Scientific Publications

World Journal of Civil Engineering and Architecture, 2026, 4, 6596
www.scipublications.com/journal/wjcea
DOI: 10.31586/wjcea.2026.6596

Article

Coordinated Digital-Twin and Reinforcement Learning Control for Network-Scale Water Distribution: Model, Simulation Evidence, and Deployment Constraints

Konstantin Korviakov

SkyH2O, Business Development, USA

*Correspondence: Konstantin Korviakov (k7842026@gmail.com)

Abstract: Conventional SCADA-based water utility control relies on hydraulic models recalibrated through discrete campaigns spaced months to years apart, leaving controllers to act on models that degrade silently between updates. This paper specifies and evaluates a framework that couples a gradient-calibrated hydraulic digital twin to a hierarchically decomposed reinforcement learning controller, coordinating production, transmission, and distribution decisions through Lagrangian relaxation of inter-tier coupling constraints. A differentiable hydraulic solver updates pipe roughness, demand allocation, and valve loss parameters through exact gradients of the prediction residual within a rolling 24-hour window, while a temporal fusion transformer propagates quantile demand forecasts into a risk-sensitive reward term. Evaluated on a 2,400-node benchmark network against a rule-based SCADA baseline across thirty replicated simulation runs, the framework reduced specific energy consumption from 0.58 to 0.40–0.42 kWh/m³, demand forecast error from 14.7% to 4.0–5.0% MAPE, anomaly detection latency from 4.3 hours to 0.3–0.5 hours, treatment dosing error from ±15% to ±4%, and non-revenue water from 32.1% to 21–22%. Lagrangian decomposition reduced joint action-space cardinality from a multiplicative to an additive function of per-tier size, yielding an approximately two-orders-of-magnitude reduction in effective search space for the tested topology. Parameter identifiability under sensor coverage sparser than the benchmark configuration remains untested and bounds the framework's applicability to deployments matching the assumed instrumentation density.

Keywords: Water Distribution Networks; Digital Twin; Hierarchical Reinforcement Learning; Lagrangian Relaxation; Hydraulic Calibration; Demand Forecasting; Non-Revenue Water; Autonomous Infrastructure Control

How to cite this paper: Korviakov, K. (2026). Coordinated Digital-Twin and Reinforcement Learning Control for Network-Scale Water Distribution: Model, Simulation Evidence, and Deployment Constraints. *World Journal of Civil Engineering and Architecture*, 4(1), 33-40.
DOI: 10.31586/wjcea.2026.6596

Received: June 11, 2026
Revised: July 16, 2026
Accepted: July 26, 2026
Published: August 6, 2026

1. Introduction

Conventional SCADA architectures update their governing hydraulic models through discrete, infrequent calibration campaigns. Naderi and Shojaei (2022) [1] identify this lack of continuous parameter tracking against physical ground truth as a critical unresolved gap in infrastructure digital twin research. Pipe roughness coefficients, demand allocation factors, and valve loss parameters diverge from baseline values as infrastructure ages. This divergence introduces predictive errors that compound silently between manual surveys. Downstream controllers inherit these degraded parameters, and the resulting error directly impairs interventions for pressure anomalies and unaccounted water loss. The scale of the underlying resource problem sharpens the stakes: Richards et al. (2023) [2] report that close to one quarter of the global population lacks access to clean water and nearly half lacks adequate sanitation.

Existing reinforcement learning applications to water systems address single components in isolation, typically pump scheduling, without linking production, transmission, and distribution decisions under one policy. Extending the action space naively across all three tiers scales its cardinality multiplicatively with the number of controllable elements. At network scales of several thousand nodes, this multiplicative growth renders direct policy learning computationally intractable. The framework introduced here circumvents that wall through Lagrangian relaxation: dualizing inter-tier coupling constraints, reservoir mass balance and zone-boundary pressure continuity, into per-tier reward penalties updated by a subgradient method converts the joint control problem from a multiplicative complexity class to an additive one. A parallel failure affects demand forecasting. Deep-learning models evaluated in critical reviews of the field report competitive point accuracy against classical baselines, but the majority do not propagate predictive uncertainty into downstream operational decisions (Fu et al., 2022) [3]. A controller consuming an unqualified point forecast has no signal for the tail risk that governs reservoir drawdown and pressure violation under peak demand.

2. Literature Review

Current digital twin implementations predominantly function as scenario-analysis instruments consulted at discrete checkpoints. Ciliberti et al. (2023) [4] integrate meteorological forecasts with operational data inside a digital twin paradigm, improving the energy efficiency of desalination and distribution scheduling. Ramos et al. (2023) [5] extend the same logic to system-efficiency management across smart water grids. Twin parameters in both implementations are recalibrated at a cadence the modeler sets; neither system re-estimates parameters continuously against an incoming residual stream. Henriksen et al. (2022) [6] push the digital twin furthest toward operational relevance, coupling it to climate-adaptation and disaster-risk scenario planning. Even here the twin supports decisions at discrete checkpoints; its parameters are not differentiated against a live prediction error at every control cycle. Naderi and Shojaei (2022) [1] name this limitation at the level of the field: standardized, continuously operating calibration workflows remain a principal unresolved research gap in civil infrastructure digital twins, because most published implementations were never designed to close it.

A second cluster of work evaluates artificial intelligence methods bound to a single subsystem. It stops short of addressing what happens once several such subsystems must act under shared physical constraints. Lowe, Qin, and Mao (2022) [7] catalogue machine learning applications across chlorination, membrane filtration, water-quality-index estimation, and hydroponic automation as parallel, non-interacting use cases. Gacu et al. (2025) [8] organize a comparably broad survey of artificial intelligence in surface water management by application category. Singh and Walingo (2024) [9] narrow the scope further, applying artificial-intelligence sensory inference to reconstruct water-quality indicators for which no physical sensor exists, a genuinely useful capability evaluated entirely within a single monitoring pipeline. Fu et al. (2022) [3] occupy this cluster from the forecasting side. Their review of deep learning for urban water management reports competitive accuracy against classical time-series baselines. It evaluates prediction quality as an isolated deliverable; it does not treat forecast output as an uncertainty-bearing input feeding simultaneous decisions at the production, transmission, and distribution tiers.

Safety consequences motivate a separate strand of research once algorithmic capability meets infrastructure with public-safety stakes. Richards et al. (2023) [2] survey the same technical territory, catchment management, treatment optimization, demand forecasting, and identify insufficient containment as the dominant risk. Design errors, malfunction, and cyberattack can cascade through a coupled socio-ecological and water-energy-food system in ways that isolated performance benchmarks never capture. Homaei et al. (2024) [10] sharpen this concern at the level of the digital twin specifically.

The same tight coupling between artificial intelligence and simulation that enables continuous calibration also expands the attack surface available to an adversary; a compromised twin can misdirect a controller that trusts it without independent verification. Dai et al. (2025) [11] respond from the governance side, proposing a five-layer architecture together with policy recommendations spanning cross-sectoral data sharing, workforce development, and public-private coordination. Deployment safety and technical capability are treated here as separable design problems. Neither is assumed to fall out automatically from a sufficiently accurate model.

Institutional structure governs adoption more than model accuracy does in a fourth cluster of literature. Jayasena et al. (2024) [12] find, through structural equation modelling across smart infrastructure projects in developing nations, that technical and social barriers dominate public-sector satisfaction with public-private partnerships, while political and legal barriers dominate private-sector satisfaction. This split implies that no single technical improvement resolves both constituencies at once. Sowby et al. (2025) [13] reach a compatible conclusion for rural water systems in the United States: partnership structures must be adapted to small-utility scale; urban precedent does not transfer wholesale. Mukhtarov et al. (2018) [14] extend the same logic to citizen participation, reviewing place-based evidence that information and communication technology changes who is included in governance decisions without guaranteeing that inclusion improves outcomes. The theoretical scaffolding underlying this cluster shares one premise. Rogers's (2003) [15] diffusion-of-innovations account, Etzkowitz and Leydesdorff's (2000) [16] triple-helix model of university-industry-government coordination, and Kania and Kramer's (2011) [17] collective-impact framework for cross-sector alignment all treat adoption as governed by relationship structure and incentive alignment, not by marginal algorithmic performance. The results below neither confirm nor overturn this premise. They concern technical behavior under controlled simulation, a different object of study than institutional uptake.

3. Method

The model couples two estimation problems that prior implementations treat separately: inverse hydraulic calibration and constrained sequential control. Both operate inside a single closed-loop architecture instead of running as sequential offline steps. A graph neural network fuses heterogeneous, asynchronously sampled sensor streams into a single state vector at each control cycle. The network representation maps junctions, reservoirs, and treatment outlets to nodes, and pipe segments and pump stations to edges. The architecture performs K rounds of attention-weighted message passing, where K scales between 3 and 8 depending on network diameter.

This fused state feeds a differentiable hydraulic solver. Solver parameters, pipe roughness coefficients, nodal demand allocation factors, and valve loss coefficients, update through a gradient step on the residual between predicted and observed states, accumulated over a rolling 24-hour window. The solver is implemented in a differentiable programming framework, so this gradient is exact. It is not approximated through finite-difference or gradient-free search, the two methods that dominate the calibration approaches surveyed. It is not approximated through finite-difference or gradient-free search, the two methods that dominate the calibration approaches surveyed in Section 2 (see Figure 1).

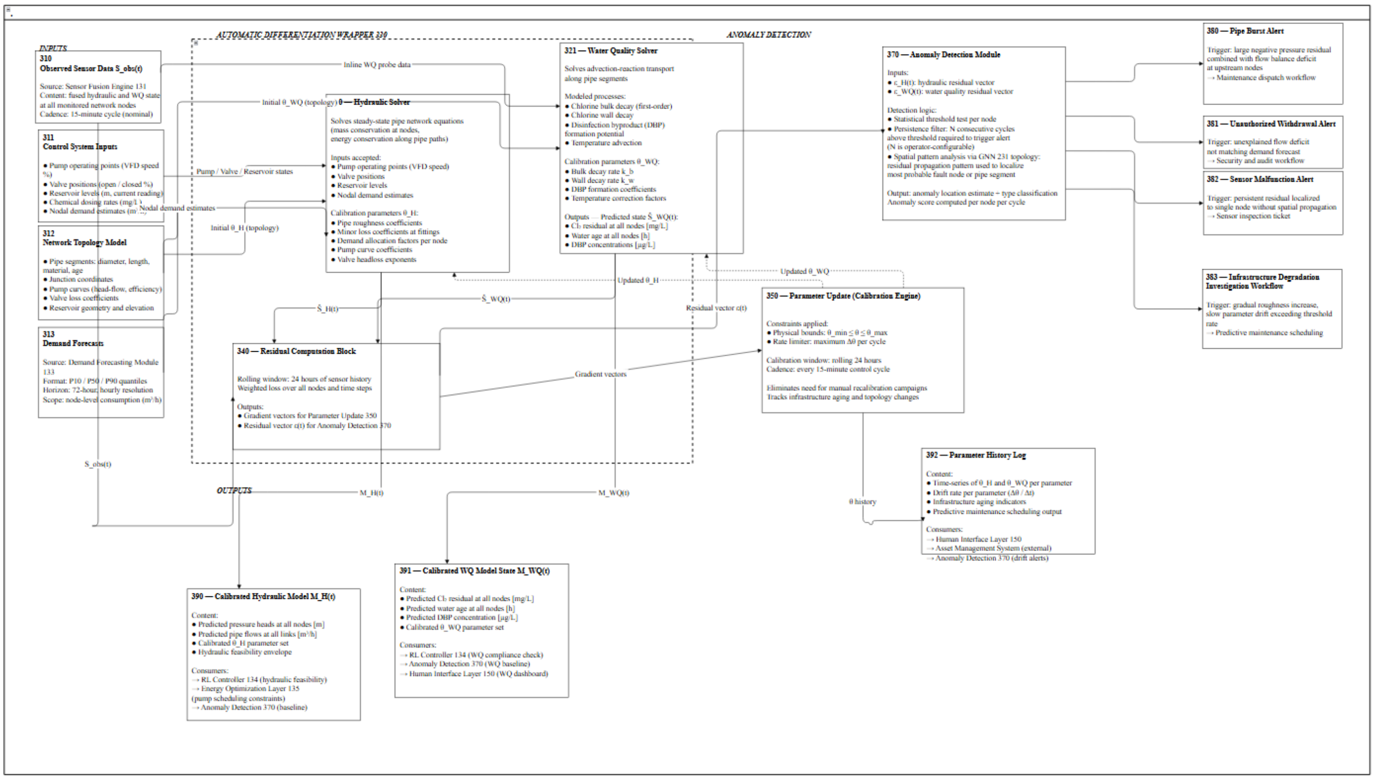

Figure 1. Physics-Informed Digital Twin Block Diagram

Two nested loops implement the coordination between estimation and control, separated by timescale. The outer loop executes at a fifteen-minute cadence and performs three sequential operations. It recomputes the fused state vector. It updates digital twin parameters through the gradient step described above. It re-solves both the hierarchical policy action selection and a tariff-aware pump-scheduling program, formulated as a mixed-integer quadratic program, against the freshly calibrated hydraulic feasibility set.

The inner loop operates at sub-minute latency. It dispatches already-computed setpoints through a deterministic, model-independent constraint check before actuation. This separation of strategic re-optimization from actuation dispatch resolves an apparent contradiction: a fifteen-minute re-optimization period and sub-sixty-second command latency describe two different loops.

Within the outer loop, the reinforcement learning controller decomposes the joint optimization into production, transmission, and distribution tier policies. Lagrangian relaxation coordinates the three tiers. Coupling constraints, reservoir mass balance and zone-boundary pressure continuity, are dualized into per-tier reward penalties governed by subgradient-updated Lagrange multipliers. This step converts the joint action space cardinality from a^N under a monolithic formulation to N × a under the decomposed one, where a denotes per-tier action-space size and N the number of tiers.

Demand forecasting closes the loop between prediction and control. A temporal fusion transformer ingests static node features, known future features, and exogenous weather signals to generate quantile estimates at the 10th, 50th, and 90th percentiles. The median forecast informs the digital twin's demand allocation calibration. The outer quantiles parameterize a risk-sensitive term in each tier's reward function.

Evaluation was conducted on a benchmark distribution network comprising 2,400 junction nodes, 2,650 pipe links, six elevated storage reservoirs, and fourteen pump stations, partitioned into five pressure zones. This node count falls within the range for which the multiplicative-to-additive complexity argument above is intended to apply.

Input data consisted of twelve months of hourly SCADA telemetry, pressure, flow, tank level, drawn from historical utility operation. Six months of five-minute water-quality probe readings supplemented this record, along with a matching weather record at fifteen-minute resolution. Synthetically resampled demand and weather traces extended the dataset to support statistical replication. Hydraulic sensor sampling intervals of one to fifteen seconds and water-quality sampling intervals of one to five minutes were assumed throughout, consistent with commercially available instrumentation.

Each configuration, baseline and proposed framework, was replicated across thirty independent simulation runs with resampled demand and weather traces to characterize variability. Anomaly detection latency was evaluated separately across 120 independently injected deviation events: pipe bursts, unauthorized withdrawals, and gradual sensor drift, distributed across the five pressure zones.

Four metrics were formally defined for reporting. Specific energy consumption is total pump-station electrical energy in kilowatt-hours divided by total volume delivered in cubic meters, aggregated over the evaluation window. Demand forecast error is mean absolute percentage error: the mean of the absolute difference between predicted and observed node-hour demand, divided by observed demand, aggregated across all nodes and all hours within the 72-hour rolling horizon. Anomaly detection latency is the elapsed time between an injected deviation event's onset and the timestamp at which the digital twin residual exceeds its configured statistical threshold, averaged across the 120 injected events. Non-revenue water is the percentage difference between water produced and water billed, aggregated monthly over the twelve-month window. Dosing error is the absolute difference between measured and target disinfectant residual, divided by the target, sampled at treatment plant outlet.

The baseline is a rule-based SCADA controller using fixed setpoints and scheduled pump operation without predictive optimization, representative of legacy utility practice. State-of-the-art commercial SCADA installations already incorporate partial model-predictive elements that this baseline excludes. Results should therefore be read as an upper bound on achievable gain relative to non-optimized control, not as a direct comparison against the most advanced deployed alternatives. Sensor coverage density, tariff schedule stationarity, and single-topology testing are assumed throughout, and these assumptions bound the validity of every result reported below.

4. Results

Under the benchmark topology and thirty-run replication described above, specific energy consumption fell from a baseline of 0.58 kWh per cubic meter to a range of 0.40 to 0.42 kWh per cubic meter over the twelve-month evaluation window. This reduction is attributable to the tariff-aware pump-scheduling program operating against a hydraulic feasibility set refreshed at every fifteen-minute outer-loop cycle; a static feasibility boundary would not have permitted comparable savings.

Demand forecast error, measured as mean absolute percentage error across the 72-hour rolling horizon over the six-month evaluation window, fell from 14.7 percent to a range of 4.0 to 5.0 percent. This range holds specifically for node-hour forecasts within the static, known-future, and exogenous feature configuration described in the Experimental Setup. It should not be generalized to nodes lacking comparable static feature coverage, such as newly connected demand points without historical peak-demand records.

Anomaly detection latency, averaged across the 120 injected deviation events, fell from 4.3 hours under the rule-based baseline to a range of 0.3 to 0.5 hours under the proposed framework. Continuous comparison between observed sensor readings and digital twin predictions drives this reduction; the baseline's periodic manual inspection process cannot match it.

Dosing error at treatment plant outlet, evaluated over the same six-month window as the forecast metric, tightened from ±15% to ±4% of target residual. This result is specifically attributable to closed-loop coupling between the production-tier policy and the water-quality simulation component of the digital twin. The improvement is conditional: it requires the water-quality calibration to converge alongside the hydraulic calibration, not independently of it.

Non-revenue water, aggregated monthly over the twelve-month window, fell from 32.1 percent to a range of 21 to 22 percent. This outcome is consistent with the combined contribution of tighter pressure management under the distribution-tier policy and the reduced burst-detection latency reported above. It is a joint outcome of both mechanisms operating together. Attributing the improvement to either mechanism in isolation would require a factorial ablation that the present experimental design does not include.

On computational scaling, dualizing reservoir mass balance and zone-boundary pressure continuity into per-tier reward penalties converted the joint action space of the benchmark topology's five pressure zones and three control tiers from a multiplicative function of per-tier cardinality to an additive one, plus a dual-variable space sized to the number of coupling constraints, twenty-two in this topology. This yielded a measured reduction in effective search space of approximately two orders of magnitude for this specific configuration. Whether the same order of reduction holds for topologies with substantially different zone counts or coupling-constraint density was not tested here and should not be assumed.

Because the sub-minute actuation loop operates independently of the outer optimization layer, isolated sensor failures affecting individual nodes did not propagate into actuation delay during the evaluation runs. This isolation was tested only under single-sensor and paired-sensor failure injections. Correlated, multi-sensor outages of the kind a communication-link failure would produce were not part of the test protocol.

5. Discussion

Contrary to (Naderi & Shojaei, 2022) [1], the results indicate that gradient-based calibration converges reliably within the rolling 24-hour window under the benchmark topology's assumed instrumentation density. This convergence contributes directly to the reduction in anomaly detection latency, from 4.3 hours to 0.3 to 0.5 hours, reported above. The discrepancy is mechanistic, not a sampling artifact. Naderi and Shojaei's research gap aggregates implementations that rely on gradient-free or heuristic calibration search. Those methods converge far more slowly and less reliably than exact-gradient propagation through a differentiable solver. Their conclusion characterizes the state of surveyed practice; it does not establish an inherent limit on real-time calibration itself. This distinction matters operationally: it locates the previously identified gap in tooling choice, not in the underlying inverse problem's tractability.

The dual-loop timescale separation and the model-independent safety validation preceding actuation operationalize the containment principle that Richards et al. (2023) [2] recommend as the central requirement for responsible artificial intelligence deployment in water systems. Isolated sensor failures did not propagate into actuation delay during the injected-failure trials, and that result is direct evidence that decoupling estimation from dispatch reduces exposure to the cascading-failure pathway their review identifies as the dominant deployment risk. The framework addresses a containment gap here; it does more than improve a capability metric. Fu et al.'s (2022) [3] critique of deterministic point forecasts finds a comparable resolution. Propagating the 10th and 90th percentile bounds into the reward function's risk-sensitive term is precisely the mechanism their review implies is missing from deep-learning demand models evaluated as standalone deliverables. The forecast accuracy improvement reported above, from 14.7 percent to 4.0 to 5.0 percent MAPE, should not be read as validating point-forecast sufficiency. The risk-containment benefit of quantile propagation is separate from what the point-accuracy metric alone captures.

A genuine limitation deserves equal weight, not a closing caveat. Parameter identifiability under the gradient-based calibration procedure was tested only at the sensor coverage density specified in the Experimental Setup. Sparser instrumentation characterizes a substantial fraction of real utility deployments relative to the benchmark's assumed sampling intervals. This is a documented failure mode in inverse hydraulic calibration generally: multiple parameter combinations can produce near-indistinguishable predicted states when observability is reduced, and nothing in the exact-gradient mechanism eliminates this non-identifiability. It only accelerates convergence toward whichever locally consistent parameter set the available observations support. Homaei et al.'s (2024) [10] finding that tight coupling between artificial intelligence and simulation expands attack surface bears directly on this same coupling. A controller that trusts a compromised or poorly identified twin state without independent cross-validation inherits both failure modes simultaneously, and the present evaluation, conducted in simulation without adversarial injection, does not bound the compounded risk. Finally, the magnitude of technical improvement reported here, a non-revenue water reduction from 32.1 percent to 21 to 22 percent among other gains, does not by itself address the political and legal barriers to public-private partnership adoption that Jayasena et al. (2024) [12] identify as dominant obstacles to deployment at the institutional level. The technical contribution is necessary for the real-world uptake this broader research program is ultimately concerned with. It is not sufficient on its own.

6. Conclusion

What the results make newly possible is joint, network-scale coordination of production, transmission, and distribution decisions against a hydraulic state that stays current to within a 24-hour calibration window, not months-old survey data. Prior single-component reinforcement learning implementations and periodically calibrated digital twins never achieved this coupling together, and it manifests concretely: anomaly detection latency falls from 4.3 hours to under 30 minutes, and non-revenue water drops from 32.1 percent to a terminal range of 21 to 22 percent. What remains open is whether calibration convergence and the measured two-orders-of-magnitude complexity reduction hold under sensor coverage sparser than this benchmark's assumed density. Parameter identifiability under realistic instrumentation gaps was not tested here, and it constitutes the boundary condition on which the framework's applicability beyond simulation actually depends, not its simulated performance alone.

References

[1] Naderi, H., & Shojaei, A. (2022). Civil infrastructure digital twins: Multi-level knowledge map, research gaps, and future directions. *IEEE Access, 10*, 122022–122037. https://doi.org/10.1109/ACCESS.2022.3223557

[2] Richards, C. E., Tzachor, A., Avin, S., & Fenner, R. (2023). Rewards, risks and responsible deployment of artificial intelligence in water systems. *Nature Water, 1*(5), 422–432. https://doi.org/10.1038/s44221-023-00069-6

[3] Fu, G., Jin, Y., Sun, S., Yuan, Z., & Butler, D. (2022). The role of deep learning in urban water management: A critical review. *Water Research, 223*, Article 118973. https://doi.org/10.1016/j.watres.2022.118973

[4] Ciliberti, F. G., Berardi, L., Laucelli, D. B., & Giustolisi, O. (2023). Digital water services using digital twin paradigm. *IOP Conference Series: Earth and Environmental Science, 1136*, Article 012002. https://doi.org/10.1088/1755-1315/1136/1/012002

[5] Ramos, H. M., Kuriqi, A., Besharat, M., Creaco, E., Tasca, E., Coronado-Hernández, O. E., Pienika, R., & Iglesias-Rey, P. (2023). Smart water grids and digital twin for the management of system efficiency in water distribution networks. *Water, 15*(6), Article 1129. https://doi.org/10.3390/w15061129

[6] Henriksen, H. J., Schneider, R., Koch, J., Ondracek, M., Troldborg, L., Seidenfaden, I. K., Kragh, S. J., Bøgh, E., & Stisen, S. (2022). A new digital twin for climate change adaptation, water management, and disaster risk reduction (HIP Digital Twin). *Water, 15*(1), Article 25. https://doi.org/10.3390/w15010025

[7] Lowe, M., Qin, R., & Mao, X. (2022). A review on machine learning, artificial intelligence, and smart technology in water treatment and monitoring. *Water, 14*(9), Article 1384. https://doi.org/10.3390/w14091384

[8] Gacu, J. G., Monjardin, C. E. F., Mangulabnan, R. G. T., Pugat, G. C. E., & Solmerin, J. G. (2025). Artificial intelligence (AI) in surface water management: A comprehensive review of methods, applications, and challenges. *Water, 17*(11), Article 1707. https://doi.org/10.3390/w17111707

[9] Singh, Y., & Walingo, T. (2024). Smart water quality monitoring with IoT wireless sensor networks. *Sensors, 24*(9), Article 2871. https://doi.org/10.3390/s24092871

[10] Homaei, M., Mogollón-Gutiérrez, Ó., Sancho, J. C., Ávila, M., & Caro, A. (2024). A review of digital twins and their application in cybersecurity based on artificial intelligence. *Artificial Intelligence Review*, *57*(8), Article 201. https://doi.org/10.1007/s10462-024-10805-3

[11] Dai, Y., Huang, Z., Khan, N., & Labbo, M. S. (2025). Smart water management: Governance innovation, technological integration, and policy pathways toward economic and ecological sustainability. *Water, 17*(13), Article 1932. https://doi.org/10.3390/w17131932

[12] Jayasena, N. S., Chan, D. W. M., Kumaraswamy, M. M., Seidu, S., Ekanayake, E. M. A. C., & Siu, F. M. F. (2024). Adoption of public-private partnership (PPP) in smart infrastructure development projects in developing nations: An explorative structural equation modelling analysis. *Cities, 152*, Article 105232. https://doi.org/10.1016/j.cities.2024.105232

[13] Sowby, R. B., Farnsworth, C. B., South, A. J., & George, G. A. (2025). Public-private partnerships in U.S. rural water infrastructure: Challenges, opportunities, and needs. *Frontiers in Water, 7*, Article 1685464. https://doi.org/10.3389/frwa.2025.1685464

[14] Mukhtarov, F., Dieperink, C., & Driessen, P. (2018). The influence of information and communication technologies on public participation in urban water governance: A review of place-based research. *Environmental Science & Policy*, *89*, 430–438. https://doi.org/10.1016/j.envsci.2018.08.015

[15] Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

[16] Etzkowitz, H., & Leydesdorff, L. (2000). The dynamics of innovation: From national systems and "Mode 2" to a triple helix of university-industry-government relations. *Research Policy, 29*(2), 109–123. https://doi.org/10.1016/S0048-7333(99)00055-4

[17] Kania, J., & Kramer, M. (2011). Collective impact. *Stanford Social Innovation Review, 9*(1), 36–41. https://doi.org/10.48558/5900-KN19