



*An Online CPD Course
brought to you by
CEDengineering.ca*

AI For Air Quality Monitoring

Course No: C01-041
Credit: 1 PDH

Brian Lisiewski, P.E.



Continuing Education and Development, Inc.

P: (877) 322-5800
info@cedengineering.ca

Table of Contents

Module 1: Introduction — Environmental Safety, Compliance & Engineer's Role	1
1.1 Why AI in Air Quality and Emissions Management	1
1.2 Engineer Accountability and Ethics	1
1.3 Compliance is Non-Negotiable	1
1.4 Integrating AI with Quality Management	2
Module 2: Air Quality & Emissions Fundamentals — Data & Processes.....	3
2.1 Pollutants and Regulatory Context.....	3
2.2 Monitoring Networks	3
2.3 Data Types and Quality Challenges	4
2.4 Integration Note.....	4
Module 3: Sensor Calibration, Data Integration & Anomaly Classification	5
3.1 Sensor Calibration & Drift Management with AI.....	5
3.2 Data Fusion and Integration	5
3.3 Emissions Anomaly Classification.....	6
3.4 Ensuring Reliability of Anomaly Alerts.....	6
Module 4: Air Quality Forecasting & Advisory Control (Human-in-the-Loop)	8
4.1 Air Quality Forecasting with AI.....	8
4.2 Operational Adjustments Based on Forecasts.....	8
4.3 Bounded Adaptive Control (Human-in-the-Loop).....	9
4.4 Ensuring Compliance & Safety in Adaptive Moves	9
Module 5: Compliance & Reporting — Monitoring, Documentation & Audit Trails.....	11
5.1 Continuous Compliance Monitoring.....	11
5.2 Documentation & Reporting	11
5.3 Professional Defensibility & Regulator Communication.....	12

5.4 Connection to Public Communication	12
Module 6: Model Validation, Uncertainty & Risk Management	13
6.1 Validating AI Models Pre-Deployment	13
6.2 Monitoring and Maintenance of Model Performance	13
6.3 Uncertainty and Decision Thresholds	14
6.4 Risk Analysis (FMEA) for AI Implementation.....	14
Module 7: Governance, Change Control & Cybersecurity	15
7.1 Configuration Management & Change Control	15
7.2 Cybersecurity Considerations	15
7.3 Ensuring Professional Defensibility	16
7.4 Continuous Improvement Outlook.....	16
References	17
Recommended Additional Reading.....	19

Module 1: Introduction — Environmental Safety, Compliance & Engineer's Role

1.1 Why AI in Air Quality and Emissions Management

Air pollution is a leading environmental health risk, causing an estimated **7 million premature deaths annually** worldwide — a WHO figure that combines both indoor and outdoor air pollution exposure [8]. Meanwhile, industries and governments operate complex emissions controls to meet regulatory standards (e.g., limits on particulate matter, NO_x, SO₂) to protect public health. **AI can strengthen these efforts** by providing **real-time analysis and early warnings** — detecting pollution spikes or predicting smog episodes ahead of time [9] — enabling more proactive responses than traditional systems. This could mean fewer people exposed to hazardous air at high levels and better adherence to ambient air quality standards or permit conditions. However, these benefits materialize only if AI is integrated with **caution and rigor**.

1.2 Engineer Accountability and Ethics

Regardless of AI tools, the **PE's duty is to the public and the environment first**. This means:

- **All critical decisions require human sign-off:** If an ML model suggests adjusting a scrubber setpoint or raising an alarm threshold, the engineer must scrutinize that suggestion, confirm it will not breach safety or environmental constraints, and explicitly authorize it. This mirrors how autopilot in an aircraft assist but the pilot remains in command.
- **No 'black box' trust:** A core ethical stance is **transparency** — if we cannot explain why an AI model is making a recommendation, we should not implement it blindly. Engineers should use interpretable models or ensure they understand key drivers (e.g., which sensor readings are triggering an alarm) before acting, so they can defend the action to regulators or stakeholders.

1.3 Compliance is Non-Negotiable

All AI implementations are subordinate to laws and permits. If a recommendation or automated change could potentially cause a violation — e.g., it reduces a control measure too much and might lead to an emission above the permitted limit — it must be **rejected or modified by the engineer**. **AI optimization can never override legal emissions limits or health standards**; it must always incorporate them as hard constraints.

1.4 Integrating AI with Quality Management

Many organizations follow ISO 9001 or ISO 14001 for environmental management. Using AI should slot into those systems:

- Document each model and its intended use, just as you would a new piece of equipment.
- **Risk assessments** — e.g., "if this model mis predicts an event, what is the impact?" — should be completed and mitigations put in place (such as backup alarms if AI fails).
- Regular **management review** of model performance — treat AI as an evolving component requiring periodic auditing, as an internal or external auditor might ask how you ensure this tool is reliable.

From the outset, we insist that **introducing AI does not diminish the engineer's legal or moral responsibility**; rather, it gives engineers a powerful assistant to meet those responsibilities more effectively. That mindset will persist throughout this course — AI helps us do our job better, but **we remain firmly at the helm** to ensure public health and environmental integrity are upheld.

Module 2: Air Quality & Emissions Fundamentals — Data & Processes

2.1 Pollutants and Regulatory Context

Common air pollutants of concern include **criteria pollutants** (e.g., PM_{2.5}, PM₁₀, NO₂, ozone, SO₂, CO, and lead) and **hazardous air pollutants (HAPs)**, such as benzene. Ambient air quality standards — like national standards or WHO guidelines — set safe concentration limits (often on time-averaged bases). For stationary sources (power plants, industries), permits enforce emission limits (mass per unit time or concentration from a stack) to ensure contributions keep ambient levels safe.

Key point: Monitoring data is compared to these standards; any AI usage must support compliance with applicable limits.

2.2 Monitoring Networks

Ambient monitoring is done via **ground-based stations** measuring pollutant concentrations in real time [10]. These may be reference-grade analyzers or **low-cost sensor networks** for finer spatial resolution. **Stationary sources** use **Continuous Emissions Monitoring Systems (CEMS)** or periodic stack tests to quantify emissions. Process data from the plant (temperature, flow, fuel composition) also correlates indirectly with emissions. **Meteorology data** (wind speed and direction, mixing height, etc.) influences pollutant dispersion — important for forecasting local air quality.



2.3 Data Types and Quality Challenges

Key data types that feed AI models include:

- **Time-Series Sensor Data:** Hourly average NO₂ at an ambient station, or continuous stack SO₂ readings. These often contain noise or calibration drift. Low-cost sensors in particular can drift in accuracy or be sensitive to humidity, requiring frequent calibration adjustments [11]. AI can help calibrate or filter these signals (see Module 3).
- **Spatial Data:** Mobile monitors (on vehicles or drones) produce spatial pollutant data which, combined with wind fields, can help attribute measured pollution to likely sources — factory plumes versus traffic — using AI pattern recognition.
- **Process & Activity Data:** Emission events often link to an activity. For instance, port emissions spike when multiple ships are simultaneously at berth, or a flare event at a refinery may occur when a process upset develops. These contextual data (ship schedules, production rates) can serve as AI model inputs to improve prediction accuracy.
- **Regulatory Data:** Permits may require computing certain metrics (e.g., 12-month rolling tons of pollutant or daily averages relative to a standard). AI can assist by continuously updating those metrics and flagging trends toward a potential violation.

2.4 Integration Note

In many cases, these data streams are siloed. **AI thrives when we integrate multi-source data:** linking an observed ambient spike with wind direction from an industrial area and that facility's production load — an ML model can incorporate all inputs to pinpoint the likely cause and advise an operator to investigate. The challenge for engineers is **ensuring data quality (calibration) and time alignment** when combining these streams — topics addressed in the next module.

Module 3: Sensor Calibration, Data Integration & Anomaly Classification

3.1 Sensor Calibration & Drift Management with AI

Low-cost sensors and even some continuous monitors require calibration to translate raw signals to pollutant concentration, and these calibrations can drift with time or environmental conditions [12]. Traditional calibration uses periodic co-location with reference monitors, but drift can occur between calibrations:

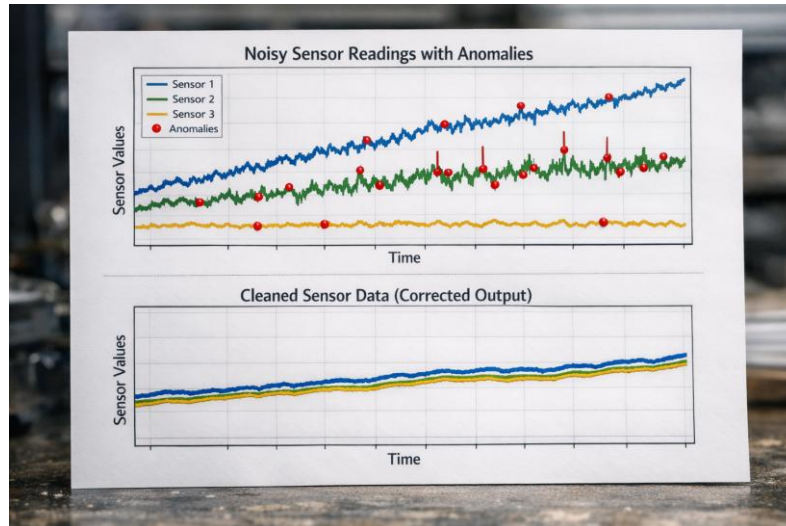
- **AIRsense Framework [3]:** A published study demonstrated an AI-based framework (referred to as AIRsense) that performs anomaly detection and correction on raw sensor readings before applying calibration models. The system used multiple algorithms to vote on whether raw data points were abnormal and then repaired those anomalies by interpolation or model prediction [3, 14], significantly improving subsequent calibration accuracy and lowering RMSE of pollutant readings.
- **Continuous Calibration Adjustment:** On-site ML models can continuously adjust calibration curves using reference comparisons or environmental correlations. For example, if one low-cost PM_{2.5} sensor in a network diverges from its neighbors, the system can flag it for recalibration. Automatic re-adjustment should be done carefully: **PE oversight** ensures we do not mis calibrate based on a false assumption (e.g., if all sensors read low due to a genuine widespread low-pollution episode, we would not want to erroneously adjust them upward).
- **Drift Detection:** Trending calibration residuals or monitoring a stable baseline (like nighttime zero readings) can feed an ML model to detect drift early. If a CO sensor baseline rises gradually over weeks, the system can compensate or alert maintenance.

3.2 Data Fusion and Integration

Multiple sensor types measuring overlapping or related variables can be integrated by ML to **improve anomaly detection**:

- A solitary reading spike might be sensor noise, but if multiple pollutants all spike in a consistent pattern, that suggests a real emissions event. AI can classify patterns across sensors, factoring in meteorology, to decide whether an anomaly is a genuine pollution incident or an outlier data point [15].
- Example: If only one sensor among ten in an area shows high NO₂ while its own NO reading is normal, AI might classify that NO₂ reading as suspect (instrument error) and initiate a calibration check. But if all ten sensors spike and meteorology indicates an upwind source, the system classifies it as a **pollution episode**.

- The **voting approach** from published research uses multiple algorithms and features to decide on anomalies [16]. In field contexts, cross-checks with external data are also valuable — e.g., did a local facility report an upset at that time, or did a dust storm occur?



3.3 Emissions Anomaly Classification

At a stationary source (e.g., a power plant stack), anomalies include sudden emission rate jumps or unusual pollutant ratios (a large CO increase without a corresponding CO₂ increase may indicate incomplete combustion). **AI can classify these anomalies:**

- **Likely operational upset vs. likely sensor error:** If all readings remain steady except one pollutant, an ML model might suspect an analyzer malfunction and raise an alert — which a PE would confirm via a backup method. Conversely, if multiple related emission parameters deviate together (CO up, O₂ down), the system classifies it as a **real process upset** needing intervention (e.g., adjusting burners).
- **Ambient anomaly classification:** For urban networks, ML might label a spike as a "fireworks event" or "dust event" versus an industrial release by examining pollutant combinations and timing — helping both the operational response and subsequent explanation to the public or regulators.

3.4 Ensuring Reliability of Anomaly Alerts

- **Hierarchical alerts:** Set "soft" anomalies for slight inconsistencies (for awareness) and "hard" anomalies for major probable issues requiring rapid action.
- **Managing false positives:** Combining evidence from multiple sensors reduces single-sensor glitch alarms. Requiring a persistence of anomaly over a few minutes before alarming avoids reactions to momentary blips.

- **Managing false negatives:** Missing an actual event is more serious; therefore, it is often prudent to tune systems to be **conservative** — some over-alerting is acceptable as long as it remains manageable — rather than missing a hazardous emissions excursion.

The outcome: by applying AI-driven data quality assurance and anomaly classification, subsequent forecasting and control modules are fed with **solid, trusted data** — ensuring the system neither cries wolf unnecessarily nor sleeps through a real event.

Module 4: Air Quality Forecasting & Advisory Control (Human-in-the-Loop)

4.1 Air Quality Forecasting with AI

Traditional air quality models (dispersion or chemical transport models) are physics-based but computationally intensive; ML offers a complementary approach that learns directly from data. Many jurisdictions now augment forecasting with ML:

- **Pollutant Concentration Forecasts:** Researchers in China and South Korea have developed ML systems that demonstrate significant improvements in forecast skill versus baseline models — demonstrating significant improvements in forecast skill versus baseline models. [17], by learning from historical sensor and meteorological data. Another system (AI-Air) specifically forecast pollutant spikes in certain regions [18].
- **Health Index Forecast (AQI):** Combining multiple pollutant forecasts into an Air Quality Index prediction for public advisory provides clearer communication of risk.
- **Confidence & Early Warnings:** ML models can produce a probability of exceeding a threshold (e.g., "40% chance PM_{2.5} reaches the Unhealthy for Sensitive Groups (USG) level tomorrow given forecasted conditions") and trigger alerts above a defined probability — creating an effective early-warning system for high-pollution days [19]. Some governments, including South Korea, have implemented AI-based early warning to enable issuing health advisories or controlling emissions proactively [20].

4.2 Operational Adjustments Based on Forecasts

Armed with a forecast, action plans triggered by pollution predictions can be more timely and geographically specific:

- **Industrial Emission Scheduling:** If stagnant air conditions (temperature inversion, low wind) are predicted the following day, an industrial facility might preemptively shift optional high-emission operations — such as a scheduled maintenance flaring — to a later date, or run pollution controls at maximum even if not strictly required under normal operating conditions.
- **Traffic Management:** City authorities might use ML to forecast when rush-hour traffic combined with weather will push pollution above thresholds. Ozone (O₃) episodes are common on hot, stagnant days; an early heads-up enables traffic flow adjustments or public advisories to reduce driving proactively.
- **Power Plant Dispatch:** An electric utility might factor air quality into dispatch decisions — running a cleaner unit more on days when an area is forecast near pollution limits — aided by AI providing a cost-versus-emission trade-off recommendation.



4.3 Bounded Adaptive Control (Human-in-the-Loop)

Moving from advisory to potential automated control requires carefully defined boundaries:

- Stationary Source Example:** A cement plant's kiln scrubber can be run at varying intensity. If an ML model detects emissions creeping upward alongside weather indicating poor dispersion, it might **advise increasing scrubber reagent feed by 10%** preemptively. That advice operates within a **bounded adjustment** (engineers define the $\pm 10\%$ range as safe and effective). The operator reviews and, if it aligns with expectations, implements the change — or the control system auto-implements after operator confirmation. AI fine-tunes control setpoints dynamically but always within pre-engineered boundaries.
- Mobile Source Example:** A smart traffic signal system could have an AI that temporarily adjusts signal timing when high pollution persists at an intersection to smooth traffic flow and reduce idling emissions. Those timing changes are limited to safe ranges to prevent congestion or safety hazards, and a traffic engineer monitors for adverse effects.

In both cases, **the human monitors outputs and retains control:** if any suggestion seems erroneous or causes unexpected effects, they override it. Safety nets — physical interlocks, backup alarms, operating limits — remain active at all times.

4.4 Ensuring Compliance & Safety in Adaptive Moves

- If an ML suggestion to ramp down operations would conflict with production or process safety, management must weigh the trade-offs; process safety typically takes precedence, but now engineers have the environmental risk insight to inform that decision.
- If an ML suggestion might itself push a secondary parameter out of compliance (e.g., increasing reagent feed creates a different waste stream issue), the engineer must evaluate and find an appropriate solution.

Final takeaway: AI forecasts and recommendations can lead to smarter, more proactive emissions management — but **engineers transform predictions into real actions** and ensure those actions are both effective and compliant. This synergy of AI foresight and human oversight can significantly curb harmful emissions and exposures while respecting all operational constraints.

Module 5: Compliance & Reporting — Monitoring, Documentation & Audit Trails

5.1 Continuous Compliance Monitoring

Modern regulations often require continuous data tracking. Some continuous emissions monitoring regulations — for example, 40 CFR Part 75 for large combustion sources — require data availability above 90–95%. Title V operating permits incorporate applicable monitoring requirements, which vary by source category and pollutant. Title V operating permits incorporate applicable monitoring requirements, which vary by source category and pollutant. Using AI does not change these obligations; in fact, it can strengthen compliance assurance by providing additional monitoring layers:

- **Non-Compliance Early Warning:** An AI model might continuously project current trends against permit limits — alerting, for example, "if current emission rate continues, the 24-hour average will exceed the limit by midnight." This gives operators time to intervene (reduce throughput, adjust controls) to avert the violation. This capability exceeds what a basic spreadsheet can do in real time.
- **Supplement to Official Monitors:** If an official CEMS fails, an AI model correlating process conditions could serve as a temporary backup estimator — helping avoid blind periods by flagging likely emissions spikes while the primary monitor is offline. Any such estimates remain subject to regulatory acceptance and must be clearly documented as model-derived, not certified measurements.

5.2 Documentation & Reporting

Regulators demand transparent records of operations, especially for deviations or emissions events. **AI-involved decisions must be explicitly recorded:**

- **Logging decisions:** Whenever AI triggers an action or alert, log the event: e.g., "14:00 — AI forecasted elevated PM for the next hour; operator reduced conveyor #2 throughput per the facility action plan; result — ambient PM remained below the permit limit." These logs become part of the compliance record.
- **Periodic reports:** Many permits require quarterly or annual reports focusing on outcomes: did the facility meet its limits? If an exceedance occurred, logs showing that AI provided a warning and the response was documented can demonstrate due diligence even when the event occurred.
- **Traceable data archiving:** All data used by AI models (raw sensor streams, model outputs) should be archived alongside normal monitoring data. If an agency audits, the response to "on what basis did you curtail operations on date X?" can be supported by the model forecast, associated uncertainty information, and the engineer's sign-off.

Compliance Log			
Date / Time	AI Alert	Operator Response	Outcome
Jan 3 10:00	Emission anomaly: CO spike detected	• Checked furnace oxygen control, • Identified issue, adjusted settings	• CO returned to normal in 15 min. • No permit exceedance.

5.3 Professional Defensibility & Regulator Communication

- **Emphasize human control:** "Our automated system can suggest adjustments, but it cannot implement them without operator confirmation, and all changes must remain within permitted operating conditions."
- **Demonstrate improvements with data:** "In the past year, our early-warning AI flagged 5 potential ozone-exceedance days, enabling proactive traffic and industrial measures — resulting in zero actual exceedances, whereas prior years averaged two per summer."
- **Be transparent about limitations:** "We use this model to support operations, but final compliance reporting is always based on calibrated EPA reference method measurements. The model helps us avoid approaching limits, not replace measurement obligations."

5.4 Connection to Public Communication

When cities or facilities publish real-time dashboards, AI outputs may indirectly feed those (such as an Air Quality Index forecast shared daily). Engineers should ensure any publicly disseminated information is vetted and includes appropriate uncertainty disclaimers — for example, "tomorrow is forecast as Moderate with a potential for Unhealthy for Sensitive Groups (USG) conditions in the afternoon" per the EPA AQI color scale. Keeping public trust means not overselling AI capabilities, but rather using it to **enhance transparency** about air quality conditions.

Module 6: Model Validation, Uncertainty & Risk Management

6.1 Validating AI Models Pre-Deployment

Before an AI model is allowed to influence a real compliance situation, it must be proven:

- **Historical Back testing:** Run the model on known past data to see if it would have predicted outcomes correctly. For example, test an air pollution forecast model on the previous year's episodes: did it detect them? Did it generate false positives? If performance falls below an acceptable threshold — perhaps a minimum 80% detection rate for high-concentration events with a manageable false-positive rate — improve the model or constrain its use.
- **Simulated Scenarios:** Create or identify edge cases — e.g., worst-case emissions combined with extreme weather — and ensure the AI responds reasonably (providing an alert with sufficient lead time). If an ML model has never encountered an extreme event in training, recognize its limits and rely on conservative manual protocols in those situations. Monte Carlo simulation can help gauge how model uncertainty grows in unusual scenarios.
- **Pilot Testing (Shadow Mode):** Run the AI tool in parallel for a period, allowing it to make "dry run" recommendations that are compared with actual operator decisions and outcomes. If the model frequently aligns with expert actions and outcomes are good, confidence grows. If not, refine or restrict the model's scope before live deployment.

6.2 Monitoring and Maintenance of Model Performance

- **Track performance metrics:** Monitor forecast errors (e.g., average difference between predicted and actual pollutant levels), anomaly detection false-alarm rates, and missed-event counts. Negative drift in these metrics over time may indicate **model drift** — underlying conditions have shifted beyond the original training data.
- **Periodic recalibration:** Retraining or updating model parameters using new data ensures the model adapts to changes, such as a new emission source in the area altering pollution patterns. Each update goes through the same change-management steps as the original deployment.
- **Cross-check with independent methods:** For crucial tasks like compliance forecasting, periodically compare ML output with a physics-based model or simpler baseline. If ML indicates no problem but a meteorological model warns of potential issues, investigate the discrepancy — the ML model may be missing an important input.

6.3 Uncertainty and Decision Thresholds

AI outputs should never be treated as certain. Providing **confidence intervals or probabilities** — and incorporating those in decisions — is essential:

- If a predicted peak NO₂ is 90% of the permit limit but carries a $\pm 15\%$ uncertainty band, it is prudent to act as if the limit could be reached.
- **Set conservative decision thresholds:** If the model indicates a 50% chance of exceedance would trigger action under a strict rule, engineers may choose 30% as their operational threshold to provide a safety margin.
- **Visualize and communicate uncertainty:** Forecast curves should show uncertainty bands (e.g., shaded ranges) so that decision-makers work with a realistic range rather than a single crisp value. Safety margins should be added, not subtracted, when acting on uncertain model outputs.

6.4 Risk Analysis (FMEA) for AI Implementation

Failure Modes and Effects Analysis (FMEA) systematically examines what-if scenarios for model errors or misuse:

- **Example FMEA Entry — Model fails to detect a major emissions event:** Consequence: possible late response, permit violation, health impact. Likelihood: moderate (if the model historically detected 19 of 20 such events). Controls: retain physical alarms and daily laboratory analyses as backstops; ensure operators are not solely reliant on AI; confirm that actual monitors will independently trigger notification if an exceedance occurs.
- **Example FMEA Entry — False positive triggers unnecessary process curtailment:** Consequence: production loss or cost. Some false positives may be accepted to ensure safety, but minimized via threshold tuning and by requiring human confirmation before action is taken.
- **Misuse risk:** If an untrained person or organization culture were to follow AI suggestions blindly, the mitigation is clear: **training** staff on what AI does and does not do, plus requiring **engineer sign-off** embedded in procedure — no direct automated control beyond the safe operating band without human authorization.

Documenting these risk evaluations — perhaps as an FMEA table appended to the facility's quality or management system documents — provides assurance for internal audits and a ready response if a regulator or auditor asks whether a specific failure scenario was considered.

By meticulously validating and monitoring models, and by formally analyzing potential failure modes, engineers create a safety net ensuring AI integration **does not inadvertently degrade environmental management reliability**. Instead, it remains an additional layer of protection and optimization, with robust human and system backups.

Module 7: Governance, Change Control & Cybersecurity

7.1 Configuration Management & Change Control

To maintain continuity and traceability, treat an AI model like any other component of the control or monitoring system:

- Each model version is **cataloged with details**: training data range, hyperparameters, performance metrics, and known limitations.
- When a model update is needed — due to drift, added data, or improved algorithms — run a **Management of Change (MOC)** process: define the change, assess impacts (e.g., will this update alter alarm frequency?), test the updated model, then obtain **approval from a senior engineer or change board** before deploying.
- This ensures no unvetted alterations slip through, just as unknown logic would not be loaded into a PLC without testing and approval. It also enables rapid rollback if something goes wrong after an update, because the prior version and change records are preserved.
- Document any operational **differences after a change** — e.g., the system became more sensitive in certain wind conditions — so operators are aware and calibrate their expectations accordingly.

7.2 Cybersecurity Considerations

AI implementations may involve new software or network connectivity (cloud or on-site edge computing), which must comply with **OT (Operational Technology) security best practices**. Applicable frameworks include **NIST SP 800-82** (Guide to Industrial Control System Security), **IEC 62443** (Industrial Automation and Control Systems Security), and **ISO/IEC 27001** for information security management:

- **Secure connectivity**: If sending data to/from cloud for AI computing, ensure encryption, authenticated channels, and compliance with applicable data-residency rules. Prefer **edge computing** (on-site model execution) to minimize unnecessary external connectivity where feasible.
- **Access control**: Only authorized engineers or data scientists should be able to access or modify the model environment. The AI system environment should be protected with the same rigor as any control system engineering station.
- **Tamper defense**: Implement **input validation** — if a sensor input suddenly reads outside its physical range, the system should question it before passing it to the model. Anomaly filtering, already discussed in Module 3, also provides a layer of cyber-defense. **Output sanity-checks** should prevent the control system from accepting any suggested action that is hazardous or physically implausible.

- **Asset inventory:** The AI system should be inventoried as an asset in the facility's ICS security register, with known network ports minimized, software kept current (subject to the same change-management testing process), and access logs maintained.

7.3 Ensuring Professional Defensibility

If the facility or engineer is ever challenged — in a regulatory hearing, audit, or legal proceeding — about decisions that involved AI:

- **Show your work:** Provide logs, performance test records, and documented oversight (e.g., engineer review sign-offs). This demonstrates that AI was used responsibly within a controlled, auditable process.
- **Characterize AI accurately:** "We employed an advanced data-analytic tool to augment our environmental monitoring, which — as our logs document — helped us detect issues faster, respond more promptly, and avoid potential permit exceedances. At all times, our licensed staff oversaw the process, ensuring we operated within safe and compliant parameters."
- **Peer comparisons:** If needed, point to industry references and published literature [1, 2, 21, 22] demonstrating that integrating AI for environmental monitoring is an emerging best practice — showing the engineer is on the forefront while still exercising appropriate due care.

7.4 Continuous Improvement Outlook

AI technology will continue to evolve, and so will the regulatory landscape. A well-governed approach ensures new improvements — better algorithms, richer sensor networks, integrated multi-media monitoring — can be adopted smoothly because a strong management system already exists. Meanwhile, **the core priority stays unchanged: protect air quality and public health.** If an AI system ever suggests something contrary to that mission, robust governance ensures a human engineer will catch it and correct course. With appropriate oversight, AI becomes a powerful ally in achieving cleaner air and regulatory excellence — amplifying the expertise and vigilance of environmental engineers rather than supplanting it.

References

- [1] World Economic Forum (2026). "Breathing digital: How AI is transforming air quality monitoring systems." WEF Stories, February 2026. <https://www.weforum.org/stories/2026/02/breathing-digital-ai-air-quality-monitoring-systems/>
- [2] World Economic Forum (2026). "Breathing digital: How AI is transforming air quality monitoring systems." (See [1] — also cited for the general framing of AI as an augmentation to engineer judgment in environmental operations.)
- [3] Zamora Zapata, M., et al. (2021). "AIrSense: AI-enhanced calibration and anomaly detection for air quality sensor networks." PMC / National Institutes of Health. <https://pmc.ncbi.nlm.nih.gov/articles/PMC9867200/>
- [4] Schneider Electric Blog (2024). "Emissions monitoring through real-time machine learning." April 18, 2024. <https://blog.se.com/industry/energies-and-chemicals/2024/04/18/emissions-monitoring-through-real-time-machine-learning/>
- [5] Schneider Electric Blog (2024). (See [4] — also cited for the framing of AI as decision-support under PE oversight.)
- [6] Zamora Zapata, M., et al. (2021). (See [3] — also cited in the context of V&V using historical data.)
- [7] World Economic Forum (2026). (See [1] — also cited for discussion of AI integration challenges including data needs, cost, and workforce.)
- [8] World Health Organization (WHO). "Air pollution." WHO Fact Sheet. Geneva: WHO. Available at: [https://www.who.int/news-room/fact-sheets/detail/ambient-\(outdoor\)-air-quality-and-health](https://www.who.int/news-room/fact-sheets/detail/ambient-(outdoor)-air-quality-and-health) (Note: the 7 million figure combines indoor and outdoor air pollution mortality estimates.)
- [9] World Economic Forum (2026). (See [1] — cited for the role of AI in real-time early warning of pollution spikes and smog episodes.)
- [10] World Economic Forum (2026). (See [1] — cited in the context of ground-based ambient monitoring station networks.)
- [11] Zamora Zapata, M., et al. (2021). (See [3] — cited for low-cost sensor drift and calibration challenges.)
- [12] Zamora Zapata, M., et al. (2021). (See [3] — cited for calibration drift in low-cost and continuous monitors.)
- [13] Zamora Zapata, M., et al. (2021). (See [3] — cited for the AIrSense anomaly detection and correction pre-processing framework.)
- [14] Zamora Zapata, M., et al. (2021). (See [3] — cited for the multi-algorithm voting and anomaly repair by interpolation or model prediction.)
- [15] Zamora Zapata, M., et al. (2021). (See [3] — cited for multi-sensor anomaly pattern classification incorporating meteorology.)
- [16] Zamora Zapata, M., et al. (2021). (See [3] — cited for the voting ensemble approach to anomaly classification.)

- [17] World Economic Forum (2026). (See [1] — cited for ML-based pollutant forecast performance improvements reported in Asian deployment studies.)
- [18] World Economic Forum (2026). (See [1] — cited for the AI-Air system for pollutant spike forecasting.)
- [19] World Economic Forum (2026). (See [1] — cited for probabilistic early-warning systems for high-pollution days.)
- [20] World Economic Forum (2026). (See [1] — cited for South Korea's implementation of AI-based air quality early-warning systems.)
- [21] World Economic Forum (2026). (See [1] — cited as evidence of industry-wide movement toward AI-assisted environmental monitoring.)
- [22] World Economic Forum (2026). (See [1] — cited alongside [21] for peer comparison supporting professional defensibility of AI usage.)

Recommended Additional Reading

- U.S. EPA (2023). Quality Assurance Handbook for Air Pollution Measurement Systems, Volume II: Ambient Air Quality Monitoring Program. EPA-454/B-17-001. Research Triangle Park, NC: U.S. EPA Office of Air Quality Planning and Standards.
- U.S. EPA (2017). 40 CFR Part 75 — Continuous Emission Monitoring. Code of Federal Regulations. Washington, DC: U.S. Government Publishing Office.
- National Institute of Standards and Technology (NIST) (2023). Guide to Industrial Control Systems (ICS) Security. NIST Special Publication 800-82, Rev. 3. Gaithersburg, MD: NIST.
- International Electrotechnical Commission (IEC). IEC 62443 Series — Security for Industrial Automation and Control Systems. Geneva: IEC.
- ISO (2015). ISO 14001:2015 — Environmental Management Systems: Requirements with Guidance for Use. Geneva: International Organization for Standardization.
- Kumar, P., et al. (2015). Real-time sensors for indoor air monitoring and challenges ahead in deploying them to urban buildings. *Science of the Total Environment*, 560–561, 150–159.
- Lary, D.J., et al. (2016). Machine learning in geosciences and remote sensing. *Geoscience Frontiers*, 7(1), 3–10.
- Shrestha, P., et al. (2021). Anomaly detection in industrial control systems using machine learning: A survey. *Sensors*, 21(24), 8210.
- NSPE (2019). NSPE Code of Ethics for Engineers. National Society of Professional Engineers. Alexandria, VA: NSPE.