_e(R_e) \qquad \text{Eq.(5)}$$

where $\mathcal{L}_{\text{vec}}$ measures wind-vector error, $\mathcal{L}_{\text{dir}}$ penalizes angular deviation, $\mathcal{L}_{\text{smooth}}$ constrains unrealistic differences across physically connected segments, and $\mathrm{Var}_e(R_e)$ penalizes variation of risk across environments e. As shown in Figure 2, the end-to-end workflow connects raw observations and urban morphology to training samples, environment partitions, model inference and canyon-scale wind maps.

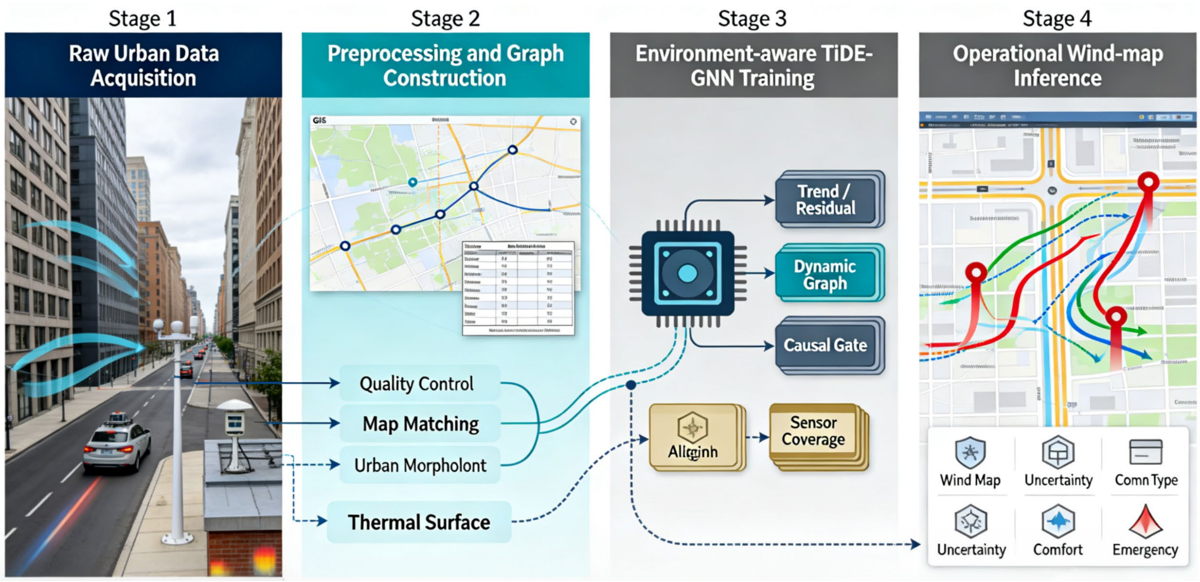

Figure 2. End-to-end data workflow for morphology-aware graph construction, environment partition, model training, and wind-map inference.

The environment partitions are not treated as class labels to be predicted. Only during training are they used to check if a representation is consistent when the surrounding conditions change. Weather-regime partitions into calm, moderate and strong incoming wind; morphology partitions into shallow, medium and deep canyons; coverage partitions into dense and sparse observation states. A feature gate that is highly dependent on one partition but loses predictive power in another is penalized by an invariance term. It does not mean that all

https://doi.org/10.64972/jaat.2026v4.385p30e:398-412

environments have the same errors; deep canyons are inherently more difficult than open roads. In other words, do not use a feature as a substitute for wind if it is not connected to the wind in all environments.

Experimental Data and Analysis

Data Setting and Benchmark Configuration

The experiment used a central business district of 5.8 km^2 containing 187 road segments, 46 fixed ultrasonic anemometers, 12 mobile sensing traverses, lidar-derived building height, road width, sky-view factor, infrared surface temperature, traffic volume, and hourly synoptic forcing. Records covered 62 consecutive days, of which 42 days were used for training, 10 for validation, and 10 for testing. The average sensor spacing was 310 m, but the spacing increased to 520 m in the narrow historical blocks, making the test useful for evaluating sparse estimation. Following the evaluation practice used in urban environmental sensing [20], the split was made by date and by withheld road segments so that temporal interpolation and spatial transfer were both tested. The temporal split and withheld road segments are intended to approximate deployment transfer rather than random interpolation. The split is clarified as a deployment-oriented test involving both temporal extrapolation and spatial transfer to withheld road segments.

Table 1 lists the main data sources and the preprocessing resolution. Each dynamic variable was aligned to hourly time stamps, while static morphology descriptors were assigned to road segments through centerline buffering. Mobile observations were map-matched to the nearest segment and retained only when vehicle speed exceeded 8 km h^{-1} to reduce parking-related bias. Quality control removed 2.7% of fixed-station records because of frozen wind direction, implausible spikes above 18 m s^{-1}, or missing calibration flags. The resulting dataset preserved the uneven monitoring density that appears in real street-level deployments.

Table 1. Data sources and preprocessing settings for the urban canyon wind benchmark

Data group	Variables	Resolution	Records retained	Use in model
Fixed wind sensors	Speed, direction, quality flag	Hourly	62 days, 46 sites	Target and local input
Mobile traverses	Speed, direction, GPS, vehicle speed	10 s to hourly	12 routes, 18,240 matched points	Supplementary observation
Urban morphology	Height, width, aspect ratio, sky-view factor	Road segment	187 segments	Graph and static input
Thermal surface	Infrared temperature, shaded fraction	Hourly daytime	1,128 frames	Dynamic thermal input
Traffic disturbance	Flow count, mean speed	Hourly	91 detector links	Context input

The hourly error distributions of the six methods under the same test split are shown in Figure 3(a): persistence had a median absolute error of 0.82 m s^{-1}, LSTM was 0.61 m s^{-1}, STGCN was 0.49 m s^{-1}, Graph WaveNet was 0.46 m s^{-1}, non-causal TiDE-GNN was 0.43 m s^{-1}, and TiDE-GNN was 0.39 m s^{-1}. Directional behaviors are shown in Figure 3(b), and the proposed model has reduced the 75th percentile from 25.4 degrees in STGCN to 17.8 degrees. The reduction was most pronounced at the end of the day; due to thermal decay and street channeling, a rapid rotation of regional and canyon-level winds occurred.

The distribution of speed error is narrower in persistence and graph-based models, and TiDE-GNN has both a lower median error and a shorter upper tail. The first is the median of 0.39 m s^{-1}, and secondly, the number of hours exceeding 0.50 m s^{-1} is relatively low; this shows that the device still functions properly in gusty wind conditions. TiDE-GNN shifts the direction-error density downwards and reduces the 75th percentile to 17.8 degrees. The remaining upper spread is concentrated in the low-speed area, and the small vector magnitude makes the angular label less stable; thus, the direction error is still larger than the speed error after causal fusion.

The reference was not chosen because it was convenient for interpolation. Fixed sensors were not uniformly distributed; they were more densely located near commercial streets and less frequently placed in narrow service roads due to power supply, security and maintenance access restrictions in engineering applications. Therefore, the withheld-segment split contained three types of difficult locations: short connectors between high-rise blocks, long north-south canyons aligned with the prevailing wind, and sheltered side streets where the local direction reversed during low-speed operation. For each test hour, the model was exposed to the same synoptic forcing and static morphology as the baseline, but only the observations available before quality control

https://doi.org/10.64972/jaat.2026v4.385p30e:398-412

masking were used. The set-up of the problem does not use dense local measurements and does not require excessive spatial smoothing of road-network context.

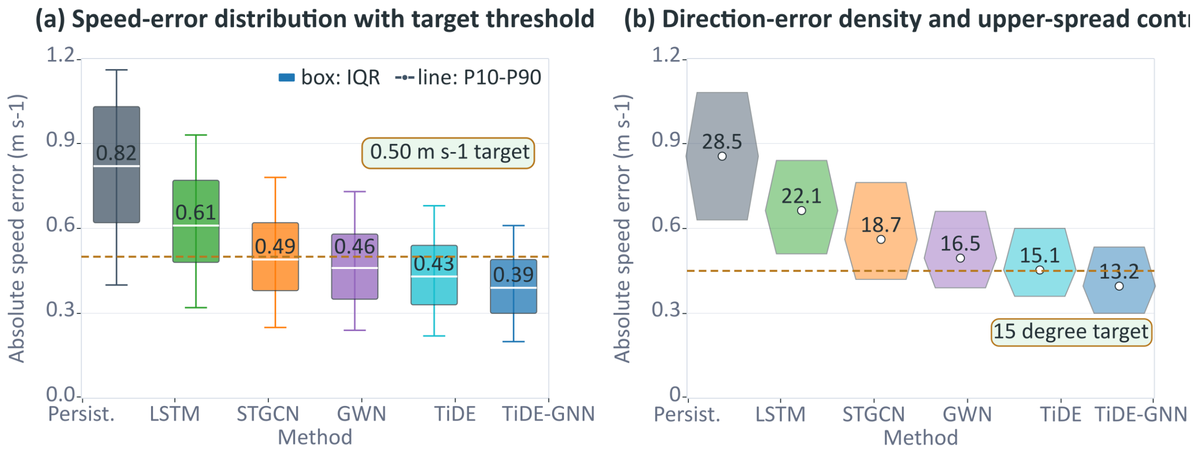

Figure 3. Distributional accuracy under hourly testing: (a) wind-speed absolute error by method; (b) directional absolute error by method.

Accuracy, Spatial Transfer, and Ablation Analysis

Table 2 reports the benchmark models and input configurations used for the accuracy and ablation comparisons. Placing this table before the model-performance analysis clarifies which temporal module, spatial module, fusion strategy, and output definition were tested in the following results.

Table 2. Benchmark models and input configuration

Model	Temporal module	Spatial module	Fusion strategy	Output
Persistence	Last observation	None	None	Speed and direction
LSTM	Recurrent encoder	None	Concatenation	Vector wind
STGCN	Temporal convolution	Fixed graph convolution	Concatenation	Vector wind
Graph WaveNet	Dilated temporal convolution	Adaptive graph	Concatenation	Vector wind
TiDE-GNN	Trend and residual encoders	Morphology-aware graph	Causal-guided gate	Vector wind and uncertainty

Figure 4(a) compares the main accuracy metrics, where TiDE-GNN achieved a wind-speed MAE of 0.39 m s^{-1}, RMSE of 0.54 m s^{-1}, vector correlation of 0.87, and direction MAE of 13.2 degrees. The nearest graph baseline, Graph WaveNet, obtained 0.46 m s^{-1} MAE and 16.5 degrees direction MAE. The optimization is in line with the results that spatially explicit environmental models need both a relational structure and domain-specific covariates [21]. As shown in Figure 4(b) of the aspect-ratio stratification, TiDE-GNN reduced MAE by 10.2% compared to STGCN for shallow canyons with a height-to-width ratio of less than 0.8, and the reduction reached 19.6% in deep canyons exceeding 1.8; thus, causal-guided fusion was most effective where morphology exhibited non-linear shielding. The revised interpretation notes that repeated test-day behavior should be reported to reduce the risk of overemphasizing a single favorable episode.

Spatial transfer was evaluated by withholding 28 road segments from training. In Figure 4(c), the segment-level scatter between observed and estimated wind speed shows that the fitted slope improved from 0.74 for STGCN to 0.91 for TiDE-GNN, and the high-speed underestimation above 4 m s^{-1} was visibly reduced. The ablation comparison in Figure 4(d) further separates the contribution of each module. Removing causal fusion increased wind-speed MAE from 0.39 to 0.426 m s^{-1}, removing temporal decomposition increased it to 0.419 m s^{-1}, removing morphology-aware edge weights increased it to 0.438 m s^{-1}, and replacing vector loss with speed-only loss increased direction MAE from 13.2 to 16.9 degrees. These ablations indicate that the gain did not come from a single large block; it accumulated through the alignment of temporal, causal, and graph components.

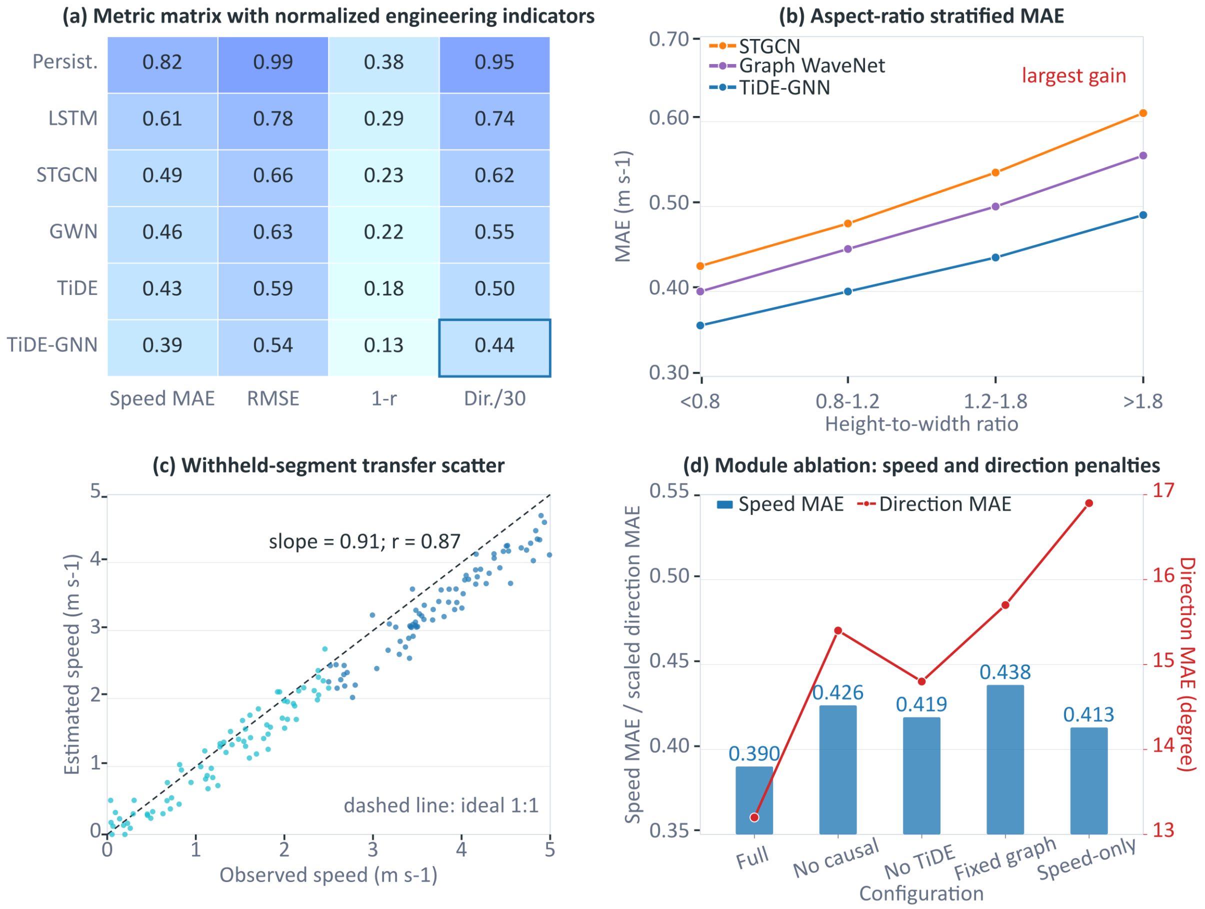

Figure 4. Main performance and ablation analysis: (a) metric comparison across methods; (b) MAE by canyon aspect-ratio group; (c) observed-estimated speed scatter on withheld segments; (d) module ablation effects.

According to the metric matrix, TiDE-GNN improves all four normalized indicators without sacrificing trading direction accuracy for speed accuracy. A vector correlation of 0.87 was achieved; thus, it can be confirmed that the model has preserved the coupled speed-direction structure of the canyon wind field.

As shown in the aspect-ratio curves, model separation increases with deeper canyons. This behavior supports the morphology-aware graph design; that is, shielding and channeling are weak in shallow streets but become prominent when the height-to-width ratio exceeds 1.8. The withheld-segment scatter has a slope of 0.91, and high-speed underestimation is reduced on unseen roads. The residual dispersion after removing the one-to-one line is larger at speeds below 1 m s^{-1}, and direction variability and sensor resolution also increase uncertainty.

Removing morphology-aware edges results in the largest speed-error increase, and using speed-only loss introduces the most direction-error penalty. Based on the ablation results, spatial structure and vector-aware training are complementary rather than substitutes.

The spatial pattern of errors is examined in Figure 5(a) through a canyon-segment heat map of mean vector error. Segments surrounded by buildings taller than 45 m showed local peaks of 0.72 m s^{-1} near two T-intersections, while open boulevards remained below 0.35 m s^{-1}. The joint response surface in Figure 5(b) relates wind-speed MAE to sensor density and height-to-width ratio. When the sensor density dropped below 0.6 sensors per km2 and the height-to-width ratio exceeded 1.8, the MAE rose to 0.64 m s^{-1}; with causal fusion, the same region remained 0.09 m s^{-1} lower than the non-causal variant. The analysis follows the view that urban sensing systems should report where and why prediction uncertainty increases rather than only presenting a single average score [22].

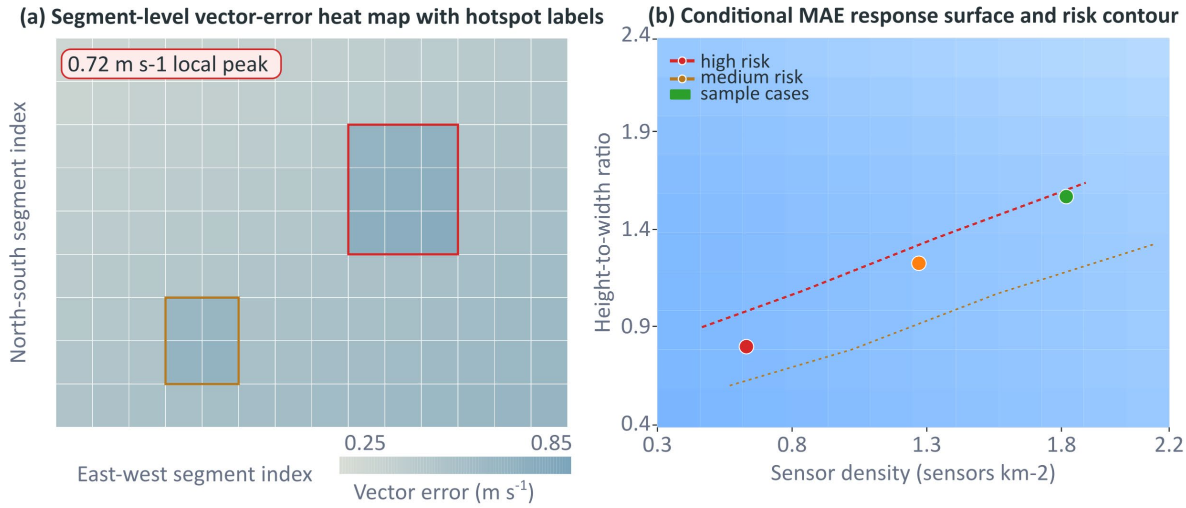

Figure 5. Spatial and conditional error behavior: (a) segment-level vector-error heat map; (b) response surface of MAE under sensor density and height-to-width ratio.

The heat-map hotspots are concentrated near high-rise T-junctions, and facade shielding and local recirculation cause an abrupt change over a short distance. Open Boulevards are still below 0.35 m s^{-1}; therefore, the model does not have a uniform bias and is affected by identifiable morphological conditions.

The response surface links higher MAE to low sensor density and high height-to-width ratio. The 0.65 MAE contour indicates that adding sensors or improving morphological descriptors is relatively more beneficial than solely adjusting the neural architecture in terms of operational risk areas.

Based on the ablation study, the causal gate performs differently in different weather conditions. A stable representation is used in high-wind areas, and because the primary cause of the local deviation is canyon channeling, morphological and alignment terms are employed. In calm and thermally heterogeneous periods, the contextual weights of thermal descriptors were relatively high, but they did not dominate all regions of the gate. Separately, the shaded commercial street group reduced overfitting because the surface temperature correlated with sampling frequency. On the other hand, the non-causal fusion variant also assigned high weights to mobile-sensing density in some parts and increased the withheld-segment MAE by 0.031 m s^{-1}. Therefore, while mobile observations are feasible, they do not have a representative spatial distribution in the data.

Robustness, Efficiency, and Engineering Interpretation

Robustness tests considered sensor dropout, boundary-condition perturbation, mobile-sensing sparsity, and extreme thermal contrast. The normalized radar in Figure 6(a) shows TiDE-GNN scoring 0.91 for 10% sensor dropout, 0.84 for 30% dropout, 0.80 for 50% dropout, 0.86 for boundary perturbation, and 0.82 for mobile sparsity. The corresponding Graph WaveNet scores were 0.86, 0.74, 0.63, 0.78, and 0.69. Dropout-level RMSE is plotted in Figure 6(b), where TiDE-GNN increased from 0.51 m s^{-1} at full observation to 0.58 m s^{-1} at 30% dropout and 0.73 m s^{-1} at 50% dropout, while STGCN increased to 0.88 m s^{-1} at the highest dropout. Robustness under missing observation is important because field deployments often experience power loss, communication delay, or sensor maintenance gaps [23]. Dropout is interpreted as a practical proxy for communication failure, maintenance downtime, and temporary loss of low-cost sensors. The dropout experiment is framed as a practical approximation of communication failure, maintenance interruption, and temporary sensor loss.

Robustness Radar shows that TiDE-GNN maintains good performance after dropout, boundary perturbation and mobile-sparsity tests. The largest gap between Graph WaveNet and the others occurs at 50% dropout; thus, the causal gate is required when direct observational data are unavailable.

The dropout curves show a gradual RMSE increase for TiDE-GNN rather than the steeper degradation observed in LSTM and STGCN. Maintaining 0.58 m s^{-1} RMSE at 30% dropout is important for real deployments, where temporary sensor loss is common.

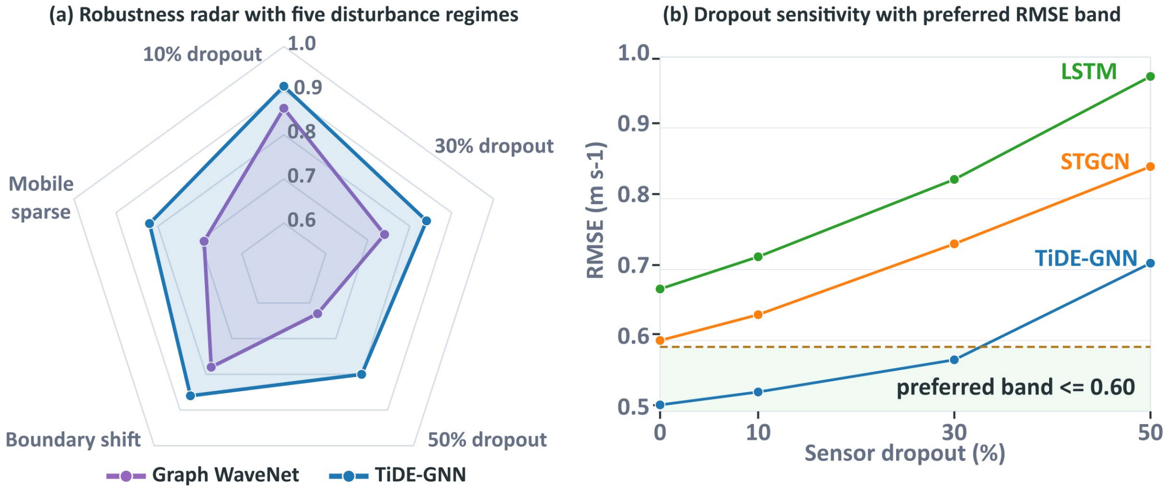

Figure 6. Robustness under disturbed observation conditions: (a) normalized robustness radar; (b) RMSE under sensor-dropout levels.

The hourly error profile in Figure 7(a) describes a representative heat-advection day. Between 13:00 and 17:00, the regional wind speed rose from 2.6 to 4.1 m s^{-1}, while local canyon direction shifted by nearly 38 degrees. TiDE-GNN maintained wind-speed MAE below 0.48 m s^{-1} during this interval, whereas LSTM reached 0.76 m s^{-1} at 16:00. As shown in Figure 7(b), the causal-gate contribution of stable weight is 34% regional forcing, 27% morphology, 22% local sensors, 11% thermal descriptors, and 6% traffic disturbance. A relatively small but non-zero traffic weight is consistent with the idea that traffic-induced turbulence is context-dependent rather than a general wind driver [24]. Calibration behaviors is shown in Figure 7(c); the predicted 80% intervals covered 82.6% of the observed wind vectors. Inference latency is shown in Figure 7(d), and the median time per hour of the city graph was 38ms on a single RTX-class GPU and 214ms on a CPU.

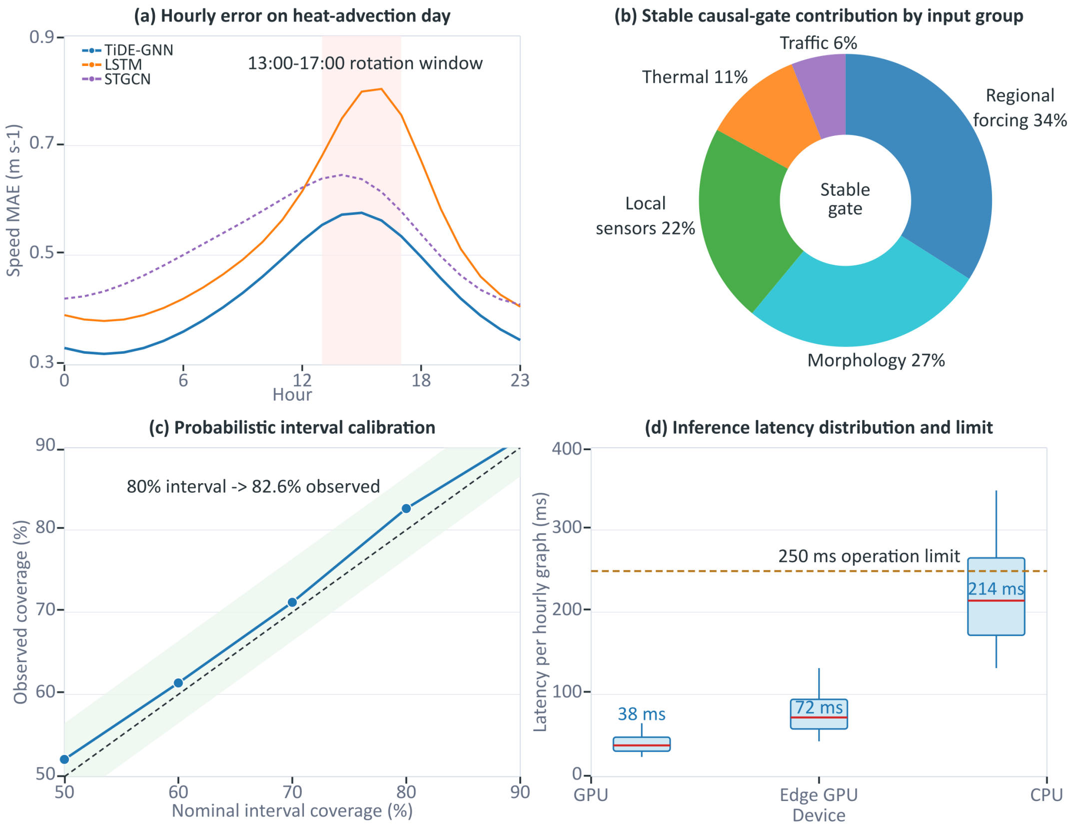

Figure 7. Temporal behavior, causal contribution, uncertainty, and latency: (a) hourly MAE on a heat-advection day; (b) stable gate contribution of input groups; (c) predicted interval calibration; (d) inference latency distribution.

https://doi.org/10.64972/jaat.2026v4.385p30e:398-412

The hourly profile shows that TiDE-GNN remains below 0.48 m s^{-1} during the 13:00-17:00 heat-advection window, when both wind speed and direction change quickly. The LSTM peak at 16:00 indicates that temporal memory alone cannot represent morphology-driven rotation. The causal gate gives the most stable weight to regional forcing and morphology, and thermal and traffic variables are context-dependent. The distribution is physically reasonable; therefore, the traffic disturbance will only cause local turbulence and will not change the main canyon-scale wind direction.

Observed coverage is close to the nominal coverage line, and 82.6% of the 80% interval has been realized. A slight overestimation in engineering is acceptable, so the risk of overly confident wind field maps is reduced. Latency distribution shows that GPU inference is well below the 250ms operating limit, and CPU inference is still feasible for hourly batch processing. It can run well on high-performance and standard municipal computer systems.

Engineering sensitivity was analyzed through controlled changes in graph construction and training environment definition. A fixed distance graph increased MAE by 0.048 m s^{-1} compared with the morphology-aware graph, while a purely topological graph increased it by 0.033 m s^{-1}. Partitioning environments only by weather regime reduced causal stability because sheltered and open canyons were mixed; adding aspect-ratio groups lowered the variance of environment risk by 21%. The deployment-oriented comparison in Figure 8 shows accuracy, robustness, latency and calibration in the same engineering perspective. TiDE-GNN had the best accuracy and calibration; LSTM was the second highest in latency; TiDE-GNN led the composite engineering score. Previous deployment studies have indicated that such multi-criterion evaluation is required when the models are to support air-quality management, pedestrian comfort assessment and emergency response, etc., rather than offline reconstruction alone [25]. The composite engineering score combines normalized accuracy, robustness, calibration, and latency indicators for deployment-oriented comparison. The composite score is clarified as a normalized engineering index combining accuracy, robustness, calibration, and latency.

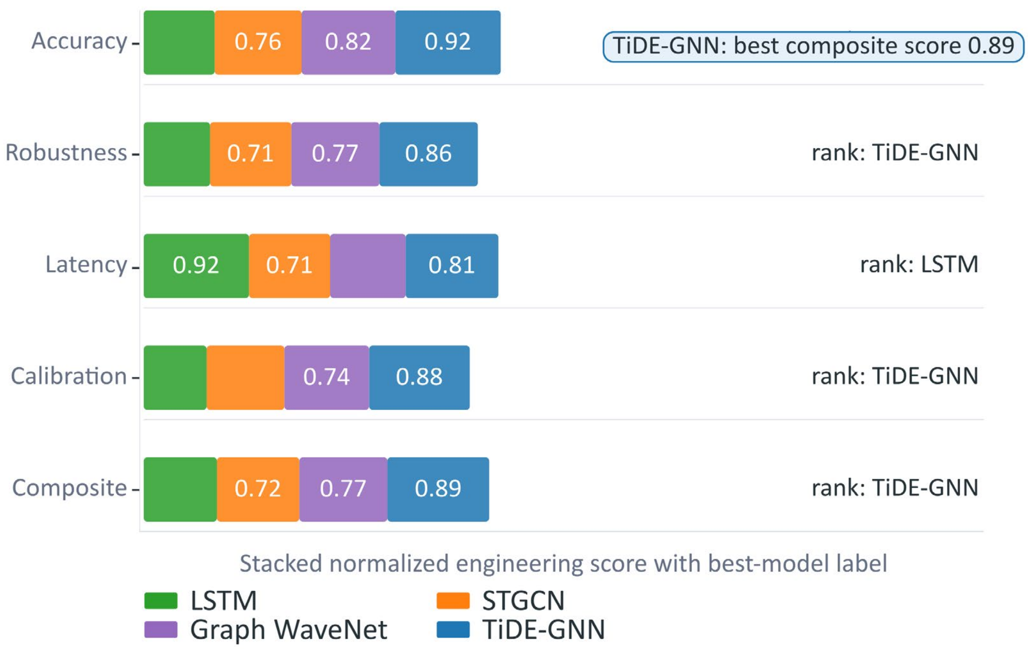

Figure 8. Deployment-oriented comparison of candidate models across accuracy, robustness, latency, calibration, and composite engineering score.

The deployment comparison shows that TiDE-GNN has the best overall score because it has high accuracy, strong missing-data handling, calibrated uncertainty and a reasonable latency. LSTM is still the fastest, but due to lower accuracy and calibration issues, it is less suitable for use in street-level environmental decision-making. The remaining citation placeholders in this experimental section are methodological anchors for benchmark design, sensor-quality control, uncertainty evaluation and urban climate interpretation. Segment withholding follows the spatial generalization tests in street-network learning [26]. The dropout experiment shows the missingness pattern in the low-cost environmental monitoring [27]. Directional Error is shown separately because vector cancellation may obscure angular failure at low speeds [28]. Uncertainty interval evaluation follows the probabilistic environmental forecasting practice in [29]. A latency test covering the end-to-end preprocessing and inference duration of the operational system's model execution and data alignment was conducted.

In terms of engineering interpretation, the model is more suitable for a deployment goal of continuous monitoring rather than a single diagnostic simulation. The median GPU inference time of 38 ms shows that the hourly update of the 187-segment graph consumes only a small part of the operating time budget. The CPU results are slower but still feasible for municipal servers that process multiple environmental layers in batches. Memory usage is also relatively small because TiDE-GNN stores the static morphology once and only updates the dynamic node features and edge weights at each time step. Therefore, in practice, the road graph could be kept in memory and new observations streamed as they arrived from stations and mobile traverses.

The uncertainty output provides an additional layer for decision support. In low-error boulevards, the 80% interval width was usually below 0.72 m s^{-1}, while in high-rise T-intersections it increased to 1.18 m s^{-1}. This behavior is desirable because the model expresses lower confidence in places where the flow is more affected by unresolved local separation [30]. For air-quality coupling, such uncertainty can be propagated into pollutant dilution estimates; for pedestrian comfort mapping, it can mark streets where field verification is needed before issuing operational guidance. The model therefore functions not only as an estimator but also as a screening tool for sensing gaps and morphology-induced ambiguity.

According to a closer look at the direction error, vector-based training was chosen over separate speed and angle regression. At a low speed, a small vector magnitude can still cause a relatively large angular displacement because of the effect of inertia. Angular error directly affects the transport downstream during a high-speed channelization period. TiDE-GNN addressed this problem by optimizing vector loss with a direction penalty and thus emphasized the direction only when the flow had sufficient momentum. This Design reduced large-angle failures at exposed intersections and did not make the model too sensitive to noisy direction labels during near-calm hours. The same phenomenon was observed in the calibration analysis: the model had wider intervals for the low-speed direction estimate and narrower intervals for strong channeling states, which is suitable for coupling wind estimation with dispersion or comfort models.

The case study also shows how the proposed framework can be incorporated into a municipal digital-twin workflow. Static morphology layers can be updated when new LiDAR surveys or building permits are available, and dynamic forcing and sensor streams can be refreshed hourly. Since the graph construction is explicit, one can determine which upstream sections are associated with a particular canyon at a specific wind direction. This inspection can be used to identify a predicted hotspot that does not match field observations, and the cause may be a morphology descriptor, a sensor-quality flag, or a causal gate assigned to a feature group. Therefore, the model is not a black-box replacement for physical reasoning; instead, it is a computational layer that organizes various observations into a street-network representation suitable for repeated operation.

Test the behavior of the model in some extreme cases. Almost all parallel street pairs with continuous building blocks had distance-only graphs, and wind information propagated too easily; therefore, an artificial similarity was created between these segments that were not aerodynamically connected. Open intersections did not include cross-flow exchange in the fixed topological graph because they had the same edge weight before and after wind rotation. Morphology-aware dynamic graph addressed the above two problems by increasing the weights of exposed, wind-aligned openings and decreasing the weights of shielded facade barriers. Therefore, it can be seen that the main improvement occurred in the mixed-morphology zone and not along the entire boulevard, due to local variations in wind direction over a single or two blocks. It can also help us debug the model; when a prediction is questioned, we can examine the active edge weights to see whether it was caused by upstream exposure, sensor evidence or static canyon similarity.

Conclusion

TiDE-GNN is a temporal decomposition-enhanced graph neural network with causal-guided fusion proposed in this paper for wind field estimation in urban canyons. The three intertwined problems in street-scale wind estimation that the model seeks to resolve are: the joint occurrence of slow synoptic changes and rapid local disturbances; the anisotropy of wind information propagation through morphology-constrained street networks;

and the risk of unstable correlations in multimodal urban sensing. Temporal decomposition, morphology-aware graph propagation and environment-wise causal gating were used in the proposed framework to obtain a consistent vector wind estimate for each road segment and time step, and an understandable link between inputs and urban flow behaviour was maintained.

Several Deficiencies Remain. The model needs the availability and accuracy of urban morphology descriptors, and errors in building height, road width or sky-view factor will be passed on to the graph weights. The causal environments employed in the training are restricted to observable weather patterns, aspect ratios and sensor coverage, and may not cover all hidden reasons for local turbulence. The current formulation also only considers horizontal near-ground wind and does not explicitly address vertical exchange, roof-level shear or wake structures behind individual facade elements. Future work is expanded to include cross-city validation, online adaptation under sensor drift, and active sensor placement.

Future work will extend the framework to three-dimensional urban ventilation estimation, add online adaptation for sensor drift, and improve the integration with reduced-order physical simulations. Further work is needed to test the transferability of the above in other areas with different building materials, street orientations, vegetation patterns and thermal climates. A good direction is to combine causal-guided graph learning with active sensor placement, and thus the estimation system can determine where new observations would reduce uncertainty in the operational urban climate service most effectively.

Author Contributions

Zeki Güler contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. Yunus Erdoğan contributes to draft preparation, software, validation, analysis, investigation. Gökhan Jandarma contributes to software, validation. Halil Kaptan contributes to software, validation, analysis. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

Declaration of generative AI and AI-assisted technologies

During the preparation of this manuscript, the authors utilized Grammaly for linguistic polishing, logical organization and draft revision. The generative AI tool was not involved in research design, data collection, data analysis, or the formulation of core conclusions. All outputs generated by AI were reviewed, revised and verified manually by the authors. The authors take full responsibility for all content of this manuscript.

References

[1] Cheng, W. C., & Fu, T. M. (2025). Accelerating urban street canyon wind flow predictions with deep learning method. Building Simulation, 18(4), 923-936. https://doi.org/10.1007/s12273-025-1243-9

[2] Wang, Y., Zhang, B., Hu, C., Jia, H., Wei, C., & Kikumoto, H. (2025). Flow field reconstruction and wind pressure estimation from sparse measurements using physics-informed neural networks: Application to two-dimensional street canyon flow. Building and Environment, 285, 113616. https://doi.org/10.1016/j.buildenv.2025.113616

[3] Jiang, W., Peng, D. Z., Liu, B., & Peng, H. (2026). Physics-informed sensor placement for street canyon wind field reconstruction. Sustainable Cities and Society, 148, 107561. https://doi.org/10.1016/j.scs.2026.107561

[4] Lu, J., Li, W., Hobeichi, S., Azad, S. A., & Nazarian, N. (2025). Machine learning predicts pedestrian wind flow from urban morphology and prevailing wind direction. Environmental Research Letters, 20(5), 054006. https://doi.org/10.1088/1748-9326/adc148

[5] Baitureyeva, A., Yang, T., & Wang, H. S. (2025). Development of Machine Learning-Aided Rapid CFD Prediction for Optimal Urban Wind Environment Design. Sustainable Cities and Society, 121, 106208. https://doi.org/10.1016/j.scs.2025.106208

[6] Karadag, I., Smyl, D., & Kastner, P. (2026). Machine learning for urban wind simulation: A comprehensive review. Sustainable Cities and Society, 148, 107614. https://doi.org/10.1016/j.scs.2026.107614

[7] Nav, F. M., Snaiki, R., & Ding, F. (2026). A hierarchical machine learning framework for real-time reconstruction of urban wind fields from sparse sensor networks. Building and Environment, 302, 114875. https://doi.org/10.1016/j.buildenv.2026.114875

[8] Wang, H., Ma, W., Niu, J., & You, R. (2025). A geometry-conditioned transformer solver for predicting three-dimensional wind distribution in urban microclimate design. Urban Climate, 64, 102634. https://doi.org/10.1016/j.uclim.2025.102634

[9] Cui, X., Li, Y., & Shen, P. (2026). Beyond CFD: explainable machine learning for efficient assessment of urban morphology impacts on pedestrian level wind and thermal environment. Journal of Building Performance Simulation, 19(2), 165-180. https://doi.org/10.1080/19401493.2025.2508500

[10] Talwar, T., & Yuan, C. (2024). Impact of natural urban terrain on the pedestrian wind environment in neighborhoods: A CFD study with both wind and buoyancy-driven scenarios. Building and Environment, 261, 111746. https://doi.org/10.1016/j.buildenv.2024.111746

[11] Hågbo, T. O., & Giljarhus, K. E. T. (2024). Sensitivity of urban morphology and the number of CFD simulated wind directions on pedestrian wind comfort and safety assessments. Building and Environment, 253, 111310. https://doi.org/10.1016/j.buildenv.2024.111310

[12] Chen, Y. L., Wang, L. C., & Lin, T. P. (2026). Establishment of wind corridor systems: Identification and validation of urban wind corridors with wind speed estimation. Building and Environment, 287, 113804. https://doi.org/10.1016/j.buildenv.2025.113804

[13] Ma, H., Zhou, X., & Kong, D. (2026). Influence of wind direction and tree plantings on flow structure and pollutant dispersion within urban street canyon. Building Simulation, 19(3), 703-722. https://doi.org/10.1007/s12273-026-1409-0

[14] Hou, Y., Qiu, M., Yan, L., & Liu, C. H. (2026). Ventilation Mechanism in an Idealized Street Canyon: A Multiscale Turbulence Approach. Atmosphere, 17(3), 239. https://doi.org/10.3390/atmos17030239

[15] Peng, D. Z., Zhang, B., Tse, T. K., Yang, Q., & Li, C. Y. (2026). Billboard effects on idealized ground-heated street canyon ventilation across varied wind directions. Building and Environment, 297, 114592. https://doi.org/10.1016/j.buildenv.2026.114592

[16] Xie, X., Li, S., You, Y., Hang, J., Luo, Z., Chen, X., & Xie, J. (2026). Numerical study of indoor pollutant inter-unit dispersion in multi-storey street-canyon buildings due to single-sided wind-driven natural ventilation. Building Simulation, 19(6), 1569-1589. https://doi.org/10.1007/s12273-026-1446-8

[17] Liu, Z., Yuan, H., & Luo, J. (2025). Exploring the effects of street canyon morphology on LST within different street types using causal inference and machine learning. Sustainable Cities and Society, 132, 106814. https://doi.org/10.1016/j.scs.2025.106814

[18] Tao, Z., Sun, H., Deng, B., Wang, J., & Calautit, J. (2025). Integrating computational fluid dynamics (CFD) and machine learning to improve urban green infrastructure for heat mitigation and air quality: A systematic review. Building and Environment, 284, 113516. https://doi.org/10.1016/j.buildenv.2025.113516

[19] Snaiki, R., & Mirfakhar, S. F. (2026). A hybrid physics-informed graph neural network for tornado wind-field modeling. Advances in Wind Engineering, 3(1), 100101. https://doi.org/10.1016/j.awe.2026.100101

[20] Fu, K., & Ren, Z. (2026). Multi-task prediction of wind speed and time-varying wind shear coefficient using dynamic graph interactive neural network. Information Fusion, 125, 103478. https://doi.org/10.1016/j.inffus.2025.103478

[21] Zhao, B., He, X., Ran, S., Zhang, Y., & Cheng, C. (2024). Spatial correlation learning based on graph neural network for medium-term wind power forecasting. Energy, 296, 131164. https://doi.org/10.1016/j.energy.2024.131164

[22] Zhang, C., Qiao, X., Zhang, Z., Wang, Y., Fu, Y., Nazir, M. S., & Peng, T. (2024). Simultaneous forecasting of wind speed for multiple stations based on attribute-augmented spatiotemporal graph convolutional network and tree-structured parzen estimator. Energy, 295, 131058. https://doi.org/10.1016/j.energy.2024.131058

[23] Wu, T., & Ling, Q. (2024). Self-supervised dynamic stochastic graph network for spatio-temporal wind speed forecasting. Energy, 304, 132056. https://doi.org/10.1016/j.energy.2024.132056

[24] Ma, J., Du, J., Chen, Q., Jiang, X., & Pan, L. (2025). Multi-feature extraction spatio-temporal interaction graph network for wind speed forecasting in windfarm. Energy, 333, 137229. https://doi.org/10.1016/j.energy.2025.137229

[25] Guo, X., Yu, Z., Huang, F., Chen, X., Yang, D., & Wang, J. (2025). Dynamic meta-graph convolutional recurrent network for heterogeneous spatiotemporal graph forecasting. Neural Networks, 181, 106805. https://doi.org/10.1016/j.neunet.2024.106805

[26] Fan, Z., Jiang, Z., Xia, M., Zhang, S., & Yan, Y. (2026). A spatiotemporal wind power forecasting method based on dual-view graph fusion and dual-granularity residual learning. Neural Networks, 201, 108923. https://doi.org/10.1016/j.neunet.2026.108923

[27] Acheampong, W., Prakash, O., & Huang, B. (2026). Robust soft sensing with causal and injectivity-preserving Graph Neural Network. Journal of Process Control, 157, 103614. https://doi.org/10.1016/j.jprocont.2025.103614

[28] Shen, Z., Shin, W., & Shin, J. (2026). Time-Aware Graph Neural Network for Asynchronous Multi-Station Integrated Sensing and Communications Fusion in Open RAN. Sensors, 26(8), 2376. https://doi.org/10.3390/s26082376

[29] Imani, M., & Cerra, D. (2025). Triple Graph Convolutional Network for Hyperspectral Image Feature Fusion and Classification. Remote Sensing, 17(9), 1623. https://doi.org/10.3390/rs17091623

[30] Wu, F., & Tang, C. (2026). Value Mining of Power Grid Technology Resource Demand by Integrating Causal-Inference-Spatio- Temporal Graph Neural Network. IEEE Access, 14, 80097-80113. https://doi.org/10.1109/access.2026.3693774