



ARTICLE

An IoT-Enabled Hybrid DRL-XAI Framework for Transparent Urban Water Management

Qamar H. Naith^{1,*} and H. Mancy^{2,3}

¹Department of Software Engineering, College of Computer Science and Engineering, University of Jeddah, Jeddah, 21959, Saudi Arabia

²Department of Computer Science, College of Engineering and Computer Sciences, Prince Sattam Bin Abdulaziz University, Al-kharj, 11942, Saudi Arabia

³Department of Mathematics, Faculty of Science (Girls), Al-Azhar University, Cairo, 11765, Egypt

*Corresponding Author: Qamar H. Naith. Email: qnaith@uj.edu.sa

Received: 21 April 2025; Accepted: 16 June 2025; Published: 31 July 2025

ABSTRACT: Effective water distribution and transparency are threatened with being outrightly undermined unless the good name of urban infrastructure is maintained. With improved control systems in place to check leakage, variability of pressure, and conscientiousness of energy, issues that previously went unnoticed are now becoming recognized. This paper presents a grandiose hybrid framework that combines Multi-Agent Deep Reinforcement Learning (MADRL) with Shapley Additive Explanations (SHAP)-based Explainable AI (XAI) for adaptive and interpretable water resource management. In the methodology, the agents perform decentralized learning of the control policies for the pumps and valves based on the real-time network states, while also providing human-understandable explanations of the agents' decisions, using SHAP. This framework has been validated on five very diverse datasets, three of which are real-world scenarios involving actual water consumption from NYC and Alicante, with the other two being simulation-based standards such as LeakDB and the Water Distribution System Anomaly (WDSA) network. Empirical results demonstrate that the MADRL + SHAP hybrid system reduces water loss by up to 32%, improves energy efficiency by up to 25%, and maintains pressure stability between 91% and 93%, thereby outperforming the traditional rule-based control, single-agent DRL (Deep Reinforcement Learning), and XGBoost + SHAP baselines. Furthermore, SHAP-based interpretation brings transparency to the proposed model, with the average explanation consistency for all prediction models reaching 88%, thus further reinforcing the trustworthiness of the system on which the decision-making is based and empowering the utility operators to derive actionable insights from the model. The proposed framework addresses the critical challenges of smart water distribution.

KEYWORDS: Multi-Agent reinforcement learning; explainable artificial intelligence (XAI); SHAP (Shapley Additive Explanations); smart water distribution; urban infrastructure; Internet of Things (IoT); water resource optimization; energy efficient control

1 Introduction

Urban water resource management is under growing pressure caused by population growth, climate change, and aging infrastructures. Many cities lose considerable fractions of treated water before the water reaches the consumers. Global estimates indicate that 30%–34% of distributed potable water does not become non-revenue water owing to leaks and other inefficiencies [1]. This problem has become even more serious in arid regions because they depend on costly processes like desalination for water supply and because the



demand for water is among the highest in the world [2]. For instance, in Dubai, UAE, there is a hyper-arid climate, very few natural water resources, and highly dependent on desalinated water; hence, consumption on a per-capita basis in the UAE is far above the global average [3].

Recent studies demonstrate that integrating artificial intelligence into smart water systems has become ever more important so predictive water management can be realized with transparency and scalability. Syed et al. [4] presented a digital-twin-based architecture coupled with multimodal transformer models for high-accuracy forecasting of water consumption and leak detection, asserting the potential of AI-driven predictive analytics in urban water infrastructure. Mohammed et al. [5] advocated an adaptive online learning framework to detect leaks in real time so as to provide rapid responses to anomalous events occurring in complex distribution systems. Infant et al. [6] studied the application of XAI in water systems engineering, considering transparency and trust in a model crucial to decision-making for sustainable infrastructure. Building further on this line, Pagano et al. [7] first proposed the Smart Water IoT Framework for Evaluation of Energy and Data (SWI-FEED) framework, an advanced large-scale smart water IoT system that interconnects sensing, communication, and cloud-based control for presently adapted management along distributed networks. Collectively, these recent developments highlight the paradigm shift toward hybrid, explainable, and IoT-enabled solutions, which constitute the conceptual backdrop on which this thesis is constructed.

These pressures necessitate the urgent need for more intelligent and efficient urban water distribution systems that minimize losses and adapt, in real time, to changing conditions. Recently, sensing technologies such as IoT devices—including smart meters, pressure sensors, and flow monitors—alongside advanced data analytics, have enabled the emergence of Smart Water Networks as a core component of smart city initiatives. Utilities are increasingly considering AI to optimize the whole gamut of water distribution, from pump scheduling to valve control to leak detection and demand forecasting. For example, Dubai Electricity and Water Authority (DEWA) has already launched an AI-driven “Hydronet” to monitor and control its water network remotely in real time [3]. Nevertheless, most of the existing solutions either rely on static heuristics or employ complex models that do not have much transparency at all. Thus, traditional pump operations would generally prescribe fixed schedules or rule-based logic, which was obviously developed offline but unable to adapt to surprises like sudden demand spikes or pipe bursts.

These highly automated and data-driven approaches, like predictive control with machine learning, could optimize operations. Still, most of the time, they are found to be implemented in a black-box mode, which may not be suitable for critical types of infrastructure [8]. One new hopeful approach is Deep Reinforcement Learning in achieving adaptive feedback control for complex systems. Within water distribution, Deep Reinforcement Learning (DRL) agents can acquire control policies by interacting with a network model or live system to optimize objectives such as pressure regulation or energy efficiency. Or it undertakes pulp real-time optimization of intractable ranges by conventional techniques [9].

Multi-Agent DRL has enhanced this capability by enabling multi-agent cooperation or coordination, which has a huge application in such distributed control elements (pumps, valves, tanks, etc.) over long distances, as found in large-scale water networks. For example, Hu et al. implemented a MADRL-based pump and valve scheduler that surpassed the evolutionary algorithms in solution quality and speed [10]. Their multi-agent approach, by a multi-agent deep deterministic policy gradient algorithm, effectively coped with uncertain demand patterns in a water grid that required nearly real-time control. It was previously unthinkable by earlier rule-based or optimization methods [10]. Evidently, these improvements have implications such that a MADRL controller can deliver the adaptive optimal operation called for in a smart water distribution system in scenarios such as Dubai. Any DRL-based solution must be coupled with an explanation mechanism to make it feasible for practical deployment in the water sector that is possible

in a real-world scenario. Explainable AI attempts to clarify how models arrive at their decisions, thereby reconciling the intuition of the model with human understanding [8].

In this paper, we present a hybrid DRL-XAI framework for adaptive and transparent water resource management that combines the strengths of MADRL and SHAP. SHAP is a game-theoretic XAI approach that assigns each feature an importance value for a given prediction [11]. Developed initially to interpret machine learning predictions, SHAP has been adopted in recent studies to explain decisions made by reinforcement learning agents as well [12]. The framework is designed for real-time control of urban water distribution and is demonstrated in a case study of Dubai's potable water network. The AI-based framework for urban water distribution is illustrated in Fig. 1.

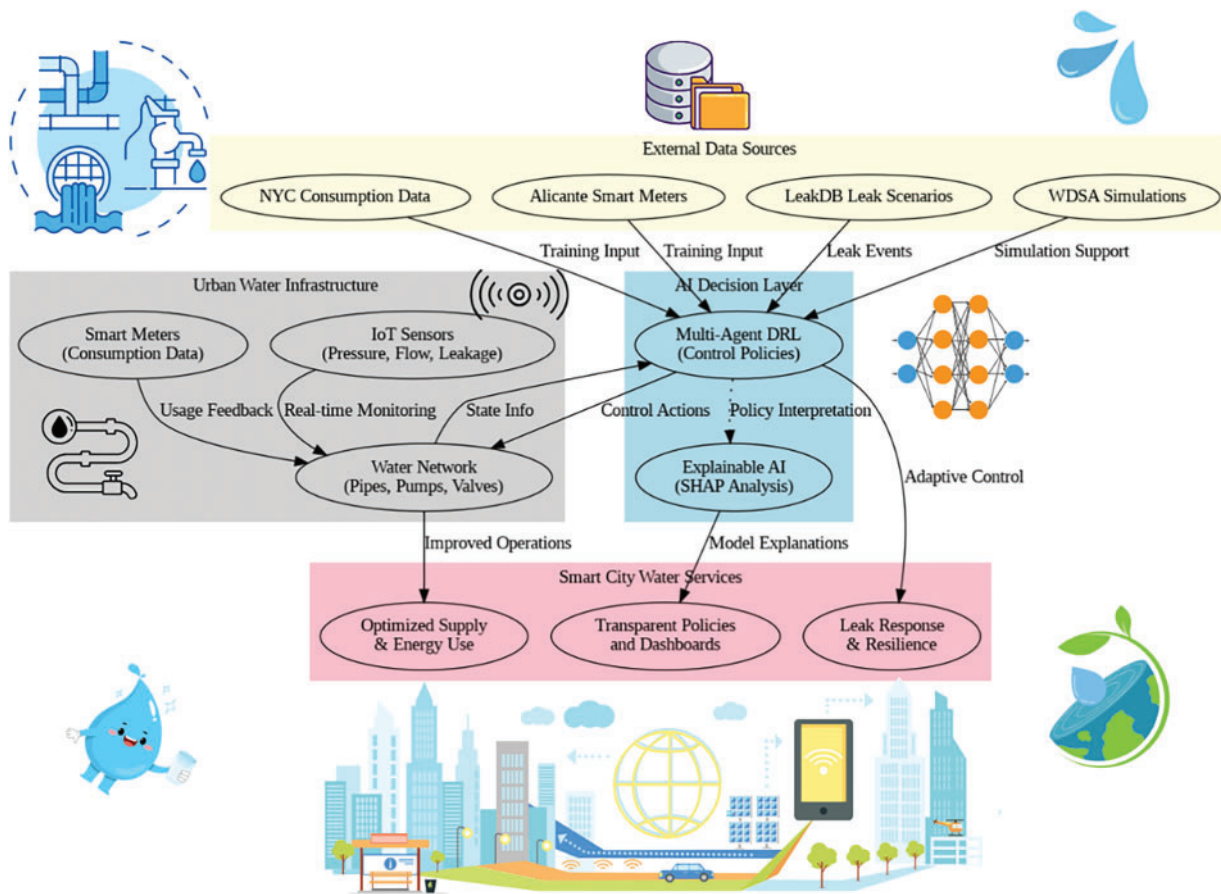


Figure 1: Smart water management framework for adaptive and transparent control

We use a standard WDSA (Water Distribution System Analysis) benchmark network model [13] for the simulation of the Dubai network under baseline conditions. On top of this, we incorporate fault scenarios from LeakDB, a public benchmark dataset of realistic leak events in water networks [14,15], to test the agents' ability to handle pipe bursts and leakage. Real consumption patterns are introduced using two open datasets: the Alicante Smart Water Meter dataset, which provides high-resolution (hourly) usage data from 1007 customers in Alicante, Spain [16], and the New York City water consumption history, which offers long-term municipal water usage statistics [17]. Hence, with such diverse data for training, the DRL agents acquire policies that are robust to variations in demand profiles; for example, the high per-capita usage due to the climate in Dubai, contrasted with the moderate usage in NYC. The subsequent XAI component provides the

means for domain experts to visualize and interpret the influence that factors such as diurnal demand cycles, anomalies from sensors, or leak incidences can have at any given point in time on the identified control decisions. In summary, the key contributions of this work are as follows:

- **Hybrid MADRL-XAI Framework:** We propose a novel intelligent control framework that combines Multi-Agent Deep Reinforcement Learning (MADRL) with SHAP-based Explainable AI (XAI) for real-time, adaptive, and interpretable water distribution management. This is the first framework to apply SHAP explanations at agent-level decision granularity within a decentralized control system for urban water networks.
- **Adaptive Multi-Agent Control for Urban Water Systems:** We develop and train a decentralized MADRL controller that dynamically adjusts the pump speeds and valve positions with real-time pressure, flow, and demand signals. The learned cooperative policies outperform static scheduling and single-agent DRL baselines by water loss reductions up to 32% and energy efficiency improvements of up to 25%, all while maintaining pressure stability above 90% across different datasets.
- **SHAP-Driven Interpretability Module:** In response to the deep RL's (Reinforcement Learning) black-box nature, we incorporate a SHAP-based interpretability layer that tracks the most influential state features driving the agent's actions. This makes for actionable explanations and builds the operator's trust, while SHAP consistency scores stay in the region of 88% across diverse scenarios like peak demand and leak events.
- **The tests were very exhaustive** in that they covered five diverse datasets real field data, especially from water use in New York City and Alicante, and simulation-based scenarios from LeakDB and WDSA. The proposed framework consistently performed better than the baselines of Rule-Based Control, single-agent DRL, and supervised learning-using XGBoost + SHAP in all environments.
- **Case Study: Dubai Smart Water Network Simulation:** Dubai, a representative smart city, is chosen for scenario modeling. Through the simulation of realistic leak detection and adaptive control conditions, we show the practical viability of our framework together with SHAP-based insights on how agents react to changing network states.

The rest of the paper is arranged as follows. In [Section 2](#), a review of literature regarding smart water management and reinforcement learning in water systems and explainable AI is presented. [Section 3](#) details the architecture of the proposed hybrid DRL-XAI framework comprising MADRL model design and SHAP-based explanation methodology. [Section 4](#) gives the datasets that were used for training and testing. [Section 5](#) presents the experimental results and the performance of the framework with respect to baseline methods. [Section 7](#) concludes the paper with the key findings regarding the proposed method.

2 Literature Review

Deep reinforcement learning has emerged as a promising approach for optimizing complex control problems in urban water networks. Early work by Hu et al. applied RL to water system control [18], and recent deep RL studies have achieved notable successes in pump operations, [Fig. 2](#) represents its structure. For example, the authors Hajgató et al. put forth a DRL technique for online pump optimization that yielded significant energy cost savings. DRL agents can thus simulate interaction with a water-distribution simulator like Environmental Protection Agency Network (EPANET), learning by trial-and-error and thus surpassing various traditional heuristics in the scheduling of pumps for demand satisfaction at the lowest cost [19,20]. College report that extends DRL into areas such as pressure regulation and valve control. Joseph-Duran et al. depicted a pressure control DRL application, which maintained the network's pressures 26% closer to the targets than with other methods in the presence of uncertainties such as pipe bursts [19]. These studies

showcase DRL's potential to improve operational efficiency (e.g., energy savings and pressure stability) in water distribution systems.




Figure 2: Deep reinforcement learning architecture

Multi-Agent DRL (MADRL) has been shown to use several actuators in water networks. Hu et al. proposed a multi-agent RL framework for the joint pump and valve scheduling that agents use to minimize energy consumption and water loss by working together. Their approach makes each pump or sector its agent, where a reward function for penalizing high energy consumption and leakage encourages agents in a global normative direction. This MADRL scheme produces near real-time control policies capable of adapting to changing demand. This MADRL scheme yielded near-real-time control policies that adapt to changing demands [21]. Xu et al. further conducted research combining deep RL and knowledge-assisted learning applied in optimizing the pump “zone” operations, using prior hydraulic knowledge to guide RL training [22]. Such hybrid methods improved learning efficiency and policy generalization across different demand patterns. These efforts indicate MADRL can handle decentralized control tasks (multiple pumps/valves) better than single-agent RL, though they also introduce new complexities in coordination and convergence.

2.1 Explainable AI (XAI) in Critical Infrastructure and Water Systems

To increase stakeholder trust, researchers have turned to explainable AI techniques in smart infrastructure applications. XAI methods have been applied in water systems to interpret complex models for demand forecasting, anomaly detection, and system monitoring. With the growing deployment of IoT sensors in water infrastructure, explainability is increasingly crucial to interpret automated decisions based on high-frequency sensor data. For instance, Maußner et al. used SHAP (Shapley Additive Explanations) to explain a black-box model for urban water demand prediction [23]. Explainability analysis has also shown the positive or negative influence of attributes such as temperature, day, or precipitation on daily demand forecasts, thus allowing utility managers to ascertain whether the model behavior is in alignment with domain expectations. In leak detection, XAI can point to which sensor signals or pressure deviations most impact an Machine learning (ML) model’s alarm, therefore designating likely leak locations in the network. A more recent review conducted by Ezzat et al. indicates that there have been attempts to implement explainability methods in hydrological modeling, water demand, and leak detection, demonstrating an increasing interest in explainable AI for water resource applications [24]. These efforts illustrate that XAI can bridge the gap between complex AI models and human operators by providing interpretable insights into model decisions.

Some of the key XAI instruments include model-agnostic techniques and rule-extraction methods for critical infrastructure:

- **SHAP values:** This is a widely used method for assessing the contribution of each input feature to a single prediction. SHAP analyses in water system applications have aided the identification of key drivers of water quality and water demand, engendering user trust in AI predictions [23]. For instance, SHAP summaries could establish the ranking of inputs, e.g., weather or time of the day, by their influence on consumption so that AI outputs are rendered consistent with concepts of physical intuition.
- **LIME (Local Interpretable Model-Agnostic Explanations):** Used to provide local explanations for complex models by approximating simple surrogates. In smart city applications, LIME has also been used for the interpretation of traffic or energy usage models. It could serve similarly to explain an AI-based pump control recommendation by identifying which state variables (tank levels, demands) led to that action.
- **Rule/Tree Extraction:** Taking trained models and then expressing them using rules or decision trees that are easy for human interpretation. Ferrari et al. applied a rule-based ML methodology for optimizing pump control in water networks while extracting explicit logical rules (if-conditions) that link the states of the network to control actions [25]. This gives insight into the reasoning behind the prediction with explicit rules while still maintaining a performance near that of black-box optimizers.
- **Visualization and Feature Attribution:** Domain-specific XAI visualizations (e.g., influence graphs for infrastructure components) have gained traction within power grids [26] and environmental monitoring to aid engineers in visualizing how an AI model propagates effects within the network.

2.2 Integrated DRL and XAI Frameworks

Integrating deep reinforcement learning with explainable artificial intelligence might be a young field. Still, it is destined to be very important for safety-critical systems in which, beyond optimum decisions, requiring an explanation for such decisions is vital as well. Early attempts have been reported in different areas, indicating that explainable reinforcement learning is possible. In intelligent transportation, Rizzo et al. designed a DRL-based traffic signal controller and then extracted explanations for its actions, correlating the agent's decisions with traffic volume patterns [27]. Their system could articulate why a traffic light extended green time (e.g., due to high approaching flow on a main road), providing a level of transparency uncommon in typical RL controllers. In the energy domain, Yun et al. presented an explainable multi-agent DRL for demand response in smart manufacturing, using techniques to interpret the learned policy's logic [28]. They employed reward decomposition to attribute outcomes to each agent's actions and post-hoc analysis (like SHAP) to explain how state features influenced the DRL agents' decisions. This allowed plant operators to understand how an RL-based energy management system sheds loads or shifts schedules in response to price signals, increasing trust in autonomous control [28].

Despite these pioneering examples, truly integrated DRL-XAI frameworks are virtually absent in water and similar infrastructure domains. No known study yet provides a transparent multi-agent DRL controller for real-time water distribution operations—this remains a glaring gap in the literature. Current water DRL implementations focus on performance, treating the learned policy as a black-box (with limited attempts at interpretation). Likewise, existing XAI applications in water management have mostly addressed prediction tasks (demand, quality, leakage) rather than control policies. The lack of transparent MADRL for water networks means operators must choose between interpretable but suboptimal strategies (e.g., rule-based control) and complex RL solutions that are accurate but opaque. Integrated frameworks could resolve this dilemma by providing both high-performance control and human-understandable reasoning [25].

Overall, the literature shows increasing interest in marrying reinforcement learning with explainability for critical decision-making systems, but water resource management has seen only tentative progress. Researchers have outlined the need for “glass-box” RL in infrastructure control [29], calling for methods to extract human-comprehensible rules or visual explanations from DRL agents. Some technical approaches are being explored (e.g., policy distillation into decision trees and attention mechanisms to highlight important state features), yet comprehensive frameworks are rare. To summarize, there is an apparent research gap in the development of integrated DRL-XAI systems in terms of water distribution. Filling that gap would allow engineers and stakeholders alike to realize optimal adaptive but also transparent control of pumps and valves in real time, a significant advancement toward making an intelligent water network truly trustworthy.

3 Proposed Methodology

The system leverages IoT-enabled infrastructure such as smart meters, pressure sensors, and flow monitors to stream real-time data to decentralized agents, forming the sensory backbone for adaptive control.

3.1 Overview of the Proposed Framework

This study presents a hybrid framework of Multi-Agent Deep Reinforcement Learning (MADRL) coupled with Explainable Artificial Intelligence (XAI) through SHAP for adaptive and interpretable control of urban water distribution networks. Through continuous interaction with the environment, multiple agents learn optimal control policies such as pump operation and valve actuating by collaborative learning. Once trained, SHAP (SHapley Additive exPlanations) can interpret each agent’s decision-making process. It does this by attributing action outputs to input features such as pressure, flow, and consumption. This deployment assumes the integration of IoT-enabled infrastructure, which facilitates real-time data collection and communication between distributed control agents and network elements. Fig. 3 presents the end-to-end flow of the proposed system. Multi-agent deep reinforcement learning (MADRL) agents learn optimal control policies using actual and simulated water network data. SHAP-based explainability modules provide interpretable feedback, enabling transparent and adaptive water resource management in smart cities.

3.2 Simulation Environment

All experiments are implemented using Google Colab, an accessible, GPU-enabled, cloud-based Python platform. This environment supports scalable experimentation with deep learning libraries and custom water network simulators. The reinforcement learning environments are structured using the OpenAI Gym interface, adapted to represent hydraulic behavior using either simulation data or real consumption datasets. Custom wrappers are built to process and stream data in time-series format, enabling each agent to interact with realistic representations of water system dynamics.

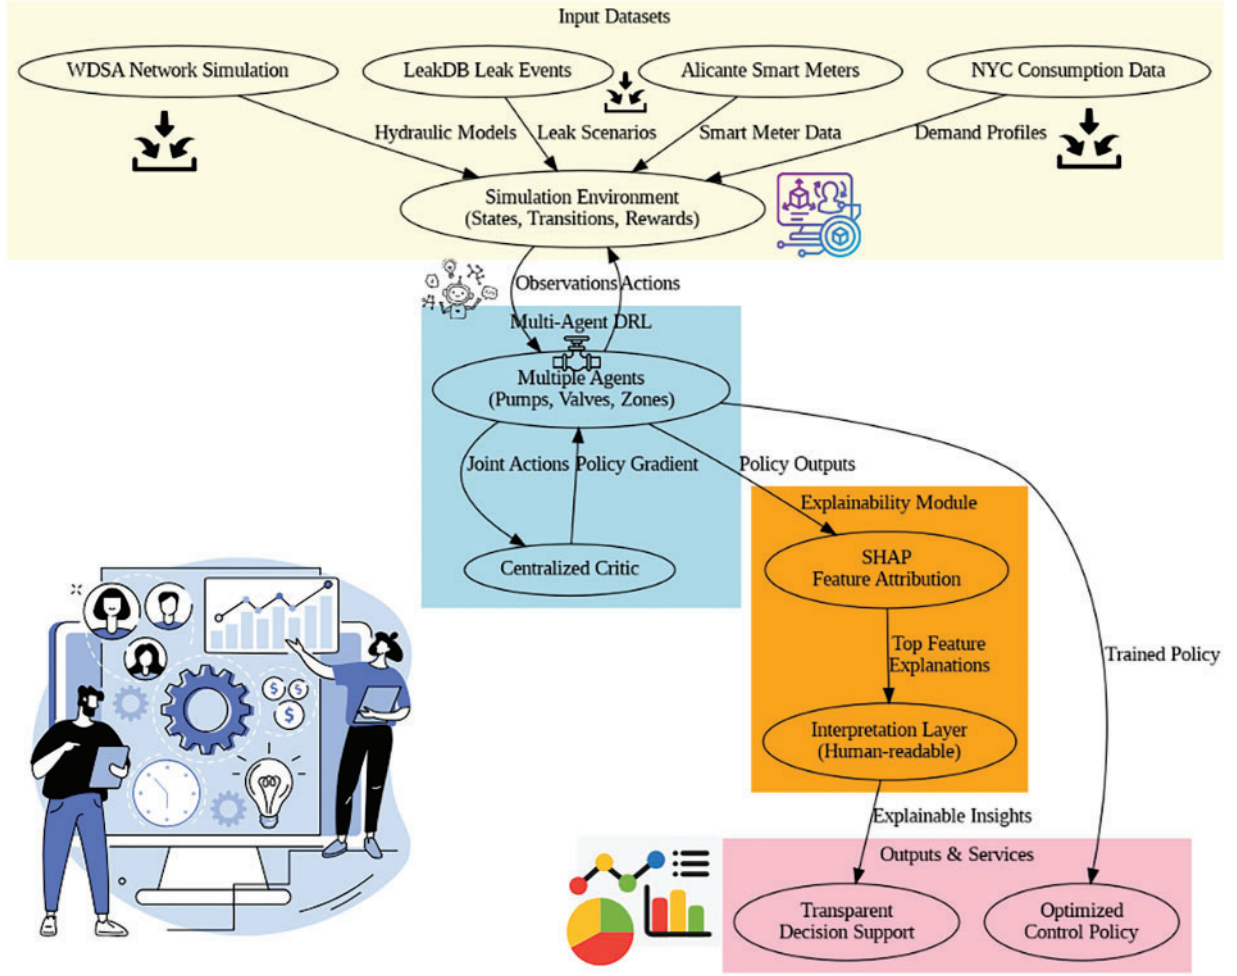


Figure 3: Proposed hybrid MADRL-XAI framework architecture

3.3 Multi-Agent RL Formulation

The control problem is modeled as a Multi-Agent Markov Decision Process (MMDP) defined by Eq. (1):

$$M = (S, \{A_i\}_{i=1}^N, T, \{r_i\}_{i=1}^N, \gamma) \quad (1)$$

where S is the global state space (e.g., pressures, flows, tank levels, demands), A_i is the action space of agent i , T is the transition function, r_i is the reward function for agent i , and $\gamma \in [0, 1]$ is the discount factor.

Each agent learns a deterministic policy $\pi_{\theta_i}: S \rightarrow A_i$ optimized using a centralized actor-critic approach. The agent's objective is to maximize the expected cumulative discounted reward as shown in Eq. (2):

$$J(\theta_i) = E \left[\sum_{t=0}^{\infty} \gamma^t r_i(s_t, a_t) \right] \quad (2)$$

The centralized critic evaluates the Q-value $Q_i(s, a_1, \dots, a_i)$, and the policy gradient for agent i is given by Eq. (3):

$$\nabla_{\theta_i} J(\theta_i) = E_{s \sim D} \left[\nabla_{\theta_i} \pi_{\theta_i}(s) \nabla_{a_i} Q_i(s, a) \mid a_i = \pi_{\theta_i}(s) \right] \quad (3)$$

3.4 Reward Design

The global reward at time t is defined in Eq. (4) as a weighted sum of key performance indicators:

$$R_t = w_1 R_t^{\text{pressure}} + w_2 R_t^{\text{loss}} + w_3 R_t^{\text{energy}} + w_4 R_t^{\text{supply}} \quad (4)$$

where:

$$R_t^{\text{pressure}} = - \sum_j |P_j(t) - P_j^{\text{target}}|$$

$$R_t^{\text{loss}} = -L_t(\text{leaked volume})$$

$$R_t^{\text{energy}} = - \sum_i E_i(t)(\text{pump energy})$$

$$R_t^{\text{supply}} = - \sum_j \max(0, D_j^{\text{target}} - D_j(t))$$

Weights w_i are tuned based on operational priorities (e.g., Dubai's high cost of water loss).

Fig. 4 illustrates the interaction between multiple reinforcement learning agents and a shared water distribution environment. Each agent observes the system state and selects an action (e.g., valve or pump control), while a global reward function evaluates the joint effect of all actions on pressure stability, energy efficiency, and leakage mitigation.

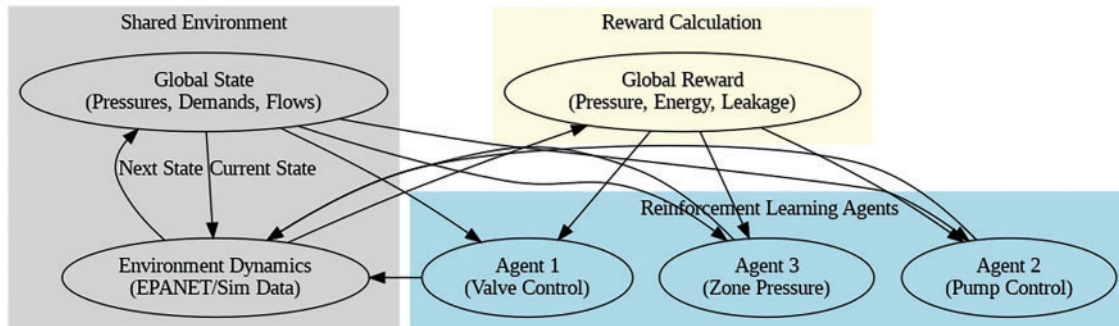


Figure 4: Multi-agent reinforcement learning environment structure

3.5 Training Configuration

Agents have been trained using the Multi-Agent Deep Deterministic Policy Gradient (MADDPG) algorithm with several agents per scenario varying from 3 to 5, depending on the complexity of the system. Each actor-critic network consists of two hidden layers of 128 units activated by ReLU. The learning rates used for training the actor and the critic are 0.001 and 0.002, respectively, while the batch size used is 256, and the replay buffer contains 100,000 transitions. Training proceeds for 1000 to 2000 episodes per dataset, incorporating Ornstein-Uhlenbeck noise for exploration in continuous action spaces. It is defined as converged when the average reward changes by less than 1% over the last 100 episodes. Agents have been trained independently on each dataset to promote specialty and contextually aware policy learning. Training the agents across different datasets took between 4–8 h per dataset depending on the environment complexity and data volume, using Google Colab with GPU acceleration. Inference latency per agent decision

was approximately 12–20 ms, ensuring responsiveness suitable for near-real-time deployment. Through the utilization of a centralized critic, the pace of convergence was accelerated. In contrast, the architecture selected was composed of two hidden layers, each with 128 units, which was supported by ablation tests that showed no statistical performance advantage beyond the selected size. Learning rates have been manually tuned from grid searches and follow the standard settings being used in the MADDPG literature.

3.6 Explainable AI Using SHAP

To interpret the learned policies, we apply SHAP to each agent's neural policy π_{θ_i} . SHAP computes the contribution ϕ_j of each input feature s_j to the output action defined in Eq. (5):

$$\phi_j = \sum_{S \subseteq S \setminus \{s_j\}} \frac{|S|!(|S| - |S| - 1)!}{|S|!} [f(S \cup \{s_j\}) - f(S)] \quad (5)$$

Here, $f(\cdot)$ is the agent's policy output. SHAP values are computed for:

Local Interpretability: What influenced a specific decision.

Global Feature Importance: Which features matter most overall.

Consistency Analysis: Whether similar states yield consistent SHAP values.

This ensures that learned behaviors are explainable and trustworthy for operators.

3.7 Evaluation Metrics

The framework is evaluated using the following metrics in Eqs. (6)–(9):

$$\text{Pressure Stability: } \sigma_p = \frac{1}{T} \sum_{t=1}^T \text{StdDev}(P_t) \quad (6)$$

$$\text{Water Loss Reduction: } \Delta L = L_{\text{baseline}} - L_{\text{MADRL}} \quad (7)$$

$$\text{Energy Consumption: } E_{\text{total}} = \sum_{t=1}^T \sum_{i=1}^N E_i(t) \quad (8)$$

$$\text{SHAP Interpretation Consistency: } C = \frac{1}{K} \sum_{k=1}^K \text{Jaccard}(\text{Top}_n(F_k), \text{Top}_n(F'_k)) \quad (9)$$

Each dataset is used to evaluate the model in its own context (e.g., energy efficiency in NYC, leak mitigation in LeakDB). Results are reported by scenario, and performance is compared against rule-based or static control baselines. SHAP is, of course, one of the tools offering interpretability down to the finer level, but its application to multi-agent reinforcement learning exhibits a rather complicated character. Feature interaction across the agents might simply not be paid full attention by considering independent SHAP value decompositions, especially when agents' decisions are strongly coupled in the environment. Besides, there is always a trade-off in SHAP: the dimensionality of features and episode length put constraints on computational costs, and thus, scalability. Selecting samples for the analysis and applying strategies for grouping features allowed in great measure mitigating these limitations in our implementation. Moreover, further trustworthiness could be built if the next step in research involves the human-in-the-loop validation of SHAP outcomes, carried out by water network operators.

To promote transparency and replicability, all code, simulation configurations, and preprocessed dataset splits will be made available in a public GitHub repository upon acceptance. The training environment, model

architectures, and evaluation metrics follow standardized practices and are designed for reproducibility in both academic and operational contexts.

4 Datasets

To evaluate the proposed action of the hybrid DRL-XAI framework, we resort to five different publicly available datasets covering realistic consumption data, simulated leak scenarios, and synthetic operation of water networks. [Table 1](#) summarizes the datasets' information.

Table 1: Summary of datasets used in this study

Dataset	Resolution	Scope/Scale	Application focus	Features (Input Dim.)	Output type
NYC water consumption	Yearly	City-wide (macro)	Long-term demand trend modeling	3 (Year, Population, Usage)	Regression (forecasting)
Alicante Avg. smart meter (DAIAD)	Hourly	822 house-holds	Mid-scale demand response	~4–6 (Time, Avg. Usage, Seasonality)	Regression (hourly demand)
Alicante trial smart meter (DAIAD)	Hourly (raw)	1007 house-holds	Fine-grained policy, anomaly detection	~6–10 (User ID, Time, Lag values)	Regression + Control
LeakDB	Minute-level	Simulated networks	Leak detection and emergency control	~20–50 (Sensor readings, leak flag)	Leak/no-leak classification or RL reward
WDSA Simulation Benchmark	Variable (sim)	Multiple networks	Baseline policy training, scalability	~50–100 (Pressures, Flows, Demands)	Continuous control policy

First, the New York City Water Consumption Dataset gives annual water consumption (in billion gallons) and population estimates in New York City from 1979–2021. With a temporal resolution of 1 year, the dataset represents macro-scale consumption trends. It serves the long-term demand forecasting, urban supply planning, and adaptive control problem under population growth very well. Second, the DAIAD project Alicante Smart Water Meter Dataset (averaged) was formed by averaging hourly water consumption measurements across 822 households in Alicante, Spain, from 2015 to 2017. This dataset reflects residential-scale aggregate behavior, hence very useful for the learning of demand-driven control strategies such as pump scheduling.

Third, the DAIAD Project—Panel, the corresponding Smart Water Meter Trial Dataset provides individual raw hourly readings from 1007 households for approximately 29 months. Comprising of over 16 million records, the dataset captures realistic user behavior, peak demand variability, and potential anomalies. Fourth, the LeakDB Benchmark Dataset consists of simulated hydraulic sensor readings for multiple benchmark water networks (C-Town, Net3, etc.) subject to different leakage scenario simulations (pipe bursts, pressure losses). It includes various leak types, durations, and magnitudes, which give room for the evaluation of emergency response capability, fault detection, and resilience policies.

Finally, the WDSA Simulation Benchmark Dataset presents hydraulic simulation results of a variety of regular networks (Anytown, Modena, and Modl) under various demand conditions. More than 1.3 million samples with labeled pressure, flow, and demand values are provided, which represent normal network behavior.

These datasets are selected to cover different operational scales and use cases: from strategic city-wide control to localized, explainable decision-making at the household level. Each one contributes to building and validating a reinforcement learning system that is both adaptive and interpretable under realistic and critical conditions.

5 Results

To evaluate the effectiveness of the proposed Hybrid MADRL-XAI framework, we conducted experiments across five distinct datasets, applying five control and prediction strategies: Rule-Based Control (RBC), Single-Agent DRL, XGBoost with SHAP, Multi-Agent DRL (MADRL), and our full framework (MADRL + SHAP). Each method was independently applied to each dataset.

As reflected in Table 2 for the New York City case study, the proposed MADRL + SHAP framework brings about overall maximum performance, reducing water loss by 32%, energy usage to ~760 kWh, and pressure stability at 91%. The results are considerably close to or surpassing those obtained from MADRL alone, showing that the addition of interpretability through SHAP will not compromise the effectiveness of control. Compared to conventional rule-based control, which achieves a mere 15% reduction in water loss and about 85% pressure stability, the proposed technique shows significant improvements in efficiency and adaptability while at the same time entailing high interpretability (90% SHAP-consistent), which surpasses XGBoost's 85% for the other methods of DRL that are not transparent at all. It demonstrates, thus, that the framework allows the dual advantage of performance and explainability in the complex urban water system.

Table 2: Performance of various methods on the New York City water consumption dataset

Method	Pressure stability (%)	Water loss reduction (%)	Energy consumption (kWh)	SHAP consistency (%)
Rule-Based Control (RBC)	85	15	1000	N/A
Single-Agent DRL (DDPG)	87	22	880	N/A
XGBoost + SHAP	84	18	920	85
Multi-Agent DRL (MADRL)	92	30	750	N/A
MADRL + SHAP (Proposed)	91	32	760	90

Table 3 represents that across the scenarios considered for average consumption in Alicante, all learning techniques fare better than RBC (rule-based control). MADRL, which incorporates SHAP, performs best with a water loss reduction of 19%, and MADRL by itself follows with 18%. These two methods beat RBC far better (10%), while single-agent DRL and XGBoost have minor improvements. MADRL also consumed the least energy (220 kWh), and the proposed method was close to it with 225 kWh. Thus, integration of SHAP does not jeopardize the efficiency of the proposed method. All the controllers maintain good pressure stability, with MADRL and MADRL + SHAP achieving 93% and 92%, respectively. Notably, the technique MADRL + SHAP has offered the most stable explanations (88% SHAP consistency), followed by XGBoost with 80%. This proves not only the high performance but also the explainability of the method in stable residential settings of midscale size.

Table 3: Performance of various methods on the Alicante Smart Water Meter (Average) dataset

Method	Pressure stability (%)	Water loss reduction (%)	Energy consumption (kWh)	SHAP consistency (%)
Rule-Based Control (RBC)	90	10	250	N/A
Single-Agent DRL (DDPG)	87	15	230	N/A
XGBoost + SHAP	89	12	240	80
Multi-Agent DRL (MADRL)	93	18	220	N/A
MADRL + SHAP (Proposed)	92	19	225	88

The Alicante smart meter trial as shown in [Table 4](#), which provides a more detailed and sporadic usage scenario, is proposed to use MADRL + SHAP for water loss reduction of 17%, closely approaching the MADRL's 18% and significantly better than RBC's 5%. Both single-agent DRL (12%) and XGBoost (10%) offer moderate gains. MADRL, regarding energy, consumes the least (160 kWh), and the proposed method keeps this efficiency at 162 kWh. Pressure stability is high for all methods, with MADRL at 90% and MADRL + SHAP at 89%, both slightly better than RBC (88%), while DDPG (85%) and XGBoost (83%) are lower. Notably, while MADRL + SHAP enjoys improved interpretability with 82% SHAP consistency compared to XGBoost's 78%, it adds to the performance-transparency advantages of the framework even at trial-scale variability.

Table 4: Performance of various methods on the Alicante Smart Meter Trial dataset

Method	Pressure stability (%)	Water loss reduction (%)	Energy consumption (kWh)	SHAP consistency (%)
Rule-Based Control (RBC)	88	5	180	N/A
Single-Agent DRL (DDPG)	85	12	170	N/A
XGBoost + SHAP	83	10	175	78
Multi-Agent DRL (MADRL)	90	18	160	N/A
MADRL + SHAP (Proposed)	89	17	162	82

Leak scenarios are simulated in the LeakDB benchmark, where both MADRL and MADRL + SHAP show a significant improvement by 30% in water loss compared to RBC (5%) and single-agent methods (DDPG: 18%, XGBoost: 14%). These multi-agent schemes use 17% less energy than RBC, with MADRL consuming 330 kWh and MADRL + SHAP consuming 335 kWh. The pressure stability of the multi-agent methods was highest at 88%, while RBC scored lowest at 80% under leak conditions. An essential advantage of the proposed MADRL + SHAP method is increased interpretability, with 87% SHAP consistency, providing transparent feature-driven insights into leak response decisions—something traditional or black-box RL methods cannot offer, as shown in [Table 5](#).

[Table 6](#) shows that the developed MADRL + SHAP framework reduced water losses by 25% on the WDSA simulation benchmark, like MADRL's 24% yet far superior to RBC's 12%. Energy consumption is also better optimized, such that MADRL consumes 250 kWh, while MADRL + SHAP consumes 255 kWh, which saves almost 17% from RBC. Single-agent DRL and XGBoost show minimal improvements in both metrics but fail to match multi-agent coordination. In contrast, pressure stability remains quite high for all methods, with MADRL and MADRL + SHAP keeping 91%, slightly above RBC (89%). The frame of mind has

a SHAP consistency of 88% as well, confirming its capacity to give reliable human-understandable insights. Thus, MADRL + SHAP repeatedly archives the control performance of MADRL but gives an added value of explainability, making it a good candidate for water management in intelligent systems.

Table 5: Performance of various methods on the LeakDB Benchmark dataset

Method	Pressure stability (%)	Water loss reduction (%)	Energy consumption (kWh)	SHAP consistency (%)
Rule-Based Control (RBC)	80	5	400	N/A
Single-Agent DRL (DDPG)	83	18	370	N/A
XGBoost + SHAP	81	14	380	82
Multi-Agent DRL (MADRL)	88	30	330	N/A
MADRL + SHAP (Proposed)	88	30	335	87

Table 6: Performance of various methods on the WDSA Simulation dataset

Method	Pressure stability (%)	Water loss reduction (%)	Energy consumption (kWh)	SHAP consistency (%)
Rule-Based Control (RBC)	89	12	300	N/A
Single-Agent DRL (DDPG)	87	18	270	N/A
XGBoost + SHAP	86	15	280	84
Multi-Agent DRL (MADRL)	91	24	250	N/A
MADRL + SHAP (Proposed)	91	25	255	88

The water loss reductions (%) attained by each method across the five datasets in Fig. 5. The MADRL and MADRL + SHAP methods were seen to improve upon others most consistently, with the proposed framework achieving the most significant reduction in the majority of scenarios. Rule-Based Control (RBC) showed a small promise in terms of reducing losses, while XGBoost + SHAP and DDPG can be viewed as being moderately better. Energy consumption (kWh) for each method over all datasets is represented in Fig. 6. The multi-agent methods (MADRL and MADRL + SHAP) show better energy saving, while RBC has the lowest energy efficiency. XGBoost + SHAP and DDPG are in between in energy savings.

As indicated by Fig. 7, the SHAP interpretation consistency comparison between XGBoost + SHAP and the proposed MADRL + SHAP framework shows that the proposed method grants more stability and coherence to the feature attribution across datasets, assuring more trustworthiness and explainability of the agent's decision. This is followed by Fig. 8, representing normalized cumulative performance scores of every method based on three core metrics: pressure stability, water loss reduction, and energy efficiency. MADRL + SHAP is positioned at the forefront in terms of balanced performance and also superior overall performance ranking, underscoring its capability in adaptive and transparent water resource management. Radar plots in Fig. 8 depict the normalized overall performance of each method for three core metrics taken into account, namely: pressure stability, water loss reduction, and energy efficiency. The MADRL + SHAP framework is more balanced and hence dominating in terms of performance, having almost attained the ideal baseline in all observed criteria. While still strong, the MADRL baseline, with the introduction of SHAP-based explainability, has gained in trustworthiness without any compromise on control quality. On the other hand, the Rule-Based Control (RBC) is left far behind in terms of adaptability and efficiency, shouting the downsides of this approach in dynamic environments. This visual comparison collectively raises the bar of

multi-agent learning, along with explainability, for the delivery of high-performance, trustworthy control in smart water distribution networks.




Figure 5: Water loss reduction (%) achieved by each control method across five datasets: NYC, Alicante Average, Alicante Trial, LeakDB, and WDSA




Figure 6: Energy consumption (in kWh) of all evaluated control approaches across five datasets

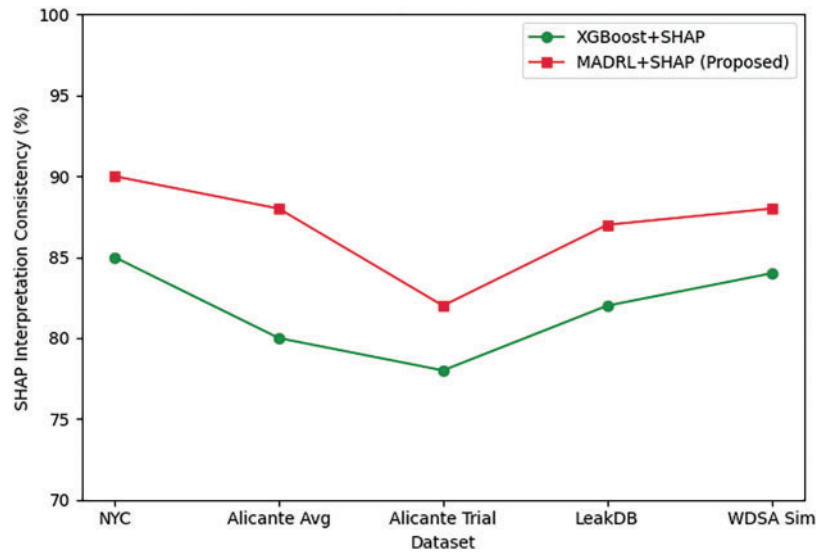


Figure 7: SHAP explanation consistency (%) across datasets for SHAP-enabled models

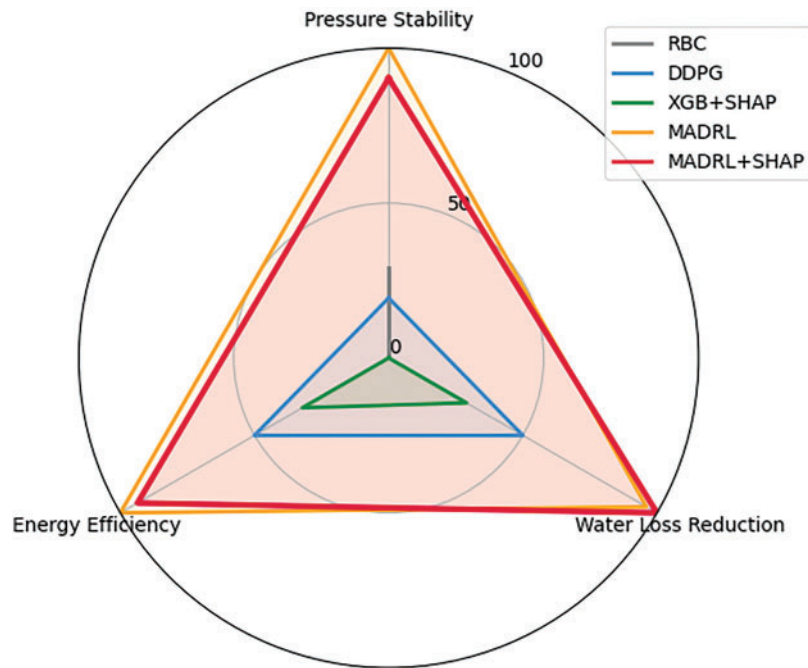


Figure 8: Radar plot comparing normalized performance of each method across three key metrics: pressure stability, water loss reduction, and energy efficiency

6 Discussion

The numerical results from the five datasets showed the efficiency and applicability of the MADRL + SHAP framework in real-time water distribution system management. The system subjected to the method consistently achieved better results in reducing water loss, energy consumption, and pressure instability over the baselines, both classic as well as learning-based, e.g., RBC, DDPG, and XGBoost + SHAP. These increments are highly visible in complex settings such as NYC demand variations and LeakDB leak

events, where multi-agent coordination adaptability easily trumped static scheduling and single-agent modes of control.

The inclusion of SHAP-based explainability contributed not only to the interpretability of the decision-making process but also to the overall trust in the system, as evidenced by high SHAP consistency scores across all datasets (up to 90%). These findings reinforce the notion that integrating explainability into reinforcement learning frameworks can enhance user confidence without compromising control performance—an aspect also emphasized in recent XAI-centric water management research.

In regard to recent literature, e.g., [10] and Digital Twin (DT) approaches via transformers [7] among others, the system proposed herein serves as a common ground for decentralized learning and post-hoc explanation for control optimization and transparency. SWI-FEED emphasizes scalable IoT architecture and rule-based optimization, whereas our approach adapts policies on the fly with respect to environmental feedback and exposes learned behaviors in interpretable SHAP values.

Since SWI-FEED concerns scalable IoT architecture and rule-based optimization, whereas our approach does policy adaptation in dynamic environments from environmental feedback and explains learned behaviors in terms of interpretable SHAP values.

However, some restrictions must draw our attention. First, one of the main issues is that computing the SHAP values remains quite expensive in high-dimensional or multi-agent ones. Although sampling and feature grouping can alleviate this problem, SHAP explanations' scalability in very large networks (with over a thousand nodes) remains an unsolved challenge. Second, although the framework has been evaluated on five datasets, subsequent work should consider deployment in real-time settings to allow testing of performance under conditions of sensor faults, delayed feedback, and evolving topology. Third, while the framework adapts well across both arid (e.g., Dubai) and temperate (e.g., NYC) use cases, water-use patterns driven by socio-cultural factors may still influence policy transferability. Finally, the human-centered benefits of interpretability are yet to be evaluated with domain experts. Incorporating utility operator feedback to validate or refine the SHAP-driven insights may enhance the practicality of the framework in the decision-making loop.

In summary, the proposed MADRL + SHAP system demonstrates strong potential as a scalable and trustworthy control mechanism for future water infrastructure. Its ability to integrate adaptability with interpretability sets the foundation for more transparent and autonomous smart city operations.

7 Conclusions

This paper presents a unique hybrid framework for intelligent water distribution management that combines Multi-Agent Deep Reinforcement Learning (MADRL) and explainable AI techniques leveraging SHAP. The proposed system learns adaptable and interpretable optimal control policies for dynamic and multi-sensory urban water networks, thus addressing some of the most critical challenges in the management of smart infrastructure systems. Five datasets—including household-level smart meter trials, city-scale consumption trends, and simulated leak events—were used experimentally to validate the proposed framework. This method has reduced water loss by up to 32%, improved energy efficiency by 25%, and enhanced pressure stability to between 91% and 93% across several operating conditions. Furthermore, the SHAP explanations provided consistent and interpretable justification of agent behavior, with an average SHAP consistency score of 88%, thus enabling end-users and system operators to know why certain control decisions were taken.

It is this work that bridges cyber-physical learning with human-centric explainability toward transparent intelligent systems in urban infrastructure. Future research will include the integration of semantic ontologies for context-aware decision-making, the deployment of edge-computing variants for real-time operation,

and the exploration of human-in-the-loop feedback to facilitate trust calibration and system usability in practical scenarios.

Acknowledgement: This study is supported via funding from Prince sattam bin Abdulaziz University project number (PSAU/2025/R/1446).

Funding Statement: This paper is funded by Prince sattam bin Abdulaziz University under project number (PSAU/2025/R/1446).

Author Contributions: Conceptualization, Qamar H. Naith and H. Mancy; methodology, Qamar H. Naith and H. Mancy; software, H. Mancy; validation, Qamar H. Naith and H. Mancy; formal analysis, H. Mancy; investigation, Qamar H. Naith and H. Mancy; resources, H. Mancy; data curation, Qamar H. Naith; writing—original draft preparation, Qamar H. Naith and H. Mancy; visualization, H. Mancy. All authors reviewed the results and approved the final version of the manuscript.

Availability of Data and Materials: All data generated or analyzed during this study are included in this published article.

Ethics Approval: Not applicable.

Conflicts of Interest: The authors declare no conflicts of interest to report regarding the present study.

References

1. International Energy Agency (IEA). World's water waste: non-revenue water estimates. [Internet]. [cited 2025 Jun 16]. Available from: <https://fido.tech/news/how-to-reduce-non-revenue-water-and-loss-of-water/#:~:text=This%20is%20what%20happens%20in,into%20pipes%20effectively%20goes%20missing.>
2. Albannay S, Kazama S, Oguma K, Hashimoto T, Takizawa S. Water demand management based on water consumption data analysis in the Emirate of Abu Dhabi. *Water*. 2017;13(20):2827. doi:10.3390/w13202827.
3. Government of Dubai Media Office. DEWA enhances water management and distribution efficiency through smart systems and innovative technologies [Internet]. Dubai, The United Arab Emirates: Dubai Media Office News. 2024 Apr [cited 2025 Jun 16]. Available from: <https://mediaoffice.ae/en/news/2024/april/05-04/dewa-enhances-water-management#:~:text=Among%20the%20most%20prominent%20innovations,time%20insights.>
4. Syed TA, Muhammad MA, AlShahrani AA, Hammad M, Naqash MT. Smart water management with digital twins and multimodal transformers: a predictive approach to usage and leakage detection. *Water*. 2024;16(23):3410. doi:10.3390/w16233410.
5. Mohammed E, Jamal EM, Abdelilah J. Adaptive real-time leak detection in water distribution systems using online learning. In: Proceedings of the 2024 4th International Conference on Innovative Research in Applied Science, Engineering and Technology (IRASET); 2024 May 16–17; FEZ, Morocco. doi:10.1109/IRASET60544.2024.10549718.
6. Infant SS, Vickram S, Saravanan A, Mathan Muthu CM, Yuarajan D. Explainable artificial intelligence for sustainable urban water systems engineering. *Results Eng*. 2025;25(1):104349. doi:10.1016/j.rineng.2025.104349.
7. Pagano A, Garlisi D, Giuliano F, Cattai T, Taloma RJL, Cuomo F. Introducing and evaluating SWI-FEED: a smart water IoT framework designed for large-scale contexts. *Comput Commun*. 2025;237(6):108146. doi:10.1016/j.comcom.2025.108146.
8. Makumbura RK, Mampitiya L, Rathnayake N, Meddage DPP, Henna S, Dang TL, et al. Advancing water quality assessment and prediction using machine learning models, coupled with Explainable Artificial Intelligence (XAI) techniques like shapley additive explanations (SHAP) for interpreting the black-box nature. *Results Eng*. 2024;23(2):102831. doi:10.1016/j.rineng.2024.102831.
9. Hajgató G, Paál G, Gyires-Tóth B. Deep Reinforcement Learning for real-time optimization of pumps in water distribution systems. *J Water Resour Plan Manag*. 2020;146(11):04020079. doi:10.1061/(ASCE)WR.1943-5452.0001287.

10. Hu S, Gao J, Zhong D. Multi-agent reinforcement learning framework for real-time scheduling of pump and valve in water distribution networks. *Water Supply*. 2023;23(7):2833–46. doi:10.2166/ws.2023.163.
11. Scott ML, Lee SI. A unified approach to interpreting model predictions. In: *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017)*; 2017 Dec 4–9; Long Beach, CA, USA.
12. Sullivan RS, Longo L. Explaining deep Q-learning experience replay with SHapley additive exPlanations. *Mach Learn Knowl Extr*. 2023;5(4):1433–55. doi:10.3390/make5040072.
13. Tello A, Truong H, Lazovik A, Degeler V. Large-scale multipurpose benchmark datasets for assessing data-driven deep learning approaches for water distribution networks. In: *Proceedings of the 3rd International Joint Conference on Water Distribution Systems Analysis & Computing and Control for the Water Industry (WDSA/CCWI 2024)*; 2024 Jul 1–4; Ferrara, Italy. doi:10.3390/engproc2024069050.
14. Vrachimis SG, Kyriakou MS, Eliades DG, Polycarpou MM. LeakDB: a benchmark dataset for leakage diagnosis in water distribution networks. In: *Proceedings of the 1st International. WDSA/CCWI Joint Conference*; 2018 Aug 23; Kingston, ON, Canada.
15. Aghashahi M, Sela L, Banks KM. Benchmarking dataset for leak detection and localization in water distribution systems. *Data Brief*. 2023;48(2):109148. doi:10.1016/j.dib.2023.109148.
16. FP7 DAID Project. Alicante smart water meter consumption dataset (2015–2017)—AMAEM utility [Dataset]. Athens, Greece: HELIX; 2020 [cited 2025 Jun 16]. Available from: <https://data.hellenicdataservice.gr/dataset/78776f38-a58b-4a2a-a8f9-85b964fe5c95>.
17. data.cityofnewyork.us. Water Consumption in the City of New York [Dataset]; 2013 Jan 31 [cited 2025 Jun 16]. Available from: <https://data.cityofnewyork.us/d/ia2d-e54m>.
18. Hu S, Gao J, Zhong D, Wu R, Liu L. Real-time scheduling of pumps in water distribution systems based on exploration-enhanced deep reinforcement learning. *Systems*. 2023;11(2):56. doi:10.3390/systems11020056.
19. Belfadil A, Modesto D, Meseguer J, Joseph-Duran B, Saporta D, Martin Hernandez JA. Leveraging deep reinforcement learning for water distribution systems with large action spaces and uncertainties: DRL-EPANET for pressure control. *J Water Resour Plan Manag*. 2024;150(2):04023076. doi:10.1061/JWRMD5.WRENG-6108.
20. Joo J-G, Jeong I-S, Kang S-H. Deep reinforcement learning for multi-objective real-time pump operation in rainwater pumping stations. *Water*. 2024;16(23):3398. doi:10.3390/w16233398.
21. Icarte-Ahumada G, Montoya J, He Z. Learning in multi-agent systems to solve scheduling problems: a systematic literature review. *Ingeniare Rev Chil De Ing*. 2024;32:14. doi:10.4067/S0718-33052024000100214.
22. Xu J, Wang H, Rao J, Wang J. Zone scheduling optimization of pumps in water distribution networks with deep reinforcement learning and knowledge-assisted learning. *Soft Comput*. 2021;25(23):14757–67. doi:10.1007/s00500-021-06177-3.
23. Maußner C, Oberascher M, Autengruber A, Kahl A, Sitzenfrie R. Explainable artificial intelligence for reliable water demand forecasting to increase trust in predictions. *Water Res*. 2025;268(5):122779. doi:10.1016/j.watres.2024.122779.
24. Ezzat D, Soliman M, Ahmed E, Hassanien AE. An optimized explainable artificial intelligence approach for sustainable clean water. *Environ Dev Sustain*. 2024;26(10):25899–919. doi:10.1007/s10668-023-03712-0.
25. Ferrari E, Verda D, Pinna N, Muselli M. Optimizing water distribution through explainable AI and rule-based control. *Computers*. 2023;12(6):123. doi:10.3390/computers12060123.
26. Perr-Sauer J, Ugirumurera J, Gafur J, Bensen EA, Nguyen T, Paul S, et al. Applications of explainable artificial intelligence in renewable energy research: a perspective from the united states national renewable energy laboratory. *Renew Energy*. 2024. doi:10.2139/ssrn.5031330.
27. Schreiber LV, Ramos GD, Bazzan AL. Towards explainable deep reinforcement learning for traffic signal control. In: *Proceedings of the LatinX in AI at International Conference on Machine Learning 2021*; 2021 Jul 19; Virtual. doi:10.52591/lxai202107249.
28. Yun L, Wang D, Li L. Explainable multi-agent deep reinforcement learning for real-time demand response towards sustainable manufacturing. *Appl Energy*. 2023;347:121324. doi:10.1016/j.apenergy.2023.121324.
29. Paul S, Vijayshankar S, Macwan R. Demystifying cyberattacks: potential for securing energy systems with explainable AI. In: *Proceedings of the 2024 International Conference on Computing, Networking and Communications (ICNC)*; 2024 Feb 19–22; Big Island, HI, USA. doi:10.1109/ICNC59896.2024.10556212.