AI Driven Predictive Monitoring and Optimization of Decentralized Water Treatment Infrastructure in Resource Constrained Urban Settlements

Ashok Khedkar, Shashikant Sahare , Sandeep Musale, Ajit Bhosale

Abstract: *Decentralized drinking water treatment in resource constrained urban settlements often depends on sparse and messy monitoring, which can delay the detection of unsafe or unreliable operation. Reactive checks and fixed thresholds may miss early warning signals and can contribute to alarm fatigue when data are intermittent. A practice-oriented framework is presented that converts low-cost sensor streams, sparse operator logs, and occasional lab results into short horizon forecasts, anomaly scores, and soft sensor estimates, and then links these outputs to human in the loop decisions for real-world use. The architecture is organized as staged sensing and logging, data quality checks and gap handling, predictive monitoring, a decision layer with persistence and capacity aware escalation rules, an action layer, and feedback logging to support threshold tuning and model update triggers. Transferability is bounded to the evaluated setting, and the scope excludes autonomous closed loop control and benchmark heavy comparisons, while prioritizing safeguards for missing data, sensor drift, and conflicting signals. The resulting monitoring to action protocol is implementation grounded and intended for WASH operators, supervisors, and dispatch teams in low-income urban settings.*

Keywords: Predictive Monitoring, Decentralized Water Treatment, Resource Constrained Operations, Operator Decision Support, Practice Oriented Frameworks, Low Income Urban Settlements

Ashok Khedkar (ashok.khedkar@cumminscollege.in), Department of E and TC, Cummins College of Engineering for Women Pune, India.
Shashikant Sahare (shashikant.sahare@cumminscollege.in), Department of E and TC, Cummins College of Engineering for Women Pune, India.
Sandeep Musale (sandeep.musale@cumminscollege.in), Department of E and TC, Cummins College of Engineering for Women Pune, India.
Ajit Bhosale (ajit.bhosale@cumminscollege.in), Department of Mechanical Engineering, Cummins College of Engineering for Women Pune, India.

Introduction

Decentralized drinking-water treatment in resource-constrained urban settlements is increasingly equipped with low-cost sensors and intermittent telemetry, opening a path for AI-enabled environmental monitoring beyond periodic spot checks (Olawade et al., 2024). However, many AI and Internet-of-Things water-quality prediction accounts still assume stable power, dense instrumentation, and reliable communications, conditions that often do not hold in decentralized operations (Alprol et al., 2024).

In practical settings, operators must decide when to verify a suspected excursion, adjust dosing, clean filters, or dispatch maintenance across multiple sites while coping with missing data, sensor drift, and delayed uploads. Under these common failure conditions, reactive threshold alarms and manual schedules can miss short-horizon degradation or create alert fatigue when data quality is poor. The central gap is therefore not prediction alone, but a monitoring-to-action design that turns uncertain signals into capacity-aware escalation rules and feasible daily and weekly routines (Alprol et al., 2024).

Health protection depends on limiting the duration of unsafe water exposure, making time-to-warning and persistence-sensitive alarms operationally meaningful endpoints (Zhou et al., 2025). Building on this operational framing, this paper contributes a practice-oriented blueprint for human-in-the-loop predictive monitoring and decision support that links minimum viable sensing and logging to data-quality safeguards, interpretable early warning outputs, and documented operator actions and feedback. By grounding these design choices in applied AI monitoring concepts while keeping deployment constraints explicit, the framework aims to support the design of improved drinking-water safety and service reliability without requiring autonomous control or curated benchmark data (Olawade et al., 2024; Zhou et al., 2025).

Definitions and scope

This study applies predictive monitoring to convert intermittent sensor and log data from decentralized drinking water treatment sites into early warnings that operators can use in practical settings. The predictive monitoring component is intentionally limited to short-horizon forecasts, anomaly scores, and soft sensor estimates that flag plausible risk before manual checks or lab results arrive.

Operational optimization is defined as selecting feasible actions under staffing, power, and connectivity constraints, rather than solving an abstract control problem. Typical levers include dosing and set point checks, alarm threshold tuning with persistence rules to reduce alert fatigue, maintenance scheduling, and dispatch prioritization across sites. Recommendations remain human-in-the-loop decision support and exclude autonomous closed loop control.

Operational water quality proxies are interpreted through basic aquatic chemistry because field sensors measure indicators rather than pathogens (chemistry, 1996). Turbidity, pH, conductivity, temperature, and disinfectant residual are treated as signals whose interpretation depends on equilibrium assumptions, including acid base balance, carbonate buffering, and chlorine speciation shifts with pH and temperature (chemistry, 1996). These assumptions are used as practical checks for calibration and spot test verification, not as lab grade calculations.

Priority operational use cases

This study prioritizes failure modes where early warning enables action: disinfectant residual decay, turbidity breakthrough from filter fouling, abrupt pH or conductivity shifts indicating source-water change or emerging contaminants, and power or telemetry loss that hides excursions. For real-world use, alerts trigger spot tests, dosing checks, maintenance dispatch, and escalation to reduce persistent contaminant risks (Samal et al., 2022).

Use-case selection logic and requirements

Use cases are selected to improve safety and reliability under sparse sensing, intermittent telemetry, and limited staffing in practical settings. Priority is given to problems with at least one low-cost signal or log proxy, a clear decision window, and a feasible response such as verification, dosing checks, maintenance, or dispatch. Because operational risk is often concentrated in a small fraction of sites, the initial portfolio targets those high leverage locations and failure modes rather than spreading attention uniformly (Moore et al., 2025).

Capacity aware triage translates skewed emitter intuition into alert rules that prioritize persistent high-risk signals and defer low impact checks (Moore et al., 2025).

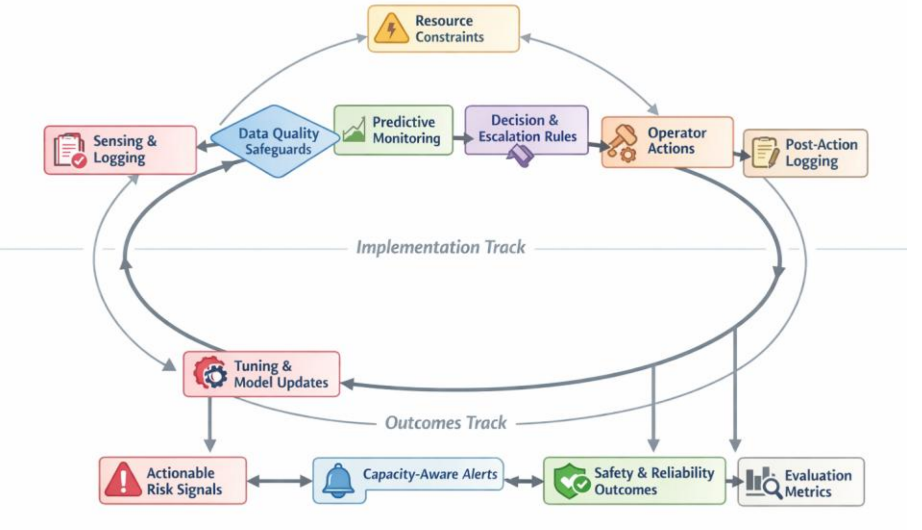

Figure 1. Use case logic model

These use-case cards translate constrained monitoring into early warnings and operator routines for decentralized drinking-water treatment, for real-world use. Fig. (1) depicts the shared sequence from signals to prediction, then to a decision hook, and finally to an action with logged follow-up.

When water-quality proxies such as turbidity, pH, conductivity, or temperature show excursions, a short-horizon risk flag or anomaly score triggers a spot test or lab sample, a review of recent dosing and set-points, and documentation of the response. Where data exist for disinfectant residual risk, a soft-sensor estimate or risk flag prompts a dosing check or adjustment and manual residual verification when available. Tab. (1) summarizes the signal-to-action mapping.

For telemetry dropout, missingness detection and uptime or power indicators support safe behaviour during gaps and schedule manual checks only when persistence rules and staffing capacity justify escalation. For suspected sensor drift, disagreement signals prompt calibration and comparison against spot tests, with outcomes logged to guide threshold tuning and model update triggers.

Table 1. Use case signals to actions matrix

Use Case	***Early Warning Signal***	***Operator Action***
Water quality proxy excursion	Short-horizon risk flag or anomaly score on turbidity, pH, conductivity, temperature	Manual verification (spot test or lab sample), check recent dosing and set-points, log findings and action taken
Disinfectant residual risk (where available)	Soft-sensor estimate or risk flag for disinfectant residual	Dosing check and adjustment per local practice, verify with manual residual test if available, record dosing event and outcome
Telemetry dropout or missing data episode	Missingness detection, telemetry uptime indicator drops, power interruption indicator	Enter safe behaviour for data gaps, schedule a manual site check based on persistence and staffing capacity, log outage window and any observed issues
Sensor drift suspected	Model-operator disagreement or drift signal separating sensor	Field calibration or sensor upkeep task, compare against spot

	error from process change	test, document drift and calibration action for threshold tuning and model update triggers

Monitoring-to-action framework

The framework links constrained monitoring streams and sparse operator logs to actionable early warnings that inform operator decisions in decentralized treatment. Fig. (2) summarizes a staged pipeline: sensing and logging, data quality and gap handling for drift and outages, predictive monitoring, capacity-aware escalation, field actions (checks, dosing, maintenance, dispatch), and feedback logging for threshold tuning.

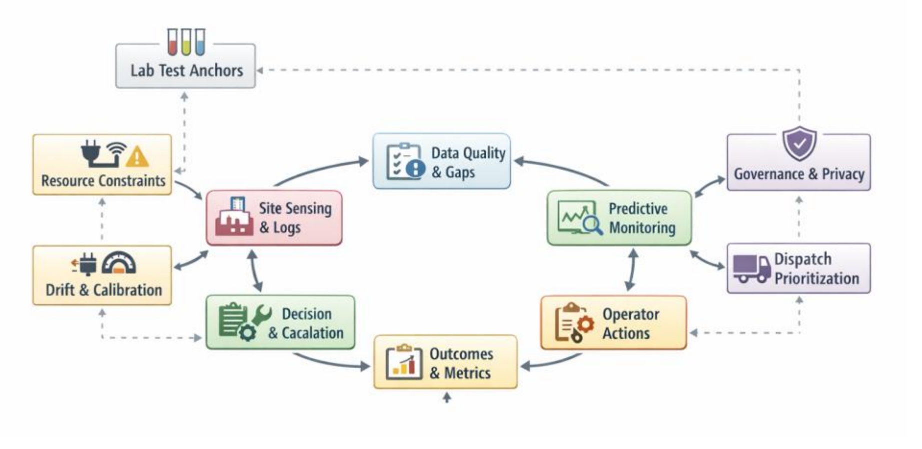

Figure 2. Monitoring-to-action system architecture

Pipeline blocks and information handoffs

The pipeline starts with sensing and logging. Low-cost water-quality proxies and operational events are time-stamped, linked to a site identifier, and buffered for intermittent synchronization. In practical settings, ingestion and storage are selected to match available capability, favouring reliable capture of raw telemetry and operator logs over complex infrastructure when power and connectivity are unstable (Gomes et al., 2020).

Next, a data-check block turns raw streams into usable inputs by identifying missingness episodes, clear sensor faults, and plausible drift signals, and by attaching data-reliability context (for example, telemetry uptime flags) to each record. This step is supported by basic visualization and quality-control dashboards built with open, maintainable tools so operators can review trends and investigate alerts without specialist software (Wang, 2012; Wessel et al., 2013).

Predictive monitoring then uses the quality-tagged series and derived features to produce interpretable outputs, including short-horizon forecasts, anomaly scores, and soft-sensor risk flags. The decision layer maps these outputs to **capacity-aware escalation rules** (including persistence checks to reduce alarm fatigue), triggering feasible actions such as verification, dosing checks, or maintenance dispatch. Actions and outcomes are logged to close the feedback loop and keep transforms reproducible across sites (Gomes et al., 2020; Wessel et al., 2013).

Safeguards and safe-mode behaviour

Operation is conservative when data are missing, sensors drift, or model outputs disagree with field checks. Inputs are labelled reliable, degraded, or unavailable, and predictions that depend on unavailable or implausible values are suppressed rather than interpolated. Persistent drift or repeated gaps move the site into safe mode, which uses validated thresholds and extra logging. State transitions are summarized in Fig. (3).

In degraded or safe mode, alerts must persist to limit alarm fatigue, and high-risk flags require manual verification before dosing changes or dispatch. Disagreements between alerts and observations are logged for review, and monitoring returns to normal only after data quality recovers and the verification outcome is documented.

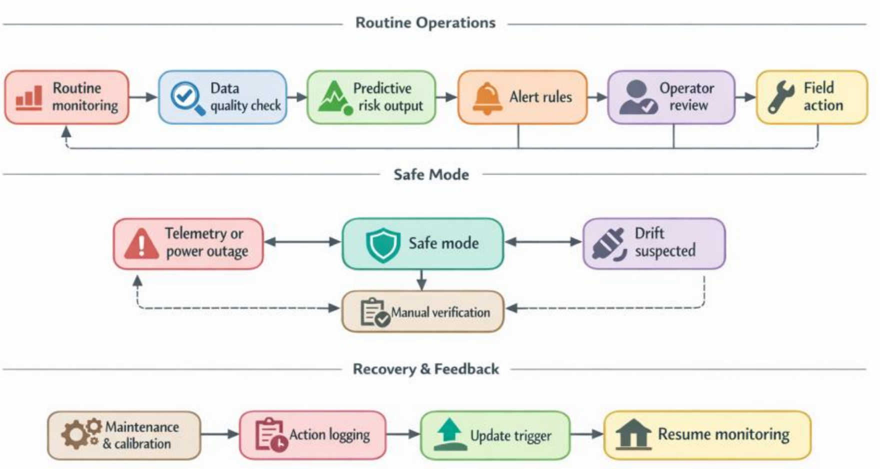

Figure 3. Safe-mode state machine

Design guidance under constraints

In practice, minimum viable deployments pair low-cost turbidity, pH, and flow proxies with dosing and logs, add lab anchors, and favour persistence alerts plus safe-mode rules for outages and drift to limit alarm fatigue.

Minimum viable sensing, logging, and data checks

A feasible monitoring foundation for decentralized treatment sites in practical settings can combine low-cost water sensors with sparse operator logs, provided that each reading or recorded event is time-stamped and linked to a site identifier. Because these systems are typically maintained with limited resources, sensor selection should balance purchase cost and measurement usefulness against the ongoing burden of calibration and drift tracking (Adjovu et al., 2023; Aigner et al., 2026). The minimum capture and linkage requirements are summarized in Tab. (2).

Data checks should treat failure regimes as primary evidence by explicitly flagging missing episodes due to power or telemetry dropouts, and by avoiding silent interpolation when synchronization fails. Telemetry availability can be tracked as the fraction of successful synchronizations out of all expected synchronizations, Eq. (1). The corresponding indicator is

$$Availability = \frac{N_{success}}{N_{total}} \tag{1}$$

Occasional lab tests and, when suitable, remote sensing can supplement sparse in situ monitoring (Adjovu et al., 2023).

Table 2. Minimum viable monitoring checklist

Item	*What To Capture*	*Minimum Check*
Low-cost water sensors	Turbidity, pH, conductivity, temperature; disinfectant residual where available	Time-stamped readings linked to site identifier and timestamp
Operator and maintenance logs	Dosing events, filter cleaning or backwash events, operator notes	Event records time-stamped and linked to site identifier and timestamp
Power and connectivity indicators	Power interruption indicator, telemetry uptime indicator	Flag and count missing data episodes; use safe behaviour when telemetry drops (no silent interpolation)
Occasional lab tests	Periodic lab results linked to time and site	Use as anchor points for water-quality risk discussion; link to site

	identifier and sample time or date

Predictive monitoring approaches that fit sparse messy data

Predictive monitoring in decentralized treatment is most reliable when models remain simple enough to withstand gaps, noise, and sensor drift. For early warning, the study uses lightweight anomaly detection and event classification based on a few water-quality proxies and operational context signals, supporting maintainability and fast review for real-world use (Liu et al., 2025).

Statistical monitoring provides a transparent baseline when failures are rare and conditions vary (Lambert & Gilbert, 1988). A common starting point is the z-score that standardizes each value against a recent baseline for rule-based alerts, shown in Eq. (2).

$$z_t = \frac{x_t - \mu_x}{\sigma_x} \tag{2}$$

During missing telemetry, baseline updates and thresholds should pause, and drift checks should be separated from abrupt excursions to limit alarm fatigue (Lambert & Gilbert, 1988; Liu et al., 2025).

Decision layer design and operational optimization levers

The decision layer translates predictive outputs into capacity-aware escalation rules for intermittent telemetry and limited staff, prioritizing dosing checks, manual verification, and dispatch over autonomous control (Jin et al., 2025; Vijayakumar et al., 2024). Thresholds on anomaly scores or soft-sensor risk flags gate escalation by persistence and response capacity. A capacity-aware decision flow is outlined in Fig. (4).

A persistence rule raises an alert only when elevated scores recur, reducing isolated alarms, Eq. (3).

$$Alert_t = I\backslash Big\left(\frac{1}{W}\sum_{i=0}^{W-1} I(s_{t-i} \geq \tau) \geq \rho\backslash Big\right) \tag{3}$$

Tab. (3) links each alert type to feasible actions, including safe behaviour during data gaps, calibration and maintenance when drift is suspected, and dosing or set-point checks when treatment risk persists. Escalation targets defined recipients with authority to act.

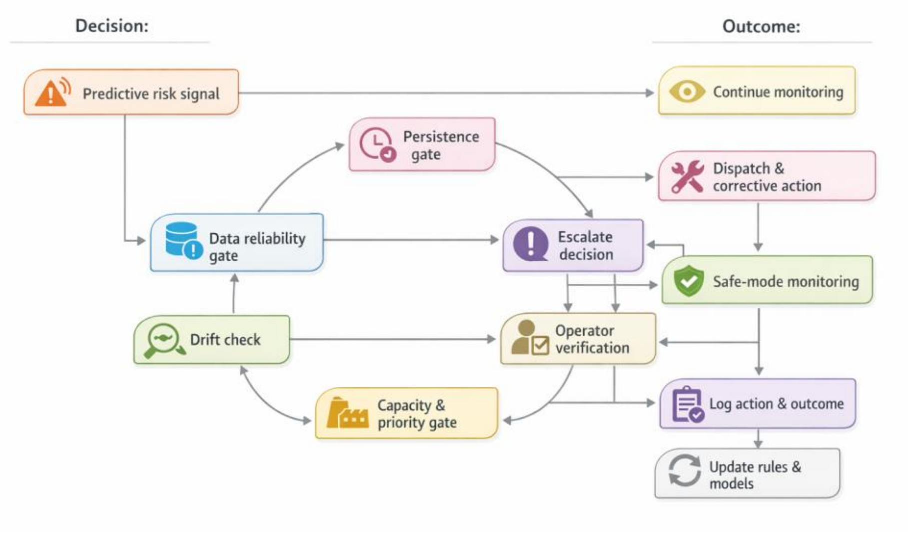

Figure 4. Capacity-aware alert decision flow

Table 3. Alert escalation and actions table

Alert Type	*Threshold Basis*	*Persistence Check*	*Action And Escalation*
Telemetry drops or missing data	Missingness detection, telemetry uptime indicator	Persists across sync cycles or repeated gaps	Enter safe behaviour for gaps (avoid silent interpolation), perform manual check, log the data outage; escalate if the gap blocks decision-making and needs dispatch prioritization
Possible sensor drift	Drift detection to separate process change from sensor error	Sustained disagreement with operator observations or spot tests	Manual verification and calibration routine, schedule sensor upkeep, log drift and

			actions; escalate to maintenance scheduling if drift recurs or affects multiple sites
Water-quality proxy anomaly	Anomaly score or short-horizon risk flag from low-cost sensor streams (turbidity, pH, conductivity, temperature) plus sparse logs	Flag repeats over a defined persistence window to reduce alarm fatigue	On-site manual verification, review recent dosing and maintenance logs, record outcome; escalate to dispatch prioritization if repeated and staffing allows response
Treatment performance risk	Soft-sensor estimate or risk flag using disinfectant residual (where available) and operational context (flow proxy, pump status, dosing logs)	Sustained risk flag after persistence rule to avoid one-off alarms	Dosing and set-point checks, review dosing event log, log actions taken; escalate to defined recipients with authorization to act if risk persists and requires coordinated response

Implementation and governance considerations

Deployment requires clear alert recipients and decision authority, with actions and outcomes logged in a living data handbook for iteration (Lowndes et al., 2017; Spector, 1957).

Operator workflow, roles, and accountability

Effective predictive monitoring for real-world use depends on clear assignment of alert ownership and on who is authorized to act. Alert recipients should include the on-shift site operator for checks, a supervisor for escalation, and a maintenance contact for dispatch. Decision rights should be limited to feasible steps such as verification, dosing checks, and maintenance scheduling, with handoff when capacity is exceeded. Recipient roles and logging duties are summarized in Fig. (5).

Documentation should link each alert to the verification step, the action taken, and the outcome at the next routine check. Routine tasks should include sensor sanity checks, review after each synchronization window, and calibration when drift is suspected or spot tests disagree.

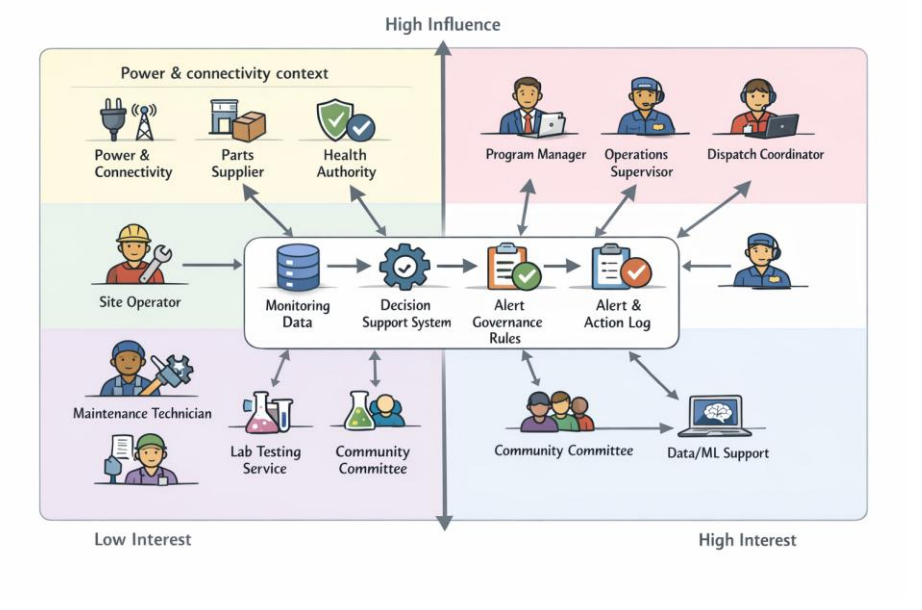

Figure 5. Operator roles and accountability map

Sensor upkeep, training burden, and sustainability

For real-world use, reliable decentralized treatment depends on routine sensor calibration, scheduled consumable replacement, and simple failure checks that operators can perform despite turnover. Monitoring and spot-testing schedules should be low-burden, supported by refresher training and clear criteria for pausing automated inferences, and should incorporate sustainability criteria in routine design (Locatelli et al., 2023). Structured greenness scoring helps compare analytic choices and logging frequency against power, material, and waste costs (Pena-Pereira et al., 2020).

Illustrative examples mapped to the framework

Illustrative mini-cases, adapted from wastewater technology surveys, link sensor logs to risk flags, capacity-aware dispatch decisions, and documented actions rather than benchmarks (Shamshad & Rehman, 2024).

End-to-end example walkthrough

A representative scenario is early warning of rising turbidity at a decentralized drinking-water site. When telemetry is available, turbidity, pH, pump runtime, and dosing logs are synchronized, and checks label missing intervals and likely sensor drift. A short-horizon model then produces a risk flag for an upcoming excursion and an anomaly score summarizing how unusual the recent pattern is.

For real-world use, an alert is sent only when the risk flag persists across windows and the escalation rule indicates that staff can respond. If telemetry is lost, the system switches to safe mode and requests a spot test instead of alarms. The operator verifies turbidity with a grab sample, checks dosing and filter cleaning, schedules maintenance or dispatch if needed, and logs the outcome to adjust thresholds and update models.

Constraint-centered scenario and response

If telemetry from a chlorination unit is unavailable for several hours, the monitoring layer detects the gap and blocks the predictor from producing risk scores. The decision layer enters safe mode: it relies on last confirmed readings, requests a manual residual check, and escalates only when abnormal patterns remain after telemetry resumes. All checks and actions are logged to support threshold tuning and drift review in practical settings.

Discussion: transferability and success criteria

Blueprint transfer is contingent on-site baselines, staffing, and data reliability, and success means earlier warnings, feasible actions, and fewer disruptions.

Limits, risks, and non-claims

This framework supports human-in-the-loop decision support for decentralized drinking-water treatment, but it depends on sparse and messy monitoring in practice. Performance may degrade with missing data, intermittent telemetry, sensor drift, and changing source-water or operating conditions. Limited staffing, turnover, and fragile supply chains for parts and calibration can delay verification and response, increase alarm fatigue or miss events. It does not claim universal model transfer, autonomous control, or guaranteed safety, and recommends cautious escalation, logging, and periodic recalibration for real-world use.

Operational success measures and scaling stance

Operational value in practical settings is indicated by earlier warnings that allow dosing or maintenance before excursions, fewer and shorter proxy exceedances, reduced downtime from faster troubleshooting, and improved dispatch prioritization across sites with manageable alert burden. Scaling should proceed in stages, requiring stable sensing, drift checks, clear escalation authority, and site-specific baselines before transfer.

Conclusion

This study presents a practice-oriented monitoring-to-action blueprint that turns sparse sensor streams and operator logs from decentralized drinking-water treatment into interpretable early warnings and feasible responses under real constraints. The architecture couples minimum viable sensing and time-stamped logging with explicit data-quality gates for missingness, sensor drift, and intermittent synchronization, and it maps forecasts, anomaly scores, or soft-sensor risk flags to persistence-based, capacity-aware escalation and verified field actions. A staged adoption path is recommended: pilot at high-leverage sites to align use cases, recipients, and manual verification, then stabilize by tuning thresholds, calibration routines, and safe-mode behavior during outages, and finally scale through standardized logging, retraining triggers, and dispatch prioritization across sites. Decision support succeeds for real-world use when alerts reliably translate into timely checks, dosing and maintenance decisions, and documented learning.

References

Adjovu, G. E., Stephen, H., James, D. E., & Ahmad, S. (2023). Overview of the application of remote sensing in effective monitoring of water quality parameters. Remote Sensing, 15(7), 1938–1938. https://doi.org/10.3390/rs15071938

Aigner, P., Chen, J., Böhm, F., Chariot, M., Emmenegger, L., Frölich, L., Grange, S. K., Kühbacher, D., Kürzinger, K., Laurent, O., Makowski, M., Rubli, P., Schmitt, A., & Wenzel, A. (2026). ACROPOLIS: Munich urban CO ₂ sensor network. Atmospheric Measurement Techniques, 19(2), 745–773. https://doi.org/10.5194/amt-19-745-2026

Alprol, A. E., Mansour, A. T., Ibrahim, E. M., & Ashour, M. (2024). Artificial intelligence technologies revolutionizing wastewater treatment: Current trends and future prospective. Water, 16(2), 314–314. https://doi.org/10.3390/w16020314

chemistry, A. (1996). Aquatic chemistry: Chemical equilibria and rates in natural waters. Choice Reviews Online, 33(11), 33–6312. https://doi.org/10.5860/choice.33-6312

Gomes, V. C. F., Queiroz, G. R. de, & Ferreira, K. R. (2020). An overview of platforms for big earth observation data management and analysis. Remote Sensing, 12(8), 1253–1253. https://doi.org/10.3390/rs12081253

Jin, J., Liu, M., Chen, B., Wu, X., Yao, L., Wang, Y., Xiong, X., Wei, L., Jiang, L., Tan, Q., Fan, D., Du, Y., Lei, Y., & Yang, N. (2025). Artificial intelligence in chemical dosing for wastewater purification and treatment: Current trends and future perspectives. Separations, 12(9), 237–237. https://doi.org/10.3390/separations12090237

Lambert, D., & Gilbert, R. O. (1988). Statistical methods for environmental pollution monitoring. Journal of the American Statistical Association, 83(404), 1226–1226. https://doi.org/10.2307/2290180

Liu, X., Bannatyne, L., Boldrini, A., Polvani, A., Cirrone, R. G., Galgani, L., Woods, S., Warner, S., & Loiselle, S. (2025). Realtime fluorescence organic matter indicators for early warning of river pollution sources. Ecological Indicators, 179, 114214–114214. https://doi.org/10.1016/j.ecolind.2025.114214

Locatelli, M., Kabir, A., Perrucci, M., Ulusoy, S., Ulusoy, H. İ., & Ali, I. (2023). Green profile tools: Current status and future perspectives. Advances in Sample Preparation, 6, 100068–100068. https://doi.org/10.1016/j.sampre.2023.100068

Lowndes, J. S., Best, B. D., Scarborough, C., Afflerbach, J. C., Frazier, M., O'Hara, C. C., Jiang, N., & Halpern, B. S. (2017). Our path to better science in less time using open data science tools. Nature Ecology & Evolution, 1(6), 160–160. https://doi.org/10.1038/s41559-017-0160

Moore, D. P., Li, N., Song, C., Zhu, J., Yi, H., Tao, L., McSpiritt, J., Sevostianov, V., Wendt, L., Rojas-Robles, N. E., Hopkins, F. M., Ren, Z. J., & Zondlo, M. A. (2025). Comprehensive assessment of the contribution of wastewater treatment to urban greenhouse gas and ammonia emissions. Nature Water, 3(10), 1114–1124. https://doi.org/10.1038/s44221-025-00490-z

Olawade, D. B., Wada, O. Z., Ige, A. O., Egbewole, B. I., Olojo, A. S., & Oladapo, B. I. (2024). Artificial intelligence in environmental monitoring: Advancements, challenges, and future directions. Hygiene and Environmental Health Advances, 12, 100114–100114. https://doi.org/10.1016/j.heha.2024.100114

Pena-Pereira, F., Wojnowski, W., & Tobiszewski, M. (2020). AGREE—analytical GREEnness metric approach and software. Analytical Chemistry, 92(14), 10076–10082. https://doi.org/10.1021/acs.analchem.0c01887

Samal, K., Mahapatra, S., & Ali, Md. H. (2022). Pharmaceutical wastewater as emerging contaminants (EC): Treatment technologies, impact on environment and human health. Energy Nexus, 6, 100076–100076. https://doi.org/10.1016/j.nexus.2022.100076

Shamshad, J., & Rehman, R. U. (2024). Innovative approaches to sustainable wastewater treatment: A comprehensive exploration of conventional and emerging technologies. Environmental Science Advances, 4(2), 189–222. https://doi.org/10.1039/d4va00136b

Spector, W. S. (1957). Handbook of biological data. AIBS Bulletin, 7(1), 41–41. https://doi.org/10.2307/1292070

Vijayakumar, G., Rajkumar, M., Chandar, N. R., Selvakumar, P., & Duraisamy, R. (2024). Environmentally friendly TDS removal from waste water by electrochemical ion exchange batch-type recirculation (EIR) technique. Environmental Science Water Research & Technology, 10(4), 826–835. https://doi.org/10.1039/d3ew00793f

Wang, Y. Q. (2012). MeteoInfo: GIS software for meteorological data visualization and analysis. Meteorological Applications, 21(2), 360–368. https://doi.org/10.1002/met.1345

Wessel, P., Smith, W. H. F., Scharroo, R., Luís, J., & Wobbe, F. (2013). Generic mapping tools: Improved version released. Eos, 94(45), 409–410. https://doi.org/10.1002/2013eo450001

Zhou, L., Ying, X., Sera, F., Vicedo-Cabrera, A. M., Abrutzky, R., Guo, Y., Tong, S., Coêlho, M. de S. Z. S., Saldiva, P. H. N., Lavigne, É., Correa, P. M., Ortega, N. V., Osorio, S., Royé, D., Kyselý, J., Orru, H., Maasikmets, M., Jaakkola, J. J. K., Ryti, N., … Kan, H. (2025). Associations of ambient exposure to benzene, toluene, ethylbenzene, and xylene with daily mortality: A multi country time-series study in 757 global locations. The Lancet Planetary Health, 9(9), 101306–101306. https://doi.org/10.1016/j.lanplh.2025.101306