



*An Online PDH Course
brought to you by
CEDengineering.com*

AI For Waste Water Treatment Optimization

Course No: C01-042
Credit: 1 PDH

Brian Lisiewski, P.E.



Continuing Education and Development, Inc.

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

Table of Contents

Module 1: Introduction — Safety-First Mindset and Engineer Accountability	1
1.1 Public Health and Environmental Protection as Top Priorities	1
1.2 The Engineer's Role in AI Integration	1
1.3 Aligning with Quality Management Systems	2
Module 2: Treatment Process Fundamentals and Data Landscape	3
2.1 Treatment Processes — A Brief Review	3
2.2 Operational Data Sources	3
2.3 Data Integration Challenges and AI's Role	4
Module 3: Process Monitoring — Anomaly and Failure Detection	5
3.1 Traditional Alarms versus AI Anomaly Detection	5
3.2 Failure Prediction for Equipment and Processes	5
3.3 Managing False Alarms and Missed Events	6
Module 4: Predictive Control and Operational Optimization.....	7
4.1 Proactive Process Adjustments	7
4.2 Adaptive versus Static Operation.....	7
4.3 Maintaining Safety Margins	8
4.4 Documenting and Validating Outcomes.....	8
Module 5: Compliance Assurance and Audit Trails	9
5.1 Compliance Monitoring and Alerts	9
5.2 Automated Reporting and Audit Trails.....	9
5.3 Defensibility to Regulators	10
Module 6: Validation, Metrics, and Uncertainty Management.....	11
6.1 Offline Testing and Commissioning of AI Models	11
6.2 Performance Metrics for AI in Operations	11

6.3 Embracing Uncertainty and Conservatism	12
Module 7: Governance, Change Control, and Professional Defensibility	14
7.1 Model Configuration Control and Updates.....	14
7.2 Cybersecurity and Access Control.....	14
7.3 Professional Defensibility and Transparency	15
7.4 Continuous Improvement and Learning	15
References	17

Module 1: Introduction — Safety-First Mindset and Engineer Accountability

Overview: Drinking water and wastewater treatment are safety-critical services — the stakes are high as they directly impact public health and the environment. This module establishes that Professional Engineers remain in responsible charge: AI/ML is a tool to assist, not a substitute for, the expert judgment of licensed operators and engineers. Public health protection and regulatory reliability are paramount, and AI must operate within those boundaries.

1.1 Public Health and Environmental Protection as Top Priorities

Water treatment plants (WTPs) and wastewater treatment plants (WWTPs) exist to ensure safe water for communities and prevent pollution of waterways. Every decision — whether made by a human or suggested by an AI — must align with those fundamental goals. For example, an algorithm might find a way to reduce chemical usage by adjusting dosing, but the engineer must ensure this will not risk microbial or chemical compliance. No optimization can violate a safe limit: delivering safe drinking water or properly treated effluent is non-negotiable. Any AI suggestions must be carefully evaluated for compliance with health standards (e.g., disinfectant residual, turbidity, nutrient limits) before implementation.

1.2 The Engineer's Role in AI Integration

Licensed PEs have the ethical and legal duty to sign off on operational changes that could affect public health. The AI's role is advisory: it can analyze large volumes of plant data and propose adjustments or flag concerns, but the human must validate and approve all significant process changes. This is consistent with overarching engineering ethics (as per NSPE) — the engineer ensures public welfare is protected and applicable code requirements (e.g., safe drinking water standards) are met.

AI recommendations are embedded into standard operating procedures so that each recommendation requires explicit human confirmation before any action is taken. For instance, an ML model might predict that a filter needs backwashing sooner than scheduled; an operator would inspect or cross-check before acting on that suggestion, just as they would with any SCADA alarm.



1.3 Aligning with Quality Management Systems

Water utilities often have robust quality management systems (such as ISO 9001 or ISO 14001) that emphasize risk management and process control. AI integration should follow the same disciplined approach: documented procedures for model training, thorough testing, and periodic review. The engineer must ensure transparency — for example, by keeping a log of when AI suggestions were implemented and their outcomes — to feed into continuous improvement cycles and to satisfy audits. If any incident were to occur, the utility can demonstrate the decision process and that an engineer oversaw every step (crucial for legal defensibility). An AI system's configuration is treated similarly to a critical piece of equipment: subject to change control, traceability, and frequent performance verification — just like a chlorination system or a SCADA logic change — all guided by the same risk-based thinking that underpins operations.

Key Takeaways — Module 1

- AI is a decision-support tool; PEs retain full legal and ethical responsibility for all process decisions
- No AI suggestion may be implemented without human review and approval
- Public health protection is the non-negotiable boundary within which all AI optimization must operate
- AI governance should follow the same disciplined change-control and documentation practices applied to SCADA and process equipment
- Transparency and audit trails are essential for regulatory defensibility

Module 2: Treatment Process Fundamentals and Data Landscape

Overview: Effective use of AI in treatment and compliance relies on understanding the processes and data intimately. This module briefly reviews key processes in water and wastewater treatment and describes the wealth of operational data (sensors, lab tests, etc.) available, which is the foundation for any AI model. The focus is on what data and process knowledge feed into AI, as AI outputs are only as good as their inputs.

2.1 Treatment Processes — A Brief Review

Water and wastewater plants have multi-step processes: coagulation, sedimentation, filtration, and disinfection for drinking water; and screening, biological treatment, clarification, and disinfection for wastewater. Each step has critical control points (e.g., clarifier effluent turbidity, aeration oxygen levels). Process reliability means avoiding upsets — such as microbial breakthrough or nitrification failure — that could lead to regulatory violations or unsafe water.

Conventional control relies on setpoints and alarms: for example, maintain chlorine residual at or above a defined minimum and backwash filters when turbidity exceeds a threshold. These setpoints are often conservative to ensure compliance, leaving room for optimization that AI might help identify — but always within safe margins.

2.2 Operational Data Sources

Modern plants generate substantial operational data from multiple sources:

- **SCADA sensor data:** Real-time flows, pressures, chemical feed rates, dissolved oxygen, pH, ORP, turbidity, and more — typically logged at intervals ranging from seconds to minutes. A large WWTP may have thousands of sensors streaming data continuously
- **Laboratory results:** Periodic analyses (e.g., daily BOD, residual chlorine, microbial counts) provide ground-truth water quality outcomes. These complement SCADA data by validating how process changes affect actual compliance metrics
- **External data:** Weather (heavy rainfall affects wastewater flows; source water changes affect drinking water processes), energy costs, and inputs from upstream processes (such as an industrial discharge composition) can all influence operations
- **Maintenance and events logs:** Operator notes and maintenance records (pump runtime, backwash frequency) contain expert knowledge of what happened and when. AI can mine these logs using natural language processing or pattern matching for clues that correlate with sensor anomalies — for example, repeated clogging events following holiday weekends



2.3 Data Integration Challenges and AI's Role

Historically, only a fraction of collected data is actively used, partly because humans cannot process it all in real time. AI excels at fusing diverse data streams. It can:

- Correlate lab results with SCADA trends (e.g., noticing that a subtle pH fluctuation predicted a compliance issue several hours later)
- Identify patterns spanning multiple processes, such as a combination of rising filter turbidity and lower UV transmittance indicating an early sign of filter breakthrough
- Handle nonlinear, multivariate influences that traditional rule-based controls may not detect (e.g., how both temperature and organic load together influence clarifier performance)

However, integration requires data cleaning (removing sensor noise or outliers) and alignment (synchronizing timestamps, bridging SCADA 5-minute data with daily lab data via interpolation or ML). Engineers must supervise these steps to ensure data quality. Traceability of data transformations is important for trust and audit — for example, documenting how missing sensor readings are handled provides a clear record for regulators and management. The payoff is that AI can then mine insights from cross-cutting data to support better decisions in ways not previously feasible in manual operation.

Key Takeaways — Module 2

- AI models are only as good as their input data — data quality and integration are foundational
- SCADA sensor data, laboratory results, external data, and maintenance logs are all valuable AI inputs
- Data cleaning, alignment, and transformation must be documented for auditability
- AI can detect nonlinear, multivariate patterns that exceed the capacity of traditional threshold-based alarms
- Engineers must supervise data pipelines, not just model outputs

Module 3: Process Monitoring — Anomaly and Failure Detection

Overview: One of the immediate benefits of AI in treatment operations is improved monitoring. This module covers how AI/ML models can detect anomalies and incipient equipment or process failures from large data sets, often earlier and more reliably than traditional alarms. The focus is on how these models provide timely, explainable alerts that engineers and operators can act on, and how to manage false alarms versus missed detections.

3.1 Traditional Alarms versus AI Anomaly Detection

In a typical plant SCADA, alarms are threshold-based — for example, blower current exceeding a high limit, or effluent chlorine falling below a low setpoint. These are straightforward but can miss nuanced patterns. AI-driven anomaly detection uses machine learning to learn the normal patterns of the plant and flag deviations from that norm.

For example, an ML model might recognize that filter turbidity normally follows a diurnal pattern; if it begins trending upward in the morning — even while still below the alarm threshold — the model flags it as an anomaly and anticipates a possible filter breakthrough. This allows operators to intervene proactively (e.g., backwash early) before an actual breach occurs. AI can consider multiple signals simultaneously: that turbidity anomaly may be significant only if paired with rising headloss and a particular raw water condition, which the ML can incorporate while a single-sensor alarm could not.

3.2 Failure Prediction for Equipment and Processes

Preventive maintenance in water facilities often relies on fixed schedules or operator intuition. AI can enhance this by detecting patterns that precede failures:

- **Pump and motor failures:** ML models can monitor vibrations, motor current signatures, and runtime patterns, learning normal signatures so as to alert when deterioration is detected (bearing wear, cavitation onset). In one study, anomaly detection applied to pump station time-series data achieved 98% accuracy in identifying pump faults ahead of breakdown. Real-world performance varies and requires site-specific validation
- **Process upsets:** Sudden changes — such as a slug of high-BOD influent entering a wastewater plant or a drop in raw water alkalinity — may not trigger immediate alarms but can lead to permit violations hours later. AI can correlate such early signals with past outcomes and generate a risk alert recommending preemptive adjustments. Early research indicates that deep learning can detect subtle multi-sensor anomalies in WWTPs with approximately 99% accuracy in controlled test scenarios; real-world results are typically lower and require careful fine-tuning before operational deployment



3.3 Managing False Alarms and Missed Events

No anomaly detector is perfect. False positives (false alarms) waste staff time or erode trust if too frequent; false negatives (missed detections) can defeat the purpose by failing to provide early notice. Tuning an AI model for a treatment plant involves establishing a decision threshold on anomaly scores to balance these outcomes.

For public health reasons, recall — capturing all true problems — is often prioritized over precision, meaning a facility would rather accept occasional false alarms than miss a serious issue. Nevertheless, alarm fatigue is a concern, so multi-tier alerts can help: an AI model might produce a watch alert when conditions are slightly abnormal and a critical alert only when confidence of a real problem is high. The PE must configure and approve these thresholds, and may incorporate a time-delay or persistence requirement to minimize reactions to momentary transients. If the AI system is offline or uncertain, the default must revert to normal conservative operation and traditional alarms.

In practice, a combined approach is emerging: AI identifies anomalies and provides context (which signals are contributing, predicted possible cause), while engineers and technicians verify on the ground. This collaboration leverages AI's pattern-spotting capability and human expertise's judgment to improve reliability and reduce unnecessary interventions. Each flagged anomaly and its response should be logged as part of an audit trail for continuous improvement and regulatory review.

Key Takeaways — Module 3

- AI anomaly detection identifies subtle or multivariate deviations that threshold-based alarms cannot detect.
- Accuracy figures from published studies represent test-scenario results; real-world performance requires site-specific validation.
- Thresholds balancing false alarms and missed detections must be configured and approved by a PE.
- A multi-tier alert structure (watch / critical) reduces alarm fatigue while preserving safety.
- Traditional alarms must remain active as a fallback whenever AI systems are offline.

Module 4: Predictive Control and Operational Optimization

Overview: AI can do more than flag issues — it can also suggest optimal operational adjustments that maintain safety while cutting costs or improving efficiency. This module explores predictive control applications, where ML models analyze current data and weather forecasts to propose setpoint changes or rescheduling actions that prevent problems and optimize performance. All such actions are reviewed by human operators and engineers, and are typically implemented gradually or in supervised mode.

4.1 Proactive Process Adjustments


Unlike conventional control (which reacts after a change is detected), predictive control attempts to anticipate changes before they occur. Examples include:

- **Storm inflow management:** A wastewater plant's AI model might predict, based on a rain forecast and catchment data, that flows will double by evening. Rather than waiting for a high-flow alarm, the system suggests preemptively lowering tank levels and starting additional pumps to create hydraulic capacity. The operator confirms these actions, which help mitigate combined sewer overflows
- **Chemical dosing optimization:** Drinking water plants often overdose coagulants or disinfectants as a safety buffer, then use other chemicals to correct the excess. ML can identify leaner dosing strategies. In at least one documented case, an AI system recommended chlorine setpoint adjustments based on influent quality predictions, achieving 10–30% reductions in chemical use while preserving full disinfection compliance (Jacobs Engineering, 2023). Such recommendations are implemented gradually and only after manual verification that they will not risk a permit violation

4.2 Adaptive versus Static Operation

The concept of adaptive control — dynamically changing setpoints or control rules — is valuable for matching operational intensity to real-time conditions. AI can predict when to ramp blowers or mixers slightly ahead of load changes, reducing energy use. However, any adaptivity must be implemented carefully:

- **Pre-validated scenarios:** Typically, the AI toggles between a few pre-approved operating modes rather than making continuous arbitrary adjustments. For instance, a peak flow mode and an average flow mode with pre-set parameter sets might be defined, with the AI selecting when to switch. Each mode must be thoroughly studied in advance by engineers for stability and compliance before being approved for use
- **Operator override and phasing:** Operators must be able to override or veto any adaptive changes at any time. Changes should also be applied incrementally — ramping rather than in abrupt steps — to avoid upsetting the biological or chemical processes, and to give operators time to observe and confirm that no adverse effects are occurring



4.3 Maintaining Safety Margins

Even with predictive control, the principle of doing no harm applies. Engineers constrain AI suggestions with documented business rules. For example, an ML model might suggest reducing aeration to minimal levels at night to save energy, but the engineer sets a hard floor: dissolved oxygen shall not drop below 2.0 mg/L, and any AI suggestion that would violate that floor is automatically rejected. This ensures that even if the model is overly aggressive, the system retains a safe operating buffer. AI proposals operate within a safe envelope defined by the engineer, often narrower than actual regulatory limits to provide an additional safety margin

4.4 Documenting and Validating Outcomes

When predictive control changes are implemented, their outcomes must be tracked and documented. If an AI-suggested change leads to improvement, that result should be recorded. If a suggestion had to be aborted — for example, because it caused a slight quality drift — that should be recorded as well. This is part of the continuous validation and model improvement cycle: the model may be retrained on these outcomes to refine future recommendations, under careful oversight. Documentation also builds an evidence base demonstrating that AI usage is controlled and safe, which is valuable for both internal quality assurance and any regulatory inquiry

Key Takeaways — Module 4

- Predictive control anticipates process conditions and allows proactive intervention before a problem develops
- Chemical savings of 10–30% have been documented in at least one case; results will vary by facility and must not be assumed
- Adaptive control must operate within pre-validated, engineer-approved operating modes
- Hard safety floors (e.g., minimum DO levels) must be enforced by the control system and cannot be overridden by AI
- All AI-driven operational changes and their outcomes must be documented for continuous validation

Module 5: Compliance Assurance and Audit Trails

Overview: Compliance with regulations — such as safe drinking water standards and NPDES (National Pollutant Discharge Elimination System) permit limits under the U.S. Clean Water Act, or equivalent frameworks in other jurisdictions — is the bedrock of water sector operations. AI can assist in monitoring and documenting compliance by predicting potential exceedances before they occur and by supporting thorough record-keeping. This module discusses how AI/ML can be integrated into compliance management without replacing regulatory accountability processes.

5.1 Compliance Monitoring and Alerts

A major concern for plant managers is an unforeseen permit violation — such as a bacteriological exceedance or an effluent nutrient spike. AI models can serve as an early warning system: by analyzing current process trends and historical conditions, an ML model might forecast that, if current trends continue, the next day's effluent ammonia could approach or exceed the permit limit. This might trigger a compliance risk alert well in advance, giving the team time to adjust operations (e.g., increase aeration or chemical dosing) to remain within limits.

This is a more proactive approach compared to traditional practice, where an issue may only be discovered when lab results return — possibly after a violation has already occurred. AI can also monitor distribution system water quality: for example, predicting a potential loss of chlorine residual in a peripheral zone days ahead (due to a heat wave or low turnover) and suggesting rechlorination at a booster station in that zone.

5.2 Automated Reporting and Audit Trails

Regulations require documentation of operations and prompt reporting of any non-compliance. AI systems should be configured to enhance, not obscure, this documentation:

- Every AI recommendation, and the operator's action (or decision not to act), should be logged automatically in the control system or a separate audit log. For example: '3:00 PM — AI predicted high turbidity in Train 2 filter; operator increased coagulant by 5%. Outcome: filter effluent remained within specification.' This creates a clear audit trail for internal review and regulators, demonstrating the facility's diligence
- Regulatory reports can be partially assisted by AI — for instance, an ML system could help identify unusual data points in monthly operating reports and confirm they were addressed or did not cause a violation. However, the PE must review and sign off on all regulatory reports; automated content is a drafting aid, not a substitute for professional review
- AI can support the trending and statistical analysis required for compliance documentation. For example, instead of a labor-intensive manual review, an AI can rapidly graph and analyze effluent trends — still with human verification — serving as a powerful tool for evidence-backed compliance narratives



5.3 Defensibility to Regulators

Using AI in operations may raise questions from regulators or the public: 'How do we know the system is not cutting corners?' The answer lies in good governance and transparency:

- **Documented model validation:** If a regulator audits the plant and encounters an AI system, they should be shown a documented validation report verifying the model's accuracy, its limitations, and the fact that a human is always in control for critical decisions. Referencing ISO 9001 risk management or AWWA/WEF best practices demonstrates due diligence
- **No black-box decisions:** Make clear that no parameter is adjusted solely by AI without operator and engineer oversight, and that all adjustments remain within proven safe ranges. If possible, provide log examples where AI gave a recommendation but the operator overrode it based on contextual knowledge — demonstrating that human expertise remains primary
- **Regular review:** Emphasize that the system is subject to regular review and fine-tuning by PEs — for example, monthly performance reviews as part of management review cycles — consistent with quality management principles. This assures regulators the technology is being actively managed and not left to run unchecked

Key Takeaways — Module 5

- AI can provide early warning of potential permit exceedances, enabling proactive corrective action
- NPDES permits are a U.S. regulatory framework; equivalent frameworks apply in other jurisdictions
- Every AI recommendation and operator action must be logged to create a defensible audit trail
- PEs must review and sign off on all regulatory reports regardless of AI assistance in drafting
- Transparency, documented validation, and regular PE review are the foundations of regulatory defensibility

Module 6: Validation, Metrics, and Uncertainty Management

Overview: Integrating AI in treatment operations introduces model risk — the chance that the model may be wrong. This module outlines how to validate AI models thoroughly before use, how to monitor their performance using appropriate metrics, and how to incorporate uncertainty considerations into decision-making to ensure safety margins are maintained. An AI model is treated similarly to a new piece of process equipment: tested, calibrated, and with known operating limits.

6.1 Offline Testing and Commissioning of AI Models

Before an AI model is trusted for operational support, it must undergo rigorous Verification and Validation (V&V):

- **Historical replay testing:** The model is run on historical data where outcomes are already known — for example, feeding it the previous year's plant data to see if it would have correctly predicted a turbidity incident or flagged a pump failure in advance. If it missed events or generated too many false alarms, it is tuned, retrained, or its scope is limited to tasks it performs reliably
- **Pilot trials in parallel with manual control:** A best practice is to initially operate the model in shadow mode — it provides recommendations or alerts, but operators do not rely on them exclusively. This allows the team to assess real-world performance and adjust thresholds and logic accordingly. For example, a plant might run an AI dosage optimizer in parallel for a month to compare its recommendations with actual dosing, verifying it would never have violated permit limits in any scenario — and that it would have achieved chemical savings. Only after this verification would acting on its recommendations be considered
- **Verification against known standards:** If the AI handles regulated metrics (such as forecasting disinfection byproduct levels), it should be calibrated against validated water quality models or pilot plant results. This ensures the ML is not overfitting noise — a significant risk in small data sets


6.2 Performance Metrics for AI in Operations

To quantify how well an AI model is contributing, several metrics are relevant:

- **Accuracy and error:** For predictions such as effluent quality forecasts, use Mean Absolute Error (MAE) — the average prediction error in the same units as the output — or Mean Absolute Percentage Error (MAPE), which expresses that error as a percentage of the true value. For example, if an ML model predicts effluent nitrate with a mean error of 0.2 mg/L and the permit margin is ± 1 mg/L, the model is sufficiently accurate for that application. If error is too high, the model may only be suitable for qualitative trending rather than direct control guidance
- **False alarm and detection rates:** Track how often anomaly detection raises flags versus actual issues. If a model raises 20 major alerts per month and only 2 prove to be real issues, precision is 10%. The team might raise the threshold or narrow the scope to improve

precision without sacrificing recall for critical events. Target metrics can be defined — for example, a false alarm rate below 5 per month and a near-miss detection rate above 90% — to guide ongoing improvement

- **Process performance impact:** Ultimately, measure how AI affects real results. Key performance indicators (KPIs) can include compliance rates (e.g., percentage of days meeting all permit limits), chemical and energy usage (reduction without increased risk), and reliability metrics (e.g., number of emergency shutdowns or process upsets per year). In documented cases, AI-based dosing control has achieved meaningful reductions in chemical use while maintaining compliance, and anomaly detection systems have measurably reduced unplanned equipment downtime



6.3 Embracing Uncertainty and Conservatism

When an AI output informs a decision, prudent engineers account for model uncertainty. ML models can often be configured to provide a confidence interval or probability range alongside each prediction. For example, if an AI predicts tomorrow's effluent turbidity will be 0.3 NTU with a ± 0.1 NTU uncertainty range at 95% confidence, and the permit limit is 0.5 NTU, the operator might make a minor process adjustment because the upper bound of 0.4 NTU is still safely below the limit. However, if the prediction is $0.45 \text{ NTU} \pm 0.1 \text{ NTU}$, placing the upper bound at 0.55 NTU — above the limit — more aggressive action would be warranted even if the central prediction is within compliance. Uncertainty ranges drive decisions alongside central predictions, preserving a conservative operating stance.

A contingency plan should always be in place: if an AI-driven action does not produce the expected result, or if data quality degrades (sensor fouling can confuse models), reversion rules must be pre-established. For example, if an ML-chosen setpoint causes effluent quality to approach a permit limit more quickly than anticipated, automated logic should revert to the last safe setpoint, or an alarm should trigger for manual override. This parallels how advanced process control systems provide bumpless transfer back to manual mode when conditions go out of the expected range. The cost of an occasional reversion to conservative operation is a small price for assured compliance.

Key Takeaways — Module 6

- All AI models must undergo V&V — including historical replay testing and shadow-mode parallel trials — before operational reliance
- MAE (Mean Absolute Error) and MAPE (Mean Absolute Percentage Error) are standard metrics for evaluating prediction accuracy
- Define target KPIs for both safety (detection rate, false alarm rate) and efficiency (chemical/energy savings, compliance rate)
- Uncertainty intervals, not just central predictions, must inform decisions to maintain conservative operating margins
- Pre-established reversion rules are essential — AI systems must have a clear fallback to safe manual operation

Module 7: Governance, Change Control, and Professional Defensibility

Overview: To sustainably use AI in operations, robust governance is needed. This final module synthesizes how to manage AI/ML models over their lifecycle to keep them safe, secure, and effective. It covers model versioning, access control, cybersecurity, and how to maintain professional defensibility — ensuring that any AI-related practice stands up to audit or legal scrutiny, just as any other engineering decision would.

7.1 Model Configuration Control and Updates


An AI model should be treated as part of the control system. If it is updated or retrained, that change should be managed like a modification to a PLC program or a control strategy:

- **Change management:** Establish a formal procedure where any model update — such as adding new data or features — is documented, tested offline, and approved by a senior engineer before deployment. This aligns with ISO 9001's requirements for control of process changes and ensures traceability. For instance, if an AI's anomaly detection threshold is adjusted, the log should capture why (e.g., false alarm rate was too high), what tests were performed to confirm improved performance, and when and by whom the change was authorized
- **Versioning and rollback:** Maintain a repository of model versions with timestamps and documented differences. If an update inadvertently causes issues, revert to the previous stable model — similar to maintaining backups of previous PLC or relay settings. This is a core risk management practice: never be in a position where a model update cannot be undone

7.2 Cybersecurity and Access Control

AI systems often run on networked servers or edge devices connected to plant control networks, and must adhere to strict operational technology (OT) cybersecurity standards. Key practices include:

- Limit access to model systems as you would for SCADA and PLC systems — for example, requiring multi-factor authentication for anyone who can view or tune the model.
- Protect data streams — especially those used for control decisions — from tampering. Use secure protocols for data ingestion to prevent an attacker from injecting false sensor readings that could mislead the AI
- Guard against adversarial ML risks: ensure that model recommendations are plausible and cross-checked. If a recommendation is wildly outside the normal operating range, require additional human confirmation before acting. This is consistent with water sector guidance from AWWA and ICS-CERT recommendations for industrial control system environments: no single compromised model component should be able to cause an unsafe action. Safety interlocks must remain in place and function independently of the AI layer



7.3 Professional Defensibility and Transparency

If ever questioned about a decision influenced by AI — for instance, why an operator changed a setpoint that led to an unusual operating condition — the engineer should be able to demonstrate:

- The decision was based on sound data analysis, with the AI's recommendation accompanied by clear rationale (e.g., 'AI flagged rising pH and predicted an effluent ammonia exceedance; alkalinity dosing was increased as a precaution')
- The entire process was documented and within design parameters, with logs showing model suggestion histories and confirmation that all adjustments remained within pre-established safe bounds
- The decision was reviewed by a competent professional (the on-duty process engineer) and was consistent with regulatory obligations, referencing the plant's operations plan as applicable

Maintaining this level of documentation and oversight produces a system that can withstand scrutiny during internal audits, third-party reviews, or legal challenges. It demonstrates that even though advanced algorithms were used, the PEs in charge exercised due care and diligence in their integration, as required by licensure laws and professional ethics.

7.4 Continuous Improvement and Learning

Any integration of AI should itself be part of a continuous improvement process. Engineers should periodically evaluate: Is the model meeting its performance goals? If not, refine or retrain. Are there new data sources or plant changes that require updating the model? This mirrors the approach taken in a quality management system: periodic management reviews examine how often AI alerts were accurate, how outcomes improved, and whether the approach should be updated.

For example, if a model's error increases after a process upgrade — due to a shift in normal operating conditions — that finding should trigger a controlled retraining using new data to realign model performance. By treating the AI with the same discipline applied to any critical piece of plant equipment or control software, it remains a reliable partner in delivering safe, compliant water and wastewater services with greater efficiency.

Key Takeaways — Module 7

- Every AI model update must follow formal change management: documented, tested offline, and approved by a senior PE before deployment
- Model version control and rollback capability are essential risk management practices
- OT cybersecurity standards apply to AI systems: limit access, secure data streams, and validate that recommendations are plausible before acting
- Professional defensibility requires complete documentation: rationale, logs, safety bounds, and PE review for every AI-influenced decision
- Continuous improvement cycles — periodic performance reviews, retraining when warranted — keep the AI aligned with evolving plant conditions

References

- [1] Jacobs Engineering Group. (2023). AI and machine learning in wastewater treatment optimization. Jacobs OneWater Reflections. <https://www.jacobs.com/newsroom/onewater-reflections/ai-and-machine-learning-wastewater-treatment-optimization>
- [2] American Water Works Association (AWWA). Cybersecurity guidance for the water sector. Available at: <https://www.awwa.org/Resources-Tools/Resource-Topics/Cybersecurity>
- [3] Water Environment Federation (WEF). Resources on process optimization and advanced controls in water resource recovery facilities. Available at: <https://www.wef.org>
- [4] International Organization for Standardization. (2015). ISO 9001:2015 — Quality management systems: Requirements. ISO.
- [5] National Society of Professional Engineers (NSPE). Code of ethics for engineers. Available at: <https://www.nspe.org/resources/ethics/code-ethics>
- [6] U.S. Environmental Protection Agency (EPA). National Pollutant Discharge Elimination System (NPDES). Available at: <https://www.epa.gov/npdes>
- [7] ICS-CERT / CISA. Recommended practice: Improving industrial control system cybersecurity with defense-in-depth strategies. U.S. Cybersecurity and Infrastructure Security Agency. Available at: <https://www.cisa.gov/ics>