RCHUB: JOURNAL OF COMPUTATIONAL INTELLIGENCE SCIENCE AND ENGINEERING (JCISE)
VOL: 2, ISSUE 4, OCTOBER - DECEMBER 2026
RCHUB PUBLISHER
ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

MACHINE LEARNING, DEEP LEARNING, AND EXPLAINABLE AI FOR MULTI-SOURCE AIR POLLUTION FORECASTING: A REVIEW

[1]*Angamuthu T*
Research scholar, Department of Information Technology,
Government Arts College, Coimbatore - 641018
Email: angamuthu.t@gmail.com
[2]*Dr.D.Devakumari*
Associate Professor, Department of Information Technology,
Government Arts College, Coimbatore - 641018
Email: ddcbe04@gmail.com

Abstract

Air pollution remains one of the most pressing environmental and public health challenges of the current decade, motivating sustained research into computational methods for forecasting pollutant concentrations. This review synthesizes recent literature on machine learning (ML), deep learning (DL), hybrid ensemble architectures, multi-source data integration, and explainable artificial intelligence (XAI) as applied to air quality forecasting. Tree-based ensemble methods such as XGBoost, LightGBM, and CatBoost, alongside recurrent architectures such as Long Short-Term Memory (LSTM) networks and their convolutional and attention-augmented variants, dominate the recent literature, frequently combined into hybrid or stacked frameworks to balance short-term responsiveness with long-term temporal dependency modeling. A growing body of work integrates heterogeneous data streams, including ground-based sensor networks, meteorological records, traffic and emission inventories, and satellite-derived aerosol and trace-gas products, to improve spatial and temporal coverage. In parallel, explainability techniques, most notably SHapley Additive exPlanations (SHAP) and Local Interpretable Model-Agnostic Explanations (LIME), are increasingly incorporated to address the opacity of high-performing black-box models and to support policy-relevant decision-making. This review consolidates findings across these strands, identifies methodological gaps relating to model transparency, heterogeneous data fusion, and generalizability across geographic contexts, and outlines directions for future hybrid explainable forecasting frameworks.

Keywords

Air pollution forecasting; machine learning; deep learning; XGBoost; LSTM; explainable artificial intelligence; SHAP; LIME; multi-source data fusion; air quality index

1. Introduction

Ambient air pollution continues to pose severe risks to human health and environmental sustainability, driving demand for forecasting systems that can anticipate pollutant concentrations with sufficient lead time to support public warnings and policy responses. Pollutants of primary concern include fine and coarse particulate matter (PM2.5 and PM10), nitrogen dioxide (NO2), sulfur dioxide (SO2), and ground-level ozone (O3), each governed by a distinct mix of meteorological, emission-related, and photochemical processes. Early forecasting approaches relied on deterministic chemical transport models and simple

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

statistical regressions, but these methods often struggled to capture the nonlinear, non-stationary behavior of urban pollution dynamics.

Over the past decade, the field has shifted decisively toward data-driven approaches. Ensemble tree-based learners and deep sequence models have repeatedly demonstrated superior predictive accuracy over classical statistical baselines, while the proliferation of heterogeneous data sources — including low-cost sensor networks, traffic cameras, and Earth-observation satellites — has expanded the feature space available to these models. At the same time, the improved accuracy of complex ensembles and deep networks has come at the cost of interpretability, prompting a parallel line of research into explainable AI methods intended to make model predictions auditable and trustworthy for environmental stakeholders and regulators.

2. Air Pollution Forecasting: A Review

This review synthesizes recent work across four interconnected themes: (i) machine learning approaches to pollutant forecasting, (ii) deep learning and temporal modeling, (iii) hybrid and ensemble architectures, (iv) multi-source data integration, and (v) explainability techniques. It concludes by identifying open research gaps that motivate the development of hybrid, explainable, multi-source forecasting frameworks.

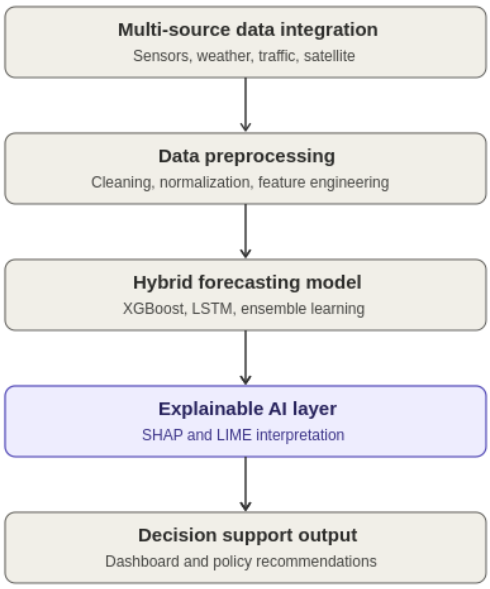


Figure 1. General block diagram of the proposed hybrid explainable AI air pollution forecasting framework.

2.1 Machine Learning Approaches for Air Quality Forecasting

Gradient-boosted tree ensembles are among the most consistently high-performing model families reported in the recent literature. Comparative evaluations have repeatedly found that XGBoost and AdaBoost perform strongly for PM2.5 prediction, while LightGBM has

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

achieved high accuracy for PM10 forecasting in South Asian urban settings [1]. A broader synthesis covering 155 studies published between 2011 and 2021 found that time-series-oriented models were particularly effective for pollutant forecasting and highlighted the growing need for explainability alongside predictive accuracy [1]. Related comparative work applying support vector regression and random forest models to multi-city air quality index (AQI) prediction found that class-balancing techniques such as SMOTE improved random forest performance in terms of root mean square error [1].

Boosting algorithms have continued to be refined for short-term forecasting tasks. In a study conducted in Ho Chi Minh City, LightGBM outperformed XGBoost, LSTM, and CNN-LSTM baselines for short-term PM2.5 forecasting by efficiently incorporating meteorological, temporal, and rolling-mean features, underscoring the computational efficiency of boosting approaches for operational deployment [2]. A long-term assessment of daily AQI prediction using meteorological and pollutant records collected over nearly a decade in eastern Türkiye similarly compared XGBoost, LightGBM, and support vector machines, reinforcing the competitiveness of boosting methods under long-horizon conditions with seasonally varying pollutant regimes [3][4]. Comparative studies applying linear regression, random forest, XGBoost, CatBoost, and LSTM to multi-year hourly PM2.5 records have found that boosting variants such as CatBoost can achieve the highest coefficient-of-determination scores among the tested models [5].

2.2. Deep Learning and Temporal Modeling Approaches

Recurrent architectures, particularly LSTM networks, remain a central tool for capturing the temporal dependencies inherent in pollutant time series. Studies applying Ridge regression, XGBoost, and LSTM across seasonal subsets of low-cost sensor data have shown that no single architecture dominates uniformly across seasons, with simpler regression baselines occasionally outperforming more complex sequence models during periods of low pollution variability [6]. To address the limitations of individual pollution episodes such as sudden concentration spikes, several works have layered attention mechanisms and optimization routines onto LSTM backbones. An architecture combining attention-based convolutional layers, quantum-inspired particle swarm optimization, and XGBoost demonstrated strong predictive power for capturing irregular AQI patterns, indicating that hybridizing optimization algorithms with deep sequence models can improve robustness for urban monitoring tasks [2].

Bidirectional LSTM variants have also been applied to high-resolution spatial estimation problems, including nationwide daily PM2.5 estimation using attention-augmented Bi-LSTM architectures, as well as PM2.5 concentration estimation using Bi-LSTM combined with metaheuristic optimization methods [7]. These developments suggest that while LSTM-based architectures remain computationally demanding relative to tree ensembles, targeted architectural augmentation continues to yield measurable forecasting gains, particularly for capturing long-range temporal dependencies and non-stationary pollution dynamics.

2.3. Hybrid and Ensemble Forecasting Frameworks

Because tree-based and deep sequence models exhibit complementary strengths — boosting methods excel at exploiting engineered tabular features, while recurrent and convolutional

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

architectures better capture sequential and spatial structure — a substantial share of recent work proposes hybrid frameworks that combine both families. A hybrid model integrating XGBoost with a multi-scale convolutional and graph-learning component (XGBoost-MSCGL) was evaluated against multilayer perceptron, CNN, LSTM, and combined baselines using hourly pollutant and meteorological data from the Fen-Wei Plain, reporting substantial reductions in root mean square error and mean absolute error alongside improved coefficient-of-determination scores relative to single-model baselines [8]. Similarly, real-time forecasting systems combining convolutional layers for spatial feature extraction, LSTM layers for short- and medium-range temporal dependencies, and lightweight mixer architectures for long-term sequence modeling have been proposed to balance predictive accuracy against the computational overhead associated with purely recurrent deep networks [9].

Stacking-based ensembles that combine gradient-boosted learners with support vector regression and Bayesian ridge regression have also been shown to enhance prediction accuracy by leveraging the complementary error characteristics of the constituent models [1]. Collectively, this body of work supports the broader observation that hybrid ensembles, rather than any single algorithm family, currently represent the state of practice for operational air pollution forecasting.

2.4. Multi-Source Environmental Data Integration

A growing share of recent research emphasizes the fusion of heterogeneous data streams to overcome the sparse spatial coverage of ground-based monitoring networks. Frameworks combining Sentinel-5P, MODIS, and Landsat-8 satellite products with meteorological variables and ground observations within attention-based spatio-temporal architectures have reported substantial improvements in forecasting accuracy relative to conventional single-source approaches, alongside explicit uncertainty quantification to support reliability assessment [10]. Related multimodal frameworks combining sensor data, meteorological information, and satellite imagery have identified traffic patterns, industrial emissions, and socioeconomic variables as promising additional inputs for improving the spatial resolution of urban air quality prediction [11].

Traffic-aware multimodal approaches have integrated closed-circuit television-derived traffic density estimates with environmental sensor data, while other studies have fused multiple satellite-derived aerosol optical depth products to construct full-coverage pollutant concentration maps [12]. At a larger operational scale, frameworks integrating emission inventories, meteorological drivers, and ground station observations across multiple megacities have used ablation analysis to quantify the marginal contribution of each data source, finding that emission inventory data and meteorological data provide largely complementary rather than redundant predictive information [13]. Real-time health-risk mapping frameworks have extended this multi-source paradigm further by incorporating fixed and mobile sensor data, meteorological inputs, satellite observations, and demographic information within a single predictive pipeline [14].

2.5. Explainable AI Techniques for Interpretability

As forecasting models have grown more complex, a parallel research stream has focused on making their predictions interpretable to environmental scientists, policymakers, and the

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

public. SHAP and LIME have emerged as the two dominant model-agnostic explanation techniques applied to air quality prediction. A systematic review of explainable AI applications in air pollution risk assessment and respiratory health outcomes found that SHAP has become the dominant interpretability technique in this domain, while noting that LIME remains comparatively underused and recommending greater exploration of causal explainability approaches in future work [15].

Applied studies have paired XGBoost with both SHAP and LIME for PM2.5 nowcasting, using SHAP-based dimensionality reduction to identify the subset of features responsible for the majority of predictive variance while preserving forecasting accuracy [16]. A physics-informed neural network framework for AQI prediction similarly combined SHAP and LIME to provide complementary global and local explanations, arguing that the joint use of both techniques enhances transparency, trust, and accountability relative to using either method alone [17]. A multi-method explainable framework applied to a major urban traffic corridor combined SHAP, LIME, partial dependence plots, and accumulated local effects with temporal decomposition and causal inference, finding that the four explanation methods yielded consistent attributions and that wind speed and lagged pollutant concentrations were dominant drivers of nitrogen oxide variability [18]. Other real-time forecasting systems combining convolutional, recurrent, and mixer-based architectures have likewise incorporated SHAP and LIME to provide feature-level interpretability for both classification and regression outputs, positioning explainability as an increasingly standard component of operational forecasting pipelines rather than an optional add-on [9].

Table 1 consolidates the studies discussed above into a comparative summary of methods, advantages, and drawbacks, to support quick cross-referencing during manuscript preparation.

Table 1 Literature Summary Table

Author(s) / Year	Methods	Advantages	Drawbacks	Notes
Özüpak, Alpsalaz & Aslan (2025) [1]	Comparative ML benchmarking (XGBoost, AdaBoost, LightGBM, SVR, Bayesian Ridge) plus stacking ensemble for AQI/PM forecasting	High accuracy from boosting methods; stacking combines complementary model strengths	Benchmark-style study; no new architecture proposed; performance is dataset/region dependent	Also synthesizes a prior 155-study review (2011–2021) on ML/DL trends
ScienceDirect review (2025) [2]	Systematic literature review of ML/DL air quality forecasting studies (LightGBM,	Broad synthesis across pollutants and regions; identifies best-	Review only — no new empirical results; conclusions	Published in Chemosphere

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

Author(s) / Year	Methods	Advantages	Drawbacks	Notes
	XGBoost, LSTM, CNN-LSTM, SVR, hybrid models)	performing model families	inherit limitations of primary studies	
Tırınk (2025) [3][4]	XGBoost, LightGBM and SVM compared for long-term daily AQI forecasting using ~8 years of meteorological and pollutant data	Long observation period supports robustness across seasons; multi-model comparison	Single-region study (eastern Türkiye); generalizability to other climates untested	-
IJLTEMAS (2025) [5]	Linear Regression, Random Forest, XGBoost, CatBoost and LSTM compared on multi-year hourly PM2.5 data (India)	LSTM achieved lowest MAE/RMSE; CatBoost achieved highest R^2	Best model differs by metric; no explainability component included	-
Engineering Proceedings (2025) [6]	Ridge regression, XGBoost and LSTM applied to low-cost sensor data across winter/summer seasons	Demonstrates feasibility of low-cost sensor networks for PM2.5 forecasting	No single model wins consistently; LSTM underestimated pollution peaks	Performance strongly season-dependent
arXiv preprint (2025) [7]	Comparative benchmarking of operational time-series models (LSTM, Bi-LSTM, ODE-based neural nets) with interpretability focus	Emphasizes operational deployment and interpretability, not accuracy alone	Preprint — not yet peer-reviewed at time of citation	Draws on Bi-LSTM/attention-augmented PM2.5 literature
ScienceDirect — XGBoost-	Hybrid XGBoost + multi-scale convolutional/graph-	Substantial RMSE/MAE reduction and R^2	Added architectural complexity;	-

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

Author(s) / Year	Methods	Advantages	Drawbacks	Notes
MSCGL (2025) [8]	learning (MSCGL) model vs MLP, CNN, LSTM baselines (Fen-Wei Plain data)	improvement over single models	validated on a single regional dataset (China)	
Discover Atmospher e (2025) [9]	Hybrid CNN + LSTM + TSMixer architecture with SHAP/LIME explainability for real-time forecasting and classification	Balances accuracy with computational efficiency; built-in explainability	Combining three architectures increases design and tuning complexity	-
Scientific Reports — MAST-Net (2025) [10]	Attention-based LSTM+CNN spatio-temporal network fusing Sentinel-5P, MODIS and Landsat-8 satellite data with ground observations	23–31% RMSE improvement over conventional single-source approaches; includes uncertainty quantification	Depends on satellite data availability and latency; complex multi-modal architecture	-
Scientific Reports (2025) [11]	Sensor + meteorological + satellite data fusion with attention, GNN and Neural-ODE components	Ablation study shows each architectural component contributes measurably to accuracy	Authors note missing traffic, industrial and socioeconomic data as a limitation	-
Scientific Reports (2023) [12]	CNN-based traffic density extraction from CCTV footage fused with environmental sensor data for AQI prediction	Novel use of visual/traffic data alongside conventional sensors	Related prior work cited within reported low accuracy and classification only, not true forecasting	-
FuXi-Air,	Emission–	Quantifies	Diminishing	Preprint

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

Author(s) / Year	Methods	Advantages	Drawbacks	Notes
arXiv (2025) [13]	meteorology–pollutant multimodal ML with ablation analysis across Beijing, Shanghai and Shenzhen	marginal contribution of each data source; large multi-city validation	returns when combining sources — meteorology's contribution shrinks once emissions and station data are present	
Scientific Reports (2025) [14]	RF, Gradient Boosting, XGBoost and LSTM fusing fixed/mobile sensors, meteorology, satellite and demographic data	Extends forecasting into health-risk mapping across urban, industrial and rural settings	Operational complexity of maintaining multiple live real-time data feeds	-
Atmosphere — systematic review (2025) [15]	Systematic review with ROBINS-I risk-of-bias assessment of ML+XAI studies in air pollution and respiratory health	Rigorous bias assessment; identifies SHAP as the dominant XAI technique	Best-performing model (extremely randomized trees) lacks explainability; LIME remains underused	Recommends future causal-XAI-ML research direction
Applied Sciences (2025) [16]	XGBoost combined with SHAP-based dimensionality reduction and LIME for PM2.5 nowcasting	SHAP-based feature selection preserves accuracy while reducing dimensionality	Focused on short-horizon nowcasting rather than long-term forecasting	-
AirSense-X, PMC (2025) [17]	Physics-Informed Neural Network (PINN) combined with SHAP and LIME for AQI regression	Combines physical-law constraints with data-driven learning; low reported error	PINN remains a black box despite physical constraints, still requiring post-hoc XAI	-

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

Author(s) / Year	Methods	Advantages	Drawbacks	Notes
		(MSE 2.86)		
Atmosphere — Istanbul Bosphorus (2026) [18]	XGBoost, LightGBM, CatBoost, RF and CNN-LSTM-Attention combined with SHAP, LIME, PDP, ALE, STL decomposition and causal inference (CCM)	Cross-validates explanations across four XAI methods; identifies causal, not just correlational, drivers	High methodological complexity — resource-intensive to replicate in full	CatBoost achieved best PM2.5 $R^2 = 0.876$
Blessington Praveen & Sivanantham (2025) [19]	AES encryption applied to Controller Area Network (CAN) bus communication in automotive embedded systems, combined with cloud technology	Protects in-vehicle data against CAN bus cyberattacks while meeting AES requirements under embedded constraints	Abstract does not report latency/bandwidth overhead benchmarks; different application domain entirely	Automotive cybersecurity chapter — included at user's request; not part of the air pollution forecasting literature base and would not normally appear in this review

3. Research Gaps and Challenges

Despite substantial progress, several gaps recur across the reviewed literature. First, although XAI techniques are increasingly applied, their use remains uneven: SHAP is applied considerably more often than LIME, and few studies systematically compare explanation methods against one another or validate explanation stability across retraining runs [15][18]. Second, multi-source data fusion studies frequently acknowledge that additional data modalities, such as traffic and demographic information, could further improve accuracy but are not yet consistently incorporated within a single unified framework, limiting comparability across studies [11][14]. Third, most reported architectures are evaluated on data from a single city or region, raising open questions about cross-regional generalizability,

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

particularly given the heterogeneous mix of emission sources and meteorological regimes across geographies [3][8]. Fourth, hybrid ensemble architectures, while accurate, tend to increase computational cost and model complexity, creating friction between real-time operational deployment requirements and the explainability techniques that themselves add further computational overhead [9][16].

4. Conclusion and Future Directions

The reviewed literature demonstrates a clear trajectory toward hybrid, multi-source, and explainable forecasting frameworks for air pollution prediction. Gradient-boosted ensembles and LSTM-based deep networks remain the dominant modeling backbones, with hybrid combinations increasingly outperforming single-model baselines. The integration of satellite, meteorological, traffic, and emission data continues to expand the feature space available to these models, while SHAP and LIME have become the standard tools for rendering these increasingly complex models interpretable. Future work would benefit from standardized cross-regional benchmarking, systematic comparison of explanation techniques for stability and faithfulness, and the development of computationally efficient hybrid architectures capable of delivering both real-time forecasts and transparent, decision-support-ready explanations for environmental policymakers.

References

[1.] Özüpak, Y., Alpsalaz, F., & Aslan, E. (2025). Air quality forecasting using machine learning: Comparative analysis and ensemble strategies for enhanced prediction. Water, Air, & Soil Pollution, 236, 464. https://doi.org/10.1007/s11270-025-08122-8

[2.] ScienceDirect (2025). Machine learning for air quality prediction and data analysis: Review on recent advancements, challenges, and outlooks. Chemosphere (in press). https://www.sciencedirect.com/science/article/pii/S0048969725022338

[3.] Tırınk, S. (2025). Machine learning-based forecasting of air quality index under long-term environmental patterns: A comparative approach with XGBoost, LightGBM, and SVM. PLOS ONE. https://doi.org/10.1371/journal.pone.0334252

[4.] PMC (2025). Machine learning-based forecasting of air quality index under long-term environmental patterns. PMC12507303.

[5.] International Journal of Latest Technology in Engineering, Management & Applied Science (2025). Machine learning approaches for PM2.5 prediction: A comparative study. https://doi.org/10.51583/IJLTEMAS.2025.1412000075

[6.] Spatio-temporal PM2.5 forecasting using machine learning and low-cost sensors: An urban perspective (2025). Engineering Proceedings. https://doi.org/10.3390/engproc2025101006

[7.] Interpretable PM2.5 forecasting for urban air quality: A comparative study of operational time-series models (2025). arXiv preprint.

[8.] A systematic study on PM2.5 and PM10 concentration prediction in air pollution using machine learning and deep learning model (2025). Results in Engineering / ScienceDirect. https://www.sciencedirect.com/science/article/pii/S2590182625000967

ISSN: 3107-8168
DOI: https://doi.org/10.65725/JCISE/2/4/009

[9.] Scalable AI-driven air quality forecasting and classification for public health applications (2025). Discover Atmosphere. https://doi.org/10.1007/s44292-025-00052-8
[10.] Air quality prediction using multi-source remote sensing data integration with hybrid deep learning framework (2025). Scientific Reports. https://doi.org/10.1038/s41598-025-32466-0
[11.] Advanced air quality prediction using multimodal data and dynamic modeling techniques (2025). Scientific Reports. https://doi.org/10.1038/s41598-025-11039-1
[12.] Deep learning based multimodal urban air quality prediction and traffic analytics (2023). Scientific Reports. https://doi.org/10.1038/s41598-023-49296-7
[13.] FuXi-Air: Urban air quality forecasting based on emission-meteorology-pollutant multimodal machine learning (2025). arXiv preprint, arXiv:2506.07616.
[14.] Machine learning-driven framework for real-time air quality assessment and predictive environmental health risk mapping (2025). Scientific Reports / PMC12328577.
[15.] Explainable artificial intelligence and machine learning for air pollution risk assessment and respiratory health outcomes: A systematic review (2025). Atmosphere, 16(10), 1154. https://doi.org/10.3390/atmos16101154
[16.] The quest for the best explanation: Comparing models and XAI methods in air quality modeling tasks (2025). Applied Sciences, 15(13), 7390. https://doi.org/10.3390/app15137390
[17.] A hybrid physics-informed neural and explainable AI approach for scalable and interpretable AQI predictions (2025). PMC12448019.
[18.] Multi-method explainable AI framework for quantifying traffic and meteorological contributions to urban air pollution: A case study of Istanbul's Bosphorus Bridge corridor (2026). Atmosphere, 17(6), 591. https://www.mdpi.com/2073-4433/17/6/591
[19.] Praveen, B., & Sivanantham, K. (2025). Ensuring safety and security in Control Area Network-based automotive embedded systems with Advanced Encryption Standard method using cloud technology. In The Intersection of 6G, AI/Machine Learning, and Embedded Systems (pp. 307–323). CRC Press. https://doi.org/10.1201/9781003540212-16
[20.] Kalimuthu, S., Ramkumar, N., Arun Prakash, N., Boorneush, M., & Dhusiyanth, M. (2026). Butterfly Optimization with Random Forest for COVID-19 Prediction Using Lung Image. Tiny Machine Learning: Fundamentals, Applications, and Security, 237-257.