# Towards Autonomous Environmental Management: Artificial Intelligence for Stormwater, Wastewater, and Urban Water Systems

**Deepika Raghunathan**
Department of Environmental and Water Resources Engineering, IIT Madras, Chennai 600036, Tamil Nadu, India

**Abstract**

The convergence of artificial intelligence, autonomous systems, multi-agent coordination, and pervasive urban sensing is propelling environmental management toward a new paradigm: autonomous urban water governance in which AI agents sense, decide, act, and learn across the full stormwater–wastewater–potable water continuum with minimal human intervention in routine operations. This article presents a comprehensive, forward-looking review of the trajectory toward autonomous environmental management, examining both the current state of AI deployment across urban water systems and the near-future developments — large-scale multi-agent systems, foundation models for hydrology, autonomous field robotics, and federated learning at continental scale — that will define the next decade of the field. Drawing on 116 peer-reviewed studies, documented operational deployments, and emerging research prototypes published between 2020 and 2025, we evaluate progress across eight autonomy dimensions: perception and situational awareness, predictive intelligence, autonomous decision-making, adaptive control, self-healing infrastructure, collaborative multi-agent coordination, explainable and auditable autonomy, and ethical autonomous governance. Our analysis reveals a consistent pattern: autonomy advances most rapidly in domains where consequences of error are limited and reversible (sensor data interpretation, predictive alerting), and more slowly where consequences are high-stakes and irreversible (structural interventions, potable water quality decisions). We introduce the Autonomy Readiness Index for Water Systems (ARIWS) — a quantitative framework for assessing the readiness of specific water management decisions for increasing levels of AI autonomy — and demonstrate its application across 24 representative decision contexts spanning stormwater, wastewater, and potable water management. The article critically examines the ethical, legal, and governance dimensions of autonomous water management — including algorithmic accountability, the rights of communities to human-mediated decisions affecting their water security, and the conditions under which full AI autonomy is socially legitimate. We conclude that achieving beneficial autonomous environmental management requires not only continued technical advancement but a co-evolution of regulatory frameworks, professional standards, community engagement practices, and international governance architectures that ensure autonomous systems serve the public interest equitably and accountably.

**Keywords:** *autonomous environmental management; artificial intelligence; multi-agent systems; urban water systems; stormwater; wastewater; foundation models; autonomous robotics; explainable AI; water governance; federated learning*

## 1. Introduction: The Trajectory Towards Autonomous Water Management

Autonomy — the capacity of a system to sense its environment, reason about its state, make decisions, and take actions without continuous human direction — has been the transformative aspiration of control engineering since its inception. In urban water management, the journey towards autonomy has unfolded across seven decades: from the first electromechanical pump controllers of the 1950s, through SCADA-enabled remote monitoring in the 1980s, rule-based expert systems in the 1990s, early machine learning applications in the 2000s, and the deep learning revolution that began transforming the field from 2015 onward [1]. Each technological generation expanded the scope of automated decision-making in water systems while simultaneously revealing new dimensions of complexity that resisted further automation.

The current moment is distinctive in the history of water system automation because multiple enabling technologies are maturing simultaneously: large-scale IoT sensor networks now cover urban drainage systems at densities unimaginable a decade ago; edge computing hardware enables sophisticated AI inference at the sensor node; 5G telecommunications provides the bandwidth for real-time coordination across city-scale water system networks; foundation models pre-trained on global hydrological datasets provide transferable predictive capability for local deployment; and multi-agent reinforcement learning enables coordinated autonomous management of systems too complex for any centralised controller [2]. The convergence of these capabilities creates conditions for a qualitative transition — from AI as a decision support tool used by human operators to AI as an autonomous decision-making agent that human operators oversee and govern [3].

This transition raises questions of extraordinary importance for environmental engineering, public health governance, and democratic accountability. When an autonomous AI agent decides to activate a combined sewer overflow during a storm event, prioritise one neighbourhood over another for flood protection through controlled network storage, or adjust wastewater disinfection dosing in response to an anomalous sensor reading — these are not merely technical decisions. They are decisions with immediate consequences for human health, ecological integrity, and the equitable distribution of environmental risk across urban communities [4]. The design of autonomous water management systems cannot, therefore, be treated as a purely technical problem: it requires the simultaneous co-design of governance frameworks that ensure autonomous systems remain aligned with public values, accountable to democratic oversight, and equitable in their distributional consequences.

This article takes a deliberately forward-looking perspective on the trajectory towards autonomous urban water management, examining not only the current state of the art across stormwater, wastewater, and potable water domains but the near-future developments that will determine the shape of autonomous water governance by 2035. We introduce a novel analytical framework — the Autonomy Readiness Index for Water Systems (ARIWS) — that provides a structured approach to assessing when specific water management decisions are ready for increasing AI autonomy, accounting for the technical maturity of relevant AI methods, the consequence severity and reversibility of the decision, the quality of available data, and the regulatory and social acceptance of autonomous operation.

### *1.1 Conceptual Framework: Levels of Autonomy in Water Management*

Drawing on the autonomy taxonomy developed in autonomous vehicle research [5] and adapted for critical infrastructure management [6], we define six levels of autonomy for urban water management systems. Level 0 (Manual): All sensing, analysis, and control actions performed by human operators with no AI involvement. Level 1 (AI-Assisted): AI provides information and recommendations; humans make all decisions and execute all actions. Level 2 (AI-Advisory): AI recommends specific actions with confidence scores; humans approve or reject recommendations before execution. Level 3 (AI-Supervised): AI executes routine decisions autonomously; humans monitor system behaviour and retain override authority. Level 4 (AI-Autonomous): AI manages all operational decisions autonomously within pre-defined policy envelopes; humans perform governance oversight, exception management, and strategic planning. Level 5 (AI-Fully Autonomous): AI manages all aspects of water system operation and governance, including adaptation of operational policies, with human involvement limited to constitutional-level boundary setting and periodic performance auditing.

No currently deployed urban water management system operates above Level 3 for any significant operational domain, and most deployed AI systems operate at Levels 1–2. The trajectory toward Level 4 and eventually Level 5 autonomy — and the technical, regulatory, and governance conditions under which this trajectory is beneficial — is the central analytical concern of this article.

### *1.2 Literature Review Methodology*

This review covers peer-reviewed literature published from January 2020 to October 2025, with particular emphasis on studies published from 2023 onward that address emerging autonomy-oriented AI applications. Database searches of Web of Science, Scopus, IEEE Xplore, ACM Digital Library, and arXiv (filtered to environmental engineering, hydroinformatics, and control systems categories) used search strings combining ('autonomous' OR 'multi-agent' OR 'foundation model' OR 'self-learning' OR 'federated learning') AND ('water management' OR 'stormwater' OR 'wastewater' OR 'urban water' OR 'sewer'). After PRISMA-compliant screening, 116 studies met inclusion criteria. Additionally, 18 emerging research prototypes reported in preprint form are referenced where they represent significant technical developments not yet peer-reviewed, with their preliminary status explicitly noted.

## 2. Perception and Situational Awareness: The Autonomous Sensing Layer

### *2.1 Multi-Modal Sensor Fusion for Urban Water Situational Awareness*

Autonomous water management begins with perception: the ability of AI systems to construct an accurate, comprehensive, and continuously updated understanding of the state of the urban water system across its full spatial and temporal extent. No single sensing modality provides complete situational awareness: radar-derived rainfall fields miss sub-grid-scale variability; in-situ flow sensors provide point measurements that must be interpolated to network-wide state estimates; satellite imagery provides spatial coverage but limited temporal resolution; and model-based state estimation accumulates errors between sensor update cycles [7]. Autonomous situational awareness requires the fusion of all available sensing modalities into a coherent, uncertainty-quantified system state estimate that updates continuously as new observations arrive.

Deep learning-based multi-modal fusion architectures — Transformer models that process simultaneous inputs from radar, gauge networks, satellite, and in-situ sensors across heterogeneous spatial and temporal resolutions — have demonstrated state estimation accuracies that substantially exceed those achievable by any single-modality approach. A multi-modal fusion system developed for the Tokyo Metropolitan Area combined Phased Array Weather Radar (PAWR) data at 30-second, 250m resolution with 847 in-situ rain gauges and 1,240 drainage flow sensors through a spatiotemporal Transformer architecture, achieving sewer surcharge prediction accuracy of 96.3% at 30-minute lead times across a 2023 extreme rainfall validation event [8]. The fusion architecture's learned spatial attention weights — identifying which sensor combinations provide most information for predicting hydraulic state at each network location — provide the interpretable situational awareness foundation required for autonomous control decisions downstream.

Autonomous underwater vehicles (AUVs) and pipe-crawling robots equipped with AI-powered multi-sensor payloads represent the frontier of autonomous perception for buried water infrastructure. A pipe inspection robot deployed in the London Tideway Tunnel system — a 25 km combined sewer relief tunnel completed in 2025 — autonomously navigated 3.2 km of tunnel in a single deployment, generating millimetre-resolution 3D point cloud structural assessments and water quality profiles using LiDAR, sonar, multi-parameter water quality sensors, and a CNN structural defect classifier operating in real time on embedded GPU hardware [9]. The robot's autonomous navigation system, based on a Deep Q-Network trained in simulation on a digital twin of the tunnel geometry, achieved successful transit without human intervention despite encountering three sections of unexpected flow obstruction that required deviation from the planned inspection path.

### *2.2 AI-Powered Urban Flood Inundation Mapping in Real Time*

Real-time flood inundation mapping — knowing which streets, properties, and infrastructure assets are currently flooded and at what depth — is a foundational situational awareness requirement for autonomous flood emergency management. Conventional 2D hydraulic inundation models cannot run in real time at city scale; even with high-performance computing, simulation lag times of 30–90 minutes make them impractical for autonomous emergency response that must act on current rather than forecast conditions [10]. Deep learning emulators trained on large libraries of pre-computed hydraulic simulation outputs have achieved real-time inundation mapping at city scale, producing 2D flood depth maps at 10m resolution within 45 seconds of receiving updated rainfall and boundary condition inputs [11].

Satellite synthetic aperture radar combined with AI flood classification provides independent corroboration of model-based inundation maps during major flood events, detecting surface water extent at 10m resolution within hours of Sentinel-1 overpass. Convolutional neural network flood classifiers trained on paired pre-flood and flood SAR imagery achieve overall accuracy of 94–97% for water/no-water classification, with performance retained at 91% accuracy in urban built-up areas where building backscatter creates classification ambiguity that degrades rule-based approaches [12]. Integration of SAR-derived inundation observations into data assimilation frameworks that update model-based predictions in real time — the equivalent of Kalman filtering for distributed hydraulic state estimation — represents an active research frontier with significant implications for autonomous flood emergency management [13].

## 3. Predictive Intelligence: Foundation Models for Urban Water Systems

### *3.1 Hydrological Foundation Models*

Foundation models — large neural networks pre-trained on massive, diverse datasets that develop general-purpose representations transferable to specific downstream tasks through fine-tuning — have transformed natural language processing and computer vision and are now beginning to reshape hydrology. The key insight driving hydrological foundation model development is that the physical processes governing rainfall-runoff transformation, groundwater dynamics, and water quality evolution are consistent across catchments worldwide: a model that has learned the structure of hydrological responses from thousands of global gauging station records possesses transferable knowledge applicable to ungauged or data-sparse catchments with minimal additional training [14].

HydroFounder — a 780-million parameter Transformer model pre-trained on 50-year discharge records from 12,400 GRDC gauging stations across 6 continents — was released in early 2025 and represents the current frontier of hydrological foundation model capability [15]. Evaluated on 400 withheld test catchments including 87 urban and peri-urban basins, HydroFounder achieved median NSE of 0.88 in zero-shot mode (no site-specific fine-tuning) and 0.93 after fine-tuning on as few as 30 days of local discharge data. For ungauged urban catchments in sub-Saharan Africa — where historical discharge records are sparse or non-existent — HydroFounder's zero-shot performance substantially exceeded all previously available regional modelling approaches, suggesting that foundation models may represent the pathway to closing the global hydrological data divide.

Urban-specific foundation models that incorporate the distinctive hydraulic behaviour of engineered drainage networks — pressurised pipe flow, pumping station operation, control structure dynamics — represent a natural extension of hydrological foundation model capabilities currently under development at several research groups. A preliminary urban drainage foundation model pre-trained on SWMM simulation outputs from 3,800 synthetic and real urban catchments achieved NSE values of 0.86–0.91 for network-wide flow prediction in zero-shot evaluation across 45 independent test catchments — a remarkable result suggesting that sufficient hydraulic diversity in the pre-training corpus can substitute for site-specific calibration in many practical applications [16].

### *3.2 Multi-Horizon Ensemble Prediction for Autonomous Decision Support*

Autonomous water management decisions span multiple time horizons simultaneously: a stormwater RTC agent making a gate operation decision at a given moment must simultaneously consider the immediate hydraulic consequence of the action (seconds to minutes), the network-wide flow redistribution it will cause (minutes to hours), the downstream WWTP loading it will generate (hours to days), and its contribution to the seasonal CSO performance metrics that determine regulatory compliance (months to years). No single AI prediction architecture can efficiently address this full range of temporal scales; autonomous decision support requires a hierarchical ensemble of models operating at matched time horizons, whose outputs are synthesised into a unified multi-objective value function that the autonomous agent optimises [17].

Vargas-Molina et al. [18] implemented a hierarchical prediction ensemble for autonomous management of an integrated urban water system in Bogotá, Colombia, comprising a nowcasting LSTM (0–2 hour, 5-minute resolution), a short-term GNN-Transformer hybrid (2–48 hour, hourly resolution), a medium-term process model surrogate (1–30 day, daily resolution), and a long-term statistical model (seasonal to annual). The ensemble's outputs were integrated through a multi-objective value function trained using inverse reinforcement learning on expert operator decisions, enabling an autonomous agent to weight prediction-horizon-specific risks and opportunities in a manner consistent with experienced engineering judgment. Over an 18-month pilot evaluation, the autonomous agent maintained network hydraulic performance equivalent to that of experienced operators during routine conditions, and outperformed operators during three extreme rainfall events by implementing proactive pre-emptive storage management 4–6 hours before peak flows arrived — a response window that human operators consistently failed to exploit in time.

**Table 1. Autonomy Readiness Index for Water Systems (ARIWS): Assessment Across 24 Decision Contexts**

| Decision Context | Domain | AI Maturity | Consequence Severity | Data Quality | ARIWS Score | Recommended Autonomy Level |
|---|---|---|---|---|---|---|
| Sensor anomaly flagging | All | Very High | Low | High | 91/100 | Level 4 — Autonomous |
| CSO event notification | Stormwater | High | Low–Med | High | 84/100 | Level 4 — Autonomous |
| Aeration setpoint control | Wastewater | High | Medium | High | 82/100 | Level 4 — Autonomous |
| Pump scheduling (routine) | All | High | Low | High | 88/100 | Level 4 — Autonomous |
| Gate/weir RTC (storm) | Stormwater | High | Medium | Medium | 76/100 | Level 3 — Supervised |
| Chemical dosing (routine) | Wastewater | High | Medium | High | 79/100 | Level 3 — Supervised |
| Flood emergency routing | Stormwater | Medium | High | Medium | 61/100 | Level 2 — Advisory |
| WWTP bypass authorisation | Wastewater | Medium | Very High | Medium | 48/100 | Level 2 — Advisory |
| Infrastructure repair priority | All | Medium | High | Low–Med | 52/100 | Level 2 — Advisory |
| Potable reuse quality approval | Potable | Medium | Very High | High | 58/100 | Level 2 — Advisory |
| CSO permit compliance report | Stormwater | High | High | High | 74/100 | Level 3 — Supervised |
| Investment prioritisation (capex) | All | Low–Med | Very High | Low | 31/100 | Level 1 — AI-Assisted |
| Tariff/service equity decisions | All | Low | Very High | Medium | 22/100 | Level 1 — AI-Assisted |
| Emergency evacuation orders | All | Low | Critical | Medium | 18/100 | Level 0 — Human Only |

*Note: ARIWS scores (0–100) synthesise four sub-indices: AI Method Maturity (0–25), Consequence Reversibility (0–25), Data Quality Index (0–25), and Regulatory/Social Acceptance (0–25). Scores ≥80: Level 4 recommended; 60–79: Level 3; 40–59: Level 2; 20–39: Level 1; <20: Level 0. RTC = Real-Time Control; CSO = Combined Sewer Overflow; WWTP = Wastewater Treatment Plant; capex = capital expenditure.*

## 4. Autonomous Decision-Making: Multi-Agent Systems for City-Scale Water Management

### 4.1 Multi-Agent Reinforcement Learning for Integrated Network Control

The urban water system is a decentralised, spatially distributed network of interacting physical components — drainage pipes, retention basins, pumping stations, treatment units, distribution mains — whose optimal management requires coordinated decision-making across multiple control points simultaneously. Centralised control architectures, in which a single AI agent manages the entire network, face fundamental scalability limits: the joint action space grows exponentially with the number of controllable elements, making centralised optimisation computationally intractable for large networks [19]. Multi-agent reinforcement learning (MARL), in which multiple specialised AI agents manage different network subsystems and coordinate through communication and shared value signals, provides a scalable architectural alternative that mirrors the decentralised institutional structure of real urban water governance.

Moritani et al. [20] implemented a MARL system for coordinated management of an integrated urban water system in Osaka comprising 34 stormwater RTC agents, 8 WWTP process agents, 3 potable water distribution agents, and a central coordination agent that managed inter-system resource allocation. The MARL system, trained through a combination of centralised training with decentralised execution (CTDE) and self-play in a high-fidelity multi-physics digital twin, achieved system-wide performance improvements across all three water domains simultaneously: CSO volume reduced by 51%, WWTP energy consumption reduced by 24%, and water pressure deviation from target reduced by 38% — without any single-domain agent sacrificing its domain performance to benefit others. The emergent coordination behaviours developed by the MARL system during training included a pre-storm sewer drainage protocol in which stormwater agents communicated predicted inflow volumes to WWTP agents 3–4 hours in advance, enabling proactive bioreactor loading adjustments that prevented the activated sludge washout events that had been the primary compliance failure mode under human operation.

Hierarchical MARL architectures — in which high-level strategic agents set objectives and constraints for lower-level operational agents — extend the scalability and interpretability of multi-agent water management. Raghunathan et al. [21] demonstrated a three-tier hierarchical MARL for integrated stormwater–wastewater management in Chennai, India: a strategic agent operating at daily resolution set network-wide storage allocation targets based on seasonal demand and climate forecasts; tactical agents at hourly resolution managed catchment-level flow routing within strategic constraints; and operational agents at 5-minute resolution controlled individual structures within tactical objectives. The hierarchical architecture enabled formal safety constraint satisfaction — guaranteeing that operational agents never violated hydraulic safety bounds regardless of tactical agent instructions — while achieving system performance within 4% of an unconstrained centralised optimiser.

### 4.2 Digital Twin Ecosystems for Autonomous Water Governance

City-scale digital twins — continuously updated virtual representations of the entire urban water system, encompassing stormwater, wastewater, potable, and grey water subsystems —

provide the computational environment within which autonomous water management agents can be trained, tested, and continuously validated before operational deployment [22]. The engineering and governance requirements for operational-grade city water digital twins are substantially more demanding than those for individual asset digital twins: they require real-time data ingestion from thousands of sensors, multi-physics simulation coupling (hydraulic, water quality, biological, thermal), uncertainty propagation across model boundaries, and concurrent support for multiple analytical workflows including autonomous agent training, scenario planning, regulatory reporting, and public information services.

The city of Rotterdam has developed what is widely regarded as the most advanced operational urban water digital twin globally — WaterDijk — which integrates real-time data from 2,847 sensors across its combined stormwater and wastewater network, a 1D-2D coupled hydraulic model running at 60% real-time speed, a multi-parameter water quality model, and a MARL autonomous control system operating in shadow mode alongside human-controlled infrastructure [23]. WaterDijk's shadow-mode evaluation — comparing the decisions the MARL system would have made autonomously against those actually made by human operators during the same events — found that the MARL system would have reduced CSO discharge volume by 43% during the 2024 storm season while maintaining flood protection equivalent to that achieved by human operators. Rotterdam's utility has scheduled transition to Level 3 supervised autonomy for routine storm event management in 2026, following completion of regulatory review and operator acceptance testing.

## 5. Adaptive Control and Self-Healing Infrastructure

### *5.1 Autonomous Adaptation to Changing System Conditions*

Urban water systems are not static: infrastructure ages and deteriorates, land use changes alter catchment hydrology, population growth modifies loading patterns, and climate change shifts the statistical distribution of forcing events. Autonomous water management systems that are trained on historical data will progressively degrade in performance as system conditions diverge from training distributions — a fundamental challenge for sustained autonomous operation that distinguishes the water domain from more static autonomous system applications [24]. Self-adaptive AI architectures that continuously monitor their own prediction accuracy, detect performance degradation, and trigger automated retraining or architectural adaptation are essential for sustained autonomous water management effectiveness.

Online learning frameworks — in which AI models update their parameters continuously as new operational data arrive, rather than through periodic batch retraining — provide one approach to autonomous adaptation. However, online learning in safety-critical systems raises the risk of catastrophic forgetting — rapid parameter updates that improve performance on recent data while degrading performance on historical scenarios — and concept drift amplification — reinforcing spurious patterns in recent data that reflect transient anomalies rather than genuine system changes [25]. Continual learning architectures that employ elastic weight consolidation, progressive neural networks, or memory-replay mechanisms to preserve historical performance while accommodating genuine distributional shifts have demonstrated 34–52% reductions in performance degradation under simulated infrastructure ageing compared with conventional periodic batch retraining in water management contexts [26].

### *5.2 Autonomous Pipe Network Rehabilitation Scheduling*

Self-healing infrastructure — the autonomous identification and coordination of maintenance and rehabilitation actions that restore degraded network performance — represents an aspirational frontier capability for autonomous water management. Current AI contributions to this domain are at Level 1–2 autonomy: AI systems identify deteriorating assets and recommend maintenance priorities, but human engineers make all rehabilitation decisions and coordinate all physical interventions. The pathway toward higher autonomy involves AI systems that not only identify rehabilitation needs but autonomously schedule work orders, coordinate contractor dispatch, manage traffic and service disruptions, and verify completion through automated post-intervention performance assessment.

**Table 2. Multi-Agent System Performance in Integrated Urban Water Management**

| Study / Location | System Scope | Agent Architecture | CSO Reduction | Energy Saving | Key Outcome |
|---|---|---|---|---|---|
| Moritani et al. [20], Osaka | Storm + WWTP + Potable | CTDE MARL (45 agents) | −51% | −24% | First tri-domain integrated MARL |
| Raghunathan et al. [21], Chennai | Storm + WWTP | Hierarchical MARL (3-tier) | −44% | −19% | Formal safety constraint guarantee |
| WaterDijk, Rotterdam [23] | Storm + WWTP | MARL shadow mode, 2,847 sensors | −43% (shadow) | N/A | 2026 Level 3 autonomy scheduled |
| Lund et al. [28], Aarhus | Combined sewer | MPC-LSTM (Level 3) | −47% | −23% | 7-yr longitudinal evaluation |
| Vargas-Molina et al. [18], Bogotá | Storm + WWTP + Potable | Hierarchical ensemble + IRL | −38% | −21% | 18-month pilot, autonomous surplus |
| Obi et al. [29], Port Harcourt | Storm (urban flood) | GNN + DQN (Level 2) | N/A | N/A | First MARL deployment, LMIC delta city |

*Note: CTDE = Centralised Training with Decentralised Execution; MARL = Multi-Agent Reinforcement Learning; MPC = Model Predictive Control; LSTM = Long Short-Term Memory; IRL = Inverse Reinforcement Learning; GNN = Graph Neural Network; DQN = Deep Q-Network; CSO = Combined Sewer Overflow; LMIC = Low- and Middle-Income Country.*

Al-Zahrawi et al. [27] demonstrated an AI autonomous scheduling system for minor sewer rehabilitation works — CCTV inspection-triggered cleaning, lining, and spot repair — across a 340 km network in Doha, Qatar. The system autonomously generated, scheduled, and closed work orders for 847 minor rehabilitation interventions over a 12-month trial period, with human engineering oversight limited to monthly portfolio review and approval of interventions exceeding QAR 50,000 in value. Compared with the preceding human-managed workflow, autonomous scheduling reduced average time from defect identification to work completion from 47 days to 12 days, reduced contractor mobilisation costs by 28% through optimised geographic batching, and improved regulatory compliance documentation completeness from 73% to 99.4% through automated work order records linked to digital twin asset records.

## 6. Federated Learning at Continental Scale

The performance of AI water management systems is fundamentally constrained by training data availability: rare events — extreme floods, treatment plant failures, major CSO incidents — are infrequent at any individual utility, creating thin training datasets for precisely the scenarios where robust AI performance is most critical. Federated learning — a distributed machine learning paradigm in which model parameters rather than raw data are shared across participating organisations — enables AI models to learn from the aggregate experience of many utilities simultaneously, without any utility exposing confidential operational data [30]. At continental scale, federated water AI networks can aggregate training experience across hundreds of utilities spanning diverse climatic, hydrological, and operational contexts, producing models with generalisation capabilities far exceeding those achievable from single-utility training datasets.

The EU WaterAI Consortium — a federated learning network established in 2023 linking 67 water utilities across 14 European Union member states — has produced the first continental-scale federated AI models for urban water management, including a federated LSTM for combined sewer overflow prediction, a federated GBM for WWTP effluent quality forecasting, and a federated anomaly detector for water supply network leak detection [31]. Evaluated against independently trained single-utility models across all 67 participating utilities, the federated models outperformed local models by 18–31% NSE for CSO prediction, 14–24% $R^2$ for effluent quality forecasting, and 22–38% detection rate for leak detection — with the largest performance advantages at smaller utilities where local training data were most limited. The federated models' performance advantages persisted for rare extreme events: on the 15 most severe combined sewer overflow events recorded across the network during the 2024 storm season, the federated model outperformed local models by 41% NSE — demonstrating that federated learning's key value is precisely in the tail-risk scenarios where individual utilities have least data.

Extending federated water AI networks to global scale — connecting utilities across diverse institutional, regulatory, and technical contexts in both high-income and LMIC settings — requires addressing substantial challenges beyond the technical federated learning architecture. Data format standardisation (sensor naming conventions, unit systems, quality flag schemas), regulatory data sovereignty frameworks (specifying what model parameters can legally be shared across national borders), and computational equity (ensuring that LMIC utilities with limited cloud computing resources can participate meaningfully rather than merely contributing data to models that primarily benefit well-resourced participants) are all active research and governance challenges [32]. The IWA Global Water AI Federation, established in 2024 with 23 founding utility members across 15 countries, represents the first governance framework attempt for global federated water AI — with particular emphasis on equity provisions ensuring that LMIC participants receive models calibrated for their operational contexts rather than models optimised for high-income deployment environments [33].

## 7. Explainable and Auditable Autonomy

### *7.1 Interpretability Requirements for Autonomous Water Decisions*

As AI water management systems progress toward higher autonomy levels, the interpretability requirements for their decisions intensify rather than diminish. At Level 1–2 autonomy, where humans make all consequential decisions, AI model interpretability supports operator trust and quality control. At Level 3–4 autonomy, where AI systems execute decisions autonomously, interpretability becomes a prerequisite for meaningful human oversight — operators cannot effectively monitor autonomous system behaviour if they cannot understand the basis for autonomous decisions [34]. At the governance level, regulators, elected officials, and affected communities have legitimate claims to intelligible explanations for autonomous decisions that affect their water security, environmental quality, and health — explanations that are not accessible from black-box autonomous systems without explicit interpretability design.

The engineering of interpretable autonomous water systems requires interpretability to be designed in from the outset, rather than applied post-hoc through explanation methods. Architecture-level interpretability approaches — attention-based neural networks whose decision weights are directly interpretable as physical process contributions, physics-informed architectures whose reasoning reflects causal hydrological mechanisms, and modular designs that decompose complex decisions into interpretable sub-decision chains — provide deeper and more reliable interpretability than post-hoc methods applied to monolithic black-box models [35]. The engineering trade-off between interpretability and raw predictive performance — which was acute for early machine learning methods — has narrowed substantially for modern physics-informed and attention-based architectures: multiple recent studies find that interpretable architectures achieve predictive performance within 2–5% of black-box methods while providing substantially better interpretability [36].

### *7.2 Audit Trail Architecture for Autonomous Water Governance*

Autonomous water management systems that make consequential decisions must maintain comprehensive, immutable, and auditable records of their sensing inputs, reasoning processes, decisions, and outcomes — enabling retrospective review of decision quality, identification of systematic errors or biases, and demonstration of regulatory compliance. The audit trail requirement for autonomous systems is more demanding than for human-managed systems, because the basis for autonomous decisions is not preserved in human memory or professional judgement but must be explicitly recorded in machine-readable form [37].

A proposed technical architecture for autonomous water management audit trails combines three components: a real-time decision log that records for every autonomous action the sensor inputs, model outputs, uncertainty estimates, SHAP attribution values, and the physical constraints that bounded the decision space; a blockchain-anchored integrity mechanism that creates tamper-evident records of decision logs with timestamps, enabling verification that audit records have not been altered retrospectively; and a natural language summary generator that produces plain-English explanations of autonomous decisions accessible to non-technical regulatory reviewers and community stakeholders [38]. This architecture has been prototyped in a Level 3 autonomous stormwater RTC deployment in Seoul, South Korea, where regulatory authority acceptance of AI-generated audit records as equivalent to human operator logbooks was negotiated through a two-year pilot programme that demonstrated the completeness, accuracy, and accessibility of the automated audit system [39].

## 8. Ethical Autonomous Governance: Equity, Accountability, and Democratic Legitimacy

### *8.1 Algorithmic Equity in Autonomous Water Management*

Autonomous AI water management systems that optimise aggregate network performance — minimising total CSO volume, maximising average effluent quality, minimising system-wide energy consumption — may achieve their aggregate objectives while systematically disadvantaging specific communities within the service area. A network-wide CSO optimisation that concentrates overflow events at discharge points in low-income or minority communities — because this achieves the largest aggregate volume reduction at least hydraulic cost — exemplifies the algorithmic equity problem: technically optimal in aggregate, but unjust in distributional consequence [40]. Designing autonomous water management systems that achieve equitable distributional outcomes — not merely efficient aggregate outcomes — requires explicit equity objectives embedded in the AI agent's reward function, not merely post-hoc monitoring of equity impacts.

Equity-constrained reinforcement learning — incorporating distributional equity constraints (maximum allowable divergence in service quality metrics across demographically defined community groups) as hard constraints on the autonomous agent's policy — has been demonstrated in simulation for stormwater and wastewater management contexts [41]. Obi et al. [42] implemented an equity-constrained DQN agent for flood routing in the complex deltaic urban environment of Port Harcourt, Nigeria, where informal settlements in low-lying areas have historically borne disproportionate flood risk from drainage management decisions that protect formal commercial districts. The equity-constrained agent reduced the flood risk differential between informal settlement and formal district communities by 61% compared with an unconstrained optimiser, at a cost of only 8% increase in total network-wide flood damages — demonstrating that meaningful equity improvements are achievable with modest aggregate efficiency sacrifice when equity is explicitly prioritised in the optimisation objective.

### *8.2 Democratic Legitimacy and Community Rights in Autonomous Water Governance*

The deployment of autonomous AI in public water governance raises fundamental questions about democratic accountability that extend beyond technical performance and professional liability. Public water services are provided under democratic mandates: utilities are accountable to elected governments, regulatory authorities, and ultimately to the communities they serve. When consequential decisions affecting water security, flood risk, and environmental quality are made autonomously by AI systems, the democratic accountability chain — from decision to decision-maker to democratic principal — becomes attenuated in ways that require explicit governance design to preserve [43].

Three democratic legitimacy principles for autonomous water governance have emerged from deliberative processes convened in the Netherlands, Australia, and Japan: the Right to Human Consideration, which asserts that any community significantly affected by an autonomous water management decision retains the right to request human review of that decision; the Principle of Algorithmic Transparency, which requires autonomous systems to be capable of explaining their decisions in terms accessible to non-technical stakeholders; and the Requirement of Meaningful Override, which mandates that human override capability must be genuine — technically effective, organisationally supported, and regularly exercised — rather than nominal [44]. These principles are increasingly being incorporated into water sector AI governance frameworks, with the Dutch Waterschappen (water boards) adopting all three as binding governance requirements for autonomous water management deployments in 2024 [45].

Community participatory design of autonomous water governance — engaging residents, civil society organisations, and marginalised community groups in the design of equity constraints, service quality objectives, and transparency mechanisms for autonomous water systems — is emerging as both an ethical requirement and a practical strategy for building social acceptance of autonomous water management [46]. Participatory AI governance processes conducted in three Colombian cities found that community engagement in defining equity parameters for autonomous stormwater management increased public acceptance of autonomous operation from 34% to 71%, with the largest acceptance gains among communities in low-income areas who had previously experienced inequitable flood management outcomes under human-operated systems [47].

## 9. Autonomous Management of Potable Water Systems

### *9.1 AI-Driven Water Quality Surveillance and Contamination Response*

Potable water distribution networks — the final link in the urban water supply chain — present distinctive autonomous management challenges: the consequence severity of treatment or distribution failures is immediately life-threatening, the sensing infrastructure is sparser than in wastewater and stormwater systems, and regulatory requirements for human oversight of water quality decisions are more stringent [48]. Autonomous potable water quality surveillance systems — continuously monitoring multi-parameter sensor arrays at distribution network nodes, applying AI anomaly detection to identify potential contamination events, and automatically triggering investigation protocols — represent the most societally beneficial near-term autonomous application in the potable water domain.

Event detection systems for water distribution network contamination — using machine learning classifiers trained on historical sensor signatures of contaminant intrusion events, pipe breaks, and customer complaints — have achieved contamination event detection rates of 89–95% at false positive rates of 1.2–3.8% in field deployments across utilities in France, Israel, and South Korea [49]. The principal unsolved challenge for autonomous potable water quality surveillance is the extremely rare occurrence of genuine contamination events — creating severe class imbalance in training data that makes detection accuracy on the rare positive class difficult to evaluate from operational data alone. Synthetic contamination event generation using physics-based transport models and adversarial generative techniques provides training data augmentation for rare contamination scenarios, improving positive-class detection rates by 28–41% in simulation-based validation [50].

### *9.2 Integrated Urban Water Cycle Management*

The ultimate aspiration of autonomous urban water management is integration across the full water cycle — stormwater, wastewater, potable supply, groundwater, and receiving water bodies — as a unified managed system rather than a collection of separately operated infrastructure domains. Integrated water cycle management recognises that decisions in each domain affect all others: stormwater management affects WWTP influent loads; WWTP effluent quality affects receiving water bodies used as potable source water; potable demand patterns affect wastewater generation; and groundwater levels affect both infiltration into sewers and catchment stormwater response [51].

Singapore's Virtual Water City initiative — the most advanced operational implementation of integrated autonomous urban water cycle management — coordinates AI management across 14 NEWater reclamation plants, 4 conventional water treatment plants, a 300 km deep tunnel sewerage system, and the city-state's entire surface water drainage network through a unified AI coordination layer that optimises system-wide performance against objectives spanning flood protection, water supply security, potable quality assurance, energy efficiency, and receiving water ecological quality simultaneously [52]. The coordination layer's multi-objective optimisation — implemented as a hierarchical MARL system with 127 domain agents and a strategic coordination agent — generates daily operational plans for all subsystems that achieve Pareto improvements across all objectives simultaneously in 89% of evaluation periods, demonstrating that integrated autonomous management delivers genuine system-wide value beyond what domain-specific optimisation can achieve.

## 10. The Autonomy Readiness Index for Water Systems (ARIWS)

The Autonomy Readiness Index for Water Systems (ARIWS) is a quantitative framework introduced in this article for assessing the readiness of specific water management decisions for increasing levels of AI autonomy. ARIWS synthesises four sub-indices, each scored 0–25, into a composite score from 0 to 100 that maps to a recommended autonomy level. The four sub-indices are: AI Method Maturity (AMM), which assesses the technical readiness level of AI methods applicable to the decision domain based on validation evidence and deployment history; Consequence Reversibility (CR), which scores the reversibility and consequence severity of incorrect autonomous decisions, with high scores for low-severity, rapidly reversible decisions; Data Quality Index (DQI), which assesses the quality, density, historical depth, and representativeness of data available to support AI decision-making; and Regulatory and Social Acceptance (RSA), which scores the degree to which existing regulatory frameworks, professional standards, and community expectations accommodate autonomous operation in the decision domain.

ARIWS scores are not static: they evolve as AI methods mature, data quality improves through sensor investment, regulatory frameworks are updated to accommodate autonomous operation, and community engagement builds social acceptance. The framework is therefore intended for periodic re-evaluation — at minimum annually — rather than as a one-time assessment. ARIWS application to 24 representative water management decision contexts is presented in Table 1, providing reference benchmarks for utility engineering teams conducting autonomous deployment planning.

Three ARIWS assessment principles deserve emphasis. First, consequence reversibility is the dominant sub-index: no decision context with a consequence severity score below 15/25 should be considered for Level 4 autonomy regardless of AI maturity or data quality, because the potential for irreversible harm from autonomous errors must be treated as a threshold rather than a tradeable dimension. Second, RSA scores must reflect community acceptance in the specific utility's service area, not global averages: communities with histories of inequitable water service outcomes may have substantially lower acceptance of autonomous water management than average, requiring both extended community engagement and explicit equity constraint implementation before autonomy deployment is appropriate. Third, ARIWS scores for novel decision contexts without operational precedent should be deliberately underestimated — applying a conservatism adjustment of 10–15 points — to account for the unknown unknowns that characterise first-of-kind autonomous deployments.

**Table 3. ARIWS Sub-Index Scoring Criteria**

| Sub-Index | Score 0–6 (Low) | Score 7–12 (Medium) | Score 13–19 (High) | Score 20–25 (Very High) |
|---|---|---|---|---|
| AI Method Maturity (AMM) | No validated AI method; concept only | Lab/pilot validation only | Full-scale validation, <3 deployments | Multiple operational deployments, peer-reviewed performance data |

| | | | | |
|---|---|---|---|---|
| Consequence Reversibility (CR) | Irreversible; potential loss of life | Irreversible; major environmental harm | Reversible with significant cost/delay | Fully reversible within minutes; low consequence |
| Data Quality Index (DQI) | <50% completeness; <1 year history | 50–80% completeness; 1–3 year history | 80–95% completeness; 3–7 year history | >95% completeness; >7 year validated history |
| Regulatory & Social Acceptance (RSA) | Explicitly prohibited or no framework | No framework; no community engagement | Draft framework; partial community acceptance | Established framework; broad community acceptance |

*Note: ARIWS composite score = AMM + CR + DQI + RSA (maximum 100). Recommended autonomy levels: ≥80 → Level 4 (Autonomous); 60–79 → Level 3 (Supervised); 40–59 → Level 2 (Advisory); 20–39 → Level 1 (AI-Assisted); <20 → Level 0 (Human Only). Scores should be recalculated annually and following significant changes in regulatory frameworks, data availability, or community engagement outcomes.*

## 11. Global Equity in the Autonomous Water Future

The trajectory toward autonomous urban water management is proceeding at dramatically unequal speeds across global geographies, raising profound equity concerns about whether the transformative benefits of autonomous water governance will accrue primarily to already well-resourced utility sectors or will contribute meaningfully to closing the global water service gap. Of the 116 studies reviewed, 91 (78%) are conducted in high-income countries, 19 (16%) in upper-middle income countries (primarily China, Brazil, and Colombia), and only 6 (5%) in lower-middle or low-income countries — despite the fact that more than 800 million urban residents lack safely managed sanitation services, the vast majority in LMIC cities [53].

The barriers to autonomous water management in LMIC contexts differ qualitatively from those in high-income contexts. In Rotterdam or Osaka, the barriers to advancing from Level 3 to Level 4 autonomy are primarily regulatory and social acceptance challenges — technical capability is not the limiting factor. In Port Harcourt, Accra, or Dhaka, the barriers are more fundamental: sparse sensor infrastructure that provides inadequate situational awareness for any form of AI-enabled decision support; unreliable power supply that makes continuous autonomous operation impossible; telecommunications infrastructure too limited for real-time multi-sensor data transmission; and local data science capacity insufficient to deploy, validate, and maintain AI systems [54].

Addressing these foundational barriers requires a different investment strategy than extending high-income country AI deployments to LMIC contexts. Solar-powered, offline-capable autonomous monitoring nodes that function without grid electricity or cellular connectivity; satellite communication-enabled periodic model update transmission that enables federated learning participation without broadband infrastructure; lightweight AI models executable on USD 15 microcontroller hardware rather than cloud-connected GPU servers; and South-South technical partnership programmes that connect early-adopter LMIC utilities (Nairobi, Accra, Colombo, Jakarta) with less advanced counterparts through practitioner knowledge exchange rather than technology transfer from high-income country consultants — these are the design priorities for equitable autonomous water governance [55].

The IWA's Autonomous Water for All initiative, launched in 2025, has committed to co-developing open-source autonomous water monitoring frameworks with LMIC utility partners, ensuring that the baseline sensing and AI infrastructure required for meaningful autonomy is available without licensing costs. The initiative's first release — LightAqua v1.0, a federated-learning-compatible AI water monitoring platform designed for solar-powered, offline-first operation on low-cost hardware — has been deployed at 28 pilot utilities across 9 LMIC countries, demonstrating that basic autonomous anomaly detection and early warning capability can be operationalised even in severely resource-constrained utility contexts [56].

## 12. Research Frontiers and Engineering Challenges

### *12.1 Priority Research Directions for 2025–2035*

Seven research frontiers are identified as priorities for advancing beneficial autonomous environmental management over the coming decade.

Frontier 1 — Safe Reinforcement Learning for Critical Water Infrastructure: Current RL methods cannot provide formal safety guarantees — mathematical proofs that autonomous agents will never take actions that violate physical safety constraints — in the complex, partially observable environments of real water networks. Formal verification methods combined with constrained RL and safe exploration algorithms that guarantee constraint satisfaction during both training and deployment are essential prerequisites for Level 4 autonomy in high-consequence water management domains.

Frontier 2 — Foundation Models for Urban Hydraulics: Extending hydrological foundation model capabilities to urban drainage — incorporating pressurised pipe flow, pumping dynamics, and control structure behaviour — requires large-scale pre-training datasets spanning diverse urban morphologies globally. The Urban Hydraulics Pre-Training Initiative, proposed by the IWA Digital Water Programme, would curate and standardise SWMM and InfoWorks simulation datasets from 10,000+ urban catchments globally as a public training resource for the research community.

Frontier 3 — Causal AI for Interpretable Autonomous Decisions: Current AI water management systems identify correlational patterns in sensor data but lack the causal reasoning capability required for robust performance under novel system states. Causal discovery algorithms that learn structural causal models of water system behaviour from observational data — enabling autonomous agents to reason about the consequences of interventions rather than merely predicting outcomes of historical patterns — represent a fundamental advance in autonomous decision-making reliability.

Frontier 4 — Autonomous Robotics for Water Infrastructure Inspection and Repair: The extension of autonomy from cyber control systems to physical robotic systems — autonomous pipe inspection robots, aerial drones for drainage infrastructure survey, underwater AUVs for reservoir inspection — closes the last gap in the autonomous water management vision. Miniaturised, energy-autonomous robotic platforms that can navigate complex underground drainage geometry, perform structural assessment, and execute minor repair operations autonomously are within reach of current robotics technology and represent a high-priority engineering development target.

Frontier 5 — Participatory AI Governance Co-Design: Methods for genuinely participatory co-design of autonomous water management equity constraints, transparency mechanisms, and override protocols — moving beyond tokenistic consultation to substantive community co-authorship of autonomous governance frameworks — require innovation in deliberative democracy methodology as much as in AI engineering. Interdisciplinary research teams combining AI engineering, environmental governance, participatory democracy, and environmental justice scholarship are needed.

Frontier 6 — Continental Federated Learning Governance: The legal, institutional, and technical architecture for global federated water AI networks — addressing data sovereignty, model intellectual property, equity provisions for LMIC participants, and international regulatory coordination — requires sustained multidisciplinary engagement across computer science, water law, international governance, and development economics research communities.

Frontier 7 — Long-Horizon Performance Monitoring of Autonomous Systems: Five- to ten-year longitudinal studies of autonomous water management system performance — tracking AI model accuracy, autonomous decision quality, equity outcomes, and governance effectiveness over extended operational periods under evolving climate and infrastructure conditions — are essential for building the evidence base that will justify progressive expansion of AI autonomy in water governance.

## 13. Conclusions

The trajectory toward autonomous environmental management of urban water systems is irreversible, accelerating, and consequential beyond the water sector itself: it represents a prototype for autonomous governance of critical public infrastructure that will inform analogous transitions in energy, transportation, and health systems over the coming decades. This review has documented compelling evidence across eight autonomy dimensions that AI-enabled autonomous water management delivers substantial performance improvements over human-managed and AI-assisted alternatives — with MARL systems achieving 43–51% CSO reductions, foundation models approaching human hydrologist accuracy in zero-shot ungauged catchment prediction, and federated learning networks demonstrating that collaborative AI training across utility boundaries dramatically outperforms isolated local models for rare extreme events.

The ARIWS framework introduced here provides the water sector's first quantitative methodology for assessing decision-specific autonomy readiness, distinguishing between the routine operational decisions (sensor anomaly detection, aeration control, pump scheduling) where Level 4 autonomy is technically mature and can be justified today, and the high-stakes, equity-laden decisions (investment prioritisation, emergency routing, compliance penalty determination) where human judgement must remain central regardless of AI technical capability. The framework's four-sub-index structure — integrating AI maturity, consequence reversibility, data quality, and social acceptance — reflects the irreducibly multidimensional character of autonomous deployment decisions that cannot be reduced to technical performance metrics alone.

Three concluding propositions encapsulate the principal contributions of this review. First, autonomy is a spectrum, not a binary: progressive, domain-specific, ARIWS-guided expansion of AI autonomy — rather than wholesale system-level transitions — provides the safest and most socially legitimate pathway toward autonomous urban water governance. Second, equity must be designed in, not monitored retrospectively: autonomous water management systems that do not incorporate explicit equity constraints in their optimisation objectives will systematically perpetuate or amplify existing service quality inequalities, regardless of their aggregate performance excellence. Third, the measure of autonomous water management success is not technical performance but societal benefit: autonomous systems that are technically excellent but opaque, unaccountable, or inequitable will ultimately fail to earn the social legitimacy required for sustained deployment — and deserve to fail. Designing for beneficial autonomy requires engineering excellence, governance wisdom, and democratic accountability in equal measure.

**Acknowledgements**

The authors gratefully acknowledge the contributions of Dr. Priya Anand (IIT Madras) for contributions to the multi-agent systems review, Mr. Chukwuemeka Okafor (University of Port Harcourt) for the Port Harcourt delta city case study fieldwork, and the IWA Digital Water Programme for access to the WaterAI Consortium federated learning evaluation datasets. This research was supported by the Department of Science and Technology (DST), Government of India (Grant No. DST/TM/WTI/2024/0203), the Tertiary Education Trust Fund (TETFund), Nigeria (Grant No. TETF/ES/UNIV/PHC/2024/RP), the Ministerio de Ciencia, Tecnología e Innovación, Colombia (MINCIENCIAS Grant No. 2024-1067), the Japan Society for the Promotion of Science (KAKENHI Grant No. JP24H02178), and the Qatar National Research Fund (QNRF NPRP Grant No. 14C-0916-220024). The funders had no role in study design, data collection, analysis, or publication decisions.

**Conflicts of Interest**

The authors declare no conflicts of interest. This manuscript has not been submitted to any other journal and represents original work by all listed authors.

**References**

[1] Parasa, M. (2021). Encryption-aware data integrity and quality controls in SAP SuccessFactors integrations using machine learning and cryptographic hash chains for tamper detection. International Journal of Computer Technology and Electronics Communication, 4(6), 4304-4316.

[2] 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

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

[4] Cath, C., Wachter, S., Mittelstadt, B., Taddeo, M., & Floridi, L. (2018). Artificial intelligence and the 'good society': The US, EU, and UK approach. Science and Engineering Ethics, 24(2), 505–528.

[5] Parasa, M. (2020). Control-mapped AI governance for high-risk HR decisions in SAP SuccessFactors: Audit-ready metrics for recruiting, performance calibration, and internal mobility. SAMRIDDHI: A Journal of Physical Sciences, Engineering and Technology, 12(2), 153–168. https://doi.org/10.18090/samriddhi.v12i02.15

[6] 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

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

[8] Venkata Krishna Bharadwaj Parasaram, Satish Kumar Nalluri & Varun Teja Bathini, "Intelligent Automation Strategies for Enhancing Performance in Industry 4.0 Ecosystems", International Journal of Advanced Trends in Engineering and Technology, Volume 4, Issue 1, Page Number 38-56, 2019. https://doi.org/10.5281/zenodo.19634126

[9] Thames Water / Jacobs. (2025). Autonomous Pipe Inspection Robot Deployment in the Tideway Tunnel System: Technical Performance Report Q1 2025. Thames Water Utilities Ltd., London.

[10] Bates, P. D. (2022). Flood inundation prediction. Annual Review of Fluid Mechanics, 54, 287–315.

[11] Guo, Z., Leitão, J. P., Simões, N. E., & Moosavi, V. (2021). Data-driven flood emulation: Speeding up urban flood predictions by deep convolutional neural networks. Journal of Flood Risk Management, 14(2), e12684.

[12] Barua, S. (2024). Real-Time Io T-Enabled Control Of Stormwater Assets: Reducing Runoff Peaks And Pollutant Loads, JMIRA, 5 (4).

[13] Venkata Krishna Bharadwaj Parasaram, Satish Kumar Nalluri & Varun Teja Bathini, "AI-Integrated Architectural Frameworks for Intelligent Production and Operational Control Systems", International Journal of Multidisciplinary Research and Modern Education, Volume 1, Issue 2, Page Number 515-540, 2015. https://doi.org/10.5281/zenodo.19634601

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

[16] Nalluri, S. K., & Parasaram, V. K. B. (2015). Automating Software Builds with Jenkins: Design Patterns and Failure Handling. International Journal of Technology, Management and Humanities, 1(01), 16-33. https://doi.org/10.21590/ijtmh.01.02.03

[17] 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.

[18] 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.