Artificial Intelligence for Climate-Resilient Stormwater and Wastewater Management

Nalini Parthasarathy[1], Kwame Asante-Boateng[2], Lucia Ferreira-Santos[3], and Tariq Al-Mansouri[4]

[1] School of Civil Engineering, Indian Institute of Technology Bombay, Mumbai 400076, India
[2] Department of Environmental Science and Engineering, Kwame Nkrumah University of Science and Technology, Kumasi, Ghana
[3] Centre for Environmental and Marine Studies (CESAM), University of Aveiro, 3810-193 Aveiro, Portugal
[4] Water Research and Technology Centre, King Abdullah University of Science and Technology (KAUST), Thuwal 23955, Saudi Arabia

Abstract

Climate change is fundamentally altering the hydrological regimes upon which stormwater and wastewater infrastructure was designed to operate, manifesting as intensified extreme precipitation events, extended drought cycles, sea-level rise-induced groundwater infiltration, and thermally stressed biological treatment processes. Conventional infrastructure design paradigms, premised on stationary climatological assumptions, are rapidly becoming obsolete. This article presents a comprehensive review of how artificial intelligence (AI) — spanning machine learning, deep learning, reinforcement learning, physics-informed neural networks, and multi-agent systems — is being deployed to build climate resilience into stormwater and wastewater management systems across diverse geographies and institutional contexts. Synthesising 112 peer-reviewed studies published between 2019 and 2025, we evaluate AI performance across five climate-critical domains: non-stationary flood frequency analysis, drought-adaptive wastewater reuse optimisation, sea-level rise and groundwater intrusion modelling, heat-stress resilience in biological wastewater treatment, and equity-weighted infrastructure investment prioritisation under climate uncertainty. Our analysis finds that physics-informed hybrid architectures consistently outperform purely data-driven models when extrapolating beyond historical climate envelopes, with Physics-Informed Neural Networks (PINNs) retaining 73–89% of in-distribution performance accuracy under 2°C and 4°C warming scenarios. We introduce the Climate-Adaptive AI Water Resilience (CAAWR) framework, which operationalises a five-stage resilience cycle — sense, anticipate, adapt, recover, and learn — using AI as the connective tissue between real-time sensing and long-term planning. Critical findings include the persistent underrepresentation of Global South case studies (23% of reviewed literature), the urgent need for climate-downscaled training datasets at sub-catchment resolution, and the transformative potential of generative AI for rapid scenario planning under deep climate uncertainty. Policy implications for national adaptation planning, infrastructure investment, and transboundary water governance are discussed.

Keywords: *climate resilience; artificial intelligence; stormwater management; wastewater treatment; non-stationarity; physics-informed neural networks; climate adaptation; urban flooding; water reuse; infrastructure planning*

1. Introduction

The assumption of climatic stationarity — that future hydrology will resemble the past — has been the silent foundation of water infrastructure design for over a century. Return period calculations, design storm specifications, and treatment plant capacity projections have all relied on the premise that historical rainfall statistics provide a reliable guide to future extremes. That assumption is now demonstrably false [1]. The Intergovernmental Panel on Climate Change Sixth Assessment Report confirms with high confidence that human-induced climate change has already intensified extreme precipitation events globally, with further intensification of 7% per degree Celsius of warming projected under all emissions scenarios [2].

The consequences for stormwater and wastewater infrastructure are profound and multifaceted. Urban drainage networks designed for 100-year return period events are experiencing hydraulic overloading with increasing frequency; a study of 42 major cities across 5 continents found that the effective return period of design storms had shortened by an average of 34% between 1980 and 2020 [3]. Wastewater treatment plants are simultaneously experiencing influent temperature increases that alter biological treatment kinetics, reduced dilution capacity in receiving water bodies during drought periods, and accelerated corrosion of concrete infrastructure driven by elevated hydrogen sulphide concentrations under warmer, more anaerobic sewer conditions [4]. Coastal utilities face the additional challenge of sea-level rise driving saline groundwater infiltration into sewer networks, diluting wastewater and increasing the energy cost of treatment [5].

Responding to this compound challenge requires management systems that can learn from non-stationary data, anticipate conditions beyond historical experience, and adapt infrastructure operation in real time to variable climatic forcing. Artificial intelligence offers precisely these capabilities. Unlike deterministic hydraulic models or statistical regression approaches, AI systems — particularly deep learning architectures trained on large multi-source datasets — can identify complex nonlinear relationships in historical data and, when appropriately constrained by physical knowledge, extrapolate to novel climate conditions with quantified uncertainty [6].

This article provides a comprehensive, evidence-based review of AI applications specifically oriented toward building climate resilience in stormwater and wastewater management. It is structured around five climate-critical challenge domains identified through a systematic literature review, followed by the introduction of an original integrative framework — the Climate-Adaptive AI Water Resilience (CAAWR) framework — and a forward-looking research and policy agenda. The review contributes to the growing literature at the intersection of AI, climate science, and water engineering by explicitly foregrounding the climate dimension often treated as peripheral in general AI-water management reviews.

1.1 Defining Climate Resilience in Urban Water Systems

For the purposes of this review, climate resilience in urban water systems is defined as the capacity of stormwater and wastewater infrastructure, institutions, and communities to anticipate, absorb, accommodate, and recover from climate hazards while maintaining essential service delivery, public health protection, and ecological integrity [7]. This definition encompasses three temporal dimensions: resistance (maintaining function during acute climate stress events), recovery (restoring function efficiently after disturbance), and adaptation (adjusting system configuration to reduce vulnerability to future stressors). AI contributes to all three dimensions but with different tools and evidence bases for each.

1.2 Scope and Search Methodology

The review covers studies published from January 2019 to July 2025, with temporal lower bound reflecting the emergence of climate-focused deep learning applications in hydrology following the publication of seminal LSTM streamflow models. Database searches of Web of Science, Scopus, Engineering Village, and the Web of Science Climate Change Collection used search strings combining ('artificial intelligence' OR 'machine learning' OR 'deep learning' OR 'neural network') AND ('climate change' OR 'climate resilience' OR 'climate adaptation' OR 'non-stationary') AND ('stormwater' OR 'wastewater' OR 'urban flooding' OR 'combined sewer'). After PRISMA-compliant screening, 112 studies met all inclusion criteria. Geographic distribution of included studies is examined as a secondary analytical dimension given equity concerns about the North–South knowledge gap in climate-water-AI research.

2. Climate Stressors and Their Impacts on Water Infrastructure

2.1 Intensified Extreme Precipitation

The thermodynamic relationship between atmospheric temperature and moisture-holding capacity — the Clausius-Clapeyron scaling — predicts approximately 7% increase in extreme precipitation intensity per degree of warming. Observational evidence from dense rain gauge networks and satellite remote sensing confirms super-Clausius-Clapeyron scaling in some convective systems, with short-duration (sub-hourly) extreme events intensifying at 14–21% per degree in several mid-latitude regions [8]. For urban stormwater systems, peak flow increases of this magnitude far exceed the design safety margins of most combined sewer networks, where residual capacity above average wet-weather flows is typically 20–40% [9].

The challenge is compounded by urbanisation-driven reductions in pervious surface area, which reduce infiltration and increase the ratio of surface runoff to total rainfall. Impervious surface coverage in rapidly growing cities across South Asia, sub-Saharan Africa, and Southeast Asia is expanding at 2–8% per decade, amplifying the stormwater volume generated by any given precipitation event [10]. The interaction between climate-intensified storms and expanding impervious surfaces creates a compounding risk trajectory that infrastructure investment planning must explicitly account for.

2.2 Drought Stress and Wastewater Reuse Imperatives

While flood risk intensifies in many regions, drought frequency and severity are simultaneously increasing across the Mediterranean basin, southern Africa, western North America, and northeastern Brazil [2]. Extended dry periods reduce streamflow in receiving water bodies, concentrating treated wastewater effluent and increasing the ambient concentration of micropollutants, nutrients, and pathogens in rivers and estuaries used as drinking water sources [11]. They also reduce groundwater recharge, depleting the aquifer storage that many utilities rely on for dry-season supply augmentation.

These pressures are accelerating interest in direct and indirect potable water reuse — the treatment of wastewater to drinking water standards for aquifer recharge or direct distribution system injection. Achieving the treatment reliability required for potable reuse mandates near-continuous process monitoring and adaptive control beyond the capability of conventional fixed-setpoint treatment protocols [12]. AI-driven adaptive treatment control is therefore not merely an efficiency opportunity in drought-stressed contexts: it is a prerequisite for safe potable reuse operation.

2.3 Sea-Level Rise and Coastal Sewer Inundation

Coastal urban areas face a distinctive climate threat: rising sea levels increase the frequency and magnitude of tidal and storm surge inundation of low-lying drainage infrastructure. Saltwater intrusion into sewer networks via overtopped manholes, cracked pipes, and backflow through outfall structures introduces chlorides that inhibit nitrification bacteria, reduce biological treatment efficiency, and accelerate corrosion of concrete and ferrous metal infrastructure [13]. A study of 23 coastal US utilities found that sea-level rise of 0.3 m — within the likely range for 2050 under intermediate emissions scenarios — would increase average annual saltwater infiltration volumes by 47–89%, with correlated increases in blending ratios at WWTPs that compromise compliance with nutrient discharge limits [5].

2.4 Thermal Stress in Biological Treatment

Activated sludge wastewater treatment relies on complex microbial communities whose metabolic rates are temperature-dependent. Increased influent temperatures driven by warming ambient conditions and urban heat island intensification alter nitrification and denitrification kinetics, changing the aeration demand patterns that treatment plants have historically managed through seasonal adjustment of operational parameters [14]. Simultaneously, reduced receiving water dissolved oxygen concentrations during warm periods tighten effluent quality requirements precisely when biological treatment performance is most challenged. AI models capable of predicting effluent quality as a function of influent characteristics, temperature, and microbial community state are essential for proactive climate-adaptive WWTP operation.

3. AI for Non-Stationary Flood Frequency Analysis

Flood frequency analysis — the statistical characterisation of extreme flow events and their recurrence probabilities — is the technical foundation of infrastructure design standards. Classical approaches assume stationarity: that the statistical moments (mean, variance, skewness) of the flood frequency distribution remain constant over time. This assumption is violated when land use change, climate change, or catchment modification alters the rainfall–runoff relationship over the design life of infrastructure [15].

AI approaches to non-stationary flood frequency analysis fall into three categories. First, covariate-based non-stationary extreme value models incorporate climate indices — the El Niño Southern Oscillation index, the Atlantic Multi-decadal Oscillation, global mean temperature anomaly — as time-varying parameters of Generalised Extreme Value (GEV) or Generalised Pareto Distribution (GPD) models [16]. Machine learning optimisation of these covariate relationships — using genetic algorithms, Bayesian optimisation, or gradient-boosted regression — has been shown to substantially outperform maximum likelihood estimation when covariate effects are nonlinear or interact with catchment characteristics.

Second, deep learning models for flood frequency analysis directly process long-period historical discharge records alongside climate model projections to generate flood frequency curves conditioned on future climate states. Studies showes Variational Autoencoder (VAE) on 80 years of daily discharge records from 156 catchments across the Indian subcontinent, using CMIP6 climate model outputs as conditioning variables. The VAE generated ensemble flood frequency curves for 2050 and 2100 under SSP2-4.5 and SSP5-8.5 scenarios, revealing that 10-year return period design flows would increase by 28–56% across the Gangetic Plain under high-emissions scenarios — a finding with direct implications for road culvert, urban drain, and detention basin design standards across one of the world's most densely populated regions.

Third, Physics-Informed Neural Networks (PINNs) that embed flood routing equations — the Saint-Venant equations for shallow water flow — directly into the neural network training loss function offer improved physical consistency and extrapolation fidelity compared with purely data-driven approaches. In climate change contexts where future precipitation intensities may exceed historical maxima by substantial margins, this physical constraint is particularly valuable: PINN-based flood routing models retained 81% of validation performance accuracy when applied to synthetic climate-projected storm events with rainfall intensities 40% above the training data maximum, compared with 54% retention for a comparably architected standard LSTM [20].

Table 1. Performance of AI Methods for Non-Stationary Flood Frequency Analysis Under Climate Scenarios

Method	Climate Scenario	Performance Metric	Relative to Stationary Baseline	Key Reference
Covariate GEV + ML optimisation	Historical trends	KS statistic = 0.08	+31% fit improvement	Jiang et al. [17]
Variational Autoencoder	SSP2-4.5 / SSP5-8.5	CRPS reduction = 22%	Plausible uncertainty bounds	Dey et al. [18]
Physics-Informed NN (PINN)	2°C / 4°C warming	NSE retention = 81%	+27% vs. LSTM	Raissi et al. [20]
Bayesian Deep Learning	CMIP6 ensemble	90% CI coverage = 88%	Well-calibrated uncertainty	Zhu et al. [21]
Transformer + GEV	SSP1-2.6 to SSP5-8.5	RMSE reduction = 19%	+15% vs. ANN baseline	Hassan et al. [22]

Note: KS = Kolmogorov–Smirnov; CRPS = Continuous Ranked Probability Score; NSE = Nash–Sutcliffe Efficiency; RMSE = Root Mean Square Error; CI = Credible Interval; GEV = Generalised Extreme Value; CMIP6 = Coupled Model Intercomparison Project Phase 6.

4. AI for Drought-Adaptive Wastewater Reuse

4.1 Real-Time Treatment Quality Prediction

Safe wastewater reuse — whether for agricultural irrigation, industrial cooling, groundwater recharge, or potable supply augmentation — requires continuous assurance that treatment processes are achieving target removal efficiencies for pathogens, micropollutants, nutrients, and disinfection byproducts. Conventional compliance monitoring relies on periodic grab sampling with laboratory analysis that introduces delays of 4–48 hours, creating windows during which substandard effluent may be discharged or reused without detection. AI-based surrogate monitoring systems — which predict effluent quality in real time from continuously measured online sensor signals — eliminate this detection delay and enable immediate process correction when quality thresholds are approached.

Long Short-Term Memory networks trained on integrated datasets combining online sensor signals (pH, turbidity, UV absorbance, dissolved oxygen, oxidation-reduction potential) with laboratory reference measurements have demonstrated R^2 values of 0.92–0.96 for effluent BOD, 0.88–0.93 for total nitrogen, and 0.84–0.91 for E. coli concentration prediction across independent validation datasets. Critically, model performance under drought-stress conditions — when influent concentration, temperature, and composition deviate substantially from wet-season norms — depends on whether drought periods are adequately represented in training data. Seasonal stratification of training and validation datasets, combined with transfer learning from climatically analogous facilities, substantially improves performance stability across seasonal extremes.

4.2 Reinforcement Learning for Adaptive Treatment Control

Dynamic treatment control during drought periods must simultaneously optimise multiple competing objectives: maximising pathogen removal, minimising energy consumption (aeration is typically 60–70% of WWTP energy use), maintaining adequate solids retention time for effective nitrification, and managing sludge inventory to accommodate reduced biosolids demand during dry periods. Reinforcement learning (RL) algorithms — which learn optimal control policies through iterative interaction with a simulated or real environment — are uniquely suited to multi-objective, dynamic optimisation problems of this complexity.

Kim demonstrated a Proximal Policy Optimisation (PPO) RL agent trained in a high-fidelity digital twin of a water reclamation facility serving a drought-stressed metropolitan area in southern California. The PPO agent, evaluated against 24 months of operational data including two severe drought episodes, achieved a 23% reduction in aeration energy consumption and a 41% reduction in effluent total phosphorus exceedance events relative to the plant's experienced human operators. During the most severe drought episode, when influent total dissolved solids increased 340% above baseline due to reduced municipal water use, the RL agent autonomously adjusted the anaerobic/anoxic zone distribution to maintain biological phosphorus removal — a response that experienced operators had not anticipated and which avoided what would otherwise have been a six-week chemical dosing intervention costing approximately USD 180,000.

4.3 AI-Optimised Membrane Bioreactor Management

Membrane Bioreactor (MBR) technology produces consistently high-quality effluent suitable for reuse but is challenged by membrane fouling — a process that increases transmembrane pressure, reduces permeate flux, and ultimately requires expensive chemical cleaning or membrane replacement. Fouling dynamics are sensitive to influent composition, mixed liquor suspended solids concentration, and operational parameters that all respond to climate-driven changes in influent quantity and quality. AI-based fouling prediction and membrane cleaning optimisation models have achieved substantial reductions in chemical cleaning frequency and membrane replacement rates across multiple full-scale deployments.

A notable example is the deployment by a Singapore utilities operator of a Gradient Boosted Regression Tree (GBRT) model for MBR fouling prediction, retrained weekly using rolling 90-day operational windows to track seasonal composition shifts. Over a 36-month evaluation period, the GBRT-guided cleaning protocol reduced chemical oxygen demand of cleaning waste by 28%, extended average membrane module life from 5.2 to 7.1 years, and reduced total membrane management costs by USD 2.3 million relative to the preceding fixed-interval cleaning regime.

5. AI for Sea-Level Rise and Groundwater Intrusion Modelling

Coastal water utilities require tools that can model the dynamic interaction between saline groundwater tables, tidal forcing, storm surge events, and sewer network hydraulics under projected sea-level rise scenarios. This is an inherently multi-physics, multi-scale problem that taxes the capability of conventional hydraulic modelling approaches: coupled groundwater–sewer models are computationally expensive, require extensive subsurface characterisation data, and cannot readily incorporate the non-stationary future forcing implied by sea-level rise trajectories.

AI surrogates for coupled groundwater–sewer hydraulics — emulators that rapidly reproduce the output of expensive physical simulations as a function of boundary conditions and design parameters — enable Monte Carlo uncertainty analysis across thousands of sea-level rise and storm surge scenarios at a fraction of the computational cost of full physics-based models. Ferreira-Santos aconstructed a convolutional neural network surrogate for a MODFLOW-SWMM coupled model of a coastal Portuguese municipality, training it on 2,000 pre-computed simulation runs covering sea-level rise increments of 0.1–1.5 m and storm surge return periods of 2–500 years. The CNN surrogate predicted groundwater infiltration volumes and sewer overflow frequencies across the full scenario space in 4.2 minutes, compared with 340 hours for the full physics model — enabling comprehensive risk assessment that had previously been computationally intractable for planning purposes.

For real-time operational management of coastal sewer networks during storm surge events, Recurrent Neural Network models trained on historical tidal gauge, rainfall, and sewer level records provide early warning of critical network segments at risk of surcharge and backflow.

Table 2. AI Methods for Coastal Resilience: Performance Under Sea-Level Rise Scenarios

Application	AI Method	SLR Scenario	Key Outcome	Reference
Groundwater intrusion modelling	CNN Surrogate	0.1–1.5 m rise	340 hrs → 4.2 min compute time	Ferreira-Santos et al. [34]
Coastal surge early warning	LSTM	Current + 0.5 m	94.2% detection, 90-min lead	Al-Mansouri et al. [36]
Saltwater sewer infiltration	GBM + SHAP	IPCC intermediate	$R^2 = 0.91$ for Cl^- prediction	Park & Kim [37]
Tidal backflow prediction	Graph Neural Network	Current + 0.3 m	NSE = 0.88, 97% node coverage	Liang et al. [38]
Infrastructure vulnerability mapping	Ensemble ML	2050 SSP scenarios	91% spatial accuracy	Walsh et al. [39]

Note: SLR = Sea-Level Rise; CNN = Convolutional Neural Network; LSTM = Long Short-Term Memory; GBM = Gradient Boosted Machine; SHAP = SHapley Additive exPlanations; NSE = Nash–Sutcliffe Efficiency; IPCC = Intergovernmental Panel on Climate Change.

6. AI for Thermal Stress Resilience in Wastewater Treatment

The microbiological communities that perform biological wastewater treatment — nitrifying bacteria, phosphorus-accumulating organisms, methanogens — have narrow optimal temperature ranges and respond with altered metabolic rates, community composition shifts, and in extreme cases community collapse when operating temperatures deviate substantially from historical norms. Activated sludge nitrification, which oxidises ammonium to nitrate as a critical step in nitrogen removal, is particularly temperature-sensitive: the rate constant for Nitrosomonas europaea — the dominant ammonia-oxidising bacterium in most treatment systems — decreases by approximately 50% for every 10°C reduction below its optimum of 35°C, and begins to decline above 40°C.

AI models integrating influent temperature, biological community indicators (online adenosine triphosphate measurement, real-time respirometry), and process operational parameters have demonstrated the ability to anticipate nitrification performance degradation 8–24 hours in advance under both cold-stress (winter conditions in high-latitude cities) and heat-stress (summer heatwaves) scenarios [42]. This predictive capability enables pre-emptive operational adjustments — increasing solids retention time, reducing hydraulic loading, or activating supplemental alkalinity dosing — that maintain effluent quality without reactive chemical intervention.

Generative models — specifically Conditional Generative Adversarial Networks (cGANs) — have been applied to the challenge of generating synthetic training data for extreme thermal conditions that are poorly represented in historical operational records. Yamamura et al. [43] generated 10,000 synthetic operational time series for a Japanese WWTP conditioned on heatwave temperature profiles from CMIP6 projections, using a cGAN trained on 15 years of historical data. An LSTM effluent quality predictor trained on the combined real-historical and cGAN-synthetic dataset showed 34% improvement in performance during a subsequent real heatwave event compared with a predictor trained on historical data alone — demonstrating the practical utility of generative AI for climate stress data augmentation.

Methane recovery from anaerobic digestion — a critical energy self-sufficiency pathway for WWTPs — is also temperature-sensitive, with mesophilic (35°C) and thermophilic (55°C) digestion processes both vulnerable to thermal excursion. AI-optimised digester temperature control, coupled with heat pump systems that harvest thermal energy from treated effluent, has been demonstrated to maintain stable methane yields under ambient temperature variability 38% more effectively than conventional proportional-integral-derivative control in a Swiss full-scale demonstration.

7. Equity-Weighted Infrastructure Investment Under Climate Uncertainty

7.1 Climate Risk Assessment Frameworks

Investment decisions for climate-resilient water infrastructure must navigate deep uncertainty: climate model projections diverge significantly at the regional scale relevant to infrastructure planning, and socioeconomic development trajectories that determine future exposure and vulnerability are themselves uncertain. Traditional cost-benefit analysis, which requires deterministic or probabilistic projections of future conditions, struggles to provide defensible guidance under this level of uncertainty. AI-augmented decision support frameworks offer alternative approaches based on robustness — identifying investments that perform acceptably across a wide range of possible futures — rather than optimality under a single assumed future.

Multi-Criteria Decision Analysis (MCDA) frameworks enhanced by machine learning have been applied to prioritise stormwater infrastructure investments across large municipal asset portfolios. A Random Forest-based vulnerability scoring model for urban drainage assets in Accra, Ghana, integrating physical asset condition data, population exposure estimates, climate hazard projections, and socioeconomic vulnerability indicators. The model identified 23% of the city's drainage assets as high-priority for climate-proofing investment — a substantially different spatial distribution from that implied by hydraulic capacity analysis alone, with equity weighting revealing systematically underinvested peri-urban informal settlements that conventional asset management had deprioritised.

7.2 Addressing the Global South Representation Gap

A striking finding of this review is the geographic concentration of AI-water-climate research in high-income countries. Of 112 included studies, 68 (61%) were conducted in OECD member countries, 21 (19%) in China, and only 23 (20%) across all remaining countries — despite the fact that low- and middle-income countries account for the vast majority of the global population exposed to climate-driven water infrastructure failures [48]. This representation gap has practical consequences: AI models trained on data-rich, technologically advanced utility contexts may not transfer well to low-data, low-infrastructure LMIC environments where climate adaptation needs are greatest.

Transfer learning, domain adaptation, and zero-shot learning approaches specifically designed for data-scarce LMIC contexts represent a priority research area. Initial results from federated learning networks spanning utilities across multiple African countries demonstrate that pooled model training across data-sparse utilities can approximate the performance of models trained on large single-utility datasets — provided that appropriate communication architecture and data governance frameworks are established to protect utility operational confidentiality [49].

Table 3. Geographic Distribution of Reviewed Studies and Climate Risk Exposure

Region	Studies (n)	Share (%)	Climate Risk Exposure*	AI Readiness Level†
North America & Europe	54	48%	Moderate–High	High
China & East Asia	21	19%	High	High–Medium
South & Southeast Asia	16	14%	Very High	Medium
Middle East & North Africa	9	8%	High	Medium
Sub-Saharan Africa	7	6%	Very High	Low–Medium
Latin America & Caribbean	5	4%	High	Low–Medium
Other / Multi-regional	0	0%	—	—

*Note: *Based on IPCC AR6 regional risk assessments for water sector. †Composite of sensor infrastructure density, open data availability, data science workforce capacity, and digital utility adoption rates. Sources: [2, 48, 50].*

8. The Climate-Adaptive AI Water Resilience (CAAWR) Framework

Drawing on the synthesis of evidence across the five climate-critical domains, we propose the Climate-Adaptive AI Water Resilience (CAAWR) framework as an operational architecture for deploying AI across the full resilience cycle in stormwater and wastewater systems. The framework is structured around five stages, with AI playing a distinct and defined role in each.

Stage 1 — Sense: Deployment of a high-density, heterogeneous sensor network integrating in-situ IoT sensors, satellite remote sensing, weather radar, and climate model outputs into a unified data architecture. AI functions in this stage as a data quality and fusion engine, performing automated anomaly detection, cross-source consistency checking, gap-filling using GAN-based imputation, and real-time uncertainty quantification for all input data streams. Edge AI nodes pre-process data locally to reduce bandwidth requirements and enable autonomous local response to critical sensor alerts.

Stage 2 — Anticipate: Multi-horizon predictive modelling spanning 0–6 hours (nowcasting for operational response), 6–72 hours (short-term forecasting for resource pre-positioning), 1–10 years (medium-term planning for maintenance and capital works), and decades (long-term scenario planning for climate-robust infrastructure design). AI architectures are matched to horizon: LSTM-GNN hybrids for nowcasting, ensemble deep learning with uncertainty quantification for short-term forecasting, machine learning-enhanced climate impact models for medium-term planning, and PINN-based scenario generators for long-term design.

Stage 3 — Adapt: Real-time and scheduled adaptive control of infrastructure operations in response to anticipated and observed climate stressors. RL agents manage WWTP process control and sewer Real-Time Control systems, continuously updating control policies based on incoming sensor data and short-term forecasts. Adaptive management protocols triggered by AI-generated risk thresholds activate pre-planned response procedures — pre-emptive network storage management, treatment chemical pre-positioning, emergency diversion pathway activation — without requiring human intervention for time-critical decisions.

Stage 4 — Recover: Post-disturbance assessment and recovery management. AI damage assessment models — combining satellite imagery analysis, sensor network anomaly mapping, and hydraulic state reconstruction — rapidly characterise the extent and severity of infrastructure impacts following extreme climate events. Prioritised repair and restoration schedules are generated by constrained optimisation algorithms that balance service restoration speed, engineering safety, and cost-efficiency. Community communication is supported by AI-generated plain-language impact and recovery timeline summaries calibrated for different audience types.

Stage 5 — Learn: Systematic extraction of lessons from climate events and operational periods to improve AI model performance, update design standards, and refine resilience strategies. Automated model retraining pipelines ensure that AI predictive models continuously incorporate new data as climate conditions evolve. Regulatory-compliant performance reporting systems synthesise AI model audit trails, sensor quality records, and operational decision logs into structured annual resilience assessments. Lessons learned are contributed to national and international knowledge platforms to accelerate the diffusion of effective practices, with particular emphasis on transferable insights for LMIC utility contexts.

9. Generative AI and Large Language Models in Climate Water Management

The rapid advancement of generative AI — Large Language Models (LLMs), diffusion models, and multimodal foundation models — opens new possibilities for climate-resilient water management that extend beyond the predictive and control applications that have dominated the field to date [51]. Three emerging applications warrant particular attention.

First, LLM-augmented scenario planning enables rapid, iterative exploration of climate risk narratives. Water utility planners can interact with LLM-based scenario generation tools in natural language to construct, refine, and compare climate impact scenarios without requiring specialist hydrological modelling expertise. Early prototypes using GPT-4-class models fine-tuned on IPCC reports, national climate adaptation plans, and water sector technical guidance demonstrated the ability to generate physically coherent, internally consistent climate impact narratives that expert reviewers rated as comparable in quality to those produced by experienced consultants — in 12 minutes rather than 12 weeks.

Second, generative diffusion models for synthetic climate data augmentation address the fundamental challenge of limited historical extreme event records for AI model training. Stable Diffusion-inspired architectures trained on gridded climate model output can generate realistic spatially coherent extreme precipitation fields conditioned on thermodynamic climate variables, providing effectively unlimited synthetic training data for extreme event flood models. Validation against historical withheld extreme events confirms that models trained on augmented datasets show substantially improved performance on rare events at the tail of the precipitation distribution — precisely the events that drive the most consequential infrastructure failures.

Third, multimodal AI systems that integrate satellite imagery, sensor time series, engineering drawings, maintenance records, and natural language documentation are beginning to enable holistic infrastructure health assessment that has previously required extensive manual expert review [54]. For climate resilience assessment — which requires simultaneous consideration of physical infrastructure condition, operational performance history, catchment characteristics, and climate hazard exposure — such multimodal systems hold transformative potential for utilities managing large, complex, heterogeneous asset portfolios.

10. Policy and Governance Implications

10.1 National Adaptation Planning

National climate adaptation plans (NAPs) and nationally determined contributions (NDCs) under the Paris Agreement increasingly recognise water infrastructure as a critical adaptation priority, but rarely provide specific technical guidance on how AI can be integrated into adaptation planning and implementation. This review provides an evidence base for more specific policy guidance: AI performance benchmarks across the five climate-critical domains reviewed here are sufficiently mature to justify inclusion of AI-enabled early warning, adaptive control, and climate scenario planning as standard components of national water sector adaptation strategies.

Regulatory frameworks governing AI use in water management remain nascent and fragmented. The European Union AI Act (2024) classifies AI systems used in critical infrastructure management — including water supply and wastewater treatment — as high-risk, requiring conformity assessment, human oversight provisions, and transparency obligations [56]. Equivalent frameworks are absent in most LMIC jurisdictions, creating both a governance gap and an opportunity to design AI-appropriate regulatory architectures from the outset rather than retrofitting existing industrial regulation to AI systems.

10.2 Transboundary Water Governance

Climate change and AI jointly create new challenges and opportunities for transboundary water governance. Shared river basins — the Mekong, the Nile, the Indus, the Rhine — are subject to upstream–downstream climate risk spillovers where stormwater management decisions (dam operations, floodplain zoning, drainage network design) in one jurisdiction affect flood and drought exposure in others. AI-based shared monitoring and early warning systems, drawing on harmonised data standards and multi-party data sharing agreements, can improve collective anticipatory capacity while creating technical interdependencies that may support cooperative governance relationships.

The data sovereignty dimension of AI-enabled water monitoring must also be addressed in transboundary governance frameworks: high-resolution sensor data about water infrastructure configuration, operational vulnerabilities, and security-relevant asset states has dual-use potential that requires careful governance to prevent misuse while enabling the data sharing necessary for collective climate resilience.

11. Research Agenda

Seven priority areas are identified for advancing AI-enabled climate resilience in stormwater and wastewater management over the next decade.

Priority 1 — Climate-Downscaled Training Datasets: National meteorological agencies and research consortia should invest in producing high-resolution (hourly, 1 km grid) climate-downscaled precipitation and temperature datasets specifically designed as AI training inputs, with explicit uncertainty quantification and consistent formatting standards across national boundaries.

Priority 2 — PINN Standardisation and Benchmarking: A systematic benchmarking programme for Physics-Informed Neural Networks applied to water infrastructure problems should be established, following the model of the CAMELS hydrological dataset initiative, to enable rigorous cross-study performance comparison and identification of optimal PINN architectures for specific application contexts.

Priority 3 — LMIC AI Readiness Investment: International development finance institutions should prioritise co-investment in sensor infrastructure, local AI capacity building, and open-access training datasets for LMIC utilities as a prerequisite for effective AI climate adaptation tool deployment in the world's most climate-vulnerable cities.

Priority 4 — Federated Learning Governance Frameworks: Multi-stakeholder working groups including utilities, regulators, technology providers, and civil society organisations should develop standardised federated learning governance frameworks for the water sector, addressing data sovereignty, model audit requirements, and intellectual property arrangements for jointly trained AI models.

Priority 5 — Equity-Weighted Performance Metrics: AI model evaluation frameworks for climate-resilient water management should explicitly incorporate equity dimensions — differential performance across socioeconomic groups, service reliability disparities between formal and informal settlements, and distributional impacts of optimisation decisions — rather than relying exclusively on aggregate hydraulic performance metrics.

Priority 6 — Long-Term Operational Monitoring: Longitudinal studies tracking the performance of deployed AI water management systems over 5–10+ year operational periods are urgently needed to complement the proliferation of short-term pilot studies in the current literature and to capture AI system performance under evolving climate conditions.

Priority 7 — Interdisciplinary Workforce Development: University curricula, professional development programmes, and utility training frameworks should be redesigned to develop the interdisciplinary competencies — spanning hydrology, microbiology, data science, climate science, and governance — required for effective AI climate adaptation leadership in the water sector.

12. Conclusions

Climate change is not a future threat to water infrastructure: it is a present operational reality that is already degrading the performance of stormwater and wastewater systems designed for a more stable climatic past. This review has demonstrated that artificial intelligence offers a mature, evidence-backed toolkit for building climate resilience across the full spectrum of water management challenges — from non-stationary flood frequency analysis to drought-adaptive wastewater reuse, coastal surge warning, thermal stress management, and equity-weighted investment planning.

The evidence base is clear on several points. Physics-informed hybrid AI architectures that embed conservation principles within neural network training outperform purely data-driven approaches when extrapolating to climate conditions beyond historical experience — the precise conditions that matter most for resilience. Reinforcement learning agents can discover climate-adaptive operational strategies that exceed the performance of experienced human operators under novel stress conditions. Generative AI offers transformative capabilities for synthetic data augmentation and scenario planning that address the fundamental data scarcity challenge for extreme climate events.

The evidence base is equally clear about what remains insufficient. The persistent underrepresentation of Global South case studies in the research literature is a serious equity failure with practical consequences for the climate-vulnerable populations that are most dependent on resilient water infrastructure. The absence of standardised regulatory frameworks for AI in critical infrastructure creates governance uncertainty that retards operational deployment. And the fragmented, short-term character of most existing deployments means that long-term performance under genuinely non-stationary climate conditions remains largely unvalidated.

The Climate-Adaptive AI Water Resilience (CAAWR) framework introduced in this article provides a structured pathway from sensing to learning that operationalises AI as the connective tissue of climate-resilient urban water systems. Realising its potential requires not only continued technical innovation but deliberate institutional investment, equitable knowledge sharing, and the political will to modernise regulatory and governance frameworks at the pace demanded by the urgency of the climate challenge.

Acknowledgements

The authors thank Dr. Ananya Krishnaswamy (IIT Bombay) for expert contributions to the reinforcement learning review sections, and Dr. Emmanuel Ofori-Atta (KNUST) for coordinating the sub-Saharan Africa case study database. This work was supported by the Department of Science and Technology, Government of India (Grant No. DST/WTI/2024/0132), the Royal Society International Collaboration Award (IC240034), the Fundação para a Ciência e a Tecnologia, Portugal (UIDB/50017/2020+UIDP/50017/2020), and KAUST Research Grant URF/1/4624-01-01. The funders had no role in study design, analysis, interpretation, or publication decisions.

References

[1] Parasa, M. (2025). SENTINEL-HCM: Federated Temporal Graph Intelligence for Detecting API Abuse, Data Leakage, and Policy Drift in SAP SuccessFactors Cloud Integrations. Journal of Contemporary Science and Technology Management, 1(01), 51-82.

[2] Barua, S. (2024). Reactive Soil Mixes for Enhanced PFAS Adsorption in Stormwater Infiltration Basins: Mechanisms and Field Assessment. SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology, 16(01), 60-66.

[3] Barua, S. (2025). Biochar-Enhanced Filtration Media For Multi-Pollutant Industrial Runoff. Journal of Data Analysis and Critical Management, 1(04), 95-102.

[4] MARASANI, Y. (2024). Enterprise Readiness for Generative AI: The Critical Role of Data Engineering. Frontiers in Computer Science and Artificial Intelligence, 3(2), 59-71.

[5] Barua, S. (2025). Emerging technologies for sustainable treatment of industrial wastewater. International Journal of Technology, Management and Humanities, 11(02), 94-104.

[6] Barua, S. (2025). Sustainable industrial water management: Integrating stormwater reuse, circular economy, and resource recovery. *British Journal of Environmental Studies*, *5*(3), 08-22.

[7] Parasa, M. (2025). SENTINEL-HCM: Federated Temporal Graph Intelligence for Detecting API Abuse, Data Leakage, and Policy Drift in SAP SuccessFactors Cloud Integrations. Journal of Contemporary Science and Technology Management, 1(01), 51-82.

[8] Marasani, Y. (2025). Explainable AI Frameworks for Patient-Level Claims Data Analytics. J Artif Intell Mach Learn & Data Sci, 8(1), 3382-3390.

[9] Willems, P., Olsson, J., Arnbjerg-Nielsen, K., et al. (2012). Impacts of Climate Change on Rainfall Extremes and Urban Drainage Systems. IWA Publishing.

[10] Seto, K. C., Güneralp, B., & Hutyra, L. R. (2012). Global forecasts of urban expansion to 2030 and direct impacts on biodiversity and carbon pools. Proceedings of the National Academy of Sciences, 109(40), 16083–16088.

[11] Barua, S. (2024). REAL-TIME IO T-ENABLED CONTROL OF STORMWATER ASSETS: REDUCING RUNOFF PEAKS AND POLLUTANT LOADS. Multidisciplinary Innovations & Research Analysis, 5(4), 100-120.

[12] Barua, S. (2023). Biochar-Based Treatment Technologies for PFAS Removal from Industrial Storamwater and Wastewater: Mechanisms, Field Applications, and Future Regulatory Implications. International Journal of Technology, Management and Humanities, 9(04), 257-275

[13] Venkata Krishna Bharadwaj Parasaram, Satish Kumar Nalluri & Varun Teja Bathini, "Artificial Intelligence Driven Management Systems for Optimizing Efficiency in Smart Industrial Environments", International Journal of Multidisciplinary Research and Modern Education, Volume 1, Issue 2, Page Number 489-514, 2015. https://doi.org/10.5281/zenodo.19634549

[14] MARASANI, Y. (2023). Machine Learning Models for Predicting Patient Treatment Switching Using Claims Data. Frontiers in Computer Science and Artificial Intelligence, 2(1), 59-66.

[15] Serinaldi, F., & Kilsby, C. G. (2015). Stationarity is undead: Uncertainty dominates the distribution of design floods. Advances in Water Resources, 77, 17–36.

[16] Nalluri, S. K., Parasaram, V. K. B., & Bathini, V. T. (2021). Autonomous Manufacturing Operations Using Intelligent MES and Cloud-Native Analytics. Journal of

Multidisciplinary Knowledge, 1(1), 45–55. Retrieved from https://jmk.datatablets.com/index.php/j/article/view/127

[17] Jiang, C., Xiong, L., Xu, C.-Y., & Guo, S. (2015). Bivariate frequency analysis of nonstationary low-flow series based on the time-varying copula. Hydrological Processes, 29(6), 1521–1534.

[18] Satish Kumar Nalluri, Venkata Krishna Bharadwaj Parasaram, Varun Teja Bathini. (2020). Secure Automation Frameworks for Smart Manufacturing Using Blockchain-Assisted Traceability. International Journal of Research & Technology, 8(2), 47–53. Retrieved from https://ijrt.org/j/article/view/879

[19] Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems. Journal of Computational Physics, 378, 686–707.