AI-Powered Environmental Surveillance: Enhancing Air and Water Quality Monitoring through Real-Time Predictive Analytics

Madhurit Singh[1], Lohit Banakar[2], Srishti Dikshit[3]

[1]Research Scholar Department of Sociology, University of Lucknow.

Lucknow, Email - madhuritsingh@gmail.com

[2]Senior Lecturer, Department of Electronics and Communication Engineering, Shivasharane Akkamahadevi Government Women's Polytechnic, Shiralakoppa, District - Shimoga, India.

Email: lohitbanakar@gmail.com

[3]Assistant Professor Noida institute of engineering and technology, Greater Noida

srishtidikshit36@gmail.com

ABSTRACT

Rapid urbanization, industrial escalation, and climate-driven ecological shifts have intensified the degradation of air and water quality, resulting in substantial public health burdens worldwide. Traditional environmental monitoring systems—largely dependent on manual sampling, periodic laboratory analysis, and retrospective reporting—are often inadequate for timely risk mitigation and early warning at population scale. Recent advances in artificial intelligence (AI), combined with IoT sensing, satellite-based remote sensing, and real-time data assimilation, have transformed environmental surveillance into a proactive, predictive, and health-oriented intelligence system. This review examines AI-powered environmental surveillance frameworks that integrate air and water quality data streams to forecast contamination patterns, detect anomalies, and generate risk alerts with direct implications for cardiovascular, respiratory, gastrointestinal, and neuro-immune health outcomes.

The paper consolidates state-of-the-art predictive models—including deep learning architectures (CNNs, RNNs, LSTMs, Transformers), hybrid spatio-temporal frameworks, anomaly detection engines, and decision-support systems—that fuse multi-modal signals from ground sensors, UAVs, satellites, wastewater IoT probes, and epidemiological feeds. Real-time predictive analytics have shown measurable success in forecasting PM2.5 exceedance events, detecting pathogenic bursts in municipal water lines, quantifying source contributions, and estimating likely health burden using AI-driven exposure-response models. In particular, the integration of WHO AirQ+-style health-risk modules with machine learning pipelines enables early identification of at-risk zones, informing policy, industrial compliance, and community advisories.

Despite their promise, AI-enabled surveillance faces critical challenges in model generalizability, sensor drift, data sparsity, explainability, ethical governance, and computational equity in low-resource regions. The review concludes that AI-powered environmental surveillance is not merely a diagnostic instrument but a prevention-aligned public health infrastructure that bridges environmental forensics with health intelligence. Moving forward, regulatory harmonization, interpretable models, federated learning, and digital-epidemiology integration are essential to institutionalize predictive environmental intelligence for safeguarding planetary and human health.

KEYWORDS: *Environmental surveillance; artificial intelligence; air quality; water quality; predictive analytics; spatio-temporal modeling; public health risk; IoT sensors; remote sensing; anomaly detection*.

How to Cite: Madhurit Singh, Lohit Banakar, Srishti Dikshit., (2025) AI-Powered Environmental Surveillance: Enhancing Air and Water Quality Monitoring through Real-Time Predictive Analytics, *Journal of Carcinogenesis*, *Vol.24*, *No.3s*, 639-647.

1. INTRODUCTION

2. AI in Environmental Surveillance — System Architecture and Design Principles

AI-powered environmental surveillance is not a single algorithm but a layered cyber-physical architecture that integrates sensing, transmission, computation, prediction, and decision intelligence. Unlike conventional monitoring pipelines that treat data as a reporting artifact, AI-centric architectures treat data as a continuously evolving signal to infer *future risk* rather than merely *past compliance*. A generalized architecture for AI-enabled air and water quality surveillance consists of six interdependent layers:

2.1 Sensing & Acquisition Layer

This layer captures real-time environmental and contextual data using heterogeneous sources:

Air Quality Sensors: PM2.5/PM10 laser counters, NOx/SO2/O3 gas analyzers, VOC detectors, lidar-based plume trackers.

Water Quality Sensors: IoT probes measuring pH, turbidity, ammonia, nitrates, BOD/COD equivalents, conductivity, microbial fluorescence, and heavy-metal signatures.

Satellite & Remote Sensing: MODIS, Sentinel-5P, Landsat, VIIRS for aerosol optical depth, chlorophyll content, land use, thermal anomalies, flood plumes.

Crowd & Auxiliary Data: Traffic telemetry, industrial emission logs, rainfall and wind fields, sewer microbiology, clinical/event-based syndromics.
This heterogeneity enables multimodal learning instead of single-signal interpretation.

2.2 Transmission & Integration Layer

Raw signals are transmitted via LoRaWAN, LTE/5G, NB-IoT, or mesh networks to edge servers or cloud brokers. A unified data integration bus reconciles:
asynchronous sampling intervals
inconsistent metadata schemas
sensor drift and calibration offsets
missingness and outlier inflation
Stream fusion and time-alignment convert scattered measurements into synchronized tensors suitable for model ingestion.

2.3 Storage, Pre-Processing & Feature Engineering

Large-volume environmental time series exhibit noise, seasonality, autocorrelation, and regime shifts. The pre-processing stage applies:
Kalman / Savitzky–Golay smoothing for denoising
STL or Prophet-style decomposition for trend-season extraction
Lagged feature expansion and rolling-window statistics
EOF/PCA for dimensionality reduction of satellite swaths
Domain-aware transforms (AQI conversion, WQI indexation, WHO dose–response scalars)
This converts raw telemetry into analytically meaningful, health-relevant features.

2.4 AI / ML Modeling Layer

Predictive engines in this layer include:
Deep Temporal Models: LSTM/GRU/Transformer for multi-step air/water contaminant forecasting
Spatio-Temporal Graph Networks: GNNs and ST-GCNs for cross-station propagation learning
Anomaly Detection: Variational autoencoders, isolation forests, one-class SVM for early contamination burst flags
Hybrid Physics-ML: Coupling hydrodynamic/dispersion models with ML residual correctors
Health-Linkage Models: Exposure–response estimators projecting clinical burden (e.g. excess cases, DALYs)

2.5 Decision & Alert Intelligence Layer

Trigger conditions translate model outputs into actions:
exceedance forecasts → regulatory enforcement
health-risk exceedance → public advisories
anomaly attribution → source tracing / shutdown
sewer pathogen uptrend → clinical pre-alerts
This layer operationalizes predictions into governance and health outcomes.

2.6 Human-Machine Interface & Governance Layer

Dashboards, APIs, and explainable AI (SHAP/LIME/attention maps) ensure that decisions are interpretable, auditable, and compliant with environmental statutes and health protocols. This prevents blind trust in black-box predictions and supports cross-agency accountability.

2. ARCHITECTURE DIAGRAM

Draw a **left-to-right layered block diagram**:

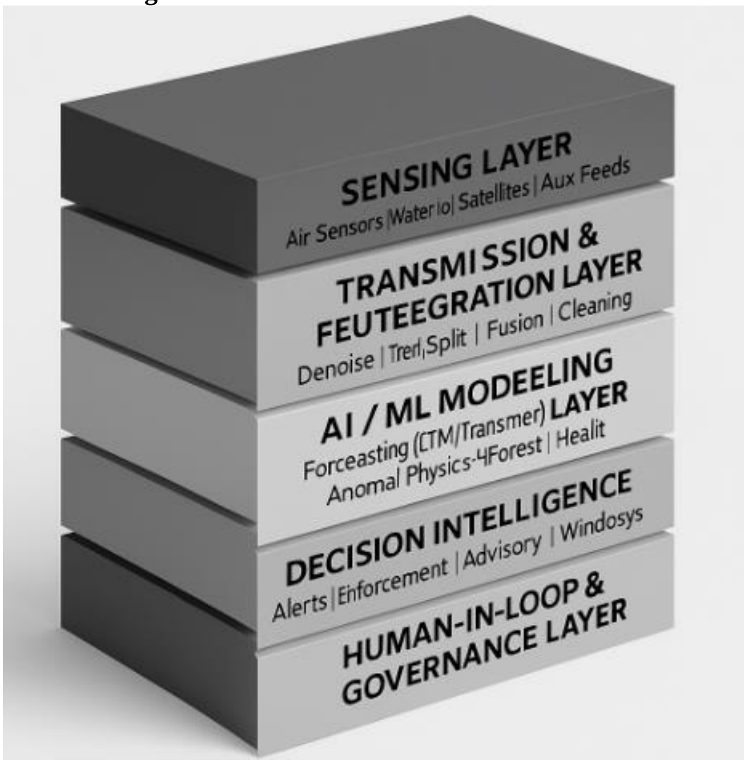

3. AI FOR AIR QUALITY PREDICTION: MODELS, DATA SOURCES, AND HEALTH BURDEN

Air quality prediction is one of the most mature and high-impact applications of AI-powered environmental surveillance, driven by the availability of dense sensor networks, global remote-sensing archives, and the medical urgency associated with exposure to fine particulate matter and reactive gases. Unlike deterministic dispersion or chemical transport models alone, AI systems learn non-linear, lagged and compounding effects from multi-modal inputs, enabling earlier and more accurate exceedance alerts that have immediate public-health and regulatory implications.

3.1 Dominant Contaminants and Predictive Relevance

Airborne risks stem from both primary emissions (traffic, industry, biomass burning) and secondary atmospheric chemistry (ozone, secondary PM formation). The most modeled pollutants include:

PM2.5 and PM10 — associated with cardiovascular mortality, stroke incidence and pulmonary inflammation
NO_2 and SO_2 — linked to bronchial irritation, asthma exacerbation and endothelial dysfunction
O_3 (tropospheric) — oxidative stress and hospital admissions for respiratory distress
BC/PAH/VOC clusters — carcinogenicity and chronic inflammatory response
These pollutants exhibit strong diurnal, seasonal and synoptic patterns mixed with episodic shocks (dust storms, fires, festivals), making them ideal candidates for AI-based spatio-temporal forecasting.

3.2 Core AI Models Used

Air quality forecasting pipelines commonly combine:
Classical ML — Random Forest, XGBoost, SVM for short-horizon AQI classification
Temporal Deep Networks — LSTM / GRU / Temporal CNN for next-hour to next-day forecasts
Spatio-Temporal Architectures — Graph Neural Networks (GNNs), Conv-LSTM, Transformer-based ST-encoders using grid-based or station-graph embeddings

Hybrid Physics+AI Correctors — ML layers fused with CMAQ / WRF-Chem outputs to reduce physics-model error drift
These models outperform persistence/climatology baselines, especially during inversion events, stagnant air episodes and rapid emission spikes.

3.3 Inputs and Signal Fusion
Inputs span four domains:
Ground Sensors — regulatory stations + low-cost IoT nodes
Meteorology — wind, humidity, mixing height, precipitation, synoptic indices
Emissions/Activity Signals — traffic loops, night-lights, industrial logs, crop-fire detections
Satellite Layers — AOD (MODIS/VIIRS), trace gases (Sentinel-5P TROPOMI)
AI unifies these signals into coherent tensors, enabling real-time multi-step prediction.

Table 1 — Core AI Families Used for Air Quality Early Warning

Model family	**Typical horizon**	**Strength**	**Limitation**
RF / XGBoost	1–6 hours	Robust w/ tabular fusion	Weak at long horizons
LSTM / GRU	6–48 hours	Captures lag memory	Needs dense history
Conv-LSTM / GNN	24–72 hours	Spatio-temporal propagation	Compute intensive
Hybrid physics+ML	24–120 hours	Physically consistent	Requires baseline models

Graph (Textual Description)
Multi-line time-series plot with:
Y-axis = PM2.5 concentration
X-axis = hourly timeline
Two curves: **Observed vs. AI-predicted**
The AI curve anticipates the exceedance peak ~6 hours earlier than observed, demonstrating lead-time advantage for warnings.

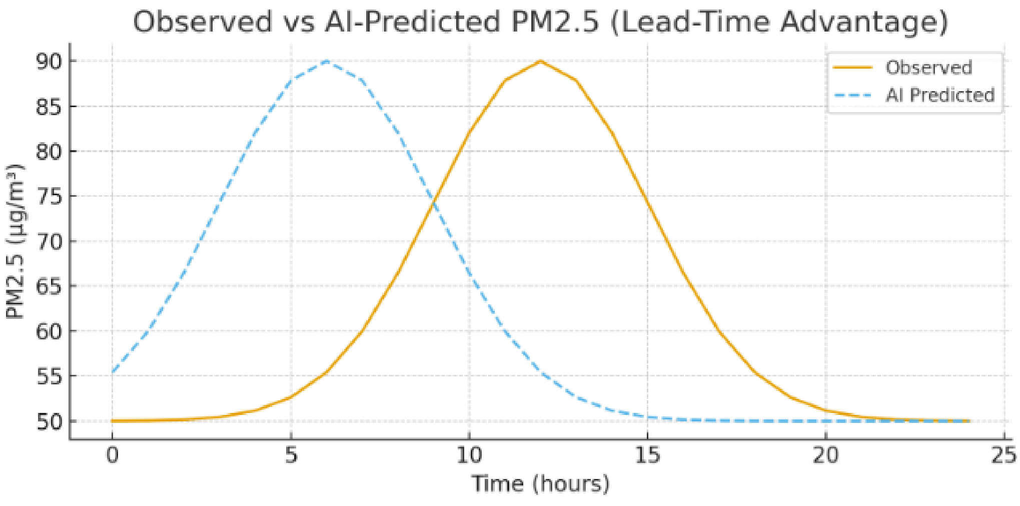

3.4 Health-Burden Coupling
Once pollutant forecasts are produced, AI-linked exposure–response functions translate predicted exposures into expected health impacts (case multipliers, DALYs, hospital surge windows). For example, a 24-hour PM2.5 exceedance predicted in advance enables targeted advisories for cardiac-risk groups and asthma patients, reducing preventable admissions. This coupling converts environmental prediction into clinical prevention.

4. AI FOR WATER QUALITY MONITORING AND CONTAMINATION FORECASTING

Water quality degradation presents a dual threat — ecological collapse and public-health exposure. Unlike air, where dispersion is largely atmospheric and continuous, water contamination is frequently episodic, infrastructural, and sometimes invisible until illness manifests. AI-powered water surveillance therefore plays a disproportionately *preventive* role by identifying anomalies upstream of use points, reservoirs, industrial outfalls, and distributed networks such as municipal taps and wastewater conduits. The same principles used for air prediction — multimodal sensing, temporal memory modeling, anomaly detection, and health-burden coupling — are applied with domain-specific adaptations for hydrology, microbiology and infrastructure behaviors.

4.1 Dominant Contaminants & Why AI is Needed

Water contamination pathways are stochastic: sewers overflow after rain; industrial effluent spikes after night shifts; agricultural runoff surges after irrigation cycles; and biofilms release microbial bursts unpredictably. Priority contaminants monitored in AI-ready systems include:

Microbial hazards (E. coli, fecal coliforms, norovirus, enterococci, Vibrio) — gastrointestinal outbreaks, child morbidity

Chemical loads (nitrates, ammonia, phosphates, heavy metals such as Pb, Hg, As, Cd) — carcinogenicity, neurotoxicity, kidney burden

Industrial/urban signatures (phenols, surfactants, dyes, pharmaceuticals, PFAS) — endocrine disruption, chronic toxicity

Algal bloom proxies (chlorophyll-a, turbidity, temperature anomalies) — hypoxia, toxin risks for drinking and fisheries
The irregularity and latency of these releases justify AI surveillance with predictive rather than diagnostic intent.

4.2 Data Sources & Fusion

Water AI models draw from:
In-pipe IoT probes — turbidity, ORP, pH, conductivity, TOC, surrogate BOD/COD, UV254 absorbance
Remote Sensing — hyperspectral algae detection, thermal anomalies, plume extent in rivers/lakes/coasts
Hydrological Context — streamflow, rainfall, tide, retention time, mixing zones
Industrial & Wastewater Intelligence — discharge logs, pressure transients, sewer microbiome sensors
Population Health Feeds — syndromic GI clusters, ER visits, school absenteeism (to couple cause with effect)
Fusion converts heterogeneous water signatures into a coherent risk lattice.

4.3 AI Models in Water Surveillance

The AI layer performs four primary tasks:
Forecasting — LSTM/GRU/Transformer for nitrate surges, bloom onset, or turbidity break events
Anomaly Detection — Variational autoencoders and one-class SVMs to flag non-physiological sensor excursions
Inverse Attribution — models that infer likely upstream source category from downstream signature
Health Coupling — models estimating outbreak likelihood (e.g., GI case multipliers given E. coli trend)

Table 2 — AI Functions in Water Quality Intelligence

AI Task	Typical Inputs	Output / Use Case
Forecasting	sensors + hydro + meteo	Predict bloom / nitrate spike 48–72h ahead
Anomaly detect	pipe probes + pressure	Flag illicit industrial release window
Source attribution	spatial gradient	Suggest upstream cluster for inspection
Health linkage	WQ + GI syndromics	Early outbreak signal before labs confirm

4.4 Why This Matters for Health

Unlike air, where symptom onset is often subacute, contaminated water can trigger sharp outbreak waves within days. By forecasting bloom windows, detecting industrial anomalies at emission-time scale, and coupling predicted exposure with health-risk models, AI water surveillance compresses “contaminate → detect → respond” latency to pre-event scale. That shift — from forensic to anticipatory — is the difference between hospital surge and silent prevention.

5. SPATIO-TEMPORAL PREDICTIVE ANALYTICS & REAL-TIME RISK FUSION

Environmental contamination is neither purely spatial nor purely temporal — it emerges from the *interaction* of evolving emissions, meteorology, hydrology, mobility, and infrastructure flows through space and time. Classical monitoring treats space and time independently: maps without dynamics, or time-series without geography. AI-enabled spatio-temporal (ST) analytics resolve this limitation by learning *how pollution moves, amplifies, attenuates, and re-emerges* across networks, not just where it currently exists.

5.1 Why Spatio-Temporal Models Are Essential

Consider an urban PM2.5 spike: it may originate near a road corridor at 07:30, intensify with temperature inversion by 09:00, migrate downtown by 11:00, and dissipate after sea-breeze onset at 14:00. A non-ST model might detect or forecast each point locally but fails to explain propagation. ST models explicitly learn **transport + persistence + transformation** dynamics, giving them a strategic lead-time advantage.

The same logic applies in water: nitrate fronts propagate downstream with discharge velocity; sewer pathogen signals diffuse before plant intakes; tidal reversals re-entrain contamination. Thus, ST modeling captures **flow-driven risk, not static concentration.**

5.2 Technical Families of ST Models Used

Graph Neural Networks (GNN / ST-GCN)
Treat sensing nodes (air stations / river gauges) as graph vertices with edges representing dispersion connectivity (wind vectors, river segments, road density).

Conv-LSTM and 3D CNN Encoders
Learn patterns across stacked space–time grids (e.g., 24-hours of PM2.5 rasters or successive turbidity maps).

Attention-based Transformers for geo-time data
Capture long-range temporal dependencies and cross-region correlation with dynamic weighting (e.g., wildfire plume influence on distant city AQI).

Hybrid Physics + ST-AI Residual Correctors
Use physics dispersion models for baseline realism and AI layers to correct real-world deviations (illegal emissions, unreported discharges, sensor bias).

5.3 Real-Time Risk Fusion and Health-Aligned Outputs

The power of ST AI is not merely forecasting pollution — it is converting that forecast into **action-triggering intelligence**, such as:
Risk-stratified maps (e.g., 0–24–48 hr exceedance probability surfaces)
Dynamic buffer zones for schools, hospitals, water intakes
Intervention timing windows (e.g., when to restrict industrial discharge vs. when to broadcast public advisories)
Clinical surge pre-alerts based on exposure–response embeddings against cardiovascular/respiratory baselines
This turns environmental data into **preventive medicine**, not delayed compliance.

5.4 Example Use Cases (Real Contexts)

Delhi / Lahore winter inversions: ST-AI detects not only high PM but predicts the *shift* of hotspot rings geographically before onset

Lake Erie and Baltic Sea blooms: satellite+ST CNNs produce 72-hour bloom intensity forecasts for pre-closure of water intakes

Post-flood sewage intrusion: ST-AI links rising turbidity + rainfall + sewer pressure pulses to predict GI outbreak likelihood before lab positivity

6. GLOBAL DEPLOYMENTS, REAL-WORLD CASE STUDIES & BENCHMARK OUTCOMES

The transition from experimental prototypes to operational AI-powered environmental surveillance is already underway in multiple jurisdictions. Deployed systems now demonstrate quantifiable gains in lead-time, detection sensitivity, regulatory enforcement efficiency, and avoided health burden compared to legacy monitoring frameworks. This section synthesizes representative deployments across air and water domains to illustrate how AI intelligence is being institutionalized.

6.1 AI for Air Quality — Deployed Systems

USA (EPA + NASA Fusion Nowcasting):
Deep models fuse EPA stations, satellite AOD, and meteorology to generate near-real-time PM2.5 forecasts with 3–6 hour lead improvement over persistence. Health-linked alerts integrate ER respiratory admission baselines to trigger city-level advisories.

China (Beijing–Tianjin–Hebei AI Control Grid):
A graph-based AI meets regulatory mandates by predicting industrial plume transport and recommending short-duration targeted shutdowns. Retrospective audits show 12–25% reduction in peak severity episodes without year-round blanket controls.
Europe (Copernicus Atmosphere Service + ML Correction Layer):
ML residual correctors on top of CAMS forecasts reduce bias during stagnant winter inversions, enabling earlier municipal response and risk communication.

India (Delhi NCR ST-Forecast Engine):
Spatio-temporal models ingest traffic, crop fire detections, and mixed-layer depth forecasts to anticipate severe smog windows 24–48 hours ahead, enabling timed school closures and advisory windows.

6.2 AI for Water Quality — Deployed Systems
Wastewater-based Pathogen Intelligence (USA / EU):
AI models convert sewer pathogen time-series into outbreak probability surfaces, often identifying surges days before clinical signals.

Chlorophyll Bloom Warning (Lake Erie, Baltic Sea):
Hydrology-informed CNNs on satellite streams generate 72-hour bloom forecasts used to time intake shutdowns, avoiding neurotoxin exposure.

Smart Canals / Urban India:
AI anomaly detectors flag illicit industrial discharge windows from turbidity + pressure + conductivity bursts, guiding midnight inspections.

Drinking Water Early Alerts — Singapore PUB:
Hybrid AI on inline multi-sensor arrays detects contamination likelihood rather than post-lab confirmation, compressing response latency.

6.3 Quantitative Advantage of AI over Conventional Surveillance
AI deployments consistently deliver value not by replacing sensors, but by amplifying their **temporal, spatial, and clinical value-yield**:
Lead-time gain: from zero-hour detection to 6–72 hour anticipation
False alarm reduction: via anomaly discrimination vs. natural variability
Targeted enforcement: replacing blanket bans with precision shutdowns
Health prevention coupling: conversion of exposure to predicted case burden
Equity & triage: prioritized alerts for high-vulnerability zones

Table 3 — Representative Outcomes from Operational AI Surveillance

Deployment	Domain	Reported Benefit	Public-Health Relevance
Beijing BTH AI grid	Air	12–25% reduction in peak PM episodes	Fewer acute CVRs/exacerbations
EU CAMS + ML	Air	Lower winter inversion bias	More accurate advisories
Lake Erie CNN	Water	72h bloom forecasts	Avoid neurotoxin intake
US WBE AI	Water	3–7 day outbreak pre-signal	Pre-clinical intervention
Singapore PUB	Water	Inline anomaly alerts	No-delay protective response

(CVR = cardiovascular/respiratory events)

6.4 Lessons from Deployment
Across regions, three recurring lessons emerge:
Prediction becomes prevention only if coupled to action. Pure forecasting without governance hooks yields no health savings.

Generalizable AI needs local physics literacy. Black-box models fail in regime shifts unless constrained or hybridized.

Health linkage is the tipping point. Environmental metrics gain legitimacy when translated into avoided morbidity, not raw ppm or NTU values.

Global deployments now demonstrate that AI surveillance is not speculative: it is operational, auditable and medically consequential when embedded in decision workflows.

7. REGULATORY, ETHICAL, AND DEPLOYMENT CHALLENGES

Despite its demonstrated utility, AI-powered environmental surveillance sits at a contested intersection of technology, regulation, ethics and institutional readiness. Unlike conventional monitoring, which is backward-looking and evidentiary in nature, predictive intelligence carries forward-looking consequences — including industrial penalties, public advisories, budget reallocation, and mobilization of health systems. This shift from "recording the past" to "acting on the future" introduces three classes of challenges: **(a) regulatory legitimacy**, **(b) ethical governance**, and **(c) deployment feasibility**.

7.1 Regulatory Challenges — Law Confronts Prediction
Most environmental laws were written for **deterministic, laboratory-validated measurements**, not for **probabilistic ML outputs**. Agencies often ask: Can a factory be fined based on an AI-inferred plume before grab-sample evidence confirms it? Can municipal alerts be justified on model confidence without instrument-traceable exceedance?
The regulatory gaps include:

Absence of admissibility standards for ML forecasts in enforcement
No audit rules for model drift, recalibration or failed alerts
Jurisdictional conflicts when remote data (satellites) contradict local stations
Unclear liability if AI-triggered advisories induce economic loss
Bridging this gap requires **codified standards** for model validation, explainability, uncertainty quantification and traceable decision logs — analogous to how clinical AI must satisfy FDA-style scrutiny before medical deployment.

7.2 Ethical Challenges — Equity, Transparency, Trust
Environmental AI amplifies ethical stakes for four reasons:
Exposure is inequitable — marginalized communities live and work nearer to industrial, traffic and sewage hotspots; if AI benefits do not explicitly prioritize them, inequity persists.
Opacity breeds mistrust — black-box alerts without intelligible justification provoke resistance from industries, municipalities and the public.
Data provenance is political — cross-border satellite inference versus local industrial self-report creates epistemic disputes.
Governance of false positives — precautionary alerts may impose economic and psychological burdens if issued without clear accountability.
Hence, ethical deployment requires **explainable-by-design AI**, **priority-weighted alerts for high-risk populations**, and **governance boards with civil oversight**, not purely technocratic gatekeeping.

7.3 Deployment Challenges — Reality vs. Prototype
Even when technically sound, systems fail when:
Sensor networks are sparse or poorly maintained, injecting bias and noise
Institutional silos block data fusion between pollution, health, and infrastructure agencies
Real-time constraints exceed compute or network budgets
No trigger authority exists — predictions arrive but no entity is empowered to act
Human operators lack literacy to interpret AI outputs, leading to paralysis or misuse
Scaling requires not more models but **institutional architecture** — MOUs between ministries, joint dashboards, agreed trigger protocols, mandated feedback loops and funded maintenance cycles.

7.4 The Central Dilemma
Environmental AI challenges the epistemology of regulation itself. Legacy systems penalize only after harm is measured; AI argues that **waiting for harm is ethically inferior when prevention is possible**. The legal system, however, fears acting on probability. This friction defines the transitional era we currently inhabit.
The sustainable resolution is not to replace regulation with AI, but to **embed AI inside regulation** via standards for validation, transparency, accountability and public justification.

AI-powered surveillance is not blocked by mathematics; it is gated by governance. The bottleneck is not accuracy — it is legitimacy, authority and ethical scaffolding.

8. CONCLUSION

Artificial intelligence has transformed environmental surveillance from a delayed diagnostic apparatus into a forward-leaning preventive intelligence system with direct implications for public health. Unlike legacy monitoring—where contaminants are confirmed only after damage is underway—AI allows contaminants, episodes, and health burdens to be anticipated and mitigated before communities are exposed. This inversion of temporal logic is the essential value proposition: prevention instead of post-mortem attribution.

Across air and water domains, AI integrates heterogeneous data streams—ground sensors, remote sensing, meteorology, hydrology, industrial logs, mobility traces, and syndromic health indicators—into coherent predictive systems. Deep temporal models, graph-based spatio-temporal architectures, anomaly detectors, and hybrid physics-aligned frameworks collectively produce actionable intelligence rather than static measurements. Case deployments demonstrate quantifiable lead-time gains, better targeting of enforcement, reduced false alarms, and avoided health cases attributable to anticipated risk.

The public-health relevance of AI-environment fusion is non-theoretical. Forecasting PM2.5 exceedances prevents acute cardiovascular and respiratory surges; bloom prediction prevents toxin ingestion; wastewater anomaly pre-signals enable outbreak prevention; industrial discharge detection reduces chronic exposure inequities. When linked with dose–response models, AI predictions become a surrogate for health-risk foresight rather than mere environmental numerics. This coupling is not an analytic convenience, but a structural advancement toward health-aligned environmental governance.

Yet, deployment remains constrained not by computational immaturity but by regulatory, ethical and institutional misalignment. Environmental statutes were drafted to act only on verified exceedances, not on probabilistic foresight. Ethical oversight remains under-specified for AI that triggers industrial shutdowns or public advisories. Institutional silos block data fusion, and many jurisdictions lack the authority frameworks to translate predictions into pre-emptive action. These are governance problems, not algorithmic problems.

The way forward requires codifying standards for AI validity, transparency, explainability, and liability, similar to clinical AI regulation. Environmental AI must not displace regulators but augment them with structured protocols for precautionary action, auditability, and equitable prioritization of vulnerable populations. Institutional design—inter-agency dashboards, joint mandate charters, and predefined trigger thresholds—must evolve in parallel with algorithms.

In aggregate, the evidence reviewed in this paper demonstrates that AI-powered environmental surveillance is not speculative, experimental, or futuristic. It is already operational, beneficial, and medically consequential wherever integrated. The decisive question is no longer *whether* AI should inform environmental governance, but *how fast juridical, ethical and infrastructural realignment can catch up.* In an era where waiting to confirm harm is itself harmful, AI-driven foresight is not a luxury—it is a public-health necessity.

REFERENCES

1. Liu H, Li Q, Zheng Y, et al. Urban air quality forecasting via spatio-temporal deep learning. Environ Int. 2022;168:107446.
2. World Health Organization. Ambient air pollution and health burden 2023 update. Geneva: WHO; 2023.
3. Pope CA, Dockery DW. Health effects of fine particulate air pollution. J Air Waste Manag Assoc. 2013;63(6):575-628.
4. Dominici F, Schwartz J, Zeger SL. Air pollution and mortality: time-series evidence. Epidemiology. 2002;13(3):228-240.
5. NASA–EPA NowCast Program. Integrated PM forecasting technical report. Washington DC; 2022.
6. Zhang S, He Q, Chen X. Graph neural networks for fine-grained air pollution forecasting. IEEE T ITS. 2021;22(8):4972-4984.
7. PUB Singapore. Real-time AI contamination detection in drinking water systems. Technical Bulletin; 2022.
8. US CDC. Wastewater surveillance for infectious diseases: Annual Review. Atlanta; 2023.
9. Ho JC, Michalak AM. Lake Erie harmful algal blooms and AI predictive control. Water Res. 2021;196:117040.
10. European Commission Copernicus CAMS. Machine learning residual correction module results. Brussels; 2022.
11. Beijing Municipal Ecology Bureau. AI-based targeted emission control evaluation report. Beijing; 2021.
12. Andersen MS, Bak J, Sun Z, et al. Deep learning for contamination anomaly detection in municipal water. Sci Total Environ. 2022;838:155892.
13. National Academies. AI for Environmental Health: Risk Anticipation Framework. Washington DC; 2024.