American Journal of Innovation in Science and Engineering (AJISE)

ISSN: 2158–7205 (ONLINE)

VOLUME 5 ISSUE 3 (2026)



PUBLISHED BY
E-PALLI PUBLISHERS, DELAWARE, USA

American Journal of Innovation in Science and Engineering (AJISE)

Volume 5 Issue 3, Year 2026
ISSN: 2158-7205 (Online)
DOI: https://doi.org/10.54536/ajise.v5i3.8048
https://journals.e-palli.com/home/index.php/ajise

Autonomous Water Infrastructure Management for High-Demand Facilities: A Hierarchical AI Framework Integrating Physics-Informed Digital Twins and Risk-Sensitive Reinforcement Learning

Konstantin Korvyakov[1*]

Article Information

Received: April 07, 2026

Accepted: June 20, 2026

Published: November 20, 2026

Keywords

Atmospheric Water Generation, Autonomous Water Infrastructure, Data Center Water Footprint, Demand Forecasting, Graph Neural Network, Hierarchical Reinforcement Learning, Hydraulic Network Modeling, Non-Revenue Water, Pump Scheduling Optimization

ABSTRACT

Municipal water distribution systems operating under conventional SCADA-based supervisory control lose between 28% and 35% of treated production volume before delivery to registered consumer meters, a failure rate attributable to static hydraulic model parameterization and decoupled subsystem architectures rather than to any fundamental physical constraint. This paper presents a unified framework coupling a continuously self-calibrating physics-informed digital twin, a hierarchical reinforcement learning controller decomposed via Lagrangian relaxation across three operational tiers, and a temporal fusion transformer whose probabilistic quantile outputs propagate directly into the controller's risk-sensitive reward function. Evaluated on a 1,847-node, 312 km distribution network serving four hyperscale data center facilities over a 12-month period, the framework achieves a 33.3% reduction in non-revenue water loss, a 29.3% reduction in specific energy consumption, and a 90.7% reduction in pipe burst detection latency relative to the SCADA baseline operating on identical sensor and tariff infrastructure.

INTRODUCTION

Municipal and industrial water distribution systems operating under conventional SCADA-based supervisory control lose between 28% and 35% of treated water volume before it reaches registered consumer meters (Siddik *et al.*, 2021), a structural failure rate that has not meaningfully declined over two decades despite incremental advances in sensor density and alarm management. Static hydraulic models calibrated on 6–12 month cycles degrade within weeks of any topology change or seasonal demand shift, leaving operators without reliable pressure-state information at the sub-hour timescale at which pipe bursts, unauthorized extractions, and pressure exceedances unfold. Hyperscale data centers consuming upward of 1.7 billion liters per day across the United States (Mytton, 2021) impose demand signatures that are poorly represented in historical consumption records, and demand forecast errors of 12–18% MAPE at the node level propagate directly into over-pressurization events at the distribution tier, accelerating leakage through the empirical leakage-pressure power law and producing under-storage conditions that compromise pressure compliance during peak computational load.

The failure is architectural rather than algorithmic. Physics-based hydraulic modeling, as deployed in EPANET-class solvers (Shehabi *et al.*, 2016), treats the simulation model as a static artifact whose non-differentiable solver prevents gradient-based parameter tracking against streaming sensor data; the model supplies feasibility checks to an external optimizer but cannot update its own parameters at control-cycle resolution. Reinforcement learning applied to pump scheduling (Raveesh *et al.*, 2021) has addressed individual components without coupling across production, transmission, and distribution tiers, and a monolithic policy network mapping the full state of a 1,000-node network to its full action space faces a joint dimensionality that renders policy learning computationally intractable at realistic infrastructure scale. Probabilistic demand forecasting surveyed by Shafeian *et al.* (2022) generates uncertainty estimates that are consumed by rule-based schedulers rather than by learning controllers, which forecloses risk-sensitive policy optimization entirely. Three individually capable subsystems remain decoupled, and the decoupling is precisely where the 28–35% NRW ceiling originates.

The framework presented here integrates a continuously self-calibrating physics-informed digital twin, a hierarchical reinforcement learning controller decomposed via Lagrangian relaxation across three operational tiers, and a temporal fusion transformer whose probabilistic quantile outputs are propagated directly into the controller's risk-sensitive reward function.

LITERATURE REVIEW

Three distinct groups of prior work converge on the problem addressed here, each sharing not a topical similarity but a structural incompleteness that prevents the subsystem it addresses from being embedded in a

[1] Business Development, United States
[*] Corresponding author's e-mail: k7842026@gmail.com

continuous closed-loop control architecture. Identifying this incompleteness precisely is necessary because the framework proposed here resolves it through architectural integration rather than algorithmic improvement, and the distinction has direct implications for how the results should be interpreted relative to each group's claims.

Physics-based hydraulic modeling and data center energy accounting share the assumption that calibration is an offline preprocessing step rather than an operational function. Shehabi *et al.* (2016) and Masanet *et al.* (2020) quantify the energy consumption trajectory of data center infrastructure and establish that demand-side water usage correlates structurally with computational load cycles at sub-15-minute timescales, but neither work provides a mechanism for updating the hydraulic model that mediates this correlation in real time. Peeters *et al.* (2020) and Salehi *et al.* (2020), examining thermodynamic constraints on atmospheric water generation, demonstrate that the water-energy nexus in high-demand settings cannot be optimized without continuous joint state estimation of both energy input and hydraulic output, and then proceed to rely on fixed thermodynamic parameter sets estimated from steady-state laboratory conditions, violating their own stated requirement. The tension between what these works identify as necessary and what their own frameworks provide is not incidental; it is the gap that a differentiable hydraulic solver directly addresses, at the cost of a computational overhead that must be benchmarked against the 0.14 bar RMSE that static calibration produces by month 12 of a deployment cycle.

Forecasting methods in the water and energy domain assume that uncertainty estimates terminate at a human operator or a deterministic rule set. Li *et al.* (2023) project AI-driven water withdrawal of 4.2–6.6 billion cubic meters by 2027, a range wide enough to reveal that node-level demand models with exogenous feature integration at the facility scale do not yet exist in operational practice. Tu and Hwang (2020) and Raveesh *et al.* (2021) survey hybrid generation and forecasting architectures and document consistent performance advantages for systems that combine thermodynamic and data-driven components, yet both works evaluate standalone generation systems in which demand is exogenous and the controller cannot act on forecast uncertainty, a context so structurally different from network distribution control that their conclusion about the primacy of thermodynamic cycle integration over forecasting quality does not transfer. Shafeian *et al.* (2022) catalog round-the-clock atmospheric water generation systems and document that hybrid cooling architectures dominate because single-mode condensation cannot sustain production across the full operational humidity range, a finding relevant here as a design constraint on decentralized supply rather than as a network control principle. Every system in this literature group applies conservative static pressure margins as a substitute for risk representation, sacrificing energy efficiency without providing quantified protection against worst-case demand scenarios.

MATERIALS AND METHODS

The Sensor Fusion Engine ingests heterogeneous telemetry from pressure transducers at 1–5 second intervals, electromagnetic flow meters at 1–15 second intervals, residual chlorine analyzers at 1–5 minute intervals, and advanced metering infrastructure at 15–60 minute intervals, aligning temporally disparate streams onto a common axis before passing them to a graph neural network that operates on the topological representation of the pipe network. Nodes correspond to junctions, reservoirs, and treatment outlets; edges correspond to pipe segments and pump stations; $K = 5$ attention-weighted message-passing iterations produce the fused state vector $S(t)$ at each 15-minute cycle. Sensor fault detection is embedded within this layer, replacing faulty readings with graph-imputed estimates and generating maintenance dispatch tickets for affected nodes.

Implementing the steady-state hydraulic solver in a differentiable programming framework allows exact gradients of the pressure prediction error at all instrumented nodes to be computed with respect to the full 2,340-parameter calibration vector, comprising pipe roughness coefficients, demand multipliers, minor loss factors, and pump curve coefficients, on a rolling 24-hour window at each control cycle. This eliminates scheduled recalibration campaigns entirely. The digital twin additionally solves advection-reaction equations for chlorine decay and disinfection byproduct formation, with water quality parameters subject to the same gradient-based update process as the hydraulic parameters, so that a single calibration mechanism maintains fidelity across both physical domains simultaneously.

Decomposing the global optimization problem into production, transmission, and distribution tier policy networks, each an actor-critic architecture with separate policy and value function networks, reduces the joint state-action dimensionality by approximately two orders of magnitude relative to a monolithic formulation for a 1,847-node network. Coupling constraints, specifically mass balance at reservoir nodes and pressure continuity at zone boundaries, are dualized into Lagrangian penalty terms, and multiplier values are updated via a subgradient method at each cycle, driving the decentralized solution toward global feasibility without requiring a centralized policy. The temporal fusion transformer generates quantile estimates at the 10th, 25th, 50th, 75th, and 90th percentiles for each demand node and forecast hour over a 72-hour horizon, with quantile outputs propagated directly into the reinforcement learning controller's reward function rather than consumed as point predictions. Pump scheduling is formulated as a mixed-integer quadratic program on a receding 24-hour horizon, incorporating hydraulic feasibility constraints from the digital twin and time-varying energy tariff signals updated every 15 minutes.

The evaluation network comprised 1,847 demand nodes, 312 km of transmission and distribution mains, three pressure zones, four pump stations with combined

installed capacity of 2.8 MW, and two balancing reservoirs totaling 18,400 m³ storage volume. Four hyperscale data center facilities appeared as demand nodes, each drawing between 180 and 420 m³/hour depending on computational load, with thermal demand telemetry fed as an exogenous feature stream into the forecasting module at 15-minute intervals. Hydraulic and energy metrics were evaluated over 12 months; demand forecast and chemical dosing metrics over 6 months, reflecting the data availability window for those subsystems. Sensor infrastructure comprised 94 pressure transducers, 61 electromagnetic flow meters, 23 residual chlorine analyzers, 12 turbidity sensors, 4 weather stations, and 1,847 AMI endpoints transmitting at 15-minute intervals via LoRaWAN. Exogenous feature streams included 72-hour numerical weather prediction outputs updated at 6-hour intervals, day-of-week and public holiday calendars, and real-time population density estimates from mobile network signaling data updated hourly.

The reinforcement learning controller's training dataset comprised 36 months of historical operational records covering 8,760 hours of pump dispatch logs, sensor readings, and energy tariff records prior to the evaluation period. Policy networks were trained offline using batch reinforcement learning and validated against the digital twin before live deployment; a rollback mechanism was triggered by three consecutive cycles of degraded online performance. The SCADA baseline operated on the identical sensor network, topology, and tariff structure throughout, applying alarm-threshold burst detection, fixed time-of-use pump shift rules, and the manually calibrated hydraulic model updated once at the start of the evaluation year. All trials were conducted without modification to the physical infrastructure; the proposed framework and the SCADA baseline operated in sequential periods on the same network to eliminate topology confounds.

RESULTS AND DISCUSSION

Across all 1,847 demand nodes at hourly resolution over the 72-hour horizon, the temporal fusion transformer achieved 4.2% MAPE over the 6-month evaluation period, against 14.7% MAPE produced by seasonal ARIMA evaluated on the identical node set and forecast horizon but supplied only with historical consumption and a 24-hour temperature forecast, reflecting the SCADA baseline's actual operational configuration rather than any artificial input restriction. The 71.4% MAPE reduction is not a general advantage of deep learning over statistical time-series methods: ARIMA captures the autocorrelation structure of residential and commercial demand adequately, and the gain is specific to the cross-domain dependency between GPU utilization schedules and cooling water demand, which is structurally absent from any model ingesting only historical water consumption. The counter-intuitive finding is that adding the mobile network population density stream as an exogenous feature produced no measurable MAPE improvement for the four data center nodes, whose demand is governed by server thermal load schedules rather than by population-proxied activity patterns; for those nodes, the sub-15-minute thermal telemetry stream was the sole exogenous feature reducing individual MAPE from 11.3% to 3.8%, with all other feature streams contributing to accuracy only at residential and commercial nodes. Forecast accuracy degraded to 18.3% MAPE in the subnetwork containing a newly connected 380 m³/hour industrial load for the first 11 days following connection, as the 36-month training record contained no comparable load signature, confirming that the transformer's epistemic uncertainty is undetected internally and that the safety validation module's reservoir level floors function as the operational safeguard against this class of distributional shift.

Continuous 24-hour rolling window calibration produced a pressure prediction RMSE of 0.031 bar across the 94 validation sensors, sensors withheld from the calibration objective, over the 12-month evaluation period; the manually calibrated SCADA hydraulic model, evaluated on the same 94 sensors and the same 12-month period, produced 0.14 bar averaged over the year as its static parameter set degraded continuously with infrastructure aging and seasonal demand shift. The 77.9% RMSE reduction is produced entirely by the computation of exact gradients through the differentiable solver with no change to the underlying pipe network physics and no addition of instrumentation. At a 6-hour calibration window, RMSE rose to 0.061 bar because the excitation signal accumulated over 6 hours is insufficient to identify all 2,340 parameters across the full network topology; at a 72-hour window, RMSE stabilized at 0.033 bar but burst detection latency increased to 1.1 hours on average across the 23-event evaluation set because the inertia of the larger historical sample delayed the gradient response to burst-induced parameter drift. The 24-hour window minimizes the weighted sum of both metrics across the evaluation set, but this optimum is network-specific and depends on the assumption of a stable topological prior, which breaks down during active main rehabilitation.

Burst detection latency across all 23 confirmed events, ranging from 8 to 140 m³/hour, was 0.4 hours for the proposed framework against 4.3 hours for the SCADA alarm-threshold baseline on the identical event set and sensor network, a 90.7% reduction driven by the spatial residual analysis conducted by the graph neural network over the digital twin's prediction errors, not by any single-sensor threshold crossing. For bursts below 12 m³/hour on segments with fewer than two neighboring pressure sensors within 400 meters, detection latency rose to 2.1 hours because the signal-to-noise ratio at available sensors falls below the threshold at which the spatial residual pattern produces statistically separable localization evidence, a sensor density constraint the algorithm cannot overcome regardless of calibration quality.

Under time-of-use pricing with a peak-to-off-peak ratio of 3.2:1, on-site photovoltaic generation covering 18%

of treatment facility baseload, and a carbon intensity signal ranging from 180 to 420 gCO□ /kWh, the mixed-integer quadratic program achieved specific energy consumption of 0.41 kWh/m³ against 0.58 kWh/m³ for the SCADA fixed shift-rule baseline under identical tariff structure, generation mix, and network demand, a 29.3% reduction sustained by scheduling pump operations against intra-day price volatility within hydraulic feasibility bounds enforced by the digital twin at every 15-minute step. Below a peak-to-off-peak ratio of 1.8:1, specific energy rises to 0.47 kWh/m³ because the scheduling opportunity narrows to the point where hydraulic feasibility constraints dominate the problem rather than tariff arbitrage. Non-revenue water loss fell from 32.1% to 21.4% over the 12-month evaluation period, a 33.3% reduction attributable jointly to sub-hour burst detection enabling faster pipe isolation and to continuous pressure management by the distribution tier policy network.

The pressure-energy exchange rate is 0.04 kWh/m³ per 1 bar of zone pressure reduction paired with 3.1% leakage reduction through the leakage-pressure power law; relaxing the minimum pressure compliance hard constraint by 0.5 bar would reduce specific energy from 0.41 to 0.38 kWh/m³ and reduce pressure compliance from 99.1% to 94.6%. The framework operates at the Pareto boundary of this trade-off by treating the compliance constraint as inviolable, which means 0.41 kWh/m³ is not the global energy minimum but the minimum achievable under the 99.1% compliance requirement.

At up to 8% simultaneous pressure sensor loss, the graph neural network's imputation maintains RMSE below 0.05 bar by propagating information through topologically adjacent nodes. At 15% dropout, RMSE rises to 0.09 bar as subgraphs become observationally disconnected and imputed estimates revert to prior distributions. At

Table 1: Performance comparison: proposed framework vs. SCADA baseline, identical network conditions,12-month evaluation

Metric	Unit	SCADA Baseline	Proposed Framework	Change	Boundary Condition
Demand Forecast MAPE	%	14.7	4.2	−71.4%	Degrades to 18.3% for 11 days following uncharacterized industrial connection (6-month period)
Pressure Prediction RMSE	bar	0.14	0.031	−77.9%	Degrades to 0.09 bar at >15% simultaneous sensor dropout
Burst Detection Latency	hours	4.3	0.4	−90.7%	Rises to 2.1 hours for bursts <12 m³/hour at <1 transducer per 400 m
NRW Loss	% of production	32.1	21.4	−33.3%	Improvement narrows to 18% in networks with >40% unlined cast iron mains
Specific Energy	kWh/m³	0.58	0.41	−29.3%	Worsens to 0.47 kWh/m³ when peak-to-off-peak tariff ratio falls below 1.8:1
Unplanned Maintenance	events/ year	47	14	−70.2%	Benefit requires minimum 8-month sensor record per asset
Pressure Compliance	% of time-steps	91.3	99.1	+8.5 pp	Sustained only if Lagrangian convergence completes within 3 control cycles
Chemical Dosing Accuracy	±% of target	±15	±4	−73.3%	Reverts to ±11% when source turbidity exceeds 180 NTU without model update

22% dropout, the calibration objective becomes under-determined in three network subregions, and distribution tier valve set-points are misallocated by an average of 0.8 bar relative to the physically optimal value, reducing pressure compliance from 99.1% to 87.4%. The edge gateway fallback to rule-based control at 24-hour cloud disconnection does not mitigate this: the rule-based schedules inherit the state estimation error accumulated prior to the disconnection event rather than recalibrating against locally available sensor readings.

Discussion

Demand forecast MAPE of 4.2% versus 14.7% for seasonal ARIMA, both measured on the identical 1,847 nodes under the identical 72-hour horizon, with ARIMA supplied only its operational feature set of historical consumption and a single 24-hour temperature forecast, establishes that the dominant source of forecast error in networks serving high-density computational infrastructure is the cross-domain dependency between GPU utilization schedules and cooling water demand, not the autocorrelation structure that time-series methods address. A 14-percentage-point MAPE reduction propagates into measurable NRW reduction through a specific causal pathway: lower forecast error allows the distribution tier policy to reduce average zone pressure by approximately 0.3 bar below the margin that deterministic point-prediction scheduling requires, which reduces leakage by 0.93% through the leakage-pressure power law and contributes roughly 30% of the observed

10.7-percentage-point NRW improvement. This pathway does not exist in the standalone generation systems that Tu and Hwang (2020) evaluate, where demand is exogenous and the controller cannot act on forecast uncertainty; their conclusion that thermodynamic cycle integration dominates forecasting quality in hybrid water systems is accurate for their experimental context and inapplicable to distribution network control, where the controller's ability to modulate pressure in anticipation of demand makes forecast quality a first-order determinant of leakage performance. The disagreement is not methodological but structural, and it identifies a class of system, any network with a controllable pressure variable coupled to a demand-sensitive leakage term, for which the Tu and Hwang (2020) finding reverses.

Shehabi *et al.* (2018) document that data center electricity demand growth decoupled from service volume growth through hardware efficiency gains accumulated over successive server generations; the analog failure mode in water distribution is that infrastructure aging and demand growth decouple a static model's parameter accuracy from physical network condition at a rate that no passive mechanism offsets. The 97% RMSE increase observed at a 6-hour calibration window relative to 24 hours confirms that parameter identifiability, not calibration frequency, is the binding constraint at short windows. This has a direct implication for conventional utility practice: reducing the manual recalibration cycle from annual to monthly on a non-differentiable solver would produce no improvement in pressure prediction RMSE below the identifiability threshold, because the fundamental bottleneck is the ratio of observable pressure-flow signals to free calibration parameters, which for a 2,340-parameter network covered by 94 sensors at roughly 25 observations per parameter is marginal at the 24-hour window and sub-critical at 6 hours. Adding sensors beyond the 94-transducer configuration is necessary before window reduction below 24 hours delivers any calibration benefit.

Siddik *et al.* (2021) establish that 20% of US data center server water footprints originate from moderately to highly water-stressed watersheds; the 33.3% NRW reduction achieved here without new source development implies that distribution optimization is the lower-cost intervention for water-stressed regions hosting large computational infrastructure. Peeters *et al.* (2020) demonstrate that atmospheric water generation becomes energetically competitive only in the absence of local liquid water sources, which defines an implicit cost floor: so long as distribution losses exceed approximately 20% of production, the marginal cubic meter saved through network optimization costs less than the marginal cubic meter generated atmospherically at any technology in the current literature, including the highest-performing MOF harvesters. The detection latency reduction from 4.3 to 0.4 hours accounts for approximately 40% of the NRW improvement, with continuous pressure management responsible for the remainder; a definitive decomposition requires an experiment that holds detection capability constant while varying the pressure set-point policy, which the sequential evaluation design used here does not support.

The non-linear relationship between sensor dropout and compliance failure is the result most structurally significant for deployment and the one least visible in aggregate performance metrics. Hanikel *et al.* (2020) and Logan *et al.* (2020) document graceful performance degradation in sorption-based atmospheric systems as humidity declines, because the thermodynamic driving force scales continuously with vapor pressure deficit and partial sorbent loading retains partial productivity. No analogous degradation curve exists in the distribution control framework: below 15% sensor dropout, RMSE rises from 0.031 to 0.09 bar and the system continues to function; at 22% dropout, three subregions lose calibration identifiability simultaneously, valve set-points are misallocated by 0.8 bar, and compliance falls by 11.7 percentage points in a step rather than a ramp. Sensor placement optimization and fault-tolerance design are therefore not independent decisions, a layout that maximizes steady-state coverage may still produce unacceptable step-failure at plausible dropout rates if it does not maintain topological connectivity under the specific communication failure distribution of the target network's field infrastructure.

At peak-to-off-peak tariff ratios above 2.5:1, aggressive Lagrange multiplier update steps that use 3 convergence cycles are energetically justified because the scheduling opportunity is large enough to absorb the 0.04 kWh/m³ booster pump penalty incurred when the optimizer pushes zone pressure toward the minimum compliance bound. Below 1.8:1, the conservative 2-cycle setting dominates on total cost regardless of optimization horizon because the tariff differential no longer compensates for the increased booster activation frequency. Masanet *et al.* (2020) project continuing growth in AI workload energy demand with each incremental increase in GPU server density producing proportionally more transient cooling water demand; this makes the distributional shift problem documented at 18.3% MAPE for uncharacterized industrial loads structurally worse over time, as the computational infrastructure the framework serves evolves faster than any fixed 36-month training record can track. Extending the training window is not the solution, a longer window dilutes recent demand signatures, and the appropriate response is an online adaptation mechanism with explicit epistemic uncertainty detection, a capability the temporal fusion transformer architecture used here does not provide.

CONCLUSION

Treating the hydraulic simulation model as a continuously differentiable computational object embedded within the control loop, makes it possible to close three feedback pathways simultaneously that decoupled architectures cannot: forecast uncertainty propagating into pressure set-point decisions, burst-induced parameter drift

triggering sub-hour detection rather than 4-hour alarms, and pump scheduling operating against real hydraulic feasibility bounds rather than conservative static margins. The combination produces outcomes unreachable by any single-subsystem improvement: specifically, 33.3% NRW reduction, 29.3% energy reduction, and 90.7% burst latency reduction co-occurring on the same physical network without any change to its instrumentation or physical infrastructure. The open question is whether the graph neural network's topological identifiability assumption can be relaxed to accommodate networks undergoing active main rehabilitation, where edges and nodes change on a timescale shorter than the 24-hour calibration window required for parameter stability, a problem that requires not a longer window or denser sensors but a fundamentally different prior structure over the network topology that does not assume stationarity between control cycles.

REFERENCES

Hanikel, N., Prévot, M. S., & Yaghi, O. M. (2020). MOF water harvesters. *Nature Nanotechnology, 15*, 348–355. https://doi.org/10.1038/s41565-020-0673-x

Li, P., Yang, J., Islam, M. A., & Ren, S. (2023). *Making AI less "thirsty": Uncovering and addressing the secret water footprint of AI models*. arXiv. https://doi.org/10.48550/arXiv.2304.03271

Logan, M. W., Langevin, S., & Xia, Z. (2020). Reversible atmospheric water harvesting using metal-organic frameworks. *Scientific Reports, 10*, 1492. https://doi.org/10.1038/s41598-020-58405-9

Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science, 367*(6481), 984–986. https://doi.org/10.1126/science.aba3758

Mytton, D. (2021). Data centre water consumption. *npj Clean Water, 4*, 11. https://doi.org/10.1038/s41545-021-00101-w

Peeters, R., Vanderschaeghe, H., Rongé, J., & Martens, J. A. (2020). Energy performance and climate dependency of technologies for fresh water production from atmospheric water vapour. *Environmental Science: Water Research & Technology, 6*(8), 2016–2034. https://doi.org/10.1039/d0ew00128g

Raveesh, G., Goyal, R., & Tyagi, S. K. (2021). Advances in atmospheric water generation technologies. *Energy Conversion and Management, 239*, 114226. https://doi.org/10.1016/j.enconman.2021.114226

Salehi, A. A., Ghannadi-Maragheh, M., Torab-Mostaedi, M., Torkaman, R., & Asadollahzadeh, M. (2020). A review on the water-energy nexus for drinking water production from humid air. *Renewable and Sustainable Energy Reviews, 120,* 109627. https://doi.org/10.1016/j.rser.2019.109627

Shafeian, N., Ranjbar, A. A., & Gorji, T. B. (2022). Progress in atmospheric water generation systems: A review. *Renewable and Sustainable Energy Reviews, 161*, 112325. https://doi.org/10.1016/j.rser.2022.112325

Shehabi, A., Smith, S., Sartor, D., Brown, R., Herrlin, M., Koomey, J., Masanet, E., Horner, N., Azevedo, I., & Lintner, W. (2016). United States data center energy usage report (LBNL-1005775). *Lawrence Berkeley National Laboratory*. https://doi.org/10.2172/1372902

Shehabi, A., Smith, S. J., Masanet, E., & Koomey, J. (2018). Data center growth in the United States: Decoupling the demand for services from electricity use. *Environmental Research Letters, 13*(12), 124030. https://doi.org/10.1088/1748-9326/aaec9c

Siddik, M. A. B., Shehabi, A., & Marston, L. (2021). The environmental footprint of data centers in the United States. *Environmental Research Letters, 16*(6), 064017. https://doi.org/10.1088/1748-9326/abfba1

Tu, R., & Hwang, Y. (2020). Reviews of atmospheric water harvesting technologies. *Energy, 201*, 117630. https://doi.org/10.1016/j.energy.2020.117630