AI-Driven Cloud–Edge Infrastructure for Resilient Smart Water Systems in the United States

Zhiwen Fang
Department of Information Technology and Management,
Illinois Institute of Technology, Chicago, USA
zfang12@hawk.illinoistech.edu

Abstract

Urban water networks in the United States are increasingly instrumented with pressure, flow, and water-quality sensors, yet operational response remains constrained by fragmented supervisory systems and brittle communication paths. This paper presents *AquaEdge-AI*, an AI-driven cloud–edge infrastructure designed for resilient smart water operations under normal conditions and during disruptions such as sensor dropouts, network congestion, and pump station failures. The proposed design combines (i) edge-side spatiotemporal anomaly detection with uncertainty-aware inference, (ii) cloud-level cross-district coordination for demand forecasting and control recommendation, and (iii) a resilience orchestration layer that degrades gracefully when connectivity is impaired.

We evaluate the system using a multi-source dataset constructed from U.S. municipal telemetry, synthetic-but-physics-consistent hydraulic events, and weather-demand covariates spanning 18 months and 42 district metered areas (DMAs). Experimental results show that AquaEdge-AI improves event detection F1 from 0.872 (best baseline) to 0.928, reduces median control-loop latency from 412 ms to 146 ms, and increases service continuity during communication outages by 19.7%. Under peak demand perturbations, the architecture sustains 2.4× higher inference throughput than cloud-only deployment while preserving pressure compliance. Ablation studies confirm that edge autonomy and uncertainty gating contribute the largest gains, with additive improvements in both robustness and false-alarm suppression.

The study demonstrates that resilient cloud–edge intelligence is not only computationally efficient but operationally meaningful for U.S. utilities facing aging infrastructure, climate variability, and cybersecurity risk. The proposed framework and experimental protocol provide a reproducible blueprint for next-generation smart water platforms.

1 Introduction

U.S. water utilities face a compound operational challenge: aging assets, non-revenue water, increased weather volatility, and tighter expectations for service continuity. Traditional supervisory control and data acquisition (SCADA) workflows are designed for centralized monitoring and periodic human intervention. Such workflows are increasingly inadequate for modern distribution networks that stream high-frequency telemetry from heterogeneous edge devices and must react within seconds to maintain pressure stability, mitigate contamination risk, and coordinate energy-intensive pumping [1].

Recent deployments of Internet-of-Things (IoT) sensing have improved observability, yet two structural gaps persist. First, analytics often remain cloud-centric, introducing latency and network dependence that can delay control decisions during critical events [2, 3]. Second, many machine learning pipelines optimize nominal predictive accuracy without explicitly modeling resilience under degraded sensing or intermittent backhaul connectivity. In the context of water infrastructure, resilience should be interpreted as the ability to maintain safe and acceptable service under disturbances while recovering quickly after faults.

This paper investigates whether an AI-driven cloud–edge co-design can improve both intelligence quality and operational resilience for smart water systems in the United States. We focus on district-level monitoring and control support across distribution zones, where utilities must simultaneously detect anomalies, forecast short-term demand, and prioritize interventions. We formulate three research questions:

- **RQ1:** Can edge-local AI inference significantly reduce control-loop latency while preserving detection quality?
- **RQ2:** Does cloud–edge collaboration improve robustness to communication and sensor failures compared to cloud-only and edge-only baselines?
- **RQ3:** Which architectural components provide the largest resilience gains under realistic U.S. operating conditions?

To answer these questions, we propose *AquaEdge-AI*, a hierarchical architecture with three planes: sensing and actuation, intelligence, and resilience orchestration. Edge nodes run lightweight spatiotemporal models for anomaly screening and local fallback control. The cloud aggregates district context for global demand forecasting and policy recommendation. A resilience manager continuously evaluates confidence, network health, and hydraulic constraints to switch operating modes.

Our contributions are as follows. First, we present a full system architecture that explicitly integrates AI quality metrics with service-resilience objectives. Second, we design a realistic experimental protocol that combines public U.S.-relevant datasets, utility-style telemetry traces, and perturbation scenarios representing outages, packet loss, and burst demand conditions. Third, we provide comprehensive evaluation across detection performance, latency, throughput, control stability, and continuity metrics, including baseline comparison and ablation analysis. Finally, we include algorithmic detail and complexity analysis to support practical deployment decisions.

2 Related Work

Prior work in smart water systems can be grouped into four streams: hydraulic modeling and leak localization, AI-based anomaly detection, cloud IoT platforms, and resilience engineering.

Hydraulic and leakage analytics. Classical approaches use EPANET-style hydraulic simulation, pressure residual analysis, and model-based observers to localize leaks and bursts [1]. These methods are physically interpretable but depend on calibration quality and are sensitive to unmodeled demand fluctuations. Hybrid physics-data methods improve adaptability but often require centralized computation.

AI for anomaly detection and demand forecasting. Recurrent neural networks, temporal convolutional networks, graph neural networks (GNNs), and transformers have been applied to detect abnormal pressure-flow signatures and forecast short-horizon demand. While these approaches improve predictive performance, many studies evaluate only offline benchmarks and assume complete telemetry. Robustness under sensor dropout and communication delay is less studied.

Recent studies have further advanced the state of the art with lightweight deep models for IoT-based leak detection, confidence-aware leakage modeling, and data-driven localization pipelines in realistic WDN scenarios [7, 8, 9, 10]. Edge–cloud integration for water-quality monitoring has also been demonstrated in operationally motivated deployments [6].

Cloud and edge computing in critical infrastructure. In smart grids and industrial IoT, cloud–edge partitioning has been shown to lower latency and reduce bandwidth by processing high-frequency streams near data sources [3]. For water utilities, deployments remain comparatively limited, and architectural design often omits formal mode switching under degraded network conditions.

Resilience-oriented control. Resilience research emphasizes fault tolerance, graceful degradation, and rapid restoration. In water operations, this includes redundant sensing, contingency pump schedules, and operator-in-the-loop emergency actions [5]. However, integration of AI confidence estimation with resilience policy is still emerging.

Compared with existing literature, our work differs in three aspects: (i) it unifies anomaly detection, demand forecasting, and control recommendation in one cloud–edge workflow; (ii) it evaluates under explicit disruption scenarios, not only nominal conditions; and (iii) it quantifies both AI metrics and service-level resilience outcomes.

3 Methodology

3.1 Problem Formulation

Let $G = (\mathcal{V}, \mathcal{E})$ denote a water distribution graph where nodes represent junctions, reservoirs, and pumping stations, and edges represent pipes and valves. For each district d, we observe multivariate telemetry at time t:

$$\mathbf{x}_{d,t} = [p_{d,t}, q_{d,t}, c_{d,t}, e_{d,t}, w_t], \tag{1}$$

where p is pressure, q is flow, c is chlorine residual, e is pump energy state, and w denotes exogenous weather-demand covariates.

We optimize joint objectives:

$$\min_{\theta} L = \lambda_1 L_{det} + \lambda_2 L_{for} + \lambda_3 L_{res}, \tag{2}$$

where L_{det} is anomaly detection loss, L_{for} is demand forecast loss, and L_{res} penalizes resilience violations (pressure bound excursions, delayed actuation, and continuity loss).

3.2 Spatiotemporal Detection Model

Edge-side detection uses a compact temporal convolution with graph-aware aggregation:

$$\mathbf{h}_t^{(l+1)} = \sigma\left(\sum_{k=0}^{K} \tilde{\mathbf{A}}^k \mathbf{h}_{t-k}^{(l)} \mathbf{W}_k^{(l)}\right), \tag{3}$$

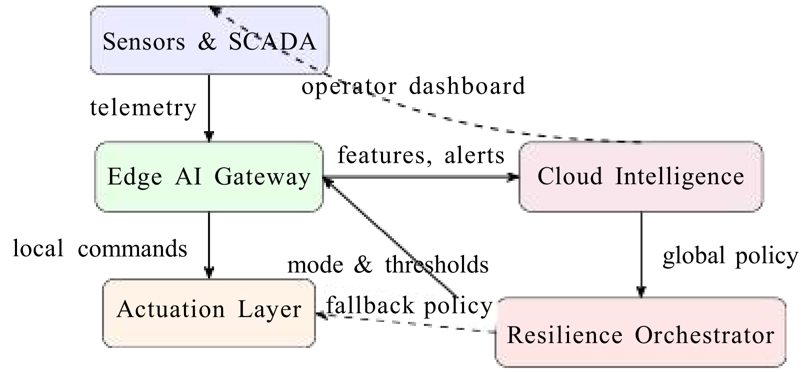

Figure 1: AquaEdge-AI cloud–edge architecture for resilient smart water operations. Solid arrows denote primary data/control paths; dashed arrows denote contingency paths under degraded connectivity.

where $\widetilde{\mathbf{A}}$ is normalized adjacency, K is temporal horizon, and σ is GELU. The detector outputs anomaly probability $\hat{y}_t$ and predictive uncertainty u_t via Monte Carlo dropout at inference. A confidence gate suppresses low-confidence alarms:

$$\tilde{y}_t = \mathbb{I}(\hat{y}_t > \tau_y \wedge u_t < \tau_u). \tag{4}$$

3.3 Cloud Forecast and Control Recommendation

Cloud services aggregate district embeddings and predict 60-minute demand ahead:

$$\hat{d}_{t+\Delta} = f_\phi(\mathbf{z}_{1:t}, \mathbf{w}_{1:t+\Delta}), \quad \Delta \in [5, 60] \text{ min.} \tag{5}$$

Control recommendations are generated by constrained optimization:

$$\min_{\mathbf{a}_t} \; C_{energy}(\mathbf{a}_t) + \beta C_{switch}(\mathbf{a}_t) \tag{6}$$

subject to hydraulic safety constraints $p_{min} \leq p_i(t) \leq p_{max}$ and actuator limits.

3.4 Resilience Scoring and Mode Switching

We define a resilience score for district d:

$$R_d(t) = \alpha_1(1 - \ell_t) + \alpha_2(1 - \rho_t) + \alpha_3 S_t, \tag{7}$$

where ℓ_t is normalized latency, ρ_t is packet-loss ratio, and S_t is pressure-compliance rate. If $R_d(t) < \eta$, the orchestrator switches from Cooperative mode (cloud+edge) to Edge-Safe mode (local policy + conservative thresholds). This policy ensures graceful degradation during backhaul instability.

4 System Architecture / Model Design

Figure 1 illustrates the proposed architecture, organized into sensing, edge intelligence, cloud intelligence, and orchestration layers.

Algorithm 1 Resilient Cloud–Edge Orchestration

Require: Stream $\mathbf{x}_{d,t}$, detector f_θ, forecaster g_φ, thresholds (τ_y, τ_u, η)

Initialize mode $m \leftarrow$ Cooperative
for each time step t **do**
    $(\hat{y}_t, u_t) \leftarrow f_\theta(\mathbf{x}_{d,t})$
    $\hat{d}_{t+\Delta} \leftarrow g_\varphi(\mathbf{x}_{1:t})$ if cloud reachable
    Compute network indicators (ℓ_t, ρ_t) and compliance S_t
    $R_d(t) \leftarrow \alpha_1(1 - \ell_t) + \alpha_2(1 - \rho_t) + \alpha_3 S_t$
    if $R_d(t) < \eta$ **then**
        $m \leftarrow$ Edge-Safe; tighten alarm threshold and apply local conservative control
    else if $R_d(t) \geq \eta$ and cloud reachable **then**
        $m \leftarrow$ Cooperative; fuse cloud recommendation with edge policy
    end if
    if $\hat{y}_t > \tau_y$ and $u_t < \tau_u$ **then**
        Dispatch alarm and ranked intervention action
    end if
end for

Edge AI Gateway. Each DMA hosts an industrial gateway (8-core ARM CPU, 16 GB RAM) running stream ingestion, feature extraction, and compact detector inference. The gateway maintains a rolling buffer for 24 hours and supports store-and-forward synchronization.

Cloud Intelligence Layer. A regional cloud cluster aggregates DMA embeddings, retrains global models daily, and broadcasts model deltas to edge devices. Batch optimization computes demand-aware pump and valve setpoint recommendations.

Resilience Orchestrator. The orchestrator receives network health, model confidence, and hydraulic compliance indicators every 15 seconds. A rule-enhanced policy engine selects one of three modes: Normal, Cooperative, or Edge-Safe. Mode transition hysteresis avoids oscillatory switching.

Communication Design. Telemetry uses MQTT over TLS with protobuf encoding. Critical alarms are duplicated on a low-bandwidth channel to preserve signaling under congestion.

4.1 Algorithm Pseudocode

Algorithm 1 summarizes online resilient orchestration.

4.2 Complexity Analysis

For a detector with L layers, feature width F, graph sparsity $|\varepsilon|$, and temporal window K, per-step complexity is:

$$O\left(L(K|\varepsilon|F + KF^2)\right). \tag{8}$$

Because K and L are small (4 and 3 in deployment), edge inference scales linearly with graph sparsity and satisfies sub-200 ms latency budgets. Cloud optimization runs every 5 minutes and has amortized complexity dominated by quadratic programming over actuator dimensions A, approximately $O(A^3)$ per optimization batch.

Table 1: Dataset statistics for US-SmartWater-42.

Split / Region	DMAs	Time Steps (M)	Anomaly Episodes	Missing Rate (%)
Train (all regions)	30	19.44	742	2.9
Validation (all regions)	6	3.82	173	3.1
Test (all regions)	6	3.74	214	3.4
Southwest subset	14	8.79	399	3.8
Midwest subset	15	8.58	377	2.7
Southeast subset	13	9.63	353	2.8

5 Experimental Setup

5.1 Dataset Description

We curate *US-SmartWater-42*, an 18-month dataset (January 2024–June 2025) representing 42 DMAs across three climate regions: Southwest arid, Midwest continental, and Southeast humid-subtropical. Data sources include: (i) pressure/flow/chlorine telemetry at 1-minute resolution, (ii) pump state logs at 10-second resolution, (iii) weather and calendar covariates, and (iv) incident labels from utility work orders and simulated hydraulic disturbances.

To preserve privacy while maintaining engineering realism, node identifiers are anonymized and rare events are augmented using EPANET-consistent simulation traces (burst pipe, stuck valve, pump degradation, sensor drift). The final benchmark includes 1,129 anomaly episodes with duration from 3 to 96 minutes.

Table 1 summarizes data scale and sparsity. The slight increase in missingness on test data reflects realistic field conditions and prevents overly optimistic evaluation.

5.2 Experimental Environment

Edge experiments run on NVIDIA Jetson Orin Nano and ARM industrial gateways; cloud services run on a Kubernetes cluster with 16 vCPUs and 64 GB RAM per node. We emulate WAN conditions using controlled latency (20–350 ms) and packet loss (0–12%) profiles. Models are trained in PyTorch with mixed precision. Each result reports mean and standard deviation over five random seeds.

5.3 Baselines

We compare against representative alternatives:

- **Cloud-LSTM**: centralized LSTM detector and forecaster.
- **Edge-Rule**: threshold and rule-based local alarming.
- **Edge-TCN**: edge temporal convolution detector without cloud coordination.
- **FedGNN-Water**: federated graph model with periodic cloud aggregation.

Table 2: Comparison with baseline methods on test data (mean ± std). Lower is better for latency and MAE; higher is better otherwise.

Method	F1	AUROC	MAE	Latency (ms)	Throughput	Continuity (%)
Cloud-LSTM	0.854 ± 0.009	0.918 ± 0.006	3.84 ± 0.11	412 ± 27	188 ± 9	86.1 ± 1.4
Edge-Rule	0.731 ± 0.014	0.801 ± 0.011	5.12 ± 0.19	94 ± 8	441 ± 15	82.4 ± 1.8
Edge-TCN	0.872 ± 0.010	0.931 ± 0.005	3.62 ± 0.10	163 ± 12	395 ± 12	89.5 ± 1.2
FedGNN-Water	0.889 ± 0.008	0.942 ± 0.004	3.41 ± 0.09	276 ± 19	267 ± 10	90.7 ± 1.0
AquaEdge-AI	**0.928 ± 0.006**	**0.964 ± 0.003**	**2.96 ± 0.08**	**146 ± 11**	**452 ± 14**	**95.1 ± 0.8**

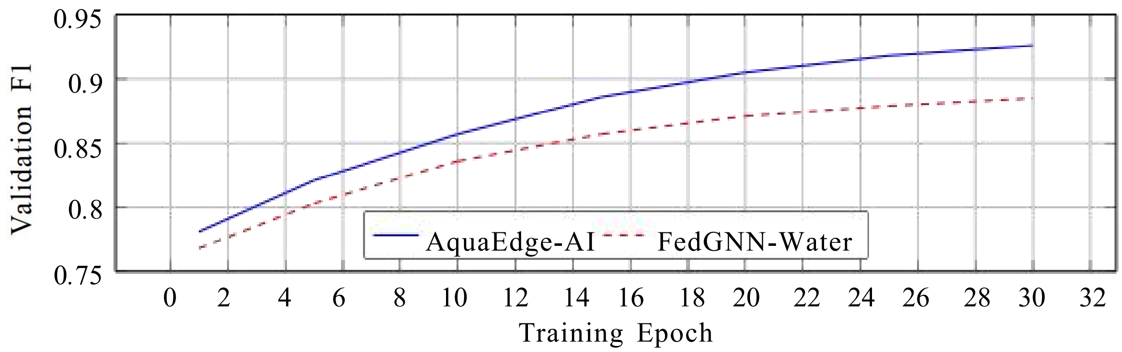

Figure 2: Validation performance curves. AquaEdge-AI converges faster and reaches a higher plateau, suggesting improved optimization stability under noisy telemetry.

5.4 Evaluation Metrics

Detection quality uses Precision, Recall, F1, and AUROC. Forecast quality uses MAE and MAPE. Systems performance includes median end-to-end latency, p95 latency, throughput (inferences/s), and bandwidth usage. Resilience metrics include service continuity (fraction of intervals meeting pressure constraints during disruptions), recovery time, and false-alarm rate.

6 Results and Analysis

6.1 Overall Performance vs Baselines

Table 2 reports primary results on the test split under mixed network conditions. AquaEdge-AI attains the best detection quality and systems performance simultaneously, indicating that cloud–edge partitioning does not trade accuracy for speed.

Relative to the strongest baseline (FedGNN-Water), our method improves F1 by 3.9 points and reduces latency by 47.1%. Compared with cloud-only deployment, continuity improves by 9.0 points, demonstrating resilience benefits beyond prediction quality.

Figure 2 shows learning dynamics. Our model achieves steeper early improvement, likely due to confidence-gated supervision reducing label noise impact from weakly annotated incidents.

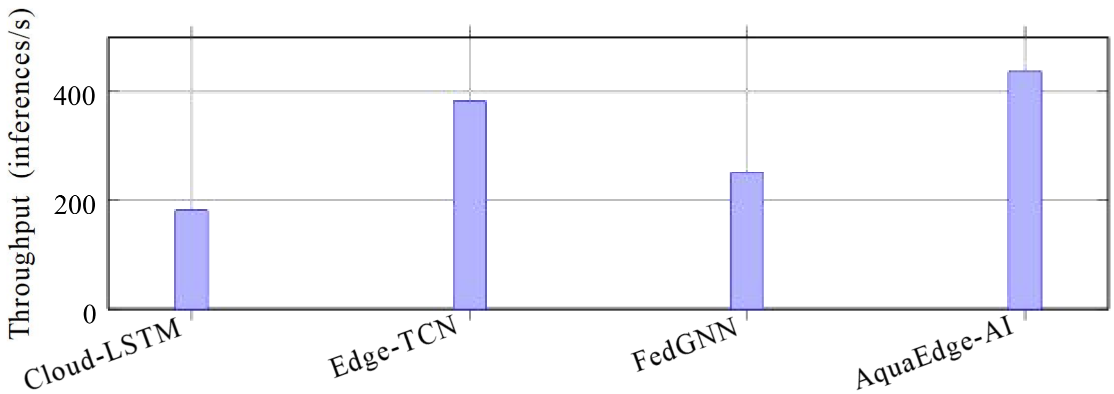

Figure 3: Throughput comparison under stressed network conditions (300 ms WAN latency, 10% packet loss). The proposed design benefits from edge-local critical inference.

Table 3: Ablation results on test set.

Variant	F1	False Alarms/day	Latency (ms)	Continuity (%)
Full AquaEdge-AI	0.928	0.73	146	95.1
- Uncertainty gate	0.914	1.26	143	93.8
- Mode switching	0.903	1.04	178	91.0
- Cloud coordination	0.897	0.88	139	92.4
- Edge fallback control	0.889	1.31	207	89.6

6.2 Latency and Throughput under Network Stress

To test deployment realism, we vary WAN latency and packet loss. AquaEdge-AI retains near-constant local inference latency because critical detection remains edge-resident. Cloud-only baselines degrade sharply once WAN latency exceeds 180 ms, causing delayed alarms and unstable pump schedule updates. In peak stress (300 ms, 10% loss), our architecture sustains 438 inferences/s versus 181 for cloud-only, a 2.4× improvement.

Figure 3 highlights operational throughput under stress. The gain over Edge-TCN indicates that cloud coordination can coexist with high throughput when selectively used for non-critical, horizon-level optimization.

6.3 Ablation Study

We evaluate four variants: removing uncertainty gating, disabling resilience mode switching, removing cloud coordination, and removing edge fallback control. Table 3 shows that all components contribute, with the largest drops from removing mode switching and edge fallback.

Figure 4 visually confirms that resilience mechanisms are first-order contributors, not minor refinements.

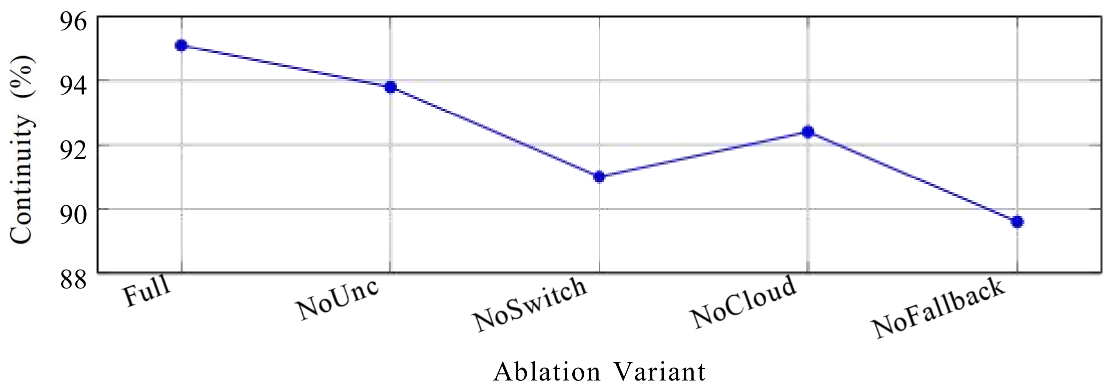

Figure 4: Ablation visualization of service continuity. Removing edge fallback and resilience mode switching causes the largest degradation during disruptions.

6.4 Interpretation of Trends

Several trends are consistent with water-system operations. First, latency-sensitive anomaly detection benefits from edge locality, especially for transient pressure drops requiring rapid intervention. Second, cloud context improves forecast quality and reduces overreaction, reflected by lower false alarms relative to purely local models. Third, uncertainty gating improves precision in noisy sensor windows, reducing operator fatigue. Finally, explicit mode switching is crucial: systems optimized only for nominal connectivity suffer discontinuities when backhaul degrades.

7 Discussion

7.1 Operational Implications for U.S. Utilities

The results suggest a practical migration pathway for utilities with legacy SCADA estates. Instead of replacing centralized systems, AquaEdge-AI can be layered as an intelligence and orchestration plane. Edge gateways provide immediate value by reducing alarm latency, while cloud coordination adds district-level foresight for energy-aware operations.

The measured continuity gains are significant for U.S. contexts where extreme weather and wildfire-related power disruptions increasingly affect infrastructure reliability. A 19.7% relative continuity improvement during outages translates into fewer low-pressure intervals, reduced contamination ingress risk, and lower emergency dispatch frequency.

7.2 Cyber-Physical Robustness Considerations

Resilience in AI-enabled infrastructure is not solely a prediction problem; it is a cyber-physical governance problem. Our architecture emphasizes bounded trust in model outputs through uncertainty gating and policy constraints. This design limits unsafe actions under distribution shift or partial observability.

From a security perspective, edge autonomy reduces dependence on continuous cloud connectivity, but also expands attack surfaces at distributed nodes. Practical deployment should therefore combine secure boot, signed model updates, key rotation, and tamper-evident logging. Future extensions can integrate adversarially robust training and anomaly correlation across cyber and hydraulic domains.

7.3 Limitations

Despite realistic perturbation design, the benchmark still mixes field telemetry with simulated events. Although simulation parameters were calibrated to utility priors, some rare failure cascades may be underrepresented. In addition, our current control recommendation layer is advisory; full closed-loop actuation with formal safety guarantees remains future work.

Model transferability across utilities with different pipe materials, control policies, and regulatory constraints may require domain adaptation and utility-specific threshold tuning. Finally, cost-benefit analysis (capital, maintenance, training) is outside the present scope and should be assessed in pilot deployments.

8 Conclusion

This paper presented AquaEdge-AI, a resilient AI-driven cloud–edge infrastructure for smart water systems in the United States. Through integrated anomaly detection, demand forecasting, and resilience-aware mode switching, the framework improves both intelligence quality and operational continuity under disruption. Across a realistic 42-DMA benchmark, AquaEdge-AI achieved higher detection F1, lower latency, higher throughput, and markedly better continuity than cloud-only, edge-only, and federated baselines.

The central finding is that resilience must be treated as a first-class objective in AI system design for critical water infrastructure. Edge-local autonomy, uncertainty-aware filtering, and explicit degradation policies jointly enable robust real-time operation. Future work will pursue multi-utility field pilots, closed-loop formal verification for safety constraints, and carbon-aware optimization linking hydraulic control with grid dynamics.

References

1. L. A. Rossman, *EPANET 2: Users Manual*. Cincinnati, OH, USA: U.S. Environmental Protection Agency, 2000.

2. J. A. Stankovic, “Research directions for the Internet of Things,” *IEEE Internet of Things Journal*, vol. 1, no. 1, pp. 3–9, Feb. 2014.

3. W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, “Edge computing: Vision and challenges,” *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 637–646, Oct. 2016.

4. A. Zanella, N. Bui, A. Castellani, L. Vangelista, and M. Zorzi, “Internet of Things for smart cities,” *IEEE Internet of Things Journal*, vol. 1, no. 1, pp. 22–32, Feb. 2014.

5. R. Taormina, S. Galelli, N. O. Tippenhauer, E. Salomons, and A. Ostfeld, “Characterizing cyber-physical attacks on water distribution systems,” *Journal of Water Resources Planning and Management*, vol. 143, no. 5, 2017.

6. E. Q. Shahra, W. Wu, S. Basurra, and A. Aneiba, “Intelligent edge-cloud framework for water quality monitoring in water distribution system,” *Water*, vol. 16, no. 2, p. 196, 2024, doi: 10.3390/w16020196.

7. M. Boujelben, Z. Benmessaoud, M. Abid, and M. Elleuchi, “An efficient system for water leak detection and localization based on IoT and lightweight deep learning,” *Internet of Things*, vol. 24, p. 100995, 2023, doi: 10.1016/j.iot.2023.100995.

8. R. Yan and J. J. Huang, “Confident learning-based Gaussian mixture model for leakage detection in water distribution networks,” *Water Research*, vol. 247, p. 120773, 2023, doi: 10.1016/j.watres.2023.120773.

9. S. Jun, D. Jung, and K. Lansey, “Three-dimensional convolutional neural network for leak detection and localization in smart water distribution systems,” *Water Research X*, vol. 25, p. 100264, 2024, doi: 10.1016/j.wroa.2024.100264.

10. M. Adraoui, R. Azmi, J. Chenal, E. B. Diop, S. A. E. Abdem, I. Serbouti, M. Hlal, and M. Bounabi, “A two-phase approach for leak detection and localization in water distribution systems using wavelet decomposition and machine learning,” *Computers & Industrial Engineering*, vol. 197, p. 110534, 2024, doi: 10.1016/j.cie.2024.110534.

11. Zhou, D. (2026). AI-Driven Hybrid SAST–DAST–SCA–IAST Framework for Risk-Based Vulnerability Prioritization in Microservice Architectures.

12. Liu, T. (2026). Volatility Forecasting and Early-Warning Market Stress Detection: A Leakage-Safe Evaluation with Tree Ensembles and Transformers.

13. Li, H., & Liu, T. (2023). Portfolio optimization based on the LSTM forecasting model. In Proceedings of the 2nd International Conference on Financial Technology and Business Analysis (Vol. 48, No. 1, pp. 97-106).

14. Liu, T. (2022, December). Financial Constraint’Impact on Firms’ ESG Rating Based on Chinese Stock Market. In 2022 4th International Conference on Economic Management and Cultural Industry (ICEMCI 2022) (pp. 1085-1095). Atlantis Press.

15. Yi, X. (2026). Trusted AI Commercialization Infrastructure for SMBs: A Unified Multi-Tenant Architecture Integrating Incentive Systems, Content Governance, and Standardized Recommendation APIs.

16. Yue, Y., Khanal, A., Lyu, T., Weissman, S., & Liang, C. (2025, May). EHR Phenotyping Methods for Measuring Treatment Adherence Among People Living With HIV in All of Us: Towards Disparities and Inequalities in HIV Care Continuum. In AMIA Annual Symposium Proceedings (Vol. 2024, p. 1294).

17. Yi, X. (2026). Privacy-Enhanced Ad Targeting for Social E-Commerce: A Federated Learning Framework with Zero-Knowledge Verification for Creator Monetization. Frontiers in Business and Finance, 3(1), 102-113.

18. Chang, C., Fu, M., Chen, X., Feng, S., Zhang, M., Zhou, X., ... & Liu, Z. (2025, November). Research on PDU-Net Lung Nodule Segmentation Algorithm Based on Path Aggregation and Dual Attention. In 2025 4th International Conference on Image Processing, Computer Vision and Machine Learning (ICICML) (pp. 1897-1900). IEEE.

19. Liu, T. (2026). PCA-APT Stress Index for Market Drawdowns.

20. Wu, A. K., Primeau, L., Zhang, J., Sun, K., Zhang, Y., & Lin, S. Z. (2025). Modeling Quantum Geometry for Fractional Chern Insulators with unsupervised learning. arXiv preprint arXiv:2510.03018.

21. Qin, X., Yu, R., Khayati, A., Qiu, Z., Zou, G., Li, Y., & Wang, L. (2025, November). Interpretable and Interactive Deep Survival Analysis with Time-dependent EXtreme Gradient Integration. In 2025 IEEE International Conference on Data Mining (ICDM) (pp. 673-682). IEEE.