

Original Research

Received 03 July 2026

Revised 30 July 2026

Accepted 12 Aug 2026

Natural Resources for Human
Health 2026, Vol. 6 (17s): 597–611

KEYWORDS:

Environmental health; artificial
intelligence; data quality;
algorithmic bias; data scarcity; low-
resource settings.

Making AI Work for Environmental Health in Low-Resource Regions: A Narrative Review

E. Rajalakshmi^{1*}, J. M. Gnanasekar², Renukadevi S.¹, Jayasri Archanadevi P.¹

¹Assistant Professor, Department of Artificial Intelligence and Data Science, Rajalakshmi Engineering College, Thandalam, Chennai, India

²Professor and Head, Department of Artificial Intelligence and Data Science, Rajalakshmi Engineering College, Thandalam, Chennai, India

*Corresponding author: E. Rajalakshmi (rajalakshmi.e@rajalakshmi.edu.in)

Email addresses: rajalakshmi.e@rajalakshmi.edu.in (E. Rajalakshmi);

jmg_sekar@yahoo.com (J. M. Gnanasekar); renukadevi.s@rajalakshmi.edu.in

(Renukadevi S.); jayasriarchanadevi.p@rajalakshmi.edu.in (Jayasri Archanadevi P.)

ABSTRACT: Artificial intelligence (AI) and machine learning (ML) are increasingly used to monitor environmental conditions, estimate exposures, and support health-related decision making. Yet the settings with substantial environmental-health burdens often have the least complete, representative, and accessible data. This narrative review examines three interconnected constraints (data quality, data bias, and data scarcity) and considers how they affect the reliability, generalisability, fairness, and practical usefulness of AI for environmental health in low-resource regions. The review synthesises 38 sources published between 2018 and 2025 covering air-quality sensing, water-quality prediction, satellite remote sensing, environmental exposure estimation, public-health surveillance, and algorithmic fairness. Across the literature, low-cost sensors can expand spatial coverage but require calibration and maintenance; monitoring and health records may under-represent exposed populations; and models developed elsewhere may transfer poorly without local validation. The review develops an integrated challenge–solution taxonomy linking each data limitation to its consequences for AI, technical mitigation strategies, and policy or community interventions. Case studies from India, Ethiopia, Uganda, and Ghana illustrate how these constraints interact in practice. The review identifies priorities for low-resource settings: shared benchmarks, local and subgroup validation, explicit uncertainty reporting, interoperable data systems, stronger environmental and health surveillance, and governance arrangements that enable responsible use of local data. The central implication is that AI can extend environmental-health monitoring, but technical innovation cannot fully substitute for representative local measurements and health data.

1. INTRODUCTION

Environmental conditions, including air pollution, unsafe water, and other exposure-related hazards, shape human health, but their distribution and measurement are highly unequal. More than 90% of deaths attributable to air pollution occur in low- and middle-income countries (LMICs) (Pinder et al., 2019). Groundwater arsenic is another major environmental-health concern: a global model estimates that 94 to 220 million people are potentially exposed to high concentrations, most of them in Asia (Podgorski and Berg, 2020). These burdens create a strong rationale for environmental monitoring that can identify exposure patterns and support prevention and public-health action.

AI and ML are increasingly used for environmental monitoring, prediction, and decision support. A bibliometric analysis of 4,762 publications from 1991 to 2024 reported rapid growth in AI research for environmental monitoring since 2010, with air and water quality among the major application areas (Alotaibi and Nassif, 2024). In principle, AI can combine measurements from regulatory monitors, low-cost sensors, satellite imagery, geospatial variables, and health information to estimate environmental conditions in places where direct measurement is limited. In practice, model performance is constrained by the quality, representativeness, and availability of these inputs.

The data these models require is often thinnest where exposure and health burdens are greatest. Of 234 countries and territories, 141 have no regular PM_{2.5} monitoring, and the global mean population distance to the nearest monitor has been estimated at 220 km (Martin et al., 2019). In 2024, 36% of countries were reported to have no government air-quality monitoring, and access to openly shared data remained uneven (OpenAQ, 2024). Health-outcome data is also incomplete: about four in ten deaths worldwide go unregistered, and in the WHO African Region only one in ten is recorded (World Health Organization, 2021). These gaps create a fundamental problem for environmental-health AI. A model may predict an environmental variable accurately while still providing limited evidence about who is exposed, who is vulnerable, and what health consequences follow.

This review therefore examines how three interconnected data problems limit AI for environmental health in low-resource regions, and which technical, governance, and community-level responses can reduce their effects. It addresses three questions:

1. What data problems limit AI models for environmental health in low-resource regions?
2. How do these problems affect model accuracy, fairness, and usefulness?
3. Which technical and policy approaches can reduce them?

Section 2 describes the applications and data sources and how environmental monitoring connects to human health. Section 3 sets out the review methodology. Section 4 presents the results and discussion: the three data challenges, the resulting challenge–solution taxonomy, three case studies, implications for practice, and limitations. Section 5 identifies open problems, and Section 6 concludes. Appendix A summarises every source surveyed.

2. BACKGROUND

2.1 Applications and Data Sources

AI for environmental health draws on four kinds of data, each with its own weaknesses in low-resource settings (Table 1).

Table 1. Data sources used by AI for environmental health and their main weaknesses in low-resource regions

Data source	Typical AI use	Main weakness in low-resource regions
Reference-grade monitoring stations	Ground truth for training and calibration	Few or none; costly to buy and maintain (Pinder et al., 2019; Martin et al., 2019)
Low-cost sensors	Dense local monitoring of air and water	Biased readings that need local calibration (Giordano et al., 2021; Concas et al., 2021)
Satellite imagery	Estimating conditions where no ground data exists	Still needs ground data to train and validate (Burke et al., 2021)
Health records and surveys	Linking exposure to disease	Incomplete registration and reporting (World Health Organization, 2021)

Air quality. ML is used to calibrate low-cost sensors, detect anomalies, and correct drift in large sensor networks (Kim and Moon, 2025). A review of 78 ML studies covers air quality monitoring and prediction more broadly (Karmoude et al., 2025). Locally built networks show what is feasible: the AirQo network in Kampala, Uganda, has released an open dataset from its low-cost monitors together with reference instrument data (Sserunjogi et al., 2022).

Water quality. A review of more than 170 studies from five years documents ML for predicting water quality indicators and indices, and lists data acquisition and model uncertainty among the main challenges (Yan et al., 2024). For groundwater, a

random forest model trained on more than 50,000 aggregated arsenic measurements and 11 geospatial predictors produced a global hazard map that flags areas with no reported measurements (Podgorski and Berg, 2020).

Remote sensing. Satellite imagery combined with ML can estimate outcomes where ground surveys are missing. A review in *Science* quantifies the shortage of ground data on human outcomes and concludes that satellite methods add to ground data collection but do not replace it (Burke et al., 2021).

Health surveillance. A scoping review of AI for public health surveillance in Africa describes applications in disease detection and prediction alongside the technical challenges of deploying them (Edet et al., 2025). Related reviews group AI health interventions in LMICs into diagnosis, risk assessment, outbreak prediction and surveillance, and health policy and planning (Schwalbe and Wahl, 2020), and point to weak infrastructure, failing health information systems, and a shortage of skilled staff as barriers in Africa (Tshimula et al., 2024).

2.2 From Environmental Monitoring to Human Health

The relevance of environmental AI to human health depends on more than accurate prediction of an environmental variable. A useful environmental-health pathway connects the measured or inferred condition to human exposure, identifies populations that may be vulnerable, and establishes how the information could support prevention, surveillance, or resource allocation. Air-quality models, for example, can improve estimates of particulate exposure. Groundwater models can identify areas where contaminant testing should be prioritised. Satellite or sensor systems can provide spatial information where conventional monitoring is sparse. However, these applications should not be interpreted as direct predictions of disease unless the underlying studies include appropriate health-outcome data.

This distinction is important in low-resource settings because the environmental and health data gaps often occur together. Sparse monitoring can obscure exposure gradients, while incomplete registration or surveillance can obscure outcomes. AI can help bridge the first gap through spatial prediction, sensor calibration, remote sensing, or transfer learning, but it cannot by itself establish an exposure-health relationship that has not been measured. Accordingly, this review treats environmental monitoring as an enabling component of environmental health, not as a substitute for health-outcome evidence.

3. REVIEW METHODOLOGY

This is a narrative review, not a systematic review, and it does not claim to be exhaustive. No PRISMA compliance is claimed. This section sets out how sources were found, selected, and synthesised, so that readers can judge the scope and limits of the evidence.

3.1 Literature Search Strategy

Sources were identified through targeted keyword searches of web and scholarly search engines and publisher and repository sites (including PubMed Central, arXiv, and publisher pages), carried out in October 2026. Search terms combined a method term with a data or domain term, for example: "low-cost sensor calibration" with "machine learning"; "air quality data gap" with "low- and middle-income countries"; "algorithmic bias" or "fairness" with "public health machine learning"; "transfer learning" or "satellite" with "air quality" and "data-poor"; "groundwater arsenic" or "fluoride" with "machine learning"; "data imputation" with "air quality"; "federated learning" with "health" or "air quality"; and "artificial intelligence" with "public health surveillance" and "Africa". These searches were supplemented by follow-up searches for related studies by the same research groups (for example, those running the AirQo network in Kampala and the sensor networks in Accra) and by candidate sources listed in the authors' preliminary drafts.

3.2 Source Selection Criteria

Sources were included if they (1) were published between 2018 and 2025; (2) were written in English; (3) addressed at least one of the three focal challenges (data quality, data bias and representativeness, or data scarcity) in environmental monitoring, exposure estimation, or health surveillance; and (4) either reported evidence from low- or middle-income settings or described methods or data infrastructure directly relevant to them. Peer-reviewed reviews and primary studies were preferred. Preprints (two sources) and conference papers were retained only where they filled a specific gap, and two institutional reports were retained because they provide global monitoring and health-data statistics that are difficult to replace with individual studies (OpenAQ, 2024; World Health Organization, 2021).

Candidate sources were excluded if their bibliographic existence could not be confirmed from a publisher, repository, or indexing record, or if they did not report data-related limitations or remedies. Several candidate sources from preliminary drafts were excluded on the first ground. The final evidence base comprises 38 sources, listed with their type and setting in Appendix A. Screening and selection were shared equally among all four authors, and the final selection was agreed by all authors.

3.3 Data Extraction and Thematic Synthesis

Each source was examined for four elements: (1) the principal data limitation; (2) its reported effect on model performance, fairness, generalisability, uncertainty, or deployment; (3) technical responses described in the source; and (4) policy, infrastructure, or community conditions required for those responses to work. Each source was then coded to one or more of the three themes (quality, bias, scarcity); the codes are shown in the final column of Table A1. A source could inform more than one theme, and findings were synthesised narratively within each theme (Sections 4.1 to 4.3).

3.4 Development of the Challenge–Solution Taxonomy

The taxonomy (Table 2) was built inductively in three steps. First, recurring sub-problems were identified within each theme, giving nine rows: four for quality, two for bias, and three for scarcity. Second, for each row, the effects on AI and the technical and policy remedies were taken from the coded extractions, and every entry drawn from a source carries its citation. Third, where the reviewed sources identified a problem but did not document a remedy, the authors added a remedy based on their own judgement, and these entries are marked with a dagger (†) in Table 2.

3.5 Distinguishing Evidence from Recommendations

Throughout the review, statements of evidence are attributed to their sources by citation. Interpretations and recommendations that go beyond what the cited studies documented are signposted as the authors' own, using phrases such as "we infer" or "we recommend". The implications in Section 4.6 and the priorities in Section 5 are the authors' recommendations, informed by the reviewed evidence but not directly tested in it.

4. RESULTS AND DISCUSSION

The findings are presented by challenge (Sections 4.1 to 4.3), brought together in a taxonomy (Section 4.4), illustrated with case studies (Section 4.5), and then discussed for practice and limitations (Sections 4.6 and 4.7).

4.1 Data Quality Challenges

Low-cost sensors make dense monitoring affordable, but their raw readings are not reliable enough for quantitative use. Low-cost particulate matter sensors are subject to biases that range from simple to correct (meteorology, sensor age) to complex (aerosol source, composition, and refractive index) (Giordano et al., 2021). A survey of ML calibration adds cross-sensitivity between pollutants, external influences such as traffic and weather, and loss of accuracy over time (Concas et al., 2021). Reported performance also varies considerably from one sensor to another (Kang et al., 2022).

Calibration improves the data but creates its own demands. Methods range from linear and quadratic corrections to complex ML models (Giordano et al., 2021; Liang, 2021). Their accuracy depends jointly on the algorithm, the amount of training data, and the predictors chosen (Liang and Daniels, 2022). The length of co-location needed differs by sensor type, and a short period can suffice only when it covers conditions similar to those of the later deployment (Levy Zamora et al., 2023).

Model choice involves a trade-off that is particularly important in data-poor settings. Traditional ML models can be comparatively lightweight but may fail to generalise across contexts, whereas deep-learning models can require larger datasets and may be harder to interpret (Kim and Moon, 2025). Consequently, an apparently accurate calibration or prediction model in one city cannot be assumed to remain accurate in another city, season, or pollution regime without local validation.

Field studies in sub-Saharan Africa confirm both the gain and the local dependence. In Kampala, a random forest calibration cut the error of low-cost PM_{2.5} sensors from an RMSE of 18.6 to 7.2 $\mu\text{g}/\text{m}^3$ and was deployed across a network of more than 120 devices (Adong et al., 2022). In Accra, linear regression and random forest gave only moderate improvement ($R^2 = 0.81$), while Gaussian mixture regression reached $R^2 = 0.88$, tolerated missing inputs, and gave an estimate of calibration certainty (McFarlane et al., 2021). A later comparison of three sensor brands in Accra found all of them strongly correlated with the reference instrument but biased high for the local mix of pollution sources (Raheja et al., 2023).

These demands are hardest to meet where reference instruments are absent. Calibration needs a reference monitor to calibrate against, yet most countries have very few (Martin et al., 2019). We infer from these sources, though none tested it directly, that calibration quality is likely to be poorest where reference monitors are scarcest; this deserves direct study in low-resource deployments.

Gaps in the record are a second quality problem. They commonly arise from sensor malfunction or communication failure (Han et al., 2023). Air quality time series often contain missing values, and the quality of imputation can significantly affect the prediction task that follows (Hua et al., 2024). One comparison tested eight imputation methods, from mean and nearest-neighbour filling to deep models such as BRITS and SAITS, on air quality datasets from several cities (Hua et al., 2024). Imputation estimates what was not measured, so it cannot repair a record that is unrepresentative to begin with.

Data also arrives from mixed sources in mixed forms. Some governments publish air quality data only in static formats that cannot be analysed directly (OpenAQ, 2024). Field studies add that incomplete, inconsistent, or non-representative data can make ML models ineffective or wrong, whatever the algorithm (Leggesse et al., 2024).

Health data has a parallel quality problem. Two-thirds of low-income countries have established a standardised system for reporting causes of death, but WHO finds that these systems need urgent strengthening (World Health Organization, 2021). For environmental-health AI, this matters: an exposure model can be technically strong while the downstream exposure–health association remains uncertain, because outcome labels are incomplete, inconsistently recorded, or unavailable for the populations of interest.

4.2 Data Bias and Representativeness

The existence of data does not mean that it represents the people, places, seasons, or exposure conditions most relevant to health. In this review, bias is considered at three connected points: measurement coverage, model development and transfer, and reporting or evaluation. These forms of bias should be distinguished from simple scarcity. Scarcity concerns how much data exists or is accessible, whereas representational bias concerns whose conditions are systematically over- or under-represented.

Coverage bias. Only 24 of 234 countries and territories have more than three PM_{2.5} monitors per million inhabitants (Martin et al., 2019). India, for example, has about one monitor for every 6.8 million people (Martin et al., 2019). Within LMICs, several megacities have invested in monitoring networks while hundreds of other cities have limited or no measurements (Pinder et al., 2019). In our interpretation, a model trained on such data learns the conditions of the monitored places and remains untested everywhere else.

Development bias. Research on AI for environmental monitoring is concentrated in a few countries (Alotaibi and Nassif, 2024). Fletcher and colleagues argue that a model must be matched to its target population before use, a criterion they call appropriateness, and that careless adoption of AI in LMICs may widen existing health disparities (Fletcher et al., 2021). They define bias as a systematic tendency of a model to favour one demographic group over another.

The global health literature makes the same point from the health side. A review in *The Lancet* found that much research on AI-driven health interventions in LMICs does not describe the ethical, regulatory, or practical considerations needed for use at scale (Schwalbe and Wahl, 2020). Asiedu and colleagues argue that fairness attributes defined in Western settings do not transfer directly, and they propose axes of disparity specific to African contexts from a scoping review, a survey of 672 people, and interviews with 28 experts (Asiedu et al., 2024).

Reporting bias. Even where data is good, fairness is rarely examined. A systematic review of 35 Dutch public health ML studies found that sampling and missing-data practices were well documented, but most studies omitted explicit fairness framing, subgroup analysis, and discussion of potential harms (Altamirano et al., 2025). This evidence comes from a single, well-resourced research system. We infer, but the reviewed sources did not test, that fairness reporting is unlikely to be more common where data is scarcer.

Predictive maps illustrate both the promise and the limit of data-driven environmental-health modelling. The global arsenic model identifies likely hazard in regions with few or no reported measurements (Podgorski and Berg, 2020). Such predictions can help prioritise testing and resource allocation, but they do not establish that an individual well is safe or unsafe. This distinction between prediction and direct measurement is essential when environmental AI is used to inform health protection.

4.3 Data Scarcity Challenges

In many low-resource regions the central problem is that the data was never collected, was collected but not digitised, or was collected but not shared. Scarcity therefore takes several forms, each with different implications for AI.

Sparse monitoring. In 2024, 36% of countries had no government air quality monitoring, and most of these were low-income or lower-middle-income countries (OpenAQ, 2024). For water, environmental data often does not exist at the spatial or temporal resolution that management decisions need (Willard et al., 2025).

Closed data. Collection does not guarantee access. A further 9% of countries monitor air quality but do not publish the results, and only 27% share data in a fully open form (OpenAQ, 2024). Health data is commonly held in separate institutions, with privacy concerns restricting access (Rieke et al., 2020).

Missing labels. Supervised models need outcomes to learn from. With four in ten deaths unregistered worldwide (World Health Organization, 2021), the health side of an exposure–outcome model is often the scarcer half. A review of satellite-based ML likewise documents a shortage of ground data on human outcomes (Burke et al., 2021).

No local data for adaptation. Transfer learning is the usual answer to scarcity, but most applications assume that some data from the target site is available for fine-tuning (Willard et al., 2025). Truly unmonitored sites, where only site characteristics and external drivers are known, remain a harder case.

The evidence on how far such methods can stretch scarce data is mixed but encouraging. A deep transfer model trained on satellite imagery of Los Angeles and New York explained up to 54% of the variation in air quality in Accra without any local labels (Yadav et al., 2024). In Rwanda, combining satellite data with local low-cost sensor readings improved surface PM_{2.5} estimates by 40% over ground sensors alone (Malings et al., 2020). In India, a random forest trained on 12,600 groundwater fluoride measurements produced a national hazard map at 1 km resolution and an estimate of about 120 million people at risk (Podgorski et al., 2018), and the approach has since been extended to a global fluoride map (Podgorski and Berg, 2022).

Scarcity compounds the other two problems. Without local reference data, sensors may not be calibratable to the required accuracy. Without local outcome data, representational bias and subgroup performance may be difficult to evaluate. Scarcity is therefore more than a lack of observations: it can prevent the quality and fairness of AI systems from being assessed in the first place.

4.4 Integrated Challenge–Solution Taxonomy

No single technique addresses all three problems, and technical fixes depend on policy, infrastructure, and community conditions. Table 2 therefore maps each challenge to its principal effect on AI, the technical responses identified in the reviewed literature, and the policy or community conditions needed to sustain those responses. The taxonomy was derived inductively from the thematic grouping described in Section 3, not imposed as a universal classification.

Table 2. Integrated challenge–solution taxonomy for AI in environmental health in low-resource regions

Group	Data challenge	Effect on AI	Technical remedies	Policy and community remedies
Quality	Sensor bias, drift, and cross-sensitivity	Errors in inputs carry into predictions; calibrations fail when moved to new sites	ML calibration with periodic recalibration (Giordano et al., 2021; Concas et al., 2021; Adong et al., 2022; McFarlane et al., 2021); co-location periods matched to deployment conditions (Levy Zamora et al., 2023); season-specific calibration (Campmier et al., 2023); ML combined with physical knowledge (Kim and Moon, 2025)	Calibration standards and regulatory benchmarks (Kim and Moon, 2025); investment in reference monitors (Pinder et al., 2019; Martin et al., 2019)

Group	Data challenge	Effect on AI	Technical remedies	Policy and community remedies
Quality	Missing values in time series	Less usable training data; poor imputation degrades later prediction	Imputation methods chosen by comparison on local data, from simple filling to deep models (Hua et al., 2024)	Reliable power, connectivity, and maintenance for monitoring sites†
Quality	Mixed sources and formats	Integration errors; data that cannot be analysed directly	Harmonised data pipelines†; networks that integrate regulatory, low-cost, and satellite sources (Martin et al., 2019)	Common data standards†; publication in machine-readable form (OpenAQ, 2024)
Quality	Unreliable health records	Noisy outcome labels weaken exposure–outcome models	Interpretable models that help expose doubtful labels†	Stronger death and cause-of-death registration (World Health Organization, 2021)
Bias	Monitoring concentrated in a few cities and countries	Models untested for under-monitored populations	Geospatial models to guide where to measure next (Podgorski and Berg, 2020; Podgorski et al., 2018; Podgorski and Berg, 2022); evaluation at both urban and rural sites (Campmier et al., 2023)	Monitoring sited by population exposure†; locally run networks (Sserunjogi et al., 2022)
Bias	Fairness seldom assessed	Disparities go unnoticed and may widen	Matching models to the target population (Fletcher et al., 2021; Asiedu et al., 2024); subgroup analysis and frameworks such as ACAR (Altamirano et al., 2025)	Fairness reporting required by journals and funders†; building on the reporting gaps found by Altamirano et al. (2025)
Scarcity	Few monitors and short records	Too little data to train, adapt, or validate	Satellite imagery with ML (Burke et al., 2021; Leggesse et al., 2024; Malings et al., 2020); transfer learning from data-rich sites (Willard et al., 2025; Yadav et al., 2024); network-wide calibration that uses existing reference monitors (Zheng et al., 2019)	Long-term funding for monitoring†; lasting local institutions and capacity (Pinder et al., 2019)
Scarcity	Closed or siloed data	Each model limited to its own small dataset	Federated learning across institutions without exchanging data (Rieke et al., 2020; Yarham et al., 2025)	Open data policies (OpenAQ, 2024); data-sharing agreements that protect privacy†
Scarcity	Missing outcome labels	Models cannot be trained or evaluated against health outcomes	Ensemble models with cross-validation when samples are few (Leggesse et al., 2024); satellite proxies validated on ground samples (Burke et al., 2021)	Health information systems and registration (World Health Organization, 2021)

Note to Table 2: entries with a citation are documented in the cited source. Entries marked † are the authors' proposals, added where the reviewed sources identified a problem but did not document a remedy.

Three points follow from the taxonomy.

First, technical remedies still require some form of local evidence. Satellite models require ground data for training or validation, and transfer learning usually requires target-site observations (Burke et al., 2021; Willard et al., 2025). Even transfer without

local labels explained only about half of the variation in one target city (Yadav et al., 2024). These methods reduce the amount of local data needed without removing the need.

Second, remedies interact. Federated learning addresses access to data held in separate institutions (Rieke et al., 2020), but we argue that it cannot by itself correct geographic coverage bias if the participating institutions all serve the same well-monitored populations. Similarly, imputation can reduce missingness in a time series but cannot make an unmeasured rural community representative of the target population.

Third, in our assessment, the policy and infrastructure environment sets a practical boundary for technical remedies. Calibration depends on reference monitors, fairness analysis depends on information about relevant subgroups, and both depend on sustained measurement, maintenance, interoperability, and responsible data access.

For environmental-health applications, the taxonomy therefore treats technical methods and governance interventions as complementary, not competing, solutions.

4.5 Case Studies

Three cases show the three problems, and the remedies, at work in specific places. Each case first summarises what the cited studies report and then gives the authors' interpretation. Because the cases come from a small number of cities and one lake, they illustrate the challenges but do not establish how widely the findings generalise across other geographical and population contexts.

4.5.1 Air quality sensing in Indian cities

India has about one PM_{2.5} monitor for every 6.8 million people (Martin et al., 2019), so low-cost sensors are an obvious supplement. Two studies show what it takes to make them usable.

In Delhi, Zheng and colleagues calibrated a network of 10 low-cost nodes against the city's 22 regulatory reference monitors using Gaussian process regression (Zheng et al., 2019). The method calibrates sensors in the field without first co-locating each one with a reference instrument, and it flags malfunctioning nodes. The authors estimate that about 20 reference stations are the optimum for this approach in Delhi.

A long-term evaluation of PurpleAir sensors at three Indian sites (Delhi, Hamirpur, and Bangalore) found that calibration using site- and season-specific information improved performance (Campmier et al., 2023). A single correction did not serve all places and seasons.

In our reading, the case links all three problems. Quality is manageable with ML calibration, but the calibration depends on reference monitors, which Delhi has in number and many other cities lack (Pinder et al., 2019). We therefore caution that a method proven in a well-monitored city may not transfer to places where data is scarcest, and where exposure estimates for health protection are most uncertain; neither study tested such transfer.

4.5.2 Water quality in Lake Tana, Ethiopia

Leggesse and colleagues estimated total nitrogen and total phosphorus in Lake Tana from Landsat 8 imagery using six ML algorithms (Leggesse et al., 2024). They note that long-term continuous water quality monitoring is lacking in most African inland waterbodies because of limited infrastructure, finance, and technical capacity.

The field data came from intermittent sampling campaigns between 2016 and 2022, at 27 to 170 locations per campaign, giving 1,101 data points for each parameter. XGBoost gave the best result for nitrogen ($R^2 = 0.80$) and random forest for phosphorus ($R^2 = 0.73$). The artificial neural network performed poorly on phosphorus ($R^2 = 0.21$).

The authors used k-fold cross-validation because field samples were limited, and they report that cloud cover in the rainy season adds uncertainty. They recommend more complete field data collection for future work.

In our reading, the case shows a satellite proxy working under scarcity, with two cautions. The model still needed several years of laboratory-analysed samples to train and validate, and simpler ensemble models outperformed the neural network on this small dataset.

4.5.3 Locally run sensor networks in Kampala and Accra

Two African cities show what locally led monitoring can achieve with few reference instruments.

In Kampala, the AirQo team at Makerere University builds its own low-cost sensors for local conditions. Calibration against a reference monitor with a random forest model reduced PM_{2.5} error from an RMSE of 18.6 to 7.2 $\mu\text{g}/\text{m}^3$, and the model was rolled out to more than 120 devices (Adong et al., 2022). The network's data has been released as an open dataset (Sserunjogi et al., 2022).

In Accra, a network of 17 low-cost monitors ran from 2018 to 2021 and was corrected with Gaussian mixture regression (McFarlane et al., 2021; Raheja et al., 2023). The corrected network average was 23.4 $\mu\text{g}/\text{m}^3$, 1.6 times the WHO daily guideline of 15 $\mu\text{g}/\text{m}^3$ (Raheja et al., 2023). Accra has also served as the test city for transfer learning from satellite imagery of US cities (Yadav et al., 2024).

In both cities, one or two reference monitors anchor a whole network (Adong et al., 2022; Raheja et al., 2023). We interpret this as supporting the view that modest investment in reference instruments, paired with local technical capacity, can yield a large return in usable exposure data; the studies did not, however, compare costs or test this directly.

4.6 Synthesis and Implications for Practice

Across the reviewed literature, data limitations emerge as a major constraint on AI for environmental health in low-resource regions. This does not mean that algorithms are unimportant. It means that algorithmic improvements cannot compensate indefinitely for missing, biased, poorly calibrated, or inaccessible observations. Studies that compare methods often report useful performance after adequate calibration or local validation (Adong et al., 2022; Leggesse et al., 2024; Zheng et al., 2019), while studies of transfer to new settings show that target-site evidence is still needed (Willard et al., 2025; Yadav et al., 2024).

The following implications are the authors' recommendations. They are informed by the reviewed evidence, cited where a specific study supports them, but their effectiveness has not been tested in the reviewed studies:

- **For researchers:** report where the training data came from and where the model was not tested. Validate on held-out sites in under-monitored areas where any exist, and report performance by subgroup (Altamirano et al., 2025).
- **For funders and governments:** a small number of reference monitors has high value because many low-cost sensors can be calibrated against them (Zheng et al., 2019). Publishing existing data in open, machine-readable form costs little and benefits every later model (OpenAQ, 2024).
- **For practitioners deploying sensors:** plan for recalibration across seasons from the start (Campmier et al., 2023), and prefer simpler models when local data is limited (Kim and Moon, 2025; Leggesse et al., 2024).
- **For public-health agencies:** treat model outputs as exposure estimates that show where to test and measure, not as direct evidence of disease (Podgorski and Berg, 2020), and strengthen death and cause-of-death registration so that exposure can be linked to outcomes (World Health Organization, 2021).

4.7 Limitations of This Review

This review has six principal limitations.

Limitations of the narrative approach. This is a narrative and purposive synthesis, not a systematic review. Relevant studies may have been missed, source selection reflects author judgement, and the searches were not designed to be reproducible in the way a systematic protocol would be. No meta-analysis was attempted, so the review cannot estimate how large or consistent the reported effects are across studies. The review is also limited to English-language sources.

Differences in study quality. The evidence base mixes peer-reviewed primary studies, reviews, institutional reports, conference papers, and two preprints. No formal quality or risk-of-bias appraisal was carried out, and the sources were not weighted by quality. Findings from single-site studies and preprints should therefore be given less weight than those from peer-reviewed reviews and multi-site evaluations.

Uneven regional data availability. The case-study evidence comes mainly from India, Ghana, Uganda, and Ethiopia, and the fairness evidence includes a review of Dutch research (Altamirano et al., 2025). Other low-resource regions, including Latin America, South-East Asia, Central Asia, and small island states, are under-represented. This partly reflects where data and

published studies exist, which is itself an example of the coverage bias discussed in Section 4.2. The findings may not generalise to settings with different pollution sources, climates, health systems, or data infrastructure.

Implementation costs. The review did not systematically extract cost data, and the remedies in Table 2 differ widely in what they would cost. Reference monitors, periodic recalibration, sensor maintenance, connectivity, data infrastructure, health registration systems, and skilled local staff all carry recurring costs that may exceed the initial purchase of sensors or software. Policymakers should not read the taxonomy as a list of interchangeable or low-cost options, and future work should evaluate the cost-effectiveness of each remedy in local conditions.

Short-term and technical evidence on mitigation. Most of the mitigation evidence reviewed here reports technical performance, such as RMSE or R^2 , over limited periods; Campmier et al. (2023) is one of the few long-term field evaluations. Few sources examine whether remedies remain effective over years, whether local institutions sustain them, or whether affected communities were involved in their design and use. Longitudinal and community-engaged evaluations are needed before any remedy can be recommended with confidence (Section 5).

Limited link to health outcomes. Relatively few included studies directly link environmental AI outputs to measured human-health outcomes. Some conclusions therefore concern the reliability of environmental exposure estimation, not demonstrated health-effect prediction.

These limitations should be considered when interpreting the taxonomy and cross-study comparisons.

5. OPEN PROBLEMS AND FUTURE DIRECTIONS

The following seven priorities are the authors' proposals, each motivated by gaps identified in the reviewed literature.

1. **Benchmarks from low-resource settings.** Models are often evaluated on data from well-monitored places. Open datasets such as the AirQo release (Sserunjogi et al., 2022) are a start, but shared benchmarks should include rural and peri-urban settings, seasonal variation, missingness, subgroup information where ethically appropriate, and uncertainty annotations.
2. **Calibration without a nearby reference.** Most sensor-calibration methods assume co-location with a reference instrument (Giordano et al., 2021; Levy Zamora et al., 2023), and network-wide methods still rely on reference monitors within the city (Zheng et al., 2019). Developing reliable reference-light or reference-free methods remains important, but such methods should be evaluated against independent reference measurements before deployment.
3. **Routine subgroup reporting.** Fairness assessment is uncommon even in well-resourced public-health ML research (Altamirano et al., 2025). Environmental-health studies should report performance across relevant geographic, demographic, socioeconomic, and exposure subgroups where such reporting is ethically and statistically justified.
4. **Honest uncertainty.** A prediction for an unmonitored place should carry an estimate of how far it can be trusted. Model uncertainty, data uncertainty, and extrapolation uncertainty should be distinguished, not collapsed into a single performance metric.
5. **Data governance.** Much existing environmental and health data is collected but not shared (OpenAQ, 2024). Federated approaches and open-data policies address different parts of this problem (Rieke et al., 2020). Future work should evaluate governance models that allow legitimate local use while protecting privacy, security, and community interests.
6. **Linking exposure to health outcomes.** A major opportunity is to connect environmental exposure estimates with longitudinal health, surveillance, or epidemiological data. Such work should make the exposure–outcome pathway explicit and assess whether improved environmental predictions translate into better identification of vulnerable populations or more effective public-health decisions. Fairness frameworks developed for clinical and public-health ML can provide a starting point, but their applicability to environmental exposure modelling requires context-specific validation (Fletcher et al., 2021; Altamirano et al., 2025).
7. **Longitudinal and community-engaged evaluation.** The mitigation strategies in Table 2 are mostly supported by short-term technical results (Section 4.7). Future studies should follow deployments over several years, report recurring costs alongside accuracy, and involve affected communities and local institutions in design, deployment, and evaluation. Locally run networks such as AirQo (Adong et al., 2022; Sserunjogi et al., 2022) offer a natural setting for such work.

6. CONCLUSION

AI can extend environmental-health monitoring to places that have historically lacked dense measurement, but its reliability remains bounded by the underlying data. This review identifies data quality, representational bias, and data scarcity as interconnected constraints that affect accuracy, generalisability, fairness, uncertainty, and the practical interpretation of AI outputs in low-resource settings. The reviewed evidence shows that calibration, satellite proxies, transfer learning, and federated learning can reduce some barriers, but they do not eliminate the need for representative local evidence. Because this is a narrative review drawing mainly on short-term technical evidence from a limited set of countries, these conclusions should be tested through longitudinal, community-engaged, and cost-aware evaluations.

The practical priority is therefore not simply to deploy more sophisticated algorithms. It is to build sustainable measurement and health-information systems, improve data interoperability and access, evaluate models across relevant populations and locations, and make uncertainty explicit. For environmental health, the ultimate test of AI is not only whether it predicts an environmental variable accurately, but whether that prediction can support a defensible understanding of exposure, vulnerability, and health risk. Future research should consequently combine technical innovation with local validation, governance, and community-informed deployment.

ACKNOWLEDGEMENTS

The authors thank the academic and research community whose published work informed this review. The authors also thank the Department of Artificial Intelligence and Data Science, Rajalakshmi Engineering College, for its support.

FUNDING

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

AUTHOR CONTRIBUTIONS

All authors contributed equally to this work.

E. Rajalakshmi: Conceptualisation, literature search and screening, analysis and synthesis, writing (original draft), and writing (review and editing).

J. M. Gnanasekar: Conceptualisation, literature search and screening, analysis and synthesis, writing (original draft), and writing (review and editing).

Renukadevi S.: Conceptualisation, literature search and screening, analysis and synthesis, writing (original draft), and writing (review and editing).

Jayasri Archanadevi P.: Conceptualisation, literature search and screening, analysis and synthesis, writing (original draft), and writing (review and editing).

All authors read and approved the final manuscript.

CONFLICT OF INTEREST

The authors declare that there is no conflict of interest.

ETHICS STATEMENT

This review used published literature and did not involve direct recruitment of human participants, collection of personal data, or animal experimentation. Ethical approval was not required.

DATA AVAILABILITY STATEMENT

No new dataset was generated for this narrative review. The evidence base consists of published and publicly available sources cited in the reference list.

APPENDIX A. SUMMARY OF SURVEYED SOURCES

Table A1 lists each source, what kind of work it is, and the point this review takes from it. Q, B, and S mark the challenge it informs: quality, bias, or scarcity.

Table A1. Sources surveyed in this review, in alphabetical order

Source	Type	Domain and setting	Point used in this review	Challenge
Adong et al., 2022	Primary study	Sensor calibration, Kampala	Random forest cut RMSE from 18.6 to 7.2 $\mu\text{g}/\text{m}^3$; over 120 devices	Q
Alotaibi and Nassif, 2024	Bibliometric review	AI in environmental monitoring	4,762 publications; output led by China, the US, and India	B
Altamirano et al., 2025	Systematic review	Public health ML, Netherlands	Most of 35 studies omit fairness framing and subgroup analysis	B
Asiedu et al., 2024	Mixed-methods study	Fairness in health ML, Africa	Fairness attributes specific to African contexts	B
Burke et al., 2021	Review	Satellite imagery and ML	Satellite methods add to ground data but do not replace it	S
Campmier et al., 2023	Primary study	PurpleAir sensors, India	Site- and season-specific calibration improves performance	Q, B
Concas et al., 2021	Survey	ML calibration of air sensors	Cross-sensitivity, external influences, and drift	Q
Edet et al., 2025	Scoping review	AI for health surveillance, Africa	Applications and deployment challenges	S
Fletcher et al., 2021	Framework	AI in global health	Appropriateness, bias, and fairness criteria	B
Giordano et al., 2021	Review	Low-cost PM sensors	Sources of bias and best practices for calibration	Q
Han et al., 2023	Survey	ML for urban air quality	Missing values arise from sensor and communication failures	Q
Hua et al., 2024	Review and experiment	Air quality imputation	Eight methods compared; imputation quality affects prediction	Q
Kang et al., 2022	Review	Low-cost air sensor performance	Performance varies considerably between sensors	Q
Karmoude et al., 2025	Review	ML for air quality	78 studies on monitoring and prediction	Q
Kim and Moon, 2025	Review	ML quality control of air sensors	Simple models generalise poorly; deep models need large datasets	Q
Leggesse et al., 2024	Primary study	Water quality, Lake Tana, Ethiopia	Landsat and ML with limited samples; R^2 0.80 for nitrogen	S
Levy Zamora et al., 2023	Primary study	Co-location calibration	Needed duration varies; conditions must match deployment	Q
Liang, 2021	Review	Sensor calibration	Techniques, trends, and challenges	Q
Liang and Daniels, 2022	Primary study	Sensor calibration	Accuracy depends on algorithm, training size, and predictors	Q
Malings et al., 2020	Primary study	Satellite and low-cost sensors, Africa	Combining both improved PM _{2.5} estimates by 40% in Rwanda	S
Martin et al., 2019	Analysis	PM _{2.5} monitoring, global	141 of 234 countries have no regular PM _{2.5} monitoring	B, S
McFarlane et al., 2021	Primary study	Sensor calibration, Accra	Gaussian mixture regression reached R^2 0.88 with certainty estimates	Q
OpenAQ, 2024	Report	Air quality data, global	36% of countries do not monitor; 27% share data fully openly	S

Source	Type	Domain and setting	Point used in this review	Challenge
Pinder et al., 2019	Perspective	Air quality, LMICs	Over 90% of air pollution deaths are in LMICs; hundreds of cities lack monitoring	B, S
Podgorski and Berg, 2020	Primary study	Groundwater arsenic, global	Random forest hazard map; 94 to 220 million people potentially exposed	S
Podgorski and Berg, 2022	Primary study	Groundwater fluoride, global	Random forest hazard map extended worldwide	S
Podgorski et al., 2018	Primary study	Groundwater fluoride, India	12,600 measurements; about 120 million people at risk	S
Raheja et al., 2023	Primary study	Sensor intercomparison, Accra	Three sensor brands biased high for local sources	Q, B
Rieke et al., 2020	Perspective	Federated learning in health	Training across institutions without exchanging data	S
Schwalbe and Wahl, 2020	Review	AI in global health	Deployment considerations are rarely described	B
Sserunjogi et al., 2022	Dataset	Air quality, Kampala	Open data from a locally run low-cost network	B, S
Tshimula et al., 2024	Review	AI for health surveillance, Africa	Infrastructure, information system, and skills barriers	Q, S
Willard et al., 2025	Survey	ML for unmonitored water sites	Transfer learning usually needs some target-site data	S
World Health Organization, 2021	Report	Health data systems, global	Four in ten deaths worldwide are unregistered	Q, S
Yadav et al., 2024	Primary study	Transfer learning, Accra	Up to 54% of variation explained without local labels	B, S
Yan et al., 2024	Review	ML for water quality	Over 170 studies; data acquisition and uncertainty are open challenges	Q, S
Yarham et al., 2025	Review	Federated learning for air quality	Collaborative training without sharing raw data	S
Zheng et al., 2019	Primary study	Sensor network, Delhi	Network-wide calibration using 22 reference monitors	Q, S

REFERENCES

- [1] Adong, P., Bainomugisha, E., Okure, D., Sserunjogi, R., 2022. *Applying machine learning for large scale field calibration of low-cost PM2.5 and PM10 air pollution sensors*. Applied AI Letters 3, e76. <https://doi.org/10.1002/aill.2.76>
- [2] Alotaibi, E., Nassif, N., 2024. *Artificial intelligence in environmental monitoring: in-depth analysis*. Discover Artificial Intelligence 4, 84. <https://doi.org/10.1007/s44163-024-00198-1>
- [3] Altamirano, S., Vreeken, A., Ghebreab, S., 2025. *Machine learning and public health: identifying and mitigating algorithmic bias through a systematic review*. Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society 8(1), 138–151. <https://doi.org/10.1609/aies.v8i1.36537>
- [4] Asiedu, M., Dieng, A., Haykel, I., Rostamzadeh, N., Pfohl, S., Nagpal, C., Nagawa, M., Oppong, A., Koyejo, S., Heller, K., 2024. *The case for globalizing fairness: a mixed methods study on colonialism, AI, and health in Africa*. Proceedings of the 4th ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO 2024), 1–24. <https://doi.org/10.1145/3689904.3694708>
- [5] Burke, M., Driscoll, A., Lobell, D.B., Ermon, S., 2021. *Using satellite imagery to understand and promote sustainable development*. Science 371(6535), eabe8628. <https://doi.org/10.1126/science.abe8628>

- [6] Campmier, M.J., Gingrich, J., Singh, S., Baig, N., Gani, S., Upadhya, A., Agrawal, P., Kushwaha, M., Mishra, H.R., Pillarisetti, A., Vakacherla, S., Pathak, R.K., Apte, J.S., 2023. *Seasonally optimized calibrations improve low-cost sensor performance: long-term field evaluation of PurpleAir sensors in urban and rural India*. Atmospheric Measurement Techniques 16(19), 4357–4374. <https://doi.org/10.5194/amt-16-4357-2023>
- [7] Concas, F., Mineraud, J., Lagerspetz, E., Varjonen, S., Liu, X., Puolamäki, K., Nurmi, P., Tarkoma, S., 2021. *Low-cost outdoor air quality monitoring and sensor calibration: a survey and critical analysis*. ACM Transactions on Sensor Networks 17(2), 1–44. <https://doi.org/10.1145/3446005>
- [8] Edet, J.E., Akyala, A.I., Omudu, E.A., Ikaidi, A.A., Epelle, J., Etimedet, E., Olatinwo, I.A., Azuka, C.G., Abah, M.I., Edem, P.A., Oluwagbohun, O., 2025. *Harnessing artificial intelligence for public health surveillance in Africa: current applications, challenges, and opportunities: a scoping review*. International Journal of Research and Innovation in Social Science 9(7), 2791–2828. <https://www.rsisinternational.org/journals/ijriss/Digital-Library/volume-9-issue-7/2791-2828.pdf>
- [9] Fletcher, R.R., Nakeshimana, A., Olubeko, O., 2021. *Addressing fairness, bias, and appropriate use of artificial intelligence and machine learning in global health*. Frontiers in Artificial Intelligence 3, 561802. <https://doi.org/10.3389/frai.2020.561802>
- [10] Giordano, M.R., Malings, C., Pandis, S.N., Presto, A.A., McNeill, V.F., Westervelt, D.M., Beekmann, M., Subramanian, R., 2021. *From low-cost sensors to high-quality data: a summary of challenges and best practices for effectively calibrating low-cost particulate matter mass sensors*. Journal of Aerosol Science 158, 105833. <https://doi.org/10.1016/j.jaerosci.2021.105833>
- [11] Han, J., Zhang, W., Liu, H., Xiong, H., 2023. *Machine learning for urban air quality analytics: a survey*. arXiv preprint arXiv:2310.09620. <https://arxiv.org/abs/2310.09620>
- [12] Hua, V., Nguyen, T., Dao, M.-S., Nguyen, H.D., Nguyen, B.T., 2024. *The impact of data imputation on air quality prediction problem*. PLOS ONE 19(9), e0306303. <https://doi.org/10.1371/journal.pone.0306303>
- [13] Kang, Y., Aye, L., Ngo, T.D., Zhou, J., 2022. *Performance evaluation of low-cost air quality sensors: a review*. Science of the Total Environment 818, 151769. <https://doi.org/10.1016/j.scitotenv.2021.151769>
- [14] Karmoude, M., Munhungewarwa, B., Chiraira, I., McKenzie, R., Kong, J., Smith, B., Ayana, G., Njara, N., Mathaha, T., Kumar, M., Mellado, B., 2025. *Machine learning for air quality prediction and data analysis: review on recent advancements, challenges, and outlooks*. Science of the Total Environment 1002, 180593. <https://doi.org/10.1016/j.scitotenv.2025.180593>
- [15] Kim, Y.-H., Moon, S.-H., 2025. *Machine learning-based quality control for low-cost air quality monitoring: a comprehensive review of the past decade*. Atmosphere 16(10), 1136. <https://doi.org/10.3390/atmos16101136>
- [16] Leggesse, E.S., Zimale, F.A., Sultan, D., Enku, T., Tilahun, S.A., 2024. *Advancing non-optical water quality monitoring in Lake Tana, Ethiopia: insights from machine learning and remote sensing techniques*. Frontiers in Water 6, 1432280. <https://doi.org/10.3389/frwa.2024.1432280>
- [17] Levy Zamora, M., Buehler, C., Datta, A., Gentner, D.R., Koehler, K., 2023. *Identifying optimal co-location calibration periods for low-cost sensors*. Atmospheric Measurement Techniques 16(1), 169–179. <https://doi.org/10.5194/amt-16-169-2023>
- [18] Liang, L., 2021. *Calibrating low-cost sensors for ambient air monitoring: techniques, trends, and challenges*. Environmental Research 197, 111163. <https://doi.org/10.1016/j.envres.2021.111163>
- [19] Liang, L., Daniels, J., 2022. *What influences low-cost sensor data calibration? A systematic assessment of algorithms, duration, and predictor selection*. Aerosol and Air Quality Research 22(9), 220076. <https://doi.org/10.4209/aaqr.220076>
- [20] Malings, C., Westervelt, D.M., Haurlyiuk, A., Presto, A.A., Grieshop, A., Bittner, A., Beekmann, M., Subramanian, R., 2020. *Application of low-cost fine particulate mass monitors to convert satellite aerosol optical depth to surface concentrations in North America and Africa*. Atmospheric Measurement Techniques 13(7), 3873–3892. <https://doi.org/10.5194/amt-13-3873-2020>
- [21] Martin, R.V., Brauer, M., van Donkelaar, A., Shaddick, G., Narain, U., Dey, S., 2019. *No one knows which city has the highest concentration of fine particulate matter*. Atmospheric Environment: X 3, 100040. <https://doi.org/10.1016/j.aeaoa.2019.100040>
- [22] McFarlane, C., Raheja, G., Malings, C., Appoh, E.K.E., Hughes, A.F., Westervelt, D.M., 2021. *Application of Gaussian mixture regression for the correction of low cost PM_{2.5} monitoring data in Accra, Ghana*. ACS Earth and Space Chemistry 5(9), 2268–2279. <https://doi.org/10.1021/acsearthspacechem.1c00217>
- [23] OpenAQ, 2024. *Open air quality data: The global landscape 2024*. OpenAQ. <https://openaq.org/about/initiatives/publications>

- [24] Pinder, R.W., Klopp, J.M., Kleiman, G., Hagler, G.S.W., Awe, Y., Terry, S., 2019. *Opportunities and challenges for filling the air quality data gap in low- and middle-income countries*. Atmospheric Environment 215, 116794. <https://doi.org/10.1016/j.atmosenv.2019.06.032>
- [25] Podgorski, J., Berg, M., 2020. *Global threat of arsenic in groundwater*. Science 368(6493), 845–850. <https://doi.org/10.1126/science.aba1510>
- [26] Podgorski, J., Berg, M., 2022. *Global analysis and prediction of fluoride in groundwater*. Nature Communications 13, 4232. <https://doi.org/10.1038/s41467-022-31940-x>
- [27] Podgorski, J.E., Labhasetwar, P., Saha, D., Berg, M., 2018. *Prediction modeling and mapping of groundwater fluoride contamination throughout India*. Environmental Science & Technology 52(17), 9889–9898. <https://doi.org/10.1021/acs.est.8b01679>
- [28] Raheja, G., Nimo, J., Appoh, E.K.-E., Essien, B., Sunu, M., Nyante, J., Amegah, M., Quansah, R., Arku, R.E., Penn, S.L., Giordano, M.R., Zheng, Z., Jack, D., Chillrud, S., Amegah, K., Subramanian, R., Pinder, R., Appah-Sampong, E., Tetteh, E.N., Borketey, M.A., Hughes, A.F., Westervelt, D.M., 2023. *Low-cost sensor performance intercomparison, correction factor development, and 2+ years of ambient PM_{2.5} monitoring in Accra, Ghana*. Environmental Science & Technology 57(29), 10708–10720. <https://doi.org/10.1021/acs.est.2c09264>
- [29] Rieke, N., Hancox, J., Li, W., Milletari, F., Roth, H.R., Albarqouni, S., Bakas, S., Galtier, M.N., Landman, B.A., Maier-Hein, K., Ourselin, S., Sheller, M., Summers, R.M., Trask, A., Xu, D., Baust, M., Cardoso, M.J., 2020. *The future of digital health with federated learning*. npj Digital Medicine 3, 119. <https://doi.org/10.1038/s41746-020-00323-1>
- [30] Schwalbe, N., Wahl, B., 2020. *Artificial intelligence and the future of global health*. The Lancet 395(10236), 1579–1586. [https://doi.org/10.1016/S0140-6736\(20\)30226-9](https://doi.org/10.1016/S0140-6736(20)30226-9)
- [31] Sserunjogi, R., Ssematimba, J., Okure, D., Ogenrwot, D., Adong, P., Muyama, L., Nsimbe, N., Bbaale, M., Bainomugisha, E., 2022. *Seeing the air in detail: hyperlocal air quality dataset collected from spatially distributed AirQo network*. Data in Brief 44, 108512. <https://doi.org/10.1016/j.dib.2022.108512>
- [32] Tshimula, J.M., Kalengayi, M., Makenga, D., Lilonge, D., Asumani, M., Madiya, D., et al., 2024. *Artificial intelligence for public health surveillance in Africa: applications and opportunities*. arXiv preprint arXiv:2408.02575. <https://arxiv.org/abs/2408.02575>
- [33] Willard, J.D., Varadharajan, C., Jia, X., Kumar, V., 2025. *Time series predictions in unmonitored sites: a survey of machine learning techniques in water resources*. Environmental Data Science 4, e7. <https://doi.org/10.1017/eds.2024.14>
- [34] World Health Organization, 2021. *SCORE for health data technical package: global report on health data systems and capacity, 2020*. World Health Organization, Geneva. ISBN 978-92-4-001870-9. <https://www.who.int/publications/i/item/9789240018709>
- [35] Yadav, N., Sorek-Hamer, M., Von Pohle, M., Asanjan, A.A., Sahasrabhojane, A., Suel, E., Arku, R.E., Lingenfelter, V., Brauer, M., Ezzati, M., Oza, N., Ganguly, A.R., 2024. *Using deep transfer learning and satellite imagery to estimate urban air quality in data-poor regions*. Environmental Pollution 342, 122914. <https://doi.org/10.1016/j.envpol.2023.122914>
- [36] Yan, X., Zhang, T., Du, W., Meng, Q., Xu, X., Zhao, X., 2024. *A comprehensive review of machine learning for water quality prediction over the past five years*. Journal of Marine Science and Engineering 12(1), 159. <https://doi.org/10.3390/jmse12010159>
- [37] Yarham, S., Behjati, M., Alobaidy, H.A.H., Majeed, A.P.P.A., Zheng, Y., 2025. *Enhancing air quality monitoring: a brief review of federated learning advances*. Lecture Notes in Networks and Systems 1316, 489–501. https://doi.org/10.1007/978-981-96-3949-6_41
- [38] Zheng, T., Bergin, M.H., Sutaria, R., Tripathi, S.N., Caldow, R., Carlson, D.E., 2019. *Gaussian process regression model for dynamically calibrating and surveilling a wireless low-cost particulate matter sensor network in Delhi*. Atmospheric Measurement Techniques 12(9), 5161–5181. <https://doi.org/10.5194/amt-12-5161-2019>