Received: March 5, 2026. Revised: April 10, 2026.

International Journal of
Intelligent Engineering & Systems

http://www.inass.org/

GeoRiskNet: A Spatio-temporal Graph Deep Learning Framework for AQI Forecasting with Regulatory Risk Interpretation

Anitha R[1]* **Shyamala Devi N[1]**

[1]*Department of Computer Applications, VISTAS, Chennai, India*
* Corresponding author's Email: ramalingam.anitha@gmail.com

Abstract: Accurate Air Quality Index (AQI) forecasting is essential for environmental monitoring, public health protection and regulatory planning in rapidly urbanizing regions. However, reliable station-level AQI prediction remains challenging due to nonlinear temporal air pollution dynamics, complex meteorological influences and cross-regional diffusion among geographically distributed monitoring stations. Conventional statistical models primarily focus on temporal patterns and often fail to capture structured spatial interdependencies and latent vulnerability characteristics inherent in urban air pollution systems. To address these limitations, this study proposes GeoRiskNet, a unified spatio-temporal graph deep learning framework for next-day AQI forecasting with integrated regulatory risk interpretation. The model formulates AQI prediction as a supervised regression problem using multivariate hourly pollutant and meteorological observations gathered from fourteen Central Pollution Control Board (CPCB) monitoring stations in the Chennai metropolitan region between 2018 and 2023. The architecture integrates a Convolutional Neural Network-Long Short-Term Memory (CNN-LSTM) network with temporal attention to capture short- and long-term pollution dynamics influencing AQI, a nonlinear autoencoder for compact latent vulnerability representation and a two-layer Graph Convolutional Network (GCN) to model spatial diffusion using geodesic distance-based adjacency. The learned representations are fused within an end-to-end regression framework to predict continuous next-day AQI values, which are subsequently mapped to CPCB-defined severity categories for regulatory interpretation. The proposed GeoRiskNet achieves strong predictive performance, with Root Mean Squared Error (RMSE) of 0.087, Mean Absolute Error (MAE) of 0.066 and coefficient of determination (R^2) of 0.953 on normalized AQI values. When mapped to the original AQI scale (0-500), this corresponds to approximately 43.5 AQI units (RMSE) and 33 AQI units (MAE), enabling practical interpretation for environmental monitoring and regulatory applications. Compared to ARIMA (RMSE = 0.124), the proposed framework reduces prediction error by 29.8%. Regulatory severity classification achieves 91.4% overall accuracy with a macro F1-score of 0.906, confirming balanced and interpretable urban air pollution forecasting.

Keywords: Air quality index forecasting, Spatio-temporal modeling, Graph convolutional networks, Deep learning, Urban air pollution, Regulatory risk assessment.

1. Introduction

Air pollution is one of the most significant public health and environmental challenges in rapidly urbanizing regions worldwide [1, 2]. Rapid industrialization, increased vehicular emissions, expanding energy consumption and urban growth have substantially elevated atmospheric pollutant concentrations in metropolitan environments [3]. Major pollutants, including particulate matter (PM2.5, PM10), nitrogen oxides (NO2), carbon monoxide (CO), sulfur dioxide (SO2) and ground-level ozone (O3), are strongly associated with respiratory and cardiovascular diseases, reduced life expectancy and broader ecological degradation [4, 5]. In addition to direct health impacts, atmospheric pollutants influence regional climate systems through complex chemical reactions and radiative processes, thereby affecting long-term environmental stability [6]. Consequently, accurate forecasting and spatial assessment of air quality risk are essential for effective urban planning, environmental monitoring and evidence-based policy formulation.

Conventional air quality forecasting approaches primarily rely on statistical time-series models and deterministic atmospheric dispersion simulations [7, 8]. While these methods provide structured analytical foundations, they often assume simplified linear relationships and may inadequately capture nonlinear interactions among meteorological variables, emission sources and atmospheric transport mechanisms. Furthermore, many traditional models treat monitoring stations independently, limiting their ability to account for spatial interdependencies and cross-regional pollutant diffusion. Since atmospheric pollutants propagate across geographically connected regions through wind-driven transport and proximity-based dispersion processes, incorporating spatial relationships is critical for comprehensive and reliable AQI forecasting.

Recent progresses in artificial intelligence and machine learning (ML) have expanded the analytical capabilities of environmental modeling systems [9, 10]. Deep learning (DL) architectures have demonstrated improved performance in modeling nonlinear temporal dynamics within complex air quality datasets [11, 12]. These methods enable automated feature extraction and adaptive learning from high-dimensional environmental observations, often outperforming classical statistical approaches. In parallel, graph-based learning paradigms have emerged as effective tools for representing spatial dependencies in interconnected systems, facilitating structured modeling of regional diffusion processes [13, 14]. Despite these developments, challenges remain in developing unified frameworks that jointly capture temporal variability and spatial diffusion within a coherent modeling architecture. Moreover, improving interpretability remains essential, as transparent AQI forecasting is necessary for regulatory implementation and public health communication.

Addressing these limitations requires an integrated analytical framework capable of simultaneously modeling nonlinear temporal dynamics, diffusion-aware spatial interdependencies and structured regional vulnerability patterns. The present study proposes a comprehensive spatio-temporal risk assessment framework for station-level AQI forecasting. The proposed approach models pollutant fluctuations across monitoring stations, incorporates graph-based spatial propagation mechanisms, extracts compact latent vulnerability representations and aligns continuous AQI predictions with CPCB-defined regulatory severity categories. While the individual components employed in this study, namely CNN-LSTM with attention, autoencoder-based latent representation and graph convolutional networks, have been explored independently in prior work, the novelty of the proposed GeoRiskNet framework lies in their structured and coordinated integration within a unified learning pipeline. Unlike conventional hybrid models that combine components in a loosely coupled or sequential manner, the proposed architecture enables complementary interaction among temporal, latent and spatial representations. Temporal attention identifies salient sequential patterns; latent compression refines these patterns into compact and denoised feature representations and graph-based propagation distributes this information across spatially connected monitoring stations. This design allows the model to simultaneously capture temporal evolution, latent feature structure and spatial diffusion, which are inherently interdependent in real-world air pollution dynamics. Furthermore, by integrating latent representation prior to spatial modeling, the framework ensures that graph-based learning operates on a structured and noise-reduced feature space, thereby enhancing generalization and predictive stability. The resulting performance improvements are therefore attributable not merely to architectural aggregation but to the coordinated interaction of these components within a unified spatio-temporal learning framework. The objectives of the study are outlined below:

- To model nonlinear temporal dynamics of urban air pollutants using sequential deep learning techniques.
- To capture spatial dependencies among monitoring stations through graph-based diffusion modeling.
- To extract latent vulnerability representations for structured AQI risk stratification.
- To develop a regulatory-aligned forecasting framework that maps continuous AQI predictions to CPCB severity categories.
- To evaluate the proposed framework using standardized performance metrics and comparative benchmarking.

The remainder of this paper is organized as follows. Section 2 reviews related literature in air pollution forecasting and spatial environmental modeling. Section 3 describes the methodological framework adopted in this study. Section 4 provides experimental results and comparative evaluation. Finally, Section 5 concludes the paper and outlines future research directions.

2. Literature review

Yenkikar et al. [15] developed a hybrid forecasting framework integrating Random Forest Regression with ARIMA-based residual correction for Air Quality Index prediction. Random Forest captured nonlinear pollutant-AQI relationships, while ARIMA modeled residual temporal structures under an expanding window cross-validation scheme. Model interpretability was ensured through SHAP analysis, identifying PM2.5, PM10 and NO2 as dominant contributors. Comparative evaluation indicated improved predictive accuracy, with the hybrid approach achieving R^2 of 0.94 and MSE of 508.46, outperforming the standalone Random Forest (RF) (MSE = 524.64; R^2 = 0.917). Diagnostic checks included Ljung-Box testing and uncertainty band assessment. Limitations included dependence on high-quality historical pollutant data, exclusion of meteorological and policy variables, limited handling of nonlinear residual dynamics and increasing uncertainty in long-term forecasts.

Karnati et al. [16] examined air quality trends in Delhi using CPCB data from 2017-2023, addressing missing observations through multiple imputation techniques, with forward-backward fill yielding the most reliable reconstruction for ozone. An 80/20 chronological split preserved temporal order for PM2.5 forecasting. Comparative evaluation involved statistical, ML and DL models, including ARIMA, SARIMA, RF, LSTM variants, GRU, TCN, WaveNet and Prophet. Daily predictions showed Bi-LSTM (RMSE = 39.56, R^2 = 0.792) and LSTM with attention (RMSE = 39.48, R^2 = 0.793) outperforming alternatives, while hourly forecasts achieved R^2 = 0.947 with RMSE ≈ 22.5. Linear models exhibited weak performance, with ARIMA recording R^2 = 0.037. Limitations included data gaps, skewed pollutant distributions, limited spatial scope to a single station and sensitivity to abrupt environmental variations.

Govande et al. [17] assessed daily PM2.5 forecasting across Chennai, Delhi, Hyderabad and Kolkata using Recurrent Neural Networks, including SimpleRNN, LSTM and GRU, trained on CPCB observations from 2018-2023. Meteorological variables and gaseous pollutants supported model training after exploratory analysis and outlier removal. Performance evaluation employed MAE, RMSE, MAPE, R^2 and correlation coefficients. GRU achieved the strongest results in Hyderabad (RMSE = 4.75, MAE = 3.57, R^2 = 0.91) and Chennai (MAE = 4.47, RMSE = 5.42, R^2 = 0.80), while Delhi recorded weaker performance, with LSTM yielding R^2 = 0.61. All models showed significant correlations (R > 0.8 in most cases). Limitations involved fixed network parameters across cities, exclusion of emission inventory data, high variability in Delhi pollution patterns and substantial computational demands associated with deep learning architectures.

Jayaraman et al. [18] developed a Time-Based-Spatial framework integrating CNN for spatial learning from normalized latitude-longitude grids and ARIMA for temporal pollutant trends using CPCB data from 22 Indian cities. The dataset was preprocessed and divided into training and testing sets to generate 6-hour AQI forecasts. Comparative evaluation against SVR, regression models, LSTM variants and ARIMA indicated that the TBS framework achieved R^2 = 0.81, MAE = 31.73, RMSE = 44.60 and MAPE = 26.02, matching the highest R^2 among competing approaches. The hybrid structure enhanced representation of spatial and temporal dependencies. Limitations included reliance on 22-city coverage, absence of detailed meteorological and industrial emission variables and performance constraints associated with ARIMA in capturing complex nonlinear temporal dynamics.

Mishra et al. [19] investigated AQI prediction for Amaravati, Bengaluru, Chennai and Hyderabad using boosting algorithms, including XGBoost, LightGBM and an ensemble stack trained on CPCB daily data from May 2019 to November 2023. Model evaluation relied on R^2 and RMSE metrics. LightGBM achieved strong accuracy in Amaravati (R2 = 0.9968, RMSE = 0.5125) and Chennai (R2 = 0.9972, RMSE = 0.9132), while XGBoost performed better in Bengaluru (R^2 = 0.9870, RMSE = 2.0229) and Hyderabad (R^2 = 0.9985, RMSE = 1.6507). The ensemble stack exhibited weaker generalization, particularly in Bengaluru (R^2 = 0.9612, RMSE = 3.4924). Limitations involved sensitivity to regional data characteristics, performance variability at higher AQI levels and reduced robustness of stacked integration due to noise and weighting issues.

Raza & Singh [20] developed an AI-driven air quality forecasting framework integrating satellite imagery, ground sensors, meteorological variables and mobile IoT data for Delhi and Mumbai. DL architectures, including Transformer, LSTM and a CNN-LSTM hybrid, were trained on hourly pollutant data from 2022-2023. Performance evaluation using MAE, RMSE and R^2 indicated superior accuracy of the hybrid model (MAE = 8.5 μg/m³, RMSE = 12.6 μg/m³, R^2 = 0.92), outperforming ARIMA (RMSE = 24.5, R^2 = 0.71) and Random Forest (RMSE = 18.9, R^2 = 0.82). A health risk assessment module translated forecasts into WHO-based exposure categories. Limitations included reduced accuracy during extreme weather events, data sparsity from

 DOI: 10.22266/ijies2026.0630.19

mobile sensors, satellite latency and high computational demands associated with Transformer training.

Ahmed et al. [21] developed a hybrid deep learning framework, CLSTM-BiGRU, for AQI prediction across seven Bangladeshi cities using CNN, LSTM and BiGRU networks optimized through Grey Wolf Optimization. The model incorporated nineteen remotely sensed meteorological predictors from MERRA-2 alongside ground-based AQI records (2014-2020). Performance evaluation indicated superior forecasting efficiency, with correlation coefficients approaching 1 and lower prediction errors compared to benchmark machine learning models. Air temperature, precipitation and wind speed exhibited moderate correlations with AQI ($r \approx 0.4$-0.5), highlighting their predictive relevance. Despite strong performance, limitations included restriction to sub-tropical humid regions, limited inclusion of broader environmental and demographic variables, seasonal variability effects and the need for deeper pollutant-specific analysis, particularly for PM2.5 and PM10 dynamics.

Rosca et al. [22] developed an integrated AQI prediction and forecasting system for the Bronx using pollutant concentrations (CO, NO_2, SO_2, PM2.5), meteorological variables and traffic data collected between 2020 and 2023. Nineteen algorithms, including twelve ML and seven DL models, were evaluated alongside a structural component forecasting approach. Customized RF Regression achieved superior predictive performance (R^2 = 99.59%, MAE = 0.22%, MAPE = 0.68%), with a validation deviation of 0.58%. The five-day AQIF forecasting model recorded R^2 = 71.62% and MAE = 0.4%. Results were integrated into an IoT-enabled web platform with automated alerts. Limitations included restricted temporal coverage, reliance on a single geographic location, default hyperparameter settings, absence of industrialization and demographic variables and comparatively lower long-term forecasting accuracy.

Rajesh et al. [23] developed a machine learning-driven framework for real-time air quality valuation and health risk mapping across industrial, urban, rural, suburban and traffic-dense zones. Multisource inputs included meteorological variables, fixed and mobile sensors, satellite data, traffic patterns and demographic overlays. Predictive modelling employed Random Forest, Gradient Boosting, XGBoost and LSTM networks, while SHAP analysis ensured interpretability of influential environmental and population factors. A cloud-based system generated five-minute risk maps and mobile alerts to identify vulnerable groups. Results demonstrated robust short-term pollutant prediction and spatially resolved Health Risk Index outputs across diverse zones. Limitations involved uneven sensor density, reliance on synthetic pilot data, reduced adaptability to extreme climatic variability, computational overhead from real-time SHAP explanations and constrained generalizability in poorly instrumented or topographically complex regions.

Patel et al. [24] evaluated ML and DL models for forecasting PM2.5 and PM10 concentrations across five Maharashtra cities using CPCB data from 2019-2023. After preprocessing and outlier removal, Linear Regression, Decision Tree, RF, XGBoost and LSTM were compared using R^2 and RMSE metrics. LSTM demonstrated superior predictive performance, achieving R^2 values between 0.9970 and 0.9999 for PM2.5 and up to 0.9980 for PM10, with low RMSE values such as 0.15 in Nashik and Pune. In contrast, traditional models recorded lower R^2 scores, ranging from 0.15 to 0.96. Limitations included sensitivity to data quality, regional dependence, limited generalizability across diverse environments, computational demands of LSTM and challenges related to complex pollutant dynamics and source variability.

Existing AQI forecasting studies have predominantly emphasized either temporal modeling or localized spatial analysis, resulting in fragmented learning of pollutant dynamics across monitoring stations [15, 18]. Several approaches achieved high predictive accuracy but relied on single-station data or lacked structured spatial dependency modeling among geographically distributed stations [16, 24]. Hybrid and ensemble models improved regression performance; however, cross-regional pollutant diffusion and latent urban vulnerability characteristics remained insufficiently represented [17, 21]. Multi-source and IoT-integrated systems enhanced real-time monitoring but often did not incorporate unified graph-based spatial reasoning within an end-to-end learning framework [20, 23]. Furthermore, limited integration of regulatory severity interpretation with continuous AQI regression constrained policy-level applicability [19], [22]. Therefore, a unified spatio-temporal framework capable of jointly modeling nonlinear temporal evolution, spatial diffusion and regulatory risk interpretation is required to address these limitations comprehensively.

3. Materials and methods

Urban air pollution exhibits nonlinear temporal variability and spatial interdependence across

Received: March 5, 2026. Revised: April 10, 2026.

geographically distributed monitoring stations, necessitating an integrated modeling framework capable of jointly capturing sequential dynamics and inter-station interactions. To address this, the proposed GeoRiskNet framework adopts a unified spatio-temporal learning architecture consisting of temporal modeling, latent feature compression, spatial graph propagation and regression-based forecasting, as shown in Fig. 1. Hourly multivariate observations of major air pollutants together with meteorological variables are first preprocessed through normalization and chronological structuring. A sliding window approach with a fixed window length of 10 hours is employed to construct sequential input samples. These sequences are processed using a hybrid CNN-LSTM architecture, where the convolutional layer extracts short-term local temporal patterns and the LSTM layer model longer-term dependencies and nonlinear pollutant interactions. A temporal attention mechanism is incorporated to assign adaptive importance weights to hidden states and generate a compact temporal embedding.

The resulting embedding is subsequently compressed using a nonlinear autoencoder to obtain a low-dimensional latent representation that preserves structured pollutant interaction characteristics while reducing redundancy. K-means clustering is then applied to the latent vectors to derive station-level vulnerability categories. In parallel, spatial relationships among monitoring stations are modeled by constructing a weighted graph based on geodesic distances between station coordinates. A two-layer GCN propagates feature information across neighboring stations to capture spatial dependencies and diffusion patterns. The temporal, latent and spatial representations are integrated through a feature concatenation to generate the final regression output corresponding to next-day AQI for each monitoring station. In this study, the target variable is the station-level next-day AQI computed according to CPCB guidelines using pollutant sub-index formulation, while pollutant and meteorological variables are used exclusively as predictive features. The predicted continuous AQI values are subsequently mapped to CPCB-defined severity categories to facilitate regulatory interpretation.

3.1 Dataset description

The dataset used in this study was obtained from the CPCB, India, through the CAAQMS network [25]. The study focuses on fourteen consistently operational monitoring stations located within the Chennai metropolitan region and adjoining districts in Tamil Nadu, representing diverse urban environments including residential, traffic-dominated and industrial zones. Each station is uniquely geo-referenced using latitude and longitude coordinates in the WGS84 geographic coordinate system, enabling spatial graph construction and geodesic distance-based adjacency modeling.

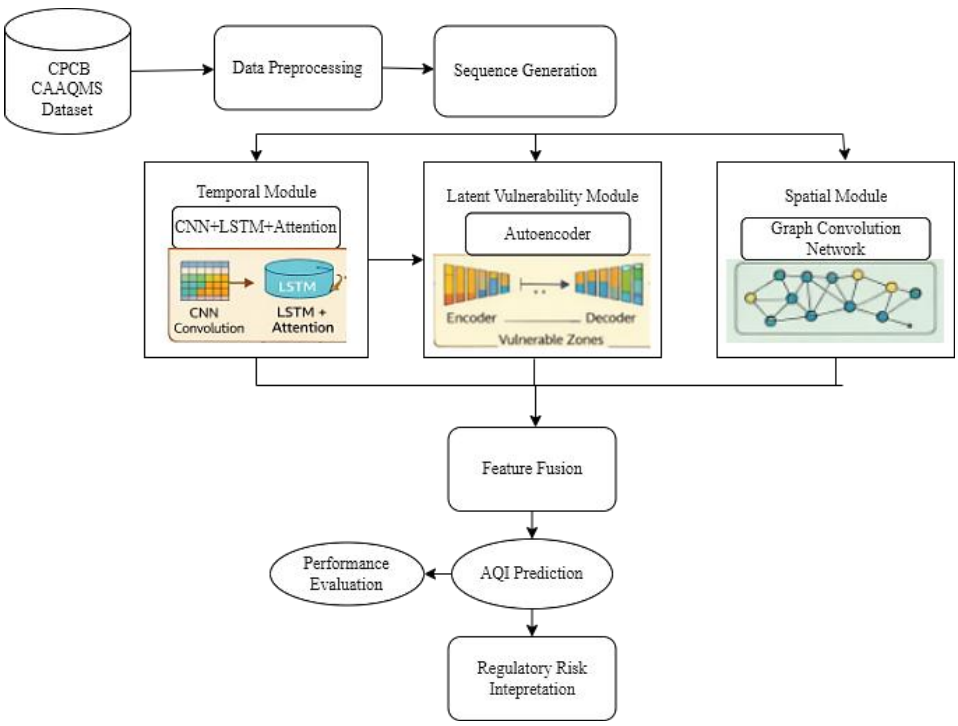

Figure. 1 GeoRiskNet framework

The dataset spans a six-year period from January 2018 to December 2023 with an hourly temporal resolution, capturing diurnal cycles, seasonal variations, episodic pollution events and long-term air quality trends. The recorded variables include major criteria pollutants regulated under the National Ambient Air Quality Standards (PM2.5, PM10, NO_2, SO_2, CO and O_3) along with meteorological parameters like relative humidity, temperature, wind speed, wind direction and atmospheric pressure. In this study, pollutant and meteorological variables are used exclusively as input features, while the prediction target corresponds to station-level AQI computed according to CPCB guidelines. Following data cleaning, time-aware imputation, normalization and sliding window construction with window length $w = 10$, the final structured dataset comprised approximately 735,000 spatio-temporal samples aggregated across all stations. To ensure temporal integrity and avoid information leakage, a chronological split strategy was adopted, wherein 80% of the samples were used for model training and the remaining 20% were reserved as a held-out test set. The final test dataset used for performance evaluation consisted of approximately 147,000 samples aggregated across all monitoring locations. This large-scale, multi-year, high-resolution dataset provides a robust and representative foundation for spatio-temporal AQI forecasting and subsequent regulatory risk interpretation. Table 1 illustrates the summary of dataset characteristics.

3.2 Data preprocessing

The station-level hourly observations obtained from the CPCB consist of multivariate pollutant and meteorological measurements recorded across the study region. Prior to analysis, the dataset was subjected to systematic preprocessing to ensure statistical consistency, temporal continuity and numerical stability. Missing observations arising from sensor downtime or transmission interruptions were corrected using time-aware imputation strategies. For short-duration gaps, linear interpolation was applied as, as shown in Eq. (1).

$$x_t = x_{t-1} + \frac{x_{t+k} - x_{t-1}}{k+1} \tag{1}$$

where k denotes the number of consecutive missing observations. For extended gaps, seasonal mean imputation was employed to preserve diurnal and monthly patterns, as shown in Eq. (2).

$$x_t = \frac{1}{N}\sum_{j=1}^{N} x_{t_j} \tag{2}$$

where x_{t_j} represents historical observations corresponding to the same hour and month. To mitigate the influence of anomalous sensor readings, statistical outlier detection was performed using Z-score normalization, as shown in Eq. (3).

$$Z = \frac{x-\mu}{\sigma} \tag{3}$$

where μ and σ denote the mean and standard deviation of the respective pollutant variable. Observations satisfying $|Z| > 3$ were treated as extreme outliers and corrected using median-based smoothing to preserve distributional integrity while eliminating measurement artifacts.

To ensure numerical stability during optimization, all continuous features were scaled using Min-Max normalization, as revealed in Eq. (4).

$$x' = \frac{x - x_{\min}}{x_{\max} - x_{\min}} \tag{4}$$

Table 1. Summary of dataset characteristics

Category	**Description**
Data Source	CPCB, India
Monitoring Network	CAAQMS
Study Region	Chennai metropolitan region and adjoining districts, Tamil Nadu, India
Temporal Coverage	January 2018 - December 2023
Temporal Resolution	Hourly observations
Monitoring Stations	14 consistently operational stations
Pollutant Variables	PM2.5, PM10, NO_2, SO_2, CO, O_3
Meteorological Variables	Relative Humidity, Temperature, Wind Speed, Wind Direction, Atmospheric Pressure
Spatial Attributes	Station ID, Latitude, Longitude, Station Type
Coordinate System	WGS84 Geographic Coordinate System
Data Structure	Timestamped station-level multivariate observations

which maps each variable to the interval [0, 1], preventing magnitude imbalance across pollutant and meteorological features. All normalization parameters were estimated using the training set only. The fitted scaler was subsequently applied to validation and test data to prevent information leakage. To capture short-term temporal dependencies, the dataset was structured into sequential samples using a sliding window formulation. For a window length w, the temporal input sequence is defined as Eq. (5).

$$S_t = \{X_{t-w+1}, X_{t-w+2}, \dots, X_t\} \tag{5}$$

where X_t signifies the multivariate feature vector at time t. In this study, a window size of $w = 10$ was adopted to preserve recent pollutant dynamics while maintaining computational efficiency. $X_t^{(i)} \in \mathbb{R}^F$ denote the multivariate feature vector of monitoring station i at time t, where F represents pollutant and meteorological variables. The model learns a mapping: $f: \{X_{t-L+1}^{(i)}, \dots, X_t^{(i)}\} \rightarrow y_{t+\Delta}^{(i)}$ where $y_{t+\Delta}^{(i)}$ denotes the AQI at forecast horizon Δ. In this study, $\Delta = 24$ hours, corresponding to next-day same-hour AQI prediction for each station. The resulting structured representation retains the intrinsic temporal patterns of the original dataset and provides a consistent foundation for subsequent spatio-temporal modeling.

3.3 Model development

3.3.1. CNN-LSTM with temporal attention for temporal feature extraction

Air pollutant concentrations exhibit nonlinear temporal behavior influenced by meteorological variability, emission cycles and urban activity patterns. To capture both short-term fluctuations and long-term dependencies in pollutant evolution, a hybrid CNN-LSTM architecture is employed for temporal feature extraction [26, 27]. The input to the temporal module entails of sequential windows of normalized pollutant features. The convolutional layer extracts localized temporal patterns according to Eq. (6).

$$C_t = \sigma(W_c * X_t + b_c) \tag{6}$$

where W_c denotes convolutional kernels, b_c represents bias, $*$ is the convolution operator and $\sigma(\cdot)$ corresponds to the ReLU activation function. This operation captures short-range temporal variations and local signal transitions. A max-pooling operation is subsequently applied to reduce dimensionality while retaining dominant temporal characteristics. The extracted convolutional features are propagated to the LSTM layer to model long-range temporal dependencies. The gating mechanism of the LSTM is defined as Eqs. (7) to (11).

$$f_t = \sigma(W_f[h_{t-1}, C_t] + b_f) \tag{7}$$

$$i_t = \sigma(W_i[h_{t-1}, C_t] + b_i) \tag{8}$$

$$o_t = \sigma(W_o[h_{t-1}, C_t] + b_o) \tag{9}$$

$$c_t = f_t \odot c_{t-1} + i_t \odot \tanh(W_c[h_{t-1}, C_t] + b_c) \tag{10}$$

$$h_t = o_t \odot \tanh(c_t) \tag{11}$$

where f_t, i_t and o_t denote the forget, input and output gates, respectively; c_t denotes the memory cell state; and h_t corresponds to the hidden representation encoding temporal dependencies. To enhance interpretability and adaptively emphasize influential time intervals, a temporal attention mechanism is incorporated. The attention weight for each time step is computed as Eq. (12).

$$\alpha_t = \frac{\exp(h_t^{\mathrm{T}} W_\alpha)}{\sum_{k=1}^{T} \exp(h_k^{\mathrm{T}} W_\alpha)} \tag{12}$$

The final aggregated temporal embedding is obtained as Eq. (13).

$$H = \sum_{t=1}^{T} \alpha_t h_t \tag{13}$$

where α_t represents normalized attention weights and H denotes the context-aware temporal representation. This module generates a compact embedding that captures nonlinear pollutant evolution patterns across sequential time intervals, serving as the foundational representation for subsequent latent and spatial modeling stages.

3.3.2. Autoencoder-based latent representation for vulnerability modeling

To identify compact and structured representations of pollutant interactions, a nonlinear autoencoder architecture is employed for latent feature extraction [28]. The objective of this model is to reduce high-dimensional input features into a lower-dimensional latent space while conserving essential structural information. The encoder performs nonlinear transformation as Eq. (14).

$$z = \tanh(W_e X + b_e) \tag{14}$$

Received: March 5, 2026. Revised: April 10, 2026.

where X denotes the input feature representation, W_e is the encoder weight matrix, b_e represents the bias term and z represents the latent vector in reduced-dimensional space. The Tanh activation function enables the capture of nonlinear interdependencies among pollutant variables. The decoder reconstructs the original input representation through Eq. (15).

$$\hat{X} = \mathrm{ReLU}(W_d z + b_d) \tag{15}$$

where W_d and b_d represent decoder parameters. The reconstruction objective is formulated using MSE, as shown in Eq. (16).

$$\mathcal{L}_{rec} = \| X - \hat{X} \|_2^2 \tag{16}$$

Minimization of this loss ensures that the latent representation retains dominant structural characteristics of pollutant variability while eliminating redundancy. To categorize latent representations into discrete vulnerability zones, K-Means clustering is applied in the reduced-dimensional space. The clustering objective minimizes within-cluster variance, as shown in Eq. (17).

$$\mathcal{J} = \sum_{i=1}^{K} \sum_{z_j \in C_i} \| z_j - \mu_i \|^2 \tag{17}$$

where μ_i denotes cluster centroids and C_i represents the corresponding clusters. In this study, four clusters are defined to represent varying levels of pollutant vulnerability. This autoencoder-based formulation provides a compact latent feature space suitable for structured vulnerability analysis and dimensionality reduction. It should be emphasized that while the latent representation generated by the encoder is incorporated into the regression pipeline, the subsequent K-Means clustering is performed solely for vulnerability stratification and interpretability analysis. The clustering process does not directly influence the regression loss or model parameter optimization during training.

Beyond interpretability, the latent representation contributes directly to the forecasting capability of the model by improving feature quality and generalization. The encoder learns a compact nonlinear embedding that captures intrinsic correlations among pollutants and meteorological variables, thereby reducing redundancy and filtering noise present in high-dimensional temporal features. This compressed representation acts as a form of feature regularization, encouraging the model to retain only the most informative structures governing pollutant dynamics.

Unlike raw temporal embeddings, which may contain correlated and noisy signals, the latent vector provides a denoised and structured feature space that enhances the stability of downstream regression. The encoded latent features are explicitly integrated into the final prediction pipeline through feature fusion with spatial representations, thereby influencing the AQI prediction process. The ablation study further demonstrates that the inclusion of latent representation consistently reduces prediction error, indicating that it contributes meaningful predictive information beyond temporal and spatial modeling alone. Thus, while clustering supports vulnerability interpretation, the autoencoder itself functions as a principled representation learning mechanism that enhances model robustness, reduces overfitting and improves generalization in spatio-temporal AQI forecasting.

3.3.3. Graph neural network for spatial dependency modelings

Air pollutant dispersion exhibits strong spatial dependency influenced by geographic proximity, meteorological drift and atmospheric transport mechanisms. To model inter-station relationships and diffusion-aware pollutant propagation, a GCN framework is employed [29, 30]. The spatial modeling pipeline implemented in this study is illustrated in Fig 2, demonstrating the sequential stages of spatial modeling, including graph construction using geodesic distance, adjacency normalization, graph convolutional propagation and spatial risk generation.

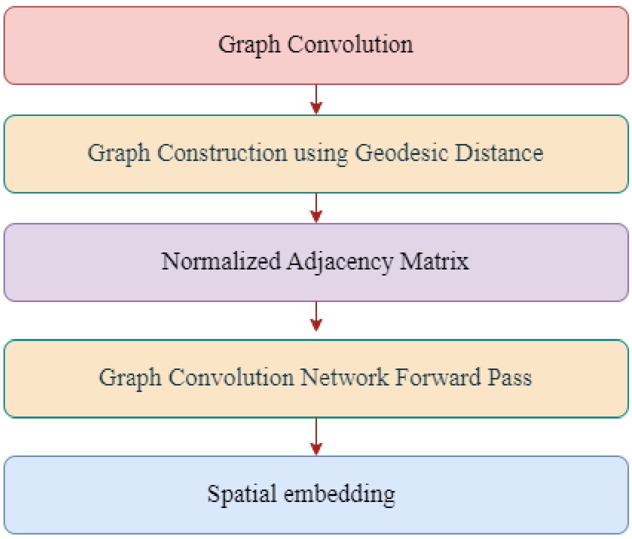

Figure. 2 Steps in spatial graph neural network modelling

Although atmospheric diffusion depends on additional factors such as wind direction and emission patterns, geodesic distance-based adjacency provides a stable and widely adopted approximation for spatial interaction modeling in station-based air quality studies.

The monitoring network is represented as a graph, as given in Eq. (18).

$$G = (V, E) \tag{18}$$

where V signifies the set of monitoring stations (nodes) and E denotes spatial connections (edges) between stations. The spatial adjacency matrix $A \in \mathbb{R}^{N\times N}$ encodes pairwise relationships between stations. Geodesic distance between stations i and j is computed using the great-circle distance formulation, as shown in Eq. (19).

$$d_{ij} = 2R \arcsin\left(\sqrt{\begin{array}{c} sin^2\left(\frac{\phi_i-\phi_j}{2}\right) + \\ cos(\phi_i)\, cos(\phi_j) \\ \cdot\, sin^2\left(\frac{\lambda_i-\lambda_j}{2}\right) \end{array}}\right) \tag{19}$$

where ϕ and λ represent latitude and longitude in radians and R is the Earth's radius. To model spatial attenuation, weighted adjacency is defined using exponential decay, as given in Eq. (20).

$$A_{ij} = \exp\left(\frac{d_{ij}}{\tau}\right) \tag{20}$$

where τ controls the spatial decay rate. Edges are retained when $d_{ij} \leq 80$km, ensuring locality-based spatial interaction. For stable graph learning, symmetric normalization is applied, as shown in Eq. (21).

$$\tilde{A} = D^{-1/2}(A + I)D^{-1/2} \tag{21}$$

where D signifies the degree matrix defined as $D_{ii} = \sum_j A_{ij}$ and I denotes self-loops to preserve node-specific information. Graph convolutional propagation is performed as Eq. (22).

$$H^{(l+1)} = \sigma\left(\tilde{A}H^{(l)}W^{(l)}\right) \tag{22}$$

where $H^{(l)}$ denotes node feature representations at layer l, $W^{(l)}$ represents learnable weight matrices and $\sigma(\cdot)$ corresponds to the ReLU activation function. The normalized adjacency matrix is processed through a two-layer Graph Convolutional Network, whose architectural flow is illustrated in Fig. 3.

The first layer aggregates pollutant information from immediate neighboring stations (local spatial influence). The second layer propagates features across higher-order neighborhoods (regional diffusion modeling). The output layer integrates spatially aggregated representations to generate final station-level risk estimates. The first graph convolution layer performs localized neighborhood aggregation, enabling each station to incorporate pollutant characteristics from directly connected spatial neighbors. This captures immediate spatial influence and localized diffusion patterns. The second graph convolution layer extends propagation to higher-order neighborhoods, modeling broader regional pollutant spread and atmospheric drift effects. Nonlinear activation functions are applied after each convolution to preserve expressive capacity and enable nonlinear spatial representation learning.

The output of the second graph convolution layer produces spatially enriched node representations that encode intrinsic pollutant behavior together with neighboring spatial influence. These representations are subsequently utilized for station-level AQI forecasting and integrated within the overall GeoRiskNet architecture. Through this structured spatial modeling approach, the GCN enables diffusion-aware learning across geographically distributed monitoring stations, ensuring that pollutant AQI forecasting reflects realistic spatial interactions rather than isolated station-level measurements.

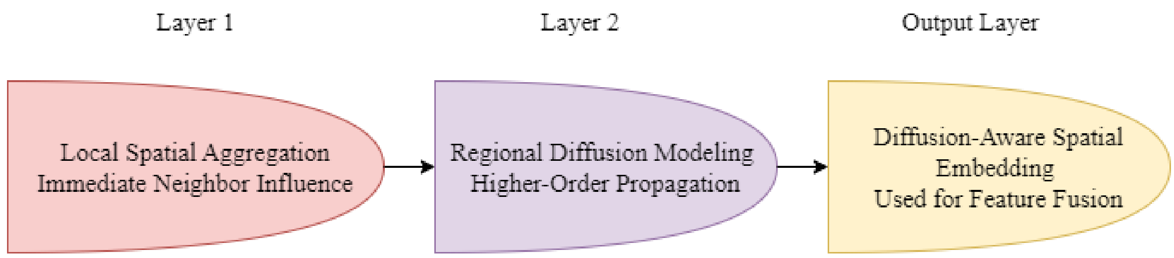

Figure. 3 Two-layer graph convolutional network forward pass

3.3.4. Proposed GeoRiskNet architecture

The proposed GeoRiskNet framework is designed as a unified end-to-end spatio-temporal DL architecture for station-level AQI risk assessment. The model integrates temporal sequence learning, latent vulnerability extraction, spatial graph propagation and regression-based AQI forecasting within a cohesive computational pipeline. The architectural configuration is designed to balance representational capacity, computational efficiency and generalization stability. The model operates on normalized hourly pollutant observations structured into sliding temporal windows of length ten. Each temporal input sample consists of multivariate pollutant features arranged as a two-dimensional tensor of shape (10 × F), where F denotes the number of pollutant and meteorological variables. These sequences are first processed through a one-dimensional convolutional layer composed of 64 filters with a kernel size of 3 and ReLU activation. This layer extracts localized short-term temporal transitions and pollutant fluctuation patterns. A max-pooling layer with pool size 2 follows the convolution stage, reducing temporal dimensionality while preserving dominant signal characteristics.

The pooled feature maps are passed into a LSTM layer consisting of 64 hidden units. The LSTM captures long-range temporal dependencies and nonlinear pollutant interactions across time steps. A dropout layer with a rate of 0.3 is applied after the LSTM to mitigate overfitting and enhance model generalization. The LSTM produces hidden state sequences that are subsequently processed through a temporal attention mechanism. The attention layer computes adaptive importance weights over the hidden states and generates a fixed-length context-aware temporal embedding of dimension 16. This embedding represents a compact summary of pollutant evolution dynamics within each temporal window.

The 16-dimensional temporal embedding is then fed into a nonlinear autoencoder module for latent vulnerability modeling. The encoder involves of a fully connected layer with 8 neurons followed by a latent compression layer of dimension 4 using Tanh activation. This four-dimensional latent vector captures intrinsic pollutant interaction structures and volatility characteristics while enforcing dimensional compactness. The decoder mirrors the encoder architecture, consisting of an 8-neuron hidden layer and a reconstruction layer restoring the 16-dimensional embedding using ReLU activation. Reconstruction regularization ensures that essential pollutant information is preserved within the compressed latent representation. The resulting latent vectors are subsequently clustered into four vulnerability zones using K-Means, enabling structured categorization of station-level pollutant risk patterns. It should be noted that the clustering stage is used exclusively for vulnerability stratification and interpretability analysis and does not directly influence the regression loss during training.

To incorporate spatial interdependencies, monitoring stations are represented as nodes within a geodesic distance-based weighted graph. Spatial connectivity is defined using an 80 km geodesic distance threshold, selected to preserve locality-driven atmospheric interaction while maintaining graph sparsity within the geographic extent of the Chennai metropolitan region. This threshold ensures that spatial influence is modeled among realistically connected monitoring stations without inducing full graph connectivity. The adjacency matrix is constructed using exponential distance-based attenuation and symmetrically normalized with self-loops to stabilize feature propagation. The spatial graph is processed through a two-layer GCN. The first graph convolution layer maps input node features to a 32-dimensional hidden representation using ReLU activation, enabling local neighborhood aggregation and short-range diffusion modeling. A dropout layer with a rate of 0.3 is applied between graph convolution layers to prevent overfitting. The second graph convolution layer reduces the representation to 16 dimensions while propagating higher-order spatial dependencies across the monitoring network. These layers collectively encode diffusion-aware pollutant interactions across geographically distributed stations.

The final predictive representation is constructed by concatenating the 4-dimensional latent vector obtained from the autoencoder encoder with the 16-dimensional spatial embedding produced by the GCN, resulting in a 20-dimensional fused feature vector for each station. This fused representation is passed through a fully connected regression layer that outputs a scalar predicted AQI value. The entire GeoRiskNet architecture is trained jointly in an end-to-end manner, allowing coordinated optimization of temporal dynamics, latent structural representation and spatial diffusion effects. Through this structured design, GeoRiskNet simultaneously captures short-term fluctuations, long-term pollutant evolution, compact vulnerability structure and realistic spatial transport mechanisms. The explicit dimensional transitions and feature fusion strategy ensure architectural transparency, modeling robustness and computational efficiency, providing a coherent multiscale framework for urban air pollution risk assessment.

3.3.5. Regulatory risk interpretation layer

While the proposed GeoRiskNet architecture is trained using a regression objective to forecast continuous next-day Air Quality Index (AQI) values, regulatory risk levels are derived through a structured post-processing layer based on official air quality standards defined by the CPCB, India. This formulation preserves the regression-based learning framework while enabling interpretable categorization of predicted air quality into policy-relevant risk classes. Let $\widehat{AQI}_{i,t+1}$ denote the predicted next-day AQI for monitoring station i at time $t+1$. The regulatory risk level $\mathcal{R}_{i,t+1}$ is obtained by mapping the continuous forecast into discrete CPCB-defined severity categories through Eq. (23).

$$\mathcal{R}_{i,t+1} = \begin{cases} \text{Low} & 0 \le \widehat{AQI}_{i,t+1} \le 100 \\ \text{Moderate} & 100 < \widehat{AQI}_{i,t+1} \le 200 \\ \text{High} & 200 < \widehat{AQI}_{i,t+1} \le 300 \\ \text{Very High} & \widehat{AQI}_{i,t+1} > 300 \end{cases} \tag{23}$$

This categorical transformation defines risk as a deterministic function of forecasted pollutant severity, as shown in Eq. (24).

$$\mathcal{R}_{i,t+1} = f(\widehat{AQI}_{i,t+1}) \tag{24}$$

where $f(\cdot)$ represents the regulatory threshold mapping operator. The CPCB air quality classification scheme used in this study is given in Table 2.

This regulatory interpretation layer does not alter the regression-based training objective, which minimizes the MSE between observed and predicted AQI values. Instead, it provides an additional decision-support interface by translating quantitative forecasts into actionable environmental risk categories aligned with national air quality standards.

By separating continuous AQI prediction from categorical risk interpretation, the proposed framework maintains numerical forecasting precision while ensuring policy-level interpretability. This dual representation enhances the practical utility of the GeoRiskNet architecture for urban environmental monitoring, public health advisories and regulatory planning. The algorithm 1 details the suggested model.

Algorithm I: **GeoRiskNet Framework**

Input: *Hourly multivariate pollutant and meteorological observations* $X_t^{(i)} \in \mathbb{R}^F$ *for station* $i = 1, \dots, N$*, Station geographic coordinates* (ϕ_i, λ_i)*, Window length* $w = 10$*, Forecast horizon* $\Delta = 24$*, Spatial threshold* $d_{th} = 80$ *km, Number of vulnerability clusters* $K = 4$

Output: *Predicted next-day AQI* $\hat{y}_{t+\Delta}^{(i)}$*, Regulatory risk category* $R_{t+\Delta}^{(i)}$

Begin

Step 1: Data Preprocessing

1. *Handle missing values:*
   - *If gap* $\le$ *short duration* $\rightarrow$ *linear interpolation*
   - *Else* $\rightarrow$ *seasonal mean imputation*
2. *Detect outliers:*
   - *Compute* $Z = (x - \mu)/\sigma$
   - *If* $|Z| > 3 \rightarrow$ *apply median smoothing*
3. *Normalize features:*
   $$x' = \frac{x - x_{min}}{x_{max} - x_{min}}$$
4. *Split dataset chronologically: 80% Training, 20% Testing*

Step 2: Sliding Window Construction

For each station i *and time* t*:*

$$S_t^{(i)} = \{X_{t-w+1}^{(i)}, \dots, X_t^{(i)}\}$$

Target: $y_{t+\Delta}^{(i)}$

Step 3: Temporal Feature Extraction (CNN-LSTM + Attention)

For each sequence $S_t^{(i)}$*:*

1. *Apply 1D Convolution:*

Table 2. CPCB AQI-based regulatory risk categories

AQI Range	Air Quality Category	Interpreted Risk Level
0-100	Good / Satisfactory	Low
101-200	Moderate	Moderate
201-300	Poor	High
>300	Very Poor / Severe	Very High

$$C_t = ReLU(W_c * S_t + b_c)$$

2. *Apply MaxPooling*
3. *Pass to LSTM:*

$$h_t = LSTM(C_t)$$

4. *Compute Attention Weights:*

$$\alpha_t = \frac{exp\,(h_t^\top W_\alpha)}{\sum_k exp\,(h_k^\top W_\alpha)}$$

5. *Compute Temporal Embedding:*

$$H^{(i)} = \sum_t \alpha_t\, h_t$$

Output:

Temporal embedding $H^{(i)} \in \mathbb{R}^{16}$

Step 4: Latent Vulnerability Encoding (Autoencoder)

1. *Encode:*

$$z^{(i)} = tanh\,(W_e H^{(i)} + b_e)$$

2. *Decode:*

$$\widehat{H}^{(i)} = ReLU(W_d z^{(i)} + b_d)$$

3. *Minimize reconstruction loss:*

$$L_{rec} = \| H^{(i)} - \widehat{H}^{(i)} \|^2$$

4. *Apply K-Means clustering on* $z^{(i)}$ *with* $K = 4 \rightarrow$ *Obtain vulnerability label* $C^{(i)}$

Step 5: Spatial Graph Construction

1. *Construct graph:*

$$G = (V, E)$$

2. *Compute geodesic distance:*

$$d_{ij}$$

3. *Define adjacency:*

$$A_{ij} = \begin{cases} exp\,(-d_{ij}/\tau), & d_{ij} \le d_{th} \\ 0, & otherwise \end{cases}$$

4. *Normalize adjacency:*

$$\tilde{A} = D^{-1/2}(A + I)D^{-1/2}$$

Step 6: Spatial Feature Propagation (GCN)

For each layer l:

$$H^{(l+1)} = ReLU(\tilde{A} H^{(l)} W^{(l)})$$

After two layers:

$$S^{(i)} \in \mathbb{R}^{16}$$

Spatial embedding obtained.

Step 7: Feature Fusion & Regression

Concatenate:

$$F^{(i)} = [z^{(i)} \parallel S^{(i)}]$$

where $F^{(i)} \in \mathbb{R}^{20}$

Predict AQI:

$$\hat{y}_{t+\Delta}^{(i)} = W_r F^{(i)} + b_r$$

Step 8: Training Objective

Minimize total loss:

$$L = L_{MSE} + \lambda L_{rec}$$

Optimize using Adam.

Step 9: Regulatory Risk Mapping

$$R_{t+\Delta}^{(i)} = \begin{cases} Low & 0 \le \hat{y} \le 100 \\ Moderate & 100 < \hat{y} \le 200 \\ High & 200 < \hat{y} \le 300 \\ Very\ High & \hat{y} > 300 \end{cases}$$

3.4 Experimental setup and reproducibility details

3.4.1. Hardware and software configuration

All experiments were conducted on a high-performance workstation configured to support spatio-temporal deep learning and graph-based computations. The system was equipped with an Intel Core i7 processor, 32 GB of RAM and an NVIDIA GPU with 8 GB of dedicated memory to accelerate model training. GPU acceleration was particularly beneficial for the CNN-LSTM temporal module and graph convolution operations, significantly reducing computational time during iterative backpropagation across the integrated architecture.

The GeoRiskNet model was implemented using the Python programming language. The deep learning components, including the CNN-LSTM, temporal attention mechanism, autoencoder, GCN and regression layer, were developed using the PyTorch framework. Spatial graph construction and graph convolution operations were implemented using the PyTorch Geometric library. Data preprocessing and temporal structuring were performed using NumPy and Pandas, while clustering and performance evaluation were carried out using Scikit-learn. All experiments were executed on a 64-bit operating system.

3.4.2. Training protocol and reproducibility settings

To ensure reproducibility, all experiments were conducted under controlled settings with fixed random seeds applied across NumPy, PyTorch and data loading procedures. The dataset was split chronologically into 80% training and 20% testing to preserve temporal consistency and prevent information leakage. Each model was trained over three independent runs, and the reported performance metrics correspond to the average results to reduce variance arising from stochastic initialization. Training for all deep learning models was conducted for a maximum of 200 epochs with a batch size of 32. Early stopping was applied based on validation loss to prevent overfitting. The Adam optimizer was used with an initial learning rate of 0.001, and weight decay of $1 \times 10^{-5}$was applied for regularization.

Hyperparameters are predefined configuration settings that control the learning behavior and structural characteristics of the model and are tuned prior to training to optimize performance and generalization. The key hyperparameters utilized in the proposed GeoRiskNet model are demonstrated in Table 3.

3.4.3. Fairness and reproducibility of comparative evaluation

Ensuring fairness and reproducibility in comparative evaluation across heterogeneous models is critical for reliable performance assessment. All models were trained using identical data preprocessing, normalization procedures, sliding window construction and chronological train-test splits to maintain consistency in input representation. Each baseline was independently configured and optimized according to model-specific best practices rather than enforcing uniform architectural settings. Sequential models such as LSTM and CNN-LSTM were designed with appropriate temporal input structures, while graph-based models were configured to effectively capture spatial dependencies.

Hyperparameter tuning was performed using a constrained grid search within a predefined search space, with model selection based on validation performance. A comparable tuning budget was allocated across all models to ensure balanced optimization effort. In addition, consistent training protocols, including batch size, epoch limits and early stopping criteria, were applied uniformly to all models. These measures ensure that performance differences arise from genuine modeling capability rather than disparities in optimization strategy or computational effort.

To account for stochastic variation in deep learning training, each model was trained over three independent runs with different random seeds. The reported performance metrics correspond to the average values across runs. The observed variance was minimal, indicating stable model convergence and robustness of the reported results. Furthermore, no model was provided with additional features, extended temporal windows or external data beyond those used in the proposed GeoRiskNet framework, ensuring that performance differences arise solely from modeling capability rather than disparities in input information.

Strict measures were implemented to prevent data leakage throughout the modeling pipeline. All preprocessing operations, including normalization, autoencoder training and latent representation learning, were performed exclusively on the training dataset. The learned transformations were subsequently applied to the test set without re-fitting. Similarly, K-means clustering for vulnerability stratification was conducted only on training embeddings. The spatial graph structure, based on geodesic distances, was constructed independently of temporal observations and remained fixed throughout training and evaluation. These measures ensure that no information from the test set influences model training or representation learning.

4. Results and discussion

4.1 Evaluation metrics

To quantitatively assess the predictive capability of the GeoRiskNet framework, standard regression-based performance metrics were employed. Since the objective involves continuous station-level AQI forecasting, regression measures are appropriate for evaluating model accuracy and generalization performance. All performance metrics reported in this study were computed on Min-Max normalized AQI values scaled to the [0, 1] interval.

Normalization was applied using parameters derived solely from the training dataset to prevent information leakage. The MSE measures the average squared difference between predicted and observed AQI values and is defined as Eq. (25).

$$\text{MSE} = \frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2 \tag{25}$$

where y_i is the observed AQI value, $\hat{y}_i$ represents the predicted AQI value and N is the total number of evaluation samples. RMSE provides error magnitude in the normalized scale and penalizes larger deviations more heavily and is computed as Eq. (26).

Table 3. Key hyperparameters of the proposed GeoRiskNet model

Hyperparameter	Value
Optimizer	Adam
Loss Function	MSE
Initial Learning Rate	0.001
Batch Size	32
Number of Epochs	200
Dropout Rate	0.3
LSTM Hidden Units	64
CNN Filters	64
CNN Kernel Size	3
Temporal Embedding Dimension	16
Latent Dimension	4
Spatial Connectivity Threshold	80 km
Weight Decay	1×10^{-5}

Received: March 5, 2026. Revised: April 10, 2026.

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(y_i - \hat{y}_i)^2} \tag{26}$$

The MAE evaluates the average absolute deviation between predicted and observed AQI values, as given by Eq. (27).

$$\text{MAE} = \frac{1}{N}\sum_{i=1}^{N} | y_i - \hat{y}_i | \tag{27}$$

MAE provides a more robust measure against outliers compared to MSE. The Coefficient of Determination quantifies the proportion of variance in the observed AQI values explained by the model, as given by Eq. (28).

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{N}(y_i - \bar{y})^2} \tag{28}$$

where $\bar{y}$ represents the mean of the observed AQI values. An R^2 value closer to 1 indicates stronger predictive performance. For evaluating the autoencoder-based latent modeling, reconstruction fidelity was assessed using Mean Squared Reconstruction Error, as shown in Eq. (29).

$$\mathcal{L}_{rec} = \| X - \hat{X} \|_2^2 \tag{29}$$

where X denotes the original feature representation and $\hat{X}$ is the reconstructed output from the decoder. To evaluate the quality of vulnerability clustering in latent space, the Silhouette Score was computed as Eq. (30).

$$S = \frac{b-a}{\max(a,b)} \tag{30}$$

where a represents the mean intra-cluster distance and b denotes the mean nearest-cluster distance. Higher values of S specify better-defined and well-separated clusters.

4.2 Temporal forecasting performance

To evaluate the ability of the CNN-LSTM with temporal attention module to capture nonlinear AQI dynamics, the temporal forecasting performance was assessed independently on the test dataset. The model was trained using sliding window sequences of length $w = 10$ and predictions were generated for all monitoring stations across the evaluation period. The quantitative results are summarized in Table 4.

The CNN-LSTM with attention achieves 0.095 RMSE and 0.074 MAE, signifying low average deviation between predicted and observed AQI values. The R2 value of 0.931 suggests that around 93.1% of the variance in temporal AQI behavior is clarified by the model.

The relatively small gap between MSE and MAE values indicates stable prediction performance without excessive influence from extreme outliers. Furthermore, the high R^2 confirms that the model efficiently captures both short-term fluctuations and longer-term temporal trends in AQI values.

Fig. 4 demonstrates a consistent decline in training loss across epochs, indicating effective parameter optimization and convergence of the Temporal CNN-LSTM model.

Table 4. Temporal forecasting performance of CNN-LSTM with attention

Metric	Value
MSE	0.0091
MAE	0.074
RMSE	0.095
R^2	0.931

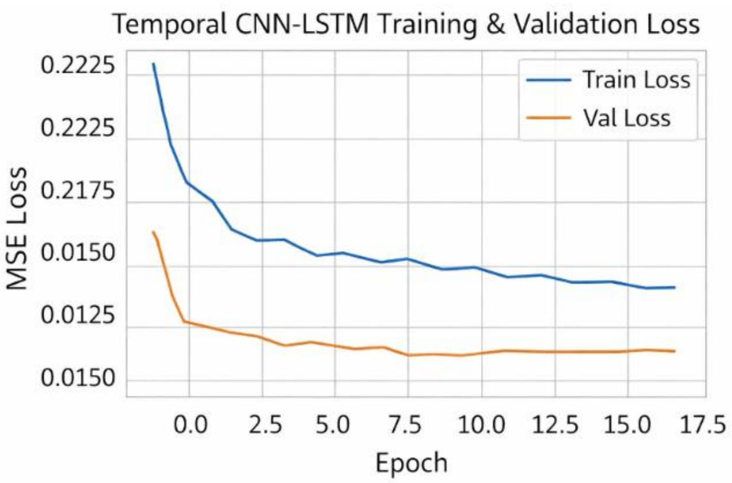

Figure. 4 Temporal CNN-LSTM training and validation loss

 DOI: 10.22266/ijies2026.0630.19

The validation loss decreases rapidly during the initial epochs and subsequently stabilizes, maintaining a lower magnitude than the training loss throughout the training process. The limited divergence between training and validation loss indicates stable convergence behavior without evidence of overfitting. The slightly lower validation loss compared to training loss is attributed to the application of dropout regularization during training, which introduces stochastic noise that is absent during validation evaluation. The stable convergence pattern confirms that the model successfully captured temporal dependencies in the AQI dataset while maintaining robust validation performance.

Fig. 5 highlights the model's ability to replicate temporal trends and short-term fluctuations in AQI values. While the actual series exhibits sharper peaks and higher variability, the forecasted curve closely follows the overall pattern and trend structure. Minor smoothing in the predicted curve indicates that the model effectively captured dominant temporal dynamics while slightly underestimating extreme spikes. The strong alignment between actual and predicted values confirms the model's forecasting reliability and its suitability for short-term air quality prediction within the study region.

4.3 Latent representation and clustering performance

To assess the effectiveness of the autoencoder-based latent modeling, reconstruction fidelity and clustering separability were evaluated. The objective of the autoencoder is to compress high-dimensional temporal embeddings into a compact latent space while preserving essential pollutant structural characteristics. Reconstruction quality was measured using the Mean Squared Reconstruction Error, while clustering quality in the reduced latent space was evaluated using the Silhouette Score. The quantitative results are given in Table 5.

The reconstruction MSE of 0.0062 indicates that the compressed four-dimensional latent space retains the majority of the original temporal information. The low reconstruction error confirms that dimensionality reduction does not significantly distort pollutant feature relationships. The Silhouette score of 0.64 suggests well-defined cluster separation in the latent space. Values above 0.5 generally indicate meaningful clustering structure, demonstrating that the vulnerability zones derived from K-means clustering are statistically coherent.

Fig. 6 demonstrates a structured clustering pattern between latent vulnerability representations and pollutant averages. As the latent feature magnitude increases, pollutant concentrations generally exhibit an upward trend, particularly in higher vulnerability zones. Lower zones display relatively compact clusters with moderate pollutant levels, whereas higher zones show greater dispersion and elevated concentrations. This indicates that the learned latent features effectively capture vulnerability gradients and encode pollutant-related risk variations across spatial clusters. The clear stratification across zones validates the representational strength of the latent embedding mechanism in differentiating environmental vulnerability profiles.

Table 5. Latent representation and clustering performance

Metric	Value
Reconstruction MSE	0.0062
Reconstruction RMSE	0.078
Silhouette Score	0.64

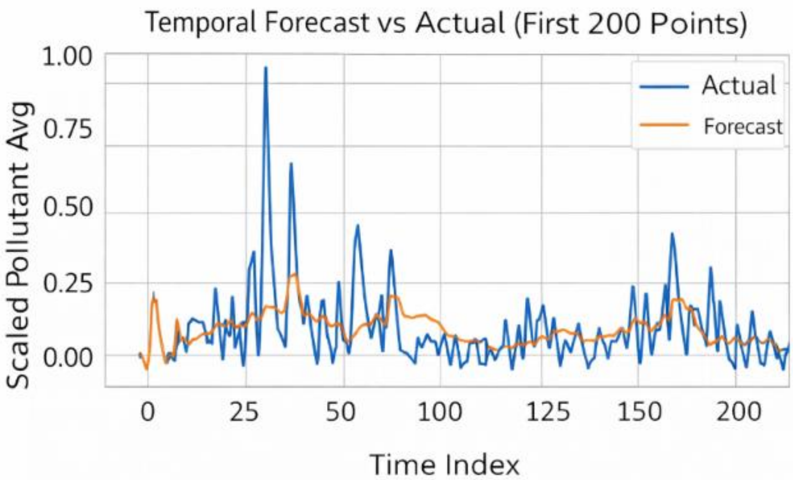

Figure. 5 Temporal forecast versus actual values (First 200 Time Steps)

 DOI: 10.22266/ijies2026.0630.19

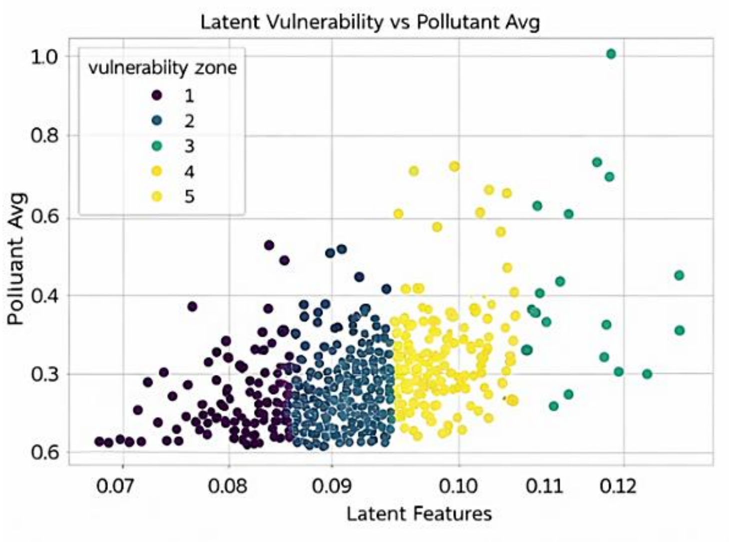

Figure. 6 Latent vulnerability vs pollutant average

4.4 Spatial modeling performance

To assess the efficacy of the Graph Convolutional Network in capturing spatial interdependencies among monitoring stations, the spatial AQI forecasting performance was assessed independently on the test dataset. The GCN aggregates information from geographically connected stations using the normalized adjacency matrix. The quantitative spatial prediction results are summarized in Table 6.

The GCN achieves an RMSE of 0.097 and MAE of 0.075, signifying accurate spatial AQI forecasting across monitoring stations. The R^2 value of 0.921 confirms that more than 92% of the spatial variance in AQI is explained by the graph-based diffusion modeling. Fig. 7 illustrates the MSE loss curves for the GNN model across 200 training epochs.The loss curves indicate rapid convergence of the GNN model during the initial training phase. Both training and validation losses exhibit a sharp decline within the first few epochs, demonstrating efficient parameter learning and fast optimization. After approximately 25-30 epochs, the curves stabilize and maintain consistently low values, indicating that the model reached convergence. The small and stable gap between training and validation loss suggests strong generalization capability and minimal overfitting. The smooth plateau observed in later epochs confirms that the GNN successfully captured spatial dependency structures without instability or divergence, thereby validating the robustness of the spatial risk modeling framework.

Table 6. Spatial AQI forecasting performance of two-layer GCN

Metric	Value
MSE	0.0094
RMSE	0.097
MAE	0.075
R^2	0.921

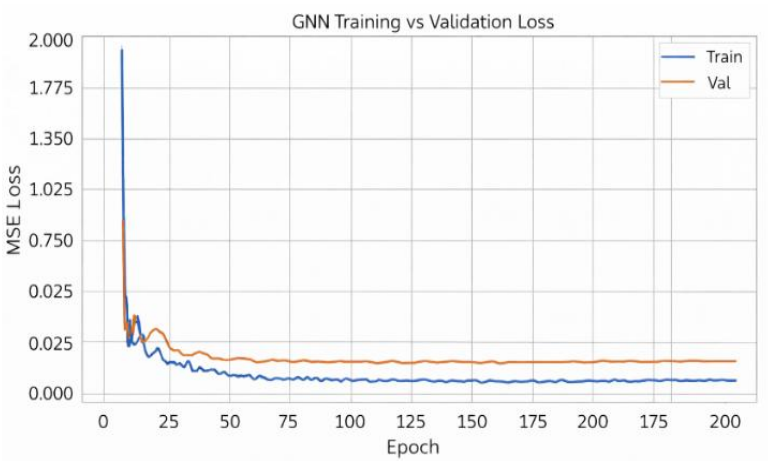

Figure. 7 Training and validation loss of GNN

Fig. 8 shows the relationship between pollutant averages (input features) and AQI-based spatial risk predicted by the GNN model. As pollutant averages increase, the AQI values are predicted by the GNN model rise correspondingly. The dispersion pattern indicates that areas with elevated pollution levels tend to exhibit higher spatial diffusion modeling effects. This confirms that the GNN framework successfully captures spatial interdependencies and neighborhood influence patterns. The upward trajectory across vulnerability zones further supports the effectiveness of graph-based learning in modeling pollutant-driven risk diffusion mechanisms.

Fig. 9 presents the Kernel Density Estimation (KDE) of spatial AQI-based risk across geographic coordinates. The color gradient represents density intensity, with brighter regions indicating higher concentration of risk estimates. The KDE visualization reveals distinct spatial hotspots concentrated in specific geographic clusters. These high-density regions correspond to urbanized and industrial zones exhibiting elevated air pollution levels influencing AQI levels. The smooth spatial gradients demonstrate continuity in spatial diffusion modeling rather than abrupt isolated spikes, confirming that the GCN effectively captures diffusion-aware spatial dependencies. The presence of localized high-intensity pockets supports the hypothesis that AQI risk is geographically heterogeneous and influenced by proximity-based interactions among monitoring stations. This spatial coherence validates the geodesic-based adjacency construction and exponential decay weighting strategy employed in the graph modeling framework.

4.5 Georisknet performance

To evaluate the effectiveness of the fully integrated GeoRiskNet architecture, the final station-level AQI predictions were assessed on the held-out test dataset. The integrated model combines temporal feature extraction, latent vulnerability encoding and spatial diffusion modeling into a unified learning framework.

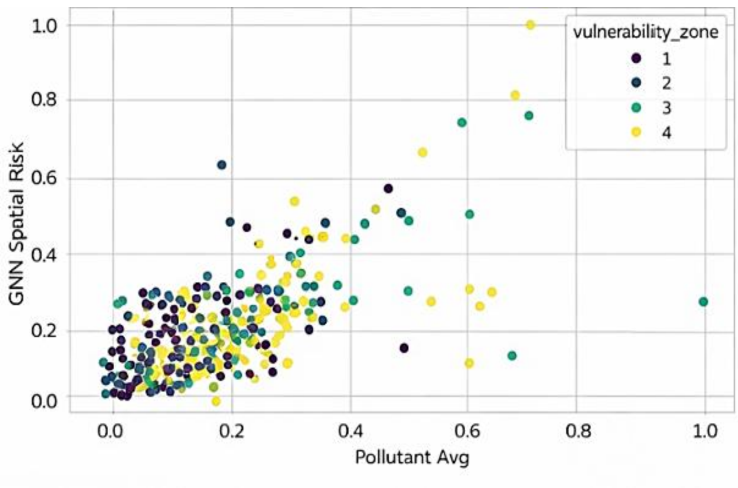

Figure. 8 Pollutant average vs AQI-Based GNN Spatial Risk

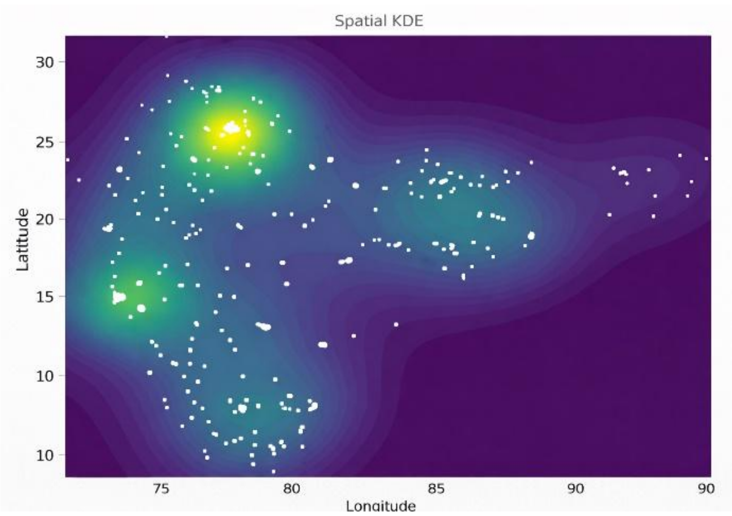

Figure. 9 Spatial KDE of forecasted AQI severity across the Chennai region

Received: March 5, 2026. Revised: April 10, 2026.

The quantitative predictive performance of the GeoRiskNet is summarized in Table 7. The proposed GeoRiskNet achieves an RMSE of 0.087 and MAE of 0.066, signifying low average prediction error across all monitoring stations and timestamps. The R^2 value of 0.953 indicates strong predictive alignment between observed and forecasted AQI values within the evaluation dataset. These values are computed on Min-Max scaled AQI values in the [0,1] interval. While normalization is applied during training to ensure numerical stability, reporting performance exclusively in normalized form limits practical interpretability. Therefore, the evaluation metrics are additionally expressed in approximate AQI units to facilitate real-world understanding.

For interpretability, the normalized error metrics are approximately mapped to AQI units using the CPCB-defined AQI range (0-500). Under this mapping, the RMSE corresponds to approximately 43.5 AQI units, and the MAE corresponds to approximately 33 AQI units. This representation enables direct assessment of prediction error in terms of standard air quality categories used in environmental monitoring and regulatory decision-making.

Compared to individual temporal and spatial modules, the integrated architecture yields improved accuracy, confirming that joint spatio-temporal learning enhances predictive robustness. The reduction in RMSE relative to standalone CNN-LSTM and GCN modules highlights the benefit of combining temporal attention, latent compression and spatial aggregation within a unified framework. To further authorize the regression behavior of the integrated GeoRiskNet framework, predicted AQI values were compared against observed test values and residual diagnostics were analyzed, as illustrated in Fig. 10.

The spatial distribution of integrated air pollution severity is illustrated in Fig. 11. The visualization reveals distinct high-risk clusters concentrated within specific geographic pockets of the Chennai metropolitan region. The spatial heterogeneity observed across latitude and longitude coordinates confirms that pollutant-driven air pollution severity is not uniformly distributed but instead influenced by localized environmental and anthropogenic factors. The gradient-based color mapping clearly identifies high-intensity corridors, supporting the effectiveness of GeoRiskNet in capturing geographically coherent risk patterns. The spatial continuity observed in the map further validates the diffusion-aware modeling capability of the GCN component.

Fig. 12 demonstrates strong positive correlations among the principal risk indicators. The high correlation between pollutant average and predicted AQI risk indicates that pollutant intensity significantly contributes to overall air pollution severity. Similarly, the near-perfect association between GNN risk and combined risk confirms that spatial modeling plays a dominant role in the integrated risk framework. These results validate the coherence and internal consistency of the multi-component risk modeling pipeline, confirming that the combined metric effectively synthesizes pollutant and spatial dependencies into a unified risk representation.

Table. 7 Overall predictive performance of the proposed GeoRiskNet

Metric	Normalized Value	Original AQI Value
MSE	0.0076	-
RMSE	0.087	43.5
MAE	0.066	33.0
R^2	0.953	0.953

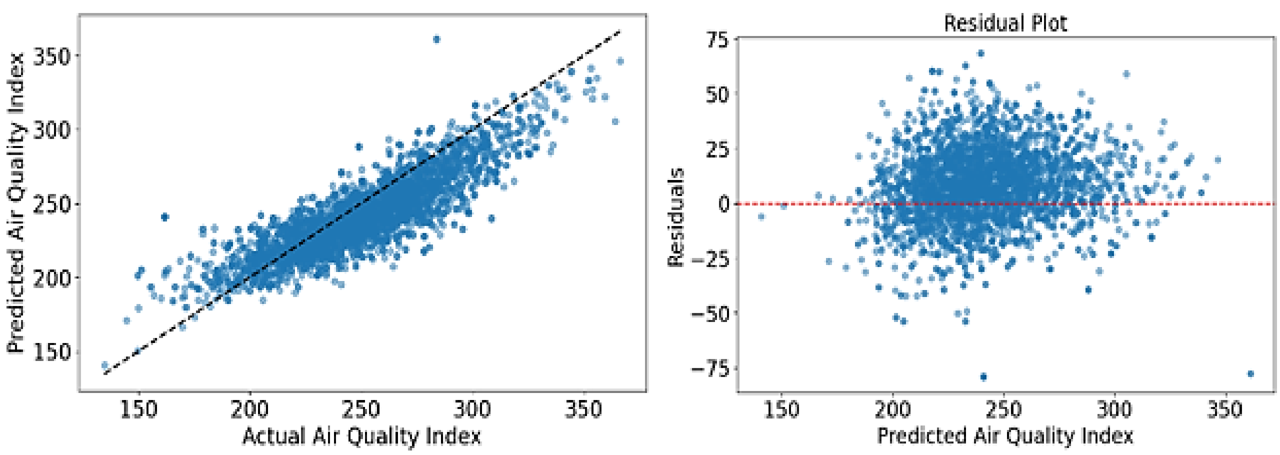

Figure. 10 Predicted vs. actual AQI and residual analysis

 DOI: 10.22266/ijies2026.0630.19

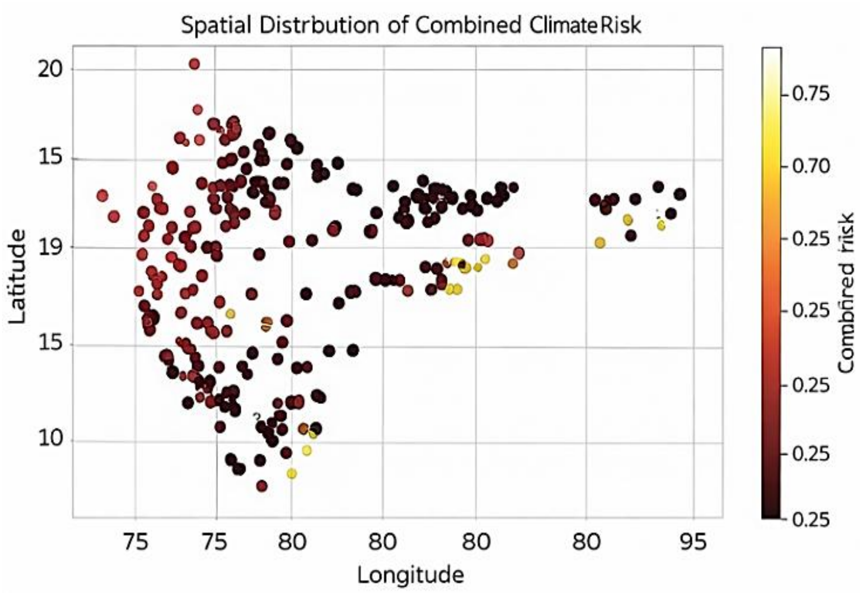

Figure. 11 Spatial distribution of combined air pollution severity

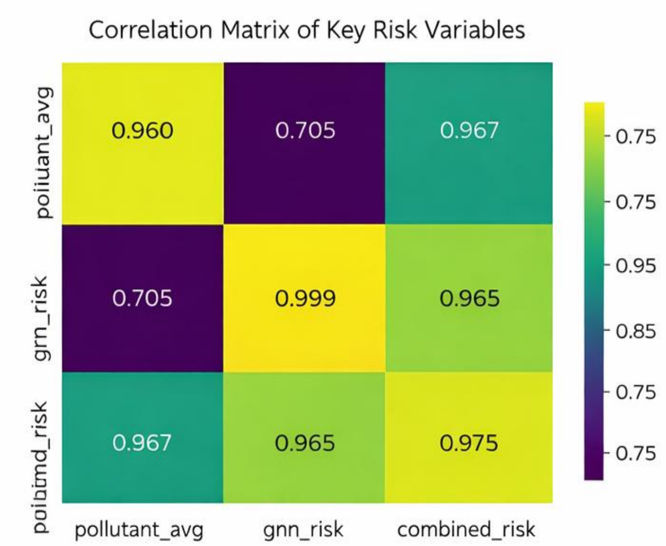

Figure. 12 Correlation matrix of key risk variables

Moreover, the high association between GNN spatial risk and combined risk confirms that spatial propagation effects are effectively incorporated within the integrated model. The consistency among these variables demonstrates the internal coherence of the multi-component modeling pipeline and validates the integration strategy employed in GeoRiskNet.

The distribution of the extracted latent vulnerability feature is shown in Fig. 13 (a), illustrating a well-structured and approximately continuous distribution without excessive skewness. This indicates stable encoding behavior in the compressed feature space. Fig. 13 (b) demonstrates a structured relationship between latent feature magnitude and average pollutant concentration. As latent values increase, pollutant concentrations exhibit a consistent upward trend across vulnerability zones. Higher zones display greater dispersion and elevated concentration levels, confirming that the latent embedding effectively captures air pollution-driven vulnerability gradients influencing AQI. The t-SNE projection in Fig. 13 (c) reveals clear cluster separation among vulnerability zones. The distinct grouping of samples indicates that the autoencoder-generated latent space preserves meaningful structural information necessary for zone differentiation. Furthermore, Fig. 13(d) shows the t-SNE embedding colored by combined risk intensity. A smooth gradient pattern is observed, where higher-risk samples cluster within specific latent regions. This alignment between latent embedding and integrated AQI forecasting confirms the internal coherence of the GeoRiskNet framework.

Received: March 5, 2026. Revised: April 10, 2026.

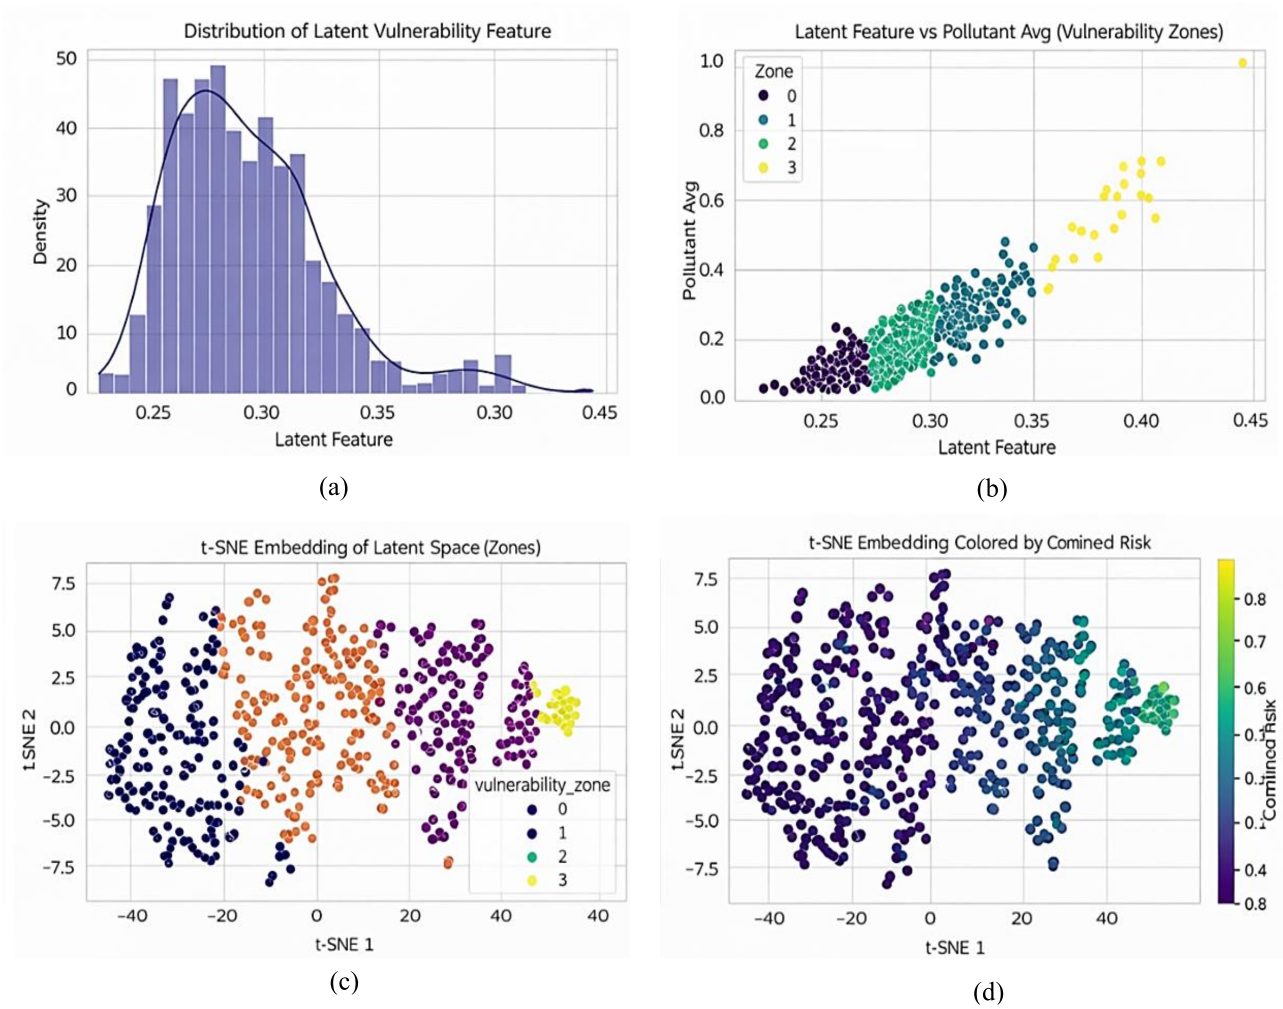

Figure. 13 Latent vulnerability representation and risk embedding in GeoRiskNet: (a) Distribution of the extracted latent feature, (b) Relationship between latent magnitude and pollutant concentration across vulnerability zones, (c) t-SNE projection showing cluster separation in latent space, and (d) t-SNE embedding colored by integrated AQI risk intensity

Number of high-pollution points in Chennai: 9

	country	state	city	station	last_update	latitude	longitude	pollutant_id	pollutant_min	pollutant_max	pollutant_avg
76	India	TamilNadu	Chennai	Arumbakkam, Chennai - TNPCB	7/1/2025 17:00	13.066400	80.211200	SO2	6.0	60.0	21.0
78	India	TamilNadu	Chennai	Perungudi, Chennai - TNPCB	7/1/2025 17:00	12.953300	80.235700	OZONE	10.0	16.0	12.0
80	India	TamilNadu	Chengalpattu	Crescent University, Chengalpattu - TNPCB	7/1/2025 17:00	12.877732	80.083461	OZONE	6.0	6.0	6.0
182	India	TamilNadu	Chennai	Manali Village, Chennai - TNPCB	7/1/2025 17:00	13.166200	80.258400	OZONE	29.0	217.0	47.0
183	India	TamilNadu	Chennai	Manali, Chennai - CPCB	7/1/2025 17:00	13.164544	80.262850	OZONE	8.0	18.0	12.0
284	India	TamilNadu	Chennai	Royapuram, Chennai - TNPCB	7/1/2025 17:00	13.103600	80.290900	OZONE	18.0	24.0	21.0
383	India	TamilNadu	Chennai	Velachery Res. Area, Chennai - CPCB	7/1/2025 17:00	13.005219	80.239812	OZONE	4.0	8.0	7.0
486	India	TamilNadu	Chennai	Alandur Bus Depot, Chennai - CPCB	7/1/2025 17:00	12.909916	80.107854	OZONE	5.0	45.0	25.0
487	India	TamilNadu	Chennai	Kodungaiyur, Chennai - TNPCB	7/1/2025 17:00	13.127800	80.264200	OZONE	7.0	43.0	30.0

Figure. 14 Top high-risk monitoring stations identified in Chennai

The GeoRiskNet model identifies nine monitoring stations within Chennai exhibiting elevated predicted risk levels, as given in Fig. 14. These stations are geographically concentrated within urban-industrial corridors, confirming spatial hotspot formation. The ranking-based node attribution demonstrates the interpretability of the graph-based risk modeling framework. Stations were ranked according to their mean predicted AQI across the test period and the top nine stations exhibiting the highest severity levels were identified as high-pollution hotspots.

4.6 Categorical regulatory risk evaluation

Although the proposed GeoRiskNet framework is optimized using a regression objective to forecast continuous next-day AQI values, regulatory air quality management requires discrete severity categorization based on CPCB standards. Therefore, an additional classification-based evaluation was conducted by mapping both predicted and observed AQI values into four regulatory categories: Low (0-100), Moderate (101-200), High (201-300) and Very High (>300), as defined by CPCB guidelines. C_i^{true} and C_i^{pred} denote the true and predicted CPCB categories for the i^{th} sample in the held-out test dataset comprising approximately 147,000 spatio-temporal instances. Classification performance was evaluated as given in Table 8.

The Low class achieved an F1-score of 0.940, indicating highly reliable identification of favorable air quality conditions. The Moderate and High categories attained F1-scores of 0.896 and 0.879, respectively, reflecting stable and consistent classification performance across mid-severity pollution levels. Importantly, the Very High category achieved an F1-score of 0.910, confirming robust detection capability for hazardous pollution events. The overall classification accuracy reached 91.4%, with a macro-averaged F1-score of 0.906, demonstrating balanced predictive capability without significant bias toward dominant classes. These results confirm that the regression-based AQI forecasts translate effectively into actionable regulatory severity categories suitable for environmental monitoring and public health applications.

4.7 Comparative performance evaluation

To evaluate the effectiveness of the proposed GeoRiskNet framework, its predictive performance was compared against multiple baseline models representing temporal-only, spatial-only and partially integrated architectures. All baseline models were implemented with clearly specified configurations to ensure reproducibility and methodological transparency. The ARIMA model was configured by selecting optimal order parameters (p, d, q) through a grid search over the range 0-3 using the Akaike Information Criterion (AIC) on the training data, ensuring systematic identification of the best temporal configuration. The LSTM baseline was implemented as a two-layer network with 64 hidden units and a dropout rate of 0.3, using a sequence length of 10 to maintain consistency with the temporal input structure.

The CNN-LSTM model employed a one-dimensional convolutional layer with 64 filters and kernel size 3, followed by an LSTM layer with 64 hidden units and dropout regularization of 0.3. The single-layer GCN serves as the baseline spatial model, while the two-layer GCN is introduced to assess the effect of deeper spatial representation learning on AQI prediction performance. The single-layer GCN model utilized a graph convolution layer with a hidden dimension of 32 and ReLU activation to capture spatial dependencies among monitoring stations.

In addition to these baseline models, two representative contemporary architectures were incorporated to reflect recent advances in AQI forecasting. The Attention-based LSTM model consists of a two-layer LSTM with 64 hidden units, followed by dropout (0.3) and a temporal attention mechanism that assigns adaptive importance weights to hidden states. The Transformer-based model includes an input embedding layer with positional encoding, followed by two encoder layers comprising multi-head self-attention (4 heads) and feedforward networks with 128 hidden units, with dropout (0.1) and layer normalization applied within each encoder block. These models are widely adopted in recent literature for capturing long-range temporal dependencies and serve as strong benchmarks for evaluating modern forecasting performance.

Hyperparameter tuning for all baseline models was conducted using a constrained grid search within a predefined search space, with model selection based on validation performance.

Table 8. Categorical regulatory classification performance of GeoRiskNet

Class	Precision	Recall	F1-score
Low	0.936	0.944	0.940
Moderate	0.902	0.891	0.896
High	0.887	0.872	0.879
Very High	0.918	0.903	0.910
Accuracy		**0.914**	
Macro Avg	0.911	0.903	0.906
Weighted Avg	0.915	0.914	0.914

For ARIMA, all combinations of (p, d, q) within the range 0-3 were evaluated. For neural network models, tuning was performed over key parameters including learning rate $(10^{-3}, 5 \times 10^{-4})$, dropout rate (0.2,0.3,0.4) and hidden layer dimensions (32, 64, 128). Each configuration was trained under identical optimization settings, and a comparable tuning budget was allocated to all models to ensure balanced optimization effort. Rather than enforcing identical architectures, each baseline was tuned according to model-specific best practices to ensure that performance differences reflect intrinsic modeling capability rather than unequal optimization.

To ensure fairness in evaluation, all models were trained using identical data preprocessing, normalization, sliding window construction and chronological train-test splitting. Consistent training protocols were applied across all models, including batch size (32), maximum epochs (200) and early stopping based on validation loss. Each model was trained over three independent runs with different random seeds, and the reported performance metrics correspond to averaged results. The variation across runs was minimal (within ±0.002 RMSE), indicating stable convergence and robustness of the comparative evaluation. In addition, statistical comparison of model performance indicated that the observed improvement of GeoRiskNet over baseline models is consistent across runs and not attributable to random initialization effects.

Strict measures were implemented to prevent data leakage throughout the modeling pipeline. All preprocessing operations, including normalization, were performed using training data statistics only and subsequently applied to the test set. The autoencoder-based latent representation was learned exclusively from the training data, and K-Means clustering was also conducted strictly within the training regime. The spatial graph, constructed using geodesic distance, was defined independently of temporal observations and remained fixed throughout training and evaluation. No information from the test set was used during model training, hyperparameter tuning or latent feature extraction, ensuring that the evaluation reflects true generalization performance. This consistent and controlled experimental protocol ensures that the reported comparative performance between GeoRiskNet and baseline models is reliable, reproducible and methodologically sound. The comparative performance results are presented in Table 9.

The results demonstrate that GeoRiskNet achieves lower RMSE compared to evaluated baseline and representative deep learning models across RMSE, MAE and R^2 metrics. Compared to classical ARIMA, the proposed model reduces RMSE by approximately 29.8%, highlighting the advantage of nonlinear deep spatio-temporal modeling. The CNN-LSTM model (RMSE = 0.102) improves upon the LSTM baseline by capturing short-term temporal patterns, but remains limited to temporal dependencies. A comparison between spatial models highlights the importance of graph depth. The single-layer GCN (RMSE = 0.104) captures only immediate neighborhood interactions and performs slightly worse than temporal models due to the absence of sequential information. In contrast, the two-layer GCN improves performance (RMSE = 0.097) by enabling higher-order spatial propagation, demonstrating that deeper graph structures better capture regional pollutant diffusion patterns.

The inclusion of Attention-LSTM and Transformer models enables a more rigorous comparison with contemporary deep learning approaches. While the Transformer model demonstrates strong performance due to its ability to capture long-range temporal dependencies, it does not explicitly account for spatial interdependencies or latent structural representations. Similarly, Attention-LSTM improves temporal feature weighting but remains limited to sequential modeling. In contrast, the proposed GeoRiskNet framework integrates temporal attention, latent feature compression and spatial graph-based diffusion within a unified architecture. This coordinated interaction allows the model to simultaneously capture temporal evolution, denoised feature structure and cross-regional air pollution dynamics influencing AQI, leading to improved predictive accuracy and generalization performance. The proposed GeoRiskNet further improves performance beyond both temporal and spatial baselines, reducing RMSE to 0.087, indicating that latent feature compression contributes to noise reduction and improved feature compactness. The improvement in R^2 to 0.953 confirms that the proposed architecture captures AQI variability more

Table 9. Comparison of GeoRiskNet with baseline models

Model	RMSE	MAE	R^2
ARIMA	0.124	0.101	0.842
LSTM	0.108	0.086	0.902
CNN-LSTM	0.102	0.081	0.914
GCN (1-layer)	0.104	0.083	0.909
GCN (2-layer)	0.097	0.075	0.921
Attention-LSTM	0.092	0.072	0.938
Transformer	0.089	0.069	0.946
Proposed GeoRiskNet	**0.087**	**0.066**	**0.953**

comprehensively than both conventional baselines and representative deep learning models.

A progressive improvement in performance is observed across model categories, starting from statistical models (ARIMA), followed by temporal models (LSTM, CNN-LSTM), spatial models (GCN), and advanced architectures (Attention-LSTM, Transformer), with GeoRiskNet achieving the best performance through unified integration of temporal, spatial and latent representations. These findings validate the effectiveness of combining temporal dynamics, latent feature compression and spatial diffusion modeling within a unified learning framework, demonstrating that GeoRiskNet is competitive with representative deep learning approaches while offering enhanced modeling capability and interpretability.

4.8 Ablation study

To investigate the influence of each architectural component in the proposed GeoRiskNet, an ablation study was conducted by systematically removing individual modules while keeping the training configuration unchanged. This analysis evaluates the impact of temporal attention, latent vulnerability modeling and spatial graph convolution on overall prediction performance. All configurations were trained under identical hyperparameters, data splits and optimization settings to ensure fair comparison. The results are presented in Table 10.

The baseline CNN-LSTM configuration achieves an RMSE of 0.102, demonstrating moderate temporal learning capability. The addition of the attention mechanism reduces RMSE to 0.095, confirming that adaptive temporal weighting improves AQI forecasting accuracy. Incorporating the spatial GCN module further reduces RMSE to 0.091, indicating that modeling inter-station spatial dependencies enhances predictive consistency. Finally, the inclusion of the autoencoder-based latent vulnerability modeling yields the best performance, reducing RMSE to 0.087 and increasing R^2 to 0.953.

The improvement observed in the final GeoRiskNet configuration compared to the CNN-LSTM + Attention + GCN model is attributable to the inclusion of the autoencoder-based latent representation, which enhances feature compactness and reduces noise, thereby improving generalization. The progressive reduction in error across configurations demonstrates that each module contributes positively to model performance. Notably, the inclusion of the latent representation leads to further reduction in RMSE and MAE beyond temporal and spatial modeling alone. This indicates that the autoencoder-based feature compression improves representation quality and contributes directly to predictive accuracy, rather than serving solely as an auxiliary interpretability component.

4.9 Robustness analysis under high pollution regimes

To evaluate the robustness of the proposed GeoRiskNet model under critical pollution conditions, an additional regime-based error analysis was conducted focusing on high-AQI scenarios, as in Table 11. Since regulatory decision-making is particularly sensitive to severe pollution events, it is important to assess whether the model maintains predictive reliability during such episodes. The test dataset was stratified based on CPCB-defined AQI categories, and model performance was analyzed specifically for higher pollution ranges (AQI > 200). This regime includes “Poor”, “Very Poor” and “Severe” air quality conditions, which represent the most critical environmental and public health scenarios.

Table 10. Ablation study of GeoRiskNet components

Model Configuration	RMSE	MAE	R^2
CNN-LSTM (baseline temporal)	0.102	0.081	0.914
CNN-LSTM + Attention	0.095	0.074	0.931
CNN-LSTM + Attention + GCN	0.091	0.071	0.941
GeoRiskNet (with latent)	**0.087**	**0.066**	**0.953**

Table 11. Regime-wise performance of GeoRiskNet across AQI levels

AQI Range	Category	RMSE	MAE
0-100	Low	0.075	0.058
101-200	Moderate	0.083	0.064
201-300	High	0.095	0.072
>300	Very High	0.108	0.081

The results indicate that prediction error increases with AQI severity, which is expected due to higher variability and nonlinear behavior during extreme pollution events. However, the increase in error remains controlled, and the model continues to capture overall trends effectively even in high-AQI regimes. Importantly, despite slightly higher RMSE in severe conditions, the model maintains strong classification performance, as reflected in the high F1-scores for the “Very High” category. This suggests that while exact numerical prediction becomes more challenging during extreme events, the model remains reliable for identifying critical pollution levels relevant to regulatory action. Overall, the GeoRiskNet framework demonstrates robust performance across varying pollution regimes, with stable behavior under moderate conditions and acceptable degradation under extreme AQI levels, confirming its suitability for real-world environmental monitoring.

5. Discussion

Recent advances in air quality forecasting have increasingly leveraged machine learning and deep learning techniques to model complex pollutant dynamics. Hybrid approaches such as Random Forest combined with ARIMA have demonstrated improved handling of nonlinear relationships and temporal residuals [15]. Attention-based LSTM architectures and recurrent models have shown strong capability in capturing temporal dependencies, achieving RMSE values around 39.48 in urban AQI prediction tasks [16]. Similarly, recurrent architectures such as GRU and LSTM have demonstrated strong predictive performance across multiple cities, though their effectiveness varies significantly depending on regional pollution characteristics [17].

Hybrid spatio-temporal models have also been explored to incorporate both spatial and temporal information. For instance, CNN-ARIMAframeworks achieved RMSE values of approximately 44.60 while modeling multi-city AQI data [18]. Ensemble and boosting-based approaches such as XGBoost and LightGBM have reported high predictive accuracy under specific regional conditions; however, their performance is often sensitive to local data characteristics and may lack robustness across varying pollution regimes [19]. Recent deep learning frameworks integrating CNN-LSTM architectures and multi-source inputs have further improved prediction accuracy, achieving RMSE values of approximately 12.6 μg/m³ for AQI prediction [20]. However, such models often rely on heterogeneous data sources and may face challenges related to data availability, scalability and consistency across regions.

Despite these advancements, several limitations remain evident in existing studies. Many approaches rely on single-station or region-specific datasets, limiting their ability to capture spatial interactions across monitoring locations [16, 24]. Hybrid and ensemble models, while improving regression performance, often lack structured mechanisms to model spatial diffusion and inter-station dependencies [17, 21]. Additionally, several studies emphasize either temporal modeling or spatial analysis independently, leading to fragmented representations of air pollution dynamics [15, 18]. Even in multi-source frameworks and IoT-integrated systems, unified graph-based spatial reasoning and end-to-end learning remain insufficiently addressed [20, 23].

The proposed GeoRiskNet framework addresses these limitations through a unified spatio-temporal modeling strategy that integrates temporal attention, latent feature compression and graph-based spatial propagation. Unlike conventional approaches that treat temporal and spatial components separately, the proposed architecture enables complementary interaction among these modules. Temporal attention enhances the modeling of dynamic pollutant fluctuations, latent representation reduces redundancy and noise, and graph convolution facilitates effective spatial information propagation across monitoring stations.

In terms of predictive performance, the proposed GeoRiskNet achieves an RMSE of approximately 43.5 AQI units, demonstrating competitive and improved performance compared to several existing approaches. For instance, the CNN-ARIMA hybrid model reported RMSE of 44.60 [18], while attention-based LSTM models reported RMSE values around 39.48 under single-region settings [16]. The proposed framework not only maintains comparable accuracy but also extends modeling capability by incorporating spatial dependencies and latent vulnerability representation, which are not explicitly addressed in many prior studies.

Furthermore, unlike models that focus solely on numerical prediction, the proposed approach aligns forecasting outputs with CPCB-defined AQI categories, enhancing its applicability for regulatory and decision-making purposes. The inclusion of latent vulnerability representation further improves robustness by providing a structured and denoised feature space, contributing to stable performance across varying pollution regimes. Overall, the GeoRiskNet framework provides a comprehensive and practically relevant solution for AQI forecasting by addressing key limitations in existing literature and demonstrating consistent predictive capability. The integration of temporal, spatial and latent representations within a unified architecture enables improved generalization and robustness, making the model suitable for real-world environmental monitoring and policy-oriented applications.

6. Conclusion

Rapid urban growth and increasing anthropogenic emissions have amplified the need for accurate, interpretable and location-aware air quality forecasting systems. Reliable AQI prediction models are essential for timely public health advisories and informed environmental governance. This study presented GeoRiskNet, a unified spatio-temporal graph deep learning framework for next-day AQI forecasting with integrated regulatory risk interpretation. The architecture jointly captures nonlinear temporal air pollution dynamics influencing AQI, latent vulnerability patterns and spatial diffusion across interconnected monitoring stations within an end-to-end learning structure. By integrating CNN-LSTM modeling with temporal attention, nonlinear latent feature compression and a two-layer GCN, the framework effectively models short-term variability, long-range dependencies and cross-regional pollutant transport dynamics. Experimental evaluation using six years of hourly data from fourteen CPCB monitoring stations in the Chennai metropolitan region demonstrates strong predictive capability. GeoRiskNet achieved an RMSE of 0.087, MSE of 0.0076, MAE of 0.066 and R2 of 0.953 on the held-out test dataset, outperforming ARIMA and other standalone deep learning baselines. The model reduced RMSE by approximately 29.8% compared to ARIMA, confirming the benefit of unified spatio-temporal integration. Regulatory severity mapping based on CPCB AQI thresholds achieved 91.4% overall accuracy with a macro F1-score of 0.906, signifying balanced and reliable classification across pollution levels. The ablation study further validated that each component, temporal attention, spatial graph modeling and latent vulnerability encoding, contributes incrementally to performance improvement. Future research can extend this framework toward detailed station-wise AQI computation, localized pollutant contribution analysis and dynamic spatial threshold adaptation to enhance micro-level environmental monitoring. Overall, GeoRiskNet offers a scalable, diffusion-aware and policy-relevant approach for accurate and interpretable urban air pollution forecasting.

Conflicts of interest

The authors declare no conflict of interest.

Author contributions

The contributions of the authors to this research are as follows: conceptualization, methodology, formal analysis by first and second author: investigation, resources, writing, draft preparation, writing, review and editing, by First author.

Acknowledgments

I would like to express my sincere gratitude to all those who contributed to the completion of this research paper. I extend my heartfelt thanks to my supervisor, my family, my colleagues and fellow researchers for their encouragement and understanding during the demanding phases of this work.

Notation list

Symbol	Description
X	Input multivariate feature matrix (pollutant + meteorological variables)
x_t	Feature vector at time step (t)
y_t	Observed AQI value at time (t)
$\hat{y}_t$	Predicted AQI value at time (t)
(T)	Total number of time steps
(w)	Sliding window length
(h)	Forecast horizon (next-day prediction)
(z)	Latent feature representation (autoencoder output)
W_e, b_e	Encoder weight matrix and bias
W_d, b_d	Decoder weight matrix and bias
L_{rec}	Reconstruction loss (MSE of autoencoder)
C_i	Cluster (i) in K-Means
μ_i	Centroid of cluster (i)
(V)	Set of nodes (monitoring stations)
(E)	Set of edges (spatial connections)
(A)	Adjacency matrix
$\tilde{A}$	Normalized adjacency matrix
(D)	Degree matrix
$H^{(l)}$	Node features at layer (l) in GCN
$W^{(l)}$	Weight matrix of GCN layer (l)
σ	Activation function (ReLU/Tanh)
d_{ij}	Geodesic distance between stations (i) and (j)
θ	Spatial decay parameter
α	Attention weight at time step (t)
h_t	Hidden state of LSTM at time (t)
c_t	LSTM cell state
$\hat{X}$	Reconstructed input from autoencoder
(N)	Number of samples

Received: March 5, 2026. Revised: April 10, 2026.

References

[1] B. Chen and H. Kan, "Air pollution and population health: a global challenge", *Environmental Health and Preventive Medicine*, Vol. 13, No. 2, pp. 94–101, 2008, doi: 10.1007/s12199-007-0018-5.

[2] Q. W. Wang, "Urbanization and global health: the role of air pollution", *Iranian Journal of Public Health*, Vol. 47, No. 11, p. 1644, 2018.

[3] S. A. Sarkodie, P. A. Owusu, and T. Leirvik, "Global effect of urban sprawl, industrialization, trade and economic development on carbon dioxide emissions", *Environmental Research Letters*, Vol. 15, No. 3, p. 034049, 2020, doi: 10.1088/1748-9326/ab7640.

[4] B. N. Aljafen, N. Shaikh, J. M. AlKhalifah, and S. A. Meo, "Effect of environmental pollutants particulate matter (PM2.5, PM10), nitrogen dioxide (NO2), sulfur dioxide (SO2), carbon monoxide (CO) and ground level ozone (O3) on epilepsy", *BMC Neurology*, Vol. 25, No. 1, p. 133, 2025, doi: 10.1186/s12883-025-04142-3.

[5] S. A. Meo, N. Shaikh, and M. Alotaibi, "Association between air pollutants particulate matter (PM2.5, PM10), nitrogen dioxide (NO2), sulfur dioxide (SO2), volatile organic compounds (VOCs), ground-level ozone (O3) and hypertension", *Journal of King Saud University-Science*, Vol. 36, No. 11, p. 103531, 2024, doi: 10.1016/j.jksus.2024.103531.

[6] I. Manisalidis, E. Stavropoulou, A. Stavropoulos, and E. Bezirtzoglou, "Environmental and health impacts of air pollution: a review", *Frontiers in Public Health*, Vol. 8, p. 14, 2020, doi: 10.3389/fpubh.2020.00014.

[7] Q. Liao, M. Zhu, L. Wu, X. Pan, X. Tang, and Z. Wang, "Deep learning for air quality forecasts: a review", *Current Pollution Reports*, Vol. 6, No. 4, pp. 399–409, 2020, doi: 10.1007/s40726-020-00159-z.

[8] Z. Zhang, C. Johansson, M. Engardt, M. Stafoggia, and X. Ma, "Improving 3-day deterministic air pollution forecasts using machine learning algorithms", *Atmospheric Chemistry and Physics*, Vol. 24, No. 2, pp. 807–851, 2024, doi: 10.5194/acp-24-807-2024.

[9] E. Alotaibi and N. Nassif, "Artificial intelligence in environmental monitoring: in-depth analysis", *Discover Artificial Intelligence*, Vol. 4, No. 1, p. 84, 2024, doi: 10.1007/s44163-024-00198-1.

[10] A. Hamdan, K. I. Ibekwe, E. A. Etukudoh, A. A. Umoh, and V. I. Ilojianya, "AI and machine learning in climate change research: a review of predictive models and environmental impact", *World Journal of Advanced Research and Reviews*, Vol. 21, No. 1, pp. 1999–2008, 2024.

[11] X. Li, L. Peng, Y. Hu, J. Shao, and T. Chi, "Deep learning architecture for air quality predictions", *Environmental Science and Pollution Research*, Vol. 23, No. 22, pp. 22408–22417, 2016.

[12] P. W. Soh, J. W. Chang, and J. W. Huang, "Adaptive deep learning-based air quality prediction model using the most relevant spatial-temporal relations", *IEEE Access*, Vol. 6, pp. 38186–38199, 2018.

[13] Z. Sha, A. Layton, B. Heydari, and D. L. Van Bossuyt, "Networks and graphs for engineering systems and design", *Journal of Computing and Information Science in Engineering*, Vol. 25, No. 6, p. 060301, 2025, doi: 10.1115/1.4068457.

[14] F. I. Hairuddin, S. Azri, and U. Ujang, "Graph-based spatio-temporal analytics: a systematic review", *Networks and Spatial Economics*, 2026, doi: 10.1007/s11067-026-09722-5.

[15] A. Yenkikar, V. P. Mishra, M. Bali, and T. Ara, "Explainable forecasting of air quality index using a hybrid random forest and ARIMA model", *MethodsX*, Vol. 15, p. 103517, 2025, doi: 10.1016/j.mex.2025.103517.

[16] H. Karnati, A. Soma, A. Alam, and B. Kalaavathi, "Comprehensive analysis of various imputation and forecasting models for predicting PM2.5 pollutant in Delhi", *Neural Computing and Applications*, Vol. 37, No. 17, pp. 11441–11458, 2025.

[17] A. Govande, R. Attada, and K. K. Shukla, "Predicting PM2.5 levels over Indian metropolitan cities using recurrent neural networks", *Earth Science Informatics*, Vol. 18, No. 1, p. 1, 2025, doi: 10.1007/s12145-024-01491-4.

[18] S. Jayaraman, N. T, A. S, and G. S, "Enhancing urban air quality prediction using time-based-spatial forecasting framework", *Scientific Reports*, Vol. 15, No. 1, p. 4139, 2025, doi: 10.1038/s41598-024-83248-z.

[19] D. Mishra, A. Sahu, and S. Naman, "Implementation of boosting algorithms in predicting air quality index of South Indian cities", *Engineering Innovations*, Vol. 14, pp. 99–115, 2025.

[20] T. Raza and K. Singh, "AI-driven multisource data fusion for real-time urban air quality forecasting and health risk assessment", *Journal of Scientific Innovation and Advanced Research*, Vol. 1, No. 2, pp. 200–206, 2025.

Received: March 5, 2026. Revised: April 10, 2026.

[21] A. A. M. Ahmed, S. J. J. Jui, E. Sharma, M. H. Ahmed, N. Raj, and A. Bose, "An advanced deep learning predictive model for air quality index forecasting with remote satellite-derived hydro-climatological variables", *Science of the Total Environment*, Vol. 906, p. 167234, 2024, doi: 10.1016/j.scitotenv.2023.167234.

[22] A. M. Rosca, M. Carbureanu, and A. Stancu, "Data-driven approaches for predicting and forecasting air quality in urban areas", *Applied Sciences*, Vol. 15, No. 8, p. 4390, 2025, doi: 10.3390/app15084390.

[23] M. Rajesh, R. G. Babu, U. Moorthy, and S. V. Easwaramoorthy, "Machine learning-driven framework for realtime air quality assessment and predictive environmental health risk mapping", *Scientific Reports*, Vol. 15, No. 1, p. 28801, 2025, doi: 10.1038/s41598-025-14214-6.

[24] P. Patel, S. Patel, K. Shah, K. Desai, S. Patel, M. Shah, *et al.*, "A systematic study on PM2.5 and PM10 concentration prediction in air pollution using machine learning and deep learning model", *Environmental Chemistry and Ecotoxicology*, Vol. 7, pp. 1401–1415, 2025, doi: 10.1016/j.enceco.2025.07.001.

[25] Central Pollution Control Board (CPCB), "National Air Quality Index", [Online]. Available: https://airquality.cpcb.gov.in/AQI_India/.

[26] M. Alhussein, K. Aurangzeb, and S. I. Haider, "Hybrid CNN-LSTM model for short-term individual household load forecasting", *IEEE Access*, Vol. 8, pp. 180544–180557, 2020, doi: 10.1109/ACCESS.2020.3028281.

[27] A. U. Rehman, A. K. Malik, B. Raza, and W. Ali, "A hybrid CNN-LSTM model for improving accuracy of movie reviews sentiment analysis", *Multimedia Tools and Applications*, Vol. 78, No. 18, pp. 26597–26613, 2019, doi: 10.1007/s11042-019-07788-7.

[28] A. Oluwasanmi, M. U. Aftab, E. Baagyere, Z. Qin, M. Ahmad, and M. Mazzara, "Attention autoencoder for generative latent representational learning in anomaly detection", *Sensors*, Vol. 22, No. 1, p. 123, 2022, doi: 10.3390/s22010123.

[29] W. Fan, Y. Ma, Q. Li, J. Wang, G. Cai, J. Tang, and D. Yin, "A graph neural network framework for social recommendations", *IEEE Transactions on Knowledge and Data Engineering*, Vol. 34, No. 5, pp. 2033–2047, 2020, doi: 10.1109/TKDE.2020.3008732.

[30] S. Li and J. Corney, "Multi-view expressive graph neural networks for 3D CAD model classification", *Computers in Industry*, Vol. 151, p. 103993, 2023, doi: 10.1016/j.compind.2023.103993.