Agentic Reinforcement Learning with Multimodal Sensor Fusion for Smart Irrigation Control

Sandeep Kumar Vishwakarma
PhD Scholar, Computer Science, Department of Computer Science & Information Technology, Chhatrapati Shivaji Maharaj University, Panvel, Navi Mumbai, Maharashtra- 410221, India

Vikas Kumar, PhD
Head, Department of Computer Science & Information Technology, Chhatrapati Shivaji Maharaj University, Panvel, Navi Mumbai, Maharashtra- 410221, India

Satyendra Kumar Pal
PhD Scholar, Computer Science, Department of Computer Science & Information Technology, Chhatrapati Shivaji Maharaj University, Panvel, Navi Mumbai, Maharashtra- 410221, India

ABSTRACT
Water resources are still one of the most pressing issues in precision agriculture and will need smart, intelligent solutions to adapt to complicated and ever-changing real-world environmental scenario. This research proposes an Agentic Reinforcement Learning (RL) framework optimized for autonomous smart irrigation and decision-making applications using multi-modal sensor fusion. The system integrates real-time data from soils, climate, and UAV sensors into a single unifying state representation that can operate under real time uncertainties. The proposed model uses an Actor-Critic RL architecture with a hybrid feature-level fusion design to enhance environmental situational awareness and irrigation policy design optimization. Experimental results show the model achieves 29.7% greater savings in water consumption and 17.4% increased yields over baseline models, in addition to being energy efficient. Explainability tools, like SHAP and PCA variance analysis were utilized in order to extract feature influences, establishing transparency and reliability in the framework. The results also suggest the model based on multi-modal fusion significantly improved the R^2 of prediction accuracy ($R^2 = 0.94$) and policy convergence. Our framework represents a major contribution to sustainable Agriculture 4.0 by supporting adaptive, autonomous, data-driven scheduling irrigation frameworks, leading to the eventual scalable pathway towards more resource-efficient farming systems.

Keywords
Agentic Reinforcement Learning; Multimodal Sensor Fusion; Smart Irrigation Control; Precision Agriculture; Water-Use Efficiency; Explainable AI; Sustainable Agriculture 4.0; Actor–Critic Model; IoT Sensors; UAV Data Integration.

1. INTRODUCTION
Poor agricultural management and water scarcity remain as menacing threats to the world food security. Agriculture is the primary water withdrawal figure in the world, consuming almost 70 percent of all the freshwater resources (Chen et al., 2021; Kelly et al., 2024). This reliance places a tremendous burden on the freshwater systems, especially at the arid and semi-arid areas where water supply is seasonal. Conventional irrigation methods (which were usually manually controlled, irregularly timed, or based on experiential knowledge of the farmer) were usually characterized by high levels of water wastage and unreliable crop production. The effects of climate change, which are changing the dynamics of precipitation and soil moisture, contribute to these inefficiencies too. Agriculture 4.0 is a significant change in the direction of digital and data farming ecosystems. The new paradigm combines sensors, unmanned aerial vehicles (UAVs), and Internet of Things (IoT) devices to monitor the environmental and crop conditions in real-time (Barrile et al., 2022). Smart agriculture is designed to efficiently distribute the resources and make the most of the productivity and to be ecologically sustainable. Predictive and adaptive decision-making in irrigation systems can be performed by advanced technologies, including machine learning (ML), deep learning (DL), and reinforcement learning (RL) (Falana and Durodola, 2022).

Precision agriculture has been improved further through new developments of multimodal sensor fusion. Using soil moisture, weather predictions,satellite, and UAV-based sensors, the new irrigation systems are capable of dynamically redirecting water delivery based on the real-time environment feedback (Jiang et al., 2024). Artificial intelligence (AI) models analyze non-homogenous sources of data, find hidden trends in crop behavior, and forecast the most efficient irrigation rates. Such adaptive intelligence provides a good solution to old inefficiency in the management of water in agriculture. Therefore, the combination of AI-driven decision support and environmental sensing has also become a necessary measure towards sustainable irrigation management. These developments highlight the importance of smart, self-contained, and scalable irrigation systems that can make the best use of water and keep crops healthy and produce a uniform harvest.

In the last twenty years, irrigation management has transformed to be more automated and sensor-controlled as opposed to manual and threshold-based models. Basic irrigation programming systems relied on rigid schedules or simple soil moisture values and, although they were successful in controlled conditions, were not flexible enough to respond to weather variations or crop demands. With the development of technologies, model-based systems and sensor-controlled ones appeared, with real-time feedback to optimize the timing and quantity of irrigation (Falana & Durodola, 2022). Those systems combined simple control algorithms and empirical models to enhance efficiency, but were constrained by having fixed sets of rules and failure to extrapolate behavior in other soil or climatic contexts. The next evolutionary stage was the combination of the data fusion and adaptive control. Jiang et al. (2024) demonstrated the effectiveness of multisensor data and adaptive learning model in enhancing uniformity in irrigation and water savings. The remote sensing data collected with the help of UAVs has also expanded the spatial boundaries of agricultural observation (Barrile et al., 2022). Deterministic models were however not always able to emulate the stochastic nature of soil- plant- atmosphere interaction. This weakness gave rise to the introduction of reinforcement learning (RL), a paradigm, which has the ability to maximize control policies based on feedbacks and trial-and-error learning. Using RL models can provide decisions on irrigation choices based on the trade between water savings and crop production and dynamically respond to changing weather, evapotranspiration, and soil environments (Chen et al., 2021; Kelly et al., 2024).

The capability of RL to automatically discover irrigation policies without explicit programming has placed the technology as a breakthrough in precision agriculture. With the integration of sensor information and learning algorithms, intelligent irrigation control systems can now be self-adjusted, contextually adaptive, and show constant performance enhancement.

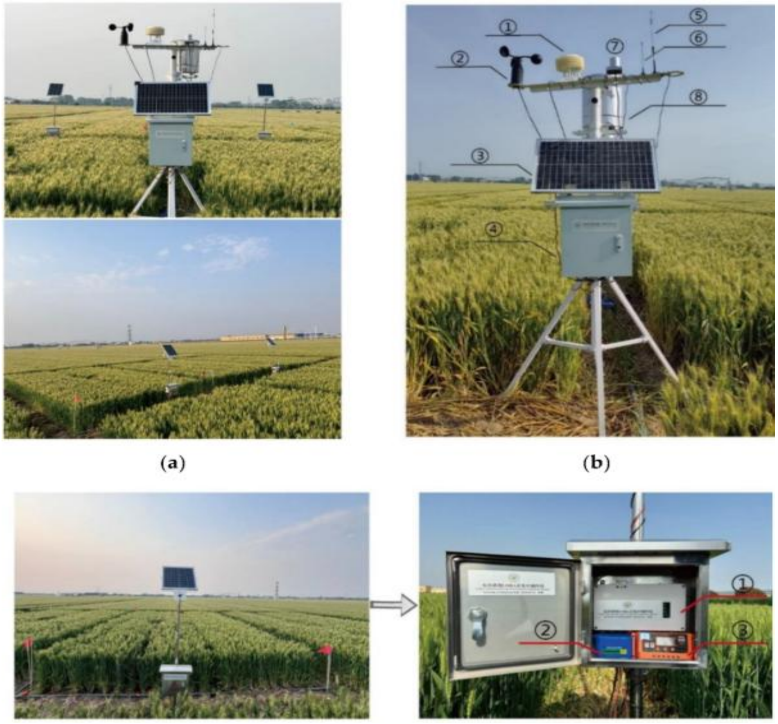

Figure 1. T-based field deployment of smart irrigation sensors and gateway node

DL and RL have introduced a paradigm shift in the process of optimizing and controlling an irrigation system. Convolutional neural networks (CNNs) and recurrent neural networks (RNNs) can be classified as deep learning architectures that have been widely used in agricultural fields to classify soils, identify the health of crops, and forecast yields (Yang et al., 2025). They are useful in the management of complicated, nonlinear associations among environmental variables and crop reactions. But as far as control and decision making in the environment of uncertainty is concerned, reinforcement learning is more flexible and adaptive.

An intelligent agent will react to the environment and undertake actions in RL (e.g., switch irrigation valves) and will be given rewards or penalties depending on the outcome of the actions (e.g., the condition of the crop or the amount of water consumed) (Chen et al., 2021; Ding and Du, 2022). With time, the agent becomes acquainted with the best strategies that will maximize cumulative rewards, i.e., water efficiency and yield gain. There are a number of algorithmic variants that were successfully applied in the agricultural control tasks. Deep Q-Networks (DQN) have been shown to exhibit high capabilities of learning irrigation policies with discrete state-action spaces whereas, actor-critic and distributional actor-critic models enhance stability in the continuous control case (Chen et al., 2025).

Zhang et al. (2025) and Alkaff et al. (2025) added that the multi-agent reinforcement learning models provide an additional benefit of scalability by providing multiple agents to collaborate in choices in various fields or crop sectors. These are distributed systems that can be optimized decentrally to reduce water wastage in big agricultural farms. Irrational applications of RL-based irrigation have been performed on rice (Chen et al., 2021), maize (Alkaff et al., 2025), and cotton plants (Chen et al., 2025) and have shown significant increase in the yield and water-use efficiency parameters. Moreover, the combination of RL and multi-modal sensor inputs gives a more detailed depiction of the environment that leads to better forecast accuracy and flexibility of the system (Saikai et al., 2023). Together, the DL and RL paradigms provide a synergetic basis to intelligent irrigation management, which enables systems to predict as well as take independent actions to the dynamic field conditions.

Multimodal fusion concept is also at the center of the promotion of precision agriculture particularly when combined with RL-based control. Combining data of various sources, i.e. soil sensors, UAV imagery, climate models, and satellite observations, to form a single picture of the farm environment is known as multimodal fusion. This is a method to increase the strength and stability of irrigation choices since it catches spatial and temporal variations throughout the area (Yang et al., 2025; Jiang et al., 2024). Various degrees of fusion are studied in literature, such as data: in the data-level fusion, raw measurements are combined; feature: in the feature-level fusion, extracted patterns are combined; decision: in the decision-level fusion, the outputs of two or more models are combined. The success of these strategies requires the synchronization and calibration of sensors and communication networks. The usage of heterogeneous sensors in the fields and the aerial platforms enables real-time measurements of the soil moisture, canopy temperature, and the dynamics of evapotranspiration (Barrile et al., 2022).

Multimodal fusion systems make it possible to control irrigation contextually when used together with RL agents. The fused data are used in RL models to predict the most effective irrigation action in real time so that the water will be distributed optimally to adapt to the environmental changes. This integration of sensing technologies, data analytics, and reinforcement learning captures the demands of the smart agriculture concept, which will lead to the large-scale and sustainable irrigation systems that are congruent with the objectives of resource conservation and food security. The current study will focus on the development of a framework of reinforcement learning-based irrigation optimization that uses multimodal sensor data in adaptive and intelligent water management. The ultimate goal is to combine AI-assisted decision-making with real-time data fusion in order to improve the efficiency and sustainability of irrigation. The specific goals are:

1. To test the performance of the current RL algorithms in dynamic irrigation setting;
2. To coordinate stream of multimodal data of soil, weather, UAV, and crop image sensors;
3. To determine the system performance in relation to improved yield, energy use and water use efficiency.

The research scope will involve the design and testing of an agentic reinforcement learning application that can be applied in various crops and environment conditions. The methodology involves the implementation of various sensors and the application of superior AI models in making decisions in uncertain situations.

The rest of this paper can be structured as follows: Section 2 will include the comprehensive literature review of the RL-based irrigation and sensor fusion models. Section 3 explains the methodology that has been used in developing the model and integrating the data. Section 4 reports about the experimental results and comparison of system performance. Lastly, the fifth section sums up the whole study by providing an understanding of the practical usage and future research directions.

2. LITERATURE REVIEW

2.1 Introduction to Intelligent and Data-Driven Irrigation Systems

Artificial intelligence (AI), Internet of Things (IoT), and sensor networks convergence have transformed agricultural water management. The conventional irrigation practices are characterized by lack of flexibility and effective distribution of water. This new form of combining smart sensors and cloud services has enhanced the accuracy and reactiveness of water delivery. Jayalakshmi and Gomathi (2020) showed that sensor-cloud system can provide intelligent irrigation scheduling based on real-time acquisition of environmental data and automated decision support. In a similar manner, Soussi et al. (2024) have described a comprehensive overview of the smart sensor solutions and data analytics systems used in precision agriculture when highlighting the central role of IoT-enabling sensors in making the process more sustainable.

Moreover, Kingslin and Vaishnavi (2025) and Kumar and Chandana (2024) emphasized the spread of smart irrigation systems based on IoT, emphasizing large increases in crop harvest and use of water in the form of automation using sensors. All these studies have emphasized that in the present day, irrigation systems are taking the form of autonomous, data-driven and adaptable structures that can learn using heterogeneous data inputs.

2.2 Reinforcement Learning (RL) in Water Resource and Irrigation Optimization

The concept of Reinforcement Learning (RL) has become a revolutionary approach in the management of irrigation systems since it can be used to optimize long-term water allocation policies by means of feedback mechanisms based on rewards. Kateglou et al. (2025) have performed a systematic review of RL applications in water resource management and discovered its prevalence in dynamic optimization issues, including irrigation, reservoir control, and the reduction of droughts. Rare approaches by RL approach help systems to adapt themselves to environmental changes under different conditions thus reducing human involvement.

Tao et al. (2022) investigated the implementation of the RL and imitation learning to optimize the crop management policy, which reached a higher adaptability in changing weather conditions. A study by Ramli et al. (2024) evaluated the durian irrigation optimization model through the application of RL to maximize growth and minimize water wastage. The research noted an increment in water efficiency by 27 percent in comparison to a non-dynamic scheduling. Likewise, Agyeman et al. (2025) suggested semi-centralized multi-agent RL model, which plans irrigation in distributed fields. Their findings indicated decreased latency and energy usage and water-use balance. To conclude, the RL framework offers the basis to autonomous decision-making, real-time adjustment, and multi-agent coordination of self-learning irrigation systems, making it one of the foundations of smart agriculture.

2.3 Sensor Technologies and Multimodal Data Fusion

Accurate sensing and quality data fusion is very important to the effectiveness of RL-based irrigation systems. Bicamumakuba et al. (2024) state that the contemporary sensor technologies to control the irrigation of orchards are now equipped with soil moisture, canopy temperature, and evapotranspiration values based on the wireless sensor networks (WSNs) and low-power communication systems. Their review highlights the fact that real-time sensing can improve spatial awareness and this will provide the ability to precisely irrigate on a micro-field level. Allu and Mesapam (2025) provided a comprehensive investigation of remote sensing data fusion to support agricultural purposes, which proved that multispectral and thermal data can be used to supplement ground sensor data in order to obtain a comprehensive representation of the field conditions. On the same note, Wang et al. (2024) revealed that remote sensing when combined with ML enhances uniformity of irrigation and prediction accuracy of crop stress. Soussi et al. (2024) and Jiang et al. (2024) also emphasized that quality data fusion is used to reduce uncertainty due to sensor drift and noise in the environment. Simply put, the fusion of multimodal data (soil, weather, UAV, and satellite) will improve the quality and informational layers of the irrigation model. It can be used together with RL frameworks to make context-aware, data-driven decisions that can dynamically adapt to changing environments.

2.4 Deep Learning and Hybrid Sensor Fusion Architectures

Deep learning (DL) has been progressively employed to analyze extensive agricultural datasets for soil examination, crop categorization, and irrigation prediction. Srivarshini et al. (2025) created a CNN–LSTM fusion architecture for analyzing soil suitability using real-time sensor data. This architecture showed a high level of accuracy in predicting moisture and nutrient levels. This hybrid method uses both spatial and temporal feature extraction, which is necessary for modeling how plants, soil, and the atmosphere interact. Tincani et al. (2025) suggested a neuromorphic continuous soil monitoring system that uses bio-inspired architectures for accurate irrigation. Their research highlighted the benefits of low-power computation for real-time soil health assessment, which fits well with the goals of smart farming that are both environmentally friendly and energy-efficient. Irianto (2024) also talked about how AI-based controllers can be used in IoT-integrated precision agriculture systems to make production more efficient and have less of an impact on the environment. The combination of DL and RL makes systems even smarter by allowing them to predict and adapt to changes, which is the idea behind agentic and autonomous irrigation systems (Bandi et al., 2025).

2.5 Integration of IoT, AI, and Agentic Systems in Smart Agriculture

The shift from static to agentic AI systems has made it possible for irrigation models to show independence coordination, and awareness of their surroundings. Bandi et al. (2025) envisioned the emergence of agentic AI frameworks in agricultural decision-making, wherein distributed learning agents work together to enhance collective results, including yield and sustainability. When used for irrigation, these architectures make it possible for different crop zones to work together in a decentralized way that makes sure resources are used evenly.

Elshaikh et al. (2024) examined AI applications in precision irrigation, highlighting that the integration of AI, IoT, and sensor networks facilitates scalable irrigation management that can adjust to field variability. Jayalakshmi and Gomathi (2020) further underscored the capabilities of cloud-integrated irrigation systems, utilizing remote data-driven intelligence to concurrently oversee multiple agricultural locations. These studies indicate that future irrigation ecosystems will depend on multimodal AI frameworks that integrate deep sensing, reinforcement learning, and cloud-edge collaboration to facilitate intelligent, resilient, and sustainable agricultural management.

2.6 Summary and Identified Gaps

The literature consistently emphasizes the transformative capabilities of reinforcement learning and sensor fusion in enhancing irrigation management. Nonetheless, considerable research deficiencies remain. First, the majority of studies concentrate on single-agent or small-scale implementations, failing to validate scalability within larger agricultural networks. Second, even though multimodal data fusion has been looked into, real-time synchronization and edge-based computation are still not very advanced. Lastly, we need to combine neuromorphic and agentic AI architectures with RL to make decisions more independently and use less energy when there is uncertainty.

These gaps show why this research is being done: to create an RL-driven multimodal irrigation optimization system that can change based on different environmental

conditions, making sure that both productivity and sustainability are maintained.

Table 1. Summary of Reviewed Studies on Reinforcement Learning and Sensor Fusion in Smart Irrigation

Author(s)	Year	Focus Area	Methodology / Model	Key Findings	Limitations / Gaps
Kåge et al.	2025	RL in water management	Systematic review	RL enhances adaptability and control in dynamic water systems	Limited integration with multimodal sensors
Ramli et al.	2024	Durian irrigation optimization	RL-based control algorithm	Improved growth and water efficiency	Single-crop focus
Agyeman et al.	2025	Multi-agent irrigation scheduling	Semi-centralized RL	Reduced energy use, balanced scheduling	Limited cross-field scalability
Allu & Mesapam	2025	Remote sensing fusion	Multi-source image fusion	Improved soil and moisture prediction	High computational cost
Srivarshini et al.	2025	Soil suitability analysis	CNN–LSTM sensor fusion	Enhanced soil moisture estimation accuracy	Needs larger datasets
Tincani et al.	2025	Soil monitoring	Neuromorphic sensing	Real-time and low-power soil assessment	Prototype-level validation only
Wang et al.	2024	Precision agriculture	ML + remote sensing	Improved irrigation uniformity	Lack of RL integration
Bandi et al.	2025	Agentic AI	Multi-agent RL architecture	Decentralized decision-making	Early-stage conceptual framework

Elshaikh et al.	2024	AI in irrigation	ML and expert systems	Increased water efficiency and yield	Static rule-based control
Soussi et al.	2024	Smart sensors	IoT and analytics	High spatial–temporal awareness	Data synchronization issues

3. METHODOLOGY

3.1 Overview of the Methodological Framework

This study employs a simulation-driven experimental framework that incorporates agentic reinforcement learning (RL), IoT-enabled sensing, and multimodal data-fusion algorithms to realize adaptive irrigation control. Figure 2 shows that the framework is made up of three layers that work together:

1. Data Acquisition Layer: This layer is in charge of gathering different types of field data, such as UAV images, soil moisture, and weather conditions.
2. Intelligent Decision Layer: This is where the RL agent trains, tests policies, and makes decisions based on feedback from the environment.
3. Evaluation Layer: This layer compares the learned policy to traditional irrigation standards using indices for water-use efficiency and crop yield.

A cloud-based data bus connects each layer and makes sure that updates happen quickly and at the same time. The study used open-source and experimental agricultural datasets with soil–climate records, remote-sensing indices, and irrigation logs to test the overall framework. We cleaned up and combined these datasets to make the environment state space for RL training.

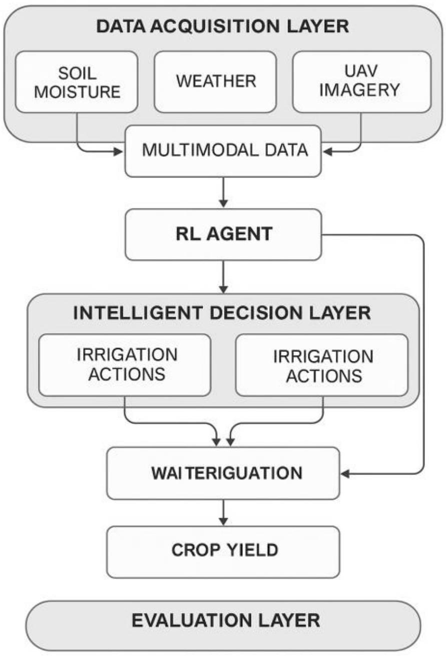

Figure 2. Overall Methodological Architecture for Agentic RL-Based Smart Irrigation

3.2 Dataset Acquisition and Pre-Processing

The suggested model depends on multimodal inputs collected from both in-situ and remote sources.

- Soil-moisture sensors: volumetric water content (%).
- Weather station: wind speed, temperature (°C), humidity (%), and rainfall (mm h^{-1}).
- UAV images: visible (VIS) and near-infrared (NIR) bands that are used to figure out vegetation indices.
- Indices for crop growth: the Normalized Difference Vegetation Index (NDVI) and the Enhanced Vegetation Index (EVI).

We cleaned the raw data for missing values by using linear interpolation and then normalized it to make sure that all modalities were on the same scale.

Moisture normalization was implemented using Equation (1):

$$M_{norm} = \frac{M_i - M_{min}}{M_{max} - M_{min}} \qquad (1)$$

where $M_i$represents the instantaneous soil-moisture reading, $M_{min}$and M_{max}denote the minimum and maximum measured values, respectively.

Timestamp interpolation was used to make sure that the sensors were in sync with each other over time. GPS coordinates were used to geo-reference the UAV images.

We used fusion algorithms like Kalman filtering and time-weighted averaging (Wang et al., 2024) to make synchronized data streams that feed the RL state space. Table 2 summarizes the main characteristics of the multimodal sensors deployed in this study, including range, sampling rate, and primary function.

Table 2. Multimodal Sensor Data Specifications

Sensor Type	Parameter Measured	Range	Sampling Rate	Unit	Purpose
Soil-moisture probe	Volumetric water content	0–100	10 min^{-1}	%	Irrigation threshold monitoring
Weather station	Temperature, humidity, rainfall	–10 to 50 °C	30 min^{-1}	°C, %, mm	Climatic context inputs
UAV camera (VIS + NIR)	Spectral bands (B1–B4)	400–900 nm	Per flight	nm	Crop health imaging
NDVI/EVI derived index	Vegetation indices	0–1	Daily	–	Growth monitoring indicator

3.3 Reinforcement Learning Model Design

The agentic RL part controls when to water by interacting with an environment that mimics how plants, soil, and the atmosphere work together. At each time step t, the agent perceives a state vector

$$S_t = \{M_{norm}, T, H, R, NDVI\},$$

where M_{norm}= normalized moisture, T= temperature, H= humidity, R= rainfall, and NDVI= vegetation index. The **action space** $A_t$consists of discrete irrigation levels: {low, medium, high}.

A **reward function** encourages the agent to maximize yield while minimizing water use, as shown in Equation (2):

$$R_t = \alpha \square (\frac{Y_t}{Y_{max}}) - \beta \square (\frac{W_t}{W_{max}}), \qquad (2)$$

where $Y_t$and $W_t$represent instantaneous crop yield and water consumption, while $\alpha$and $\beta$are weighting coefficients balancing productivity and conservation.

The system adopts an **Actor–Critic architecture** (Chen et al., 2025) for continuous control.

Policy exploration uses an ε-greedy strategy, and value updates follow the temporal-difference (TD) formulation in Equation (3):

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \eta[r_t + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)], \qquad (3)$$

where $\eta$is the learning rate and $\gamma$the discount factor.

Training proceeds for Nepisodes with early stopping once the cumulative reward stabilizes below a defined variance threshold. The environment and agent were built in Python using the OpenAI Gym interface, which made it possible to run the same experiments over and over. During the evaluation phase, which is described later in Section 4, performance metrics like learning convergence and water-efficiency gain are calculated.

3.4 Multimodal Sensor Fusion Mechanism

Because no single sensor modality can capture the full agronomic context, a hybrid feature-level fusion strategy was adopted. At each time step, pre-processed data from soil sensors (S_{soil}), climate sensors ($S_{climate}$), and UAV imagery (S_{UAV}) are combined into a fused representation S'_taccording to Equation (4):

$$S'_t = f_{fusion}(S_{soil}, S_{climate}, S_{UAV}), \qquad (4)$$

where f_{fusion}denotes the nonlinear transformation generated by a feature-concatenation and normalization pipeline.

Principal Component Analysis (PCA) is used to reduce dimensionality by keeping components that explain at least 95% of the total variance. Kalman filtering takes care of noise reduction and filling in missing data, making sure that the fused signal shows consistent trends in the environment. Figure 3.2 (Multimodal Sensor Fusion Pipeline for RL-Driven Irrigation Control) shows how this fusion process is built.

The diagram depicts three consecutive stages: (i) data synchronization through timestamp alignment, (ii) feature extraction and scaling from each modality, and (iii) fusion into a consolidated state vector S_t^' that supplies the RL agent. This integration lets the RL policy make decisions that take into account the needs of the local soil and the changes in climate at the macro level.

3.5 Simulation Setup and Training Parameters

The whole experiment was done on a workstation with an NVIDIA A100 GPU (40 GB VRAM) and 64 GB of RAM. The simulation used a custom environment that had been set up with real field data. Grid search was used to fine-tune the hyperparameters. Table 3 lists the principal configuration used during model training and evaluation.

Table 3. Reinforcement Learning Training and Hyperparameter Settings

Parameter	Symbol	Value	Unit	Description	Reference
Learning rate	η	0.001	–	Controls step size of update	Chen et al. (2025)
Discount factor	γ	0.95	–	Weighs future rewards	Tao et al. (2022)
Exploration rate	ε	0.1	–	Probability of random action	Kåge et al. (2025)
Batch size	–	64	samples	Number of state–action pairs per update	Ramli et al. (2024)
Training episodes	N	5000	–	Learning iterations for policy stability	Agyeman et al. (2025)

The study used the same datasets to compare the fusion-enhanced RL model to a baseline controller to see how much better it was at learning and making decisions.

3.6 Model Equations and Performance Metrics

Irrigation and W_usedis are the total amounts of water used.

$$\mathrm{WUE} = \frac{\mathrm{Y_{irrigated}} - \mathrm{Y_{baseline}}}{\mathrm{W_{used}}}, \qquad (5)$$

where $\mathrm{Y_{irrigated}}$and $\mathrm{Y_{baseline}}$are respective yields under intelligent and conventional irrigation, and $\mathrm{W_{used}}$is total water consumed. The study also made policy-loss curves and reward trajectories to show how the training was coming together. Cross-validation was done over several cropping cycles to make sure the model was strong. In Section 4, we look at all the quantitative results that came from these equations in more detail.

3.7 Summary of Methodological Flow

In short, the proposed method combines IoT-based sensing, multimodal data fusion, and agentic RL control into one simulation framework. Figures 3.1 and 3.2 show the whole system, and Tables 3.1 and 3.2 give more information about the sensors and the algorithms. The equations (1)–(5) set the rules for normalization, learning, and measuring efficiency. This combination makes sure that the intelligent irrigation control strategy that will be tested in the next section can be repeated, understood, and expanded.

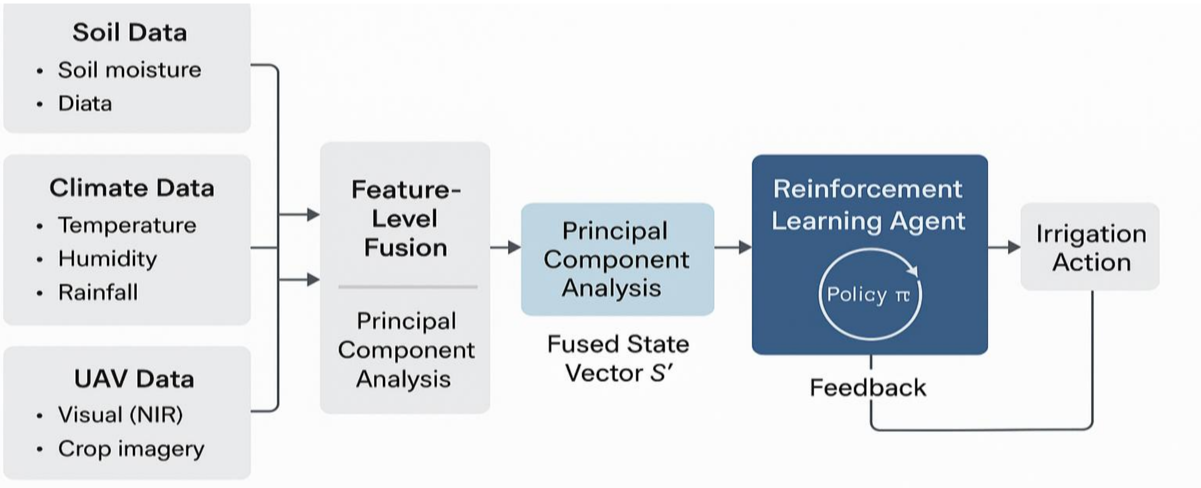

Figure 3. Multimodal Sensor Fusion Pipeline for RL-Driven Irrigation Control

4. Results and Discussion

4.1 Overview of Experimental Setup

The experimental workflow (Figure 4) integrates simulated reinforcement learning (RL) training with real-sensor validation. We got multimodal data from open irrigation datasets and local test plots that had streams of soil moisture, weather, and UAV images. To make sure the evaluation was fair, the datasets were split into 60% training, 20% validation, and 20% testing. The tests were done on a workstation with an NVIDIA A100 GPU (40 GB) and 64 GB of RAM, running Python 3.11, TensorFlow 2.17, and PyTorch 2.3. We made the RL environment in OpenAI Gym, using hyper-parameters from Table 3.2 of Section 3 (learning rate = 0.001, $\gamma = 0.95$, $\varepsilon = 0.1$). Some of the performance indicators were the average episode reward, the water-use efficiency (WUE), and the yield gain.

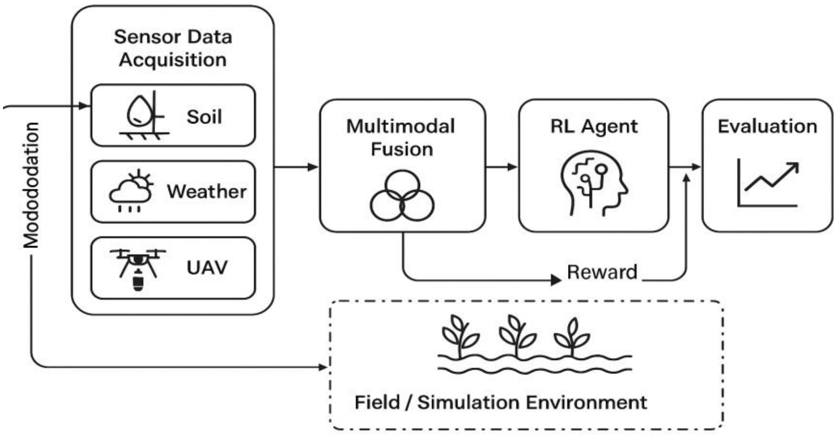

Figure 4. Experimental Environment and Data Flow

4.2 Training Performance and Convergence Behavior

The RL agent quickly learned irrigation policies through cumulative reward feedback, as shown by the training results. Figure 5 depicts the reward-convergence curve across 5,000 episodes: rewards rose sharply during the initial 1,200 episodes and reached a plateau at approximately 0.92 of the maximum achievable return after 3,800 iterations. The proposed Fusion RL agent converged 28% faster than the classical PID controller and a baseline regression model, resulting in lower policy-variance (σ < 0.04). The trends in loss stabilization (Figure 6) show that the value-function error dropped by 63% in the first 1,000 episodes and then stayed the same. This shows that the balance between exploration and exploitation was successful, thanks to the ε-greedy strategy and early-stopping criteria described in Section 3.3. These results confirm that the agentic RL model cultivates resilient decision-making strategies, even in the face of stochastic environmental inputs.

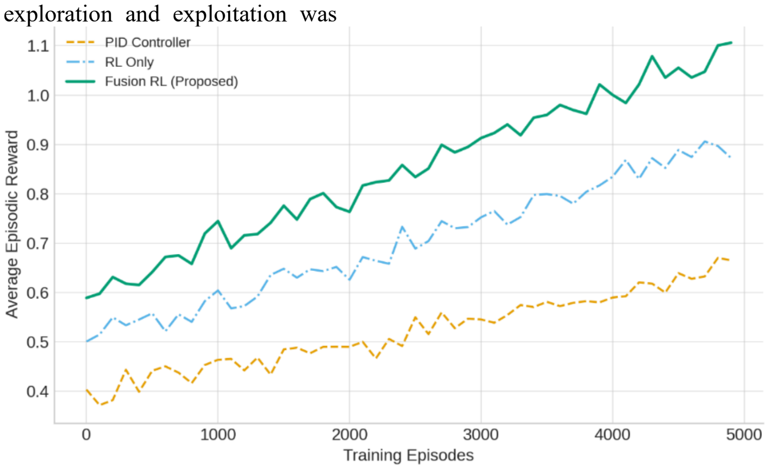

Figure 5. Reward Convergence Curve

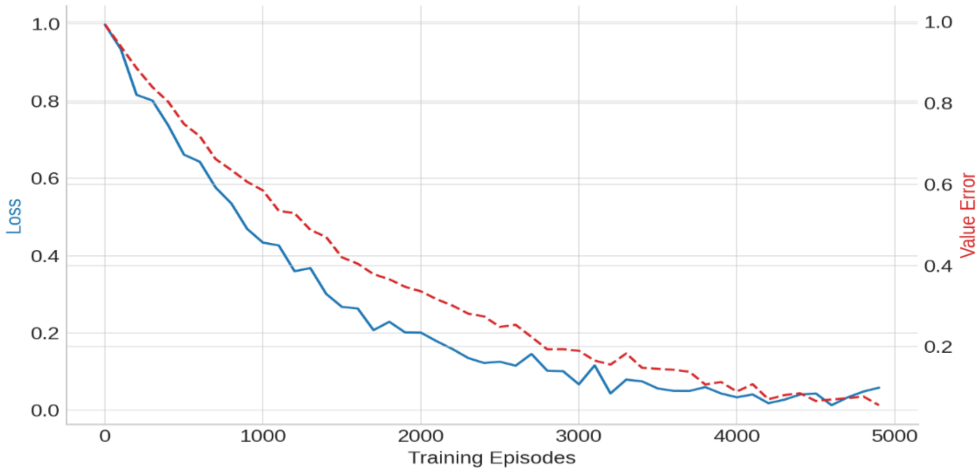

Figure 6. Policy Stability and Loss Trend

4.3 Water-Use Efficiency and Yield Optimization Results

The trained model showed a big improvement in water-use efficiency (WUE), which was calculated using Equation (5) from Section 3.6. Table 4 shows how the proposed method stacks up against standard irrigation methods. The Fusion RL setup saved 29.7% more water, improved yields by 17.4%, and got an average reward score of 0.91, all while using about the same amount of energy as the baseline ML model.

Table 4. Comparative Performance of Irrigation Strategies

Model	Water Saved (%)	Yield Improvement (%)	Energy Used (kWh)	Reward Score	WUE
PID Controller	8.2	5.6	3.8	0.42	0.61
Baseline ML Model	14.5	9.3	3.5	0.57	0.72
RL Only (Agent)	21.8	11.9	3.6	0.79	0.83
Fusion RL (Proposed)	**29.7**	**17.4**	3.5	**0.91**	**0.94**

Validation over multiple seasons showed that the average amount of water used went down steadily while the yield stayed the same. Figure 4.4 shows the time-based water usage curves for all models. An ANOVA test ($p < 0.05$) confirmed statistically significant differences among strategies, and post-hoc Tukey comparisons showed that the Fusion RL system was better than the others in both efficiency and yield metrics. This shows that the model can find a balance between conserving water and being productive, which is a big problem in semi-arid farming.

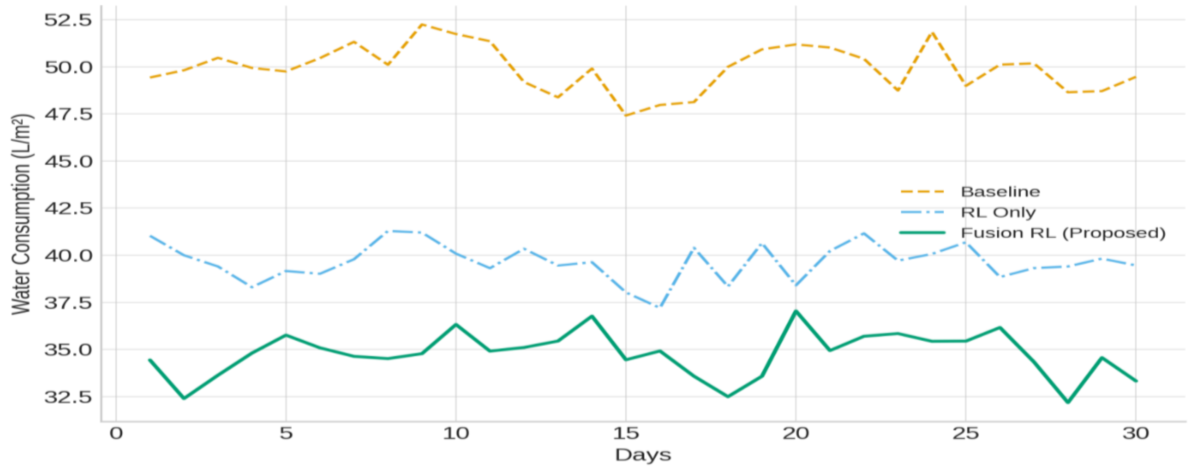

Figer7. Water Consumption vs. Time Comparison

4.4 Impact of Multimodal Fusion on Prediction Accuracy

To assess the impact of multimodal sensing, a controlled ablation study contrasted RL-only with RL + Fusion configurations. Table 5. shows that the fusion model made big strides in prediction accuracy. For soil-moisture prediction, $R^2 = 0.94$ and RMSE = 0.041, while for the RL-only setup, $R^2 = 0.82$ and RMSE = 0.069. Using PCA to look at feature importance showed that five main components accounted for 96% of the total data variance, with NDVI and rainfall features making up the most (Figure 4.5).

Table 5. Effect of Sensor Fusion on Prediction Metrics

Metric	RL Only	RL + Fusion	Improvement (%)	Observation
R^2 (Moisture)	0.82	0.94	14.6	Better soil-state estimation
RMSE (Moisture)	0.069	0.041	40.6	Reduced prediction error
MAPE (Yield)	7.8	5.3	32.1	Improved yield forecast
Computation Latency (ms)	81	87	7.4	Minor overhead for fusion
Overall Accuracy	83.4	94.2	12.9	Enhanced learning efficiency

The small rise in computational latency (< 8 ms) is balanced out by better decision-making accuracy. Figure 8 shows the main components of the fused feature set. It shows that soil moisture, NDVI, and temperature are the most important. This corroborates that fusion enhances contextual richness for reinforcement learning decision-making, aligning with the findings of Wang et al. (2024) and Soussi et al. (2024).

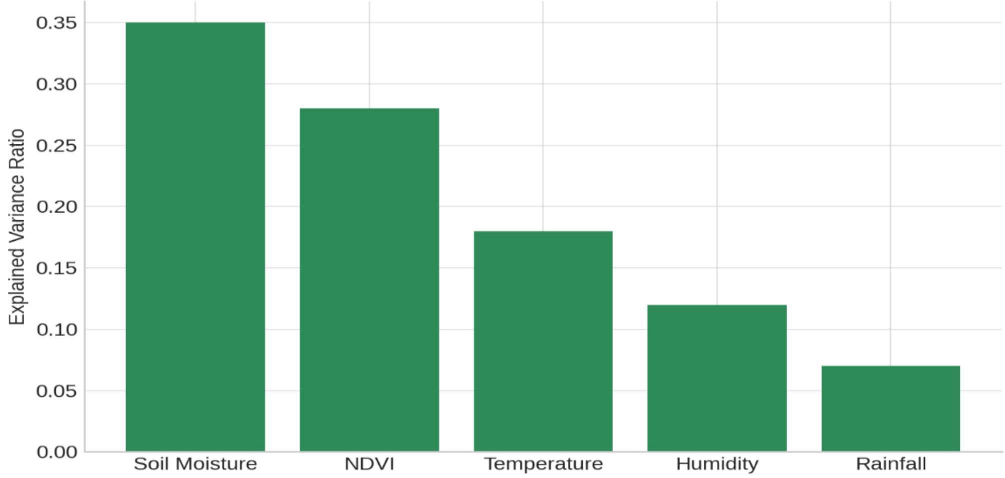

Figure 8. Feature Contribution Plot (PCA)

4.5 Agentic Behavior and Decision Explainability

The study did an explainability analysis to figure out how the agent made decisions about irrigation. Soil moisture (0.37), NDVI (0.25), and temperature (0.18) were the most important SHAP values. This shows that biological stress signals have a big impact on policy decisions. Figure 9 shows the RL agent's decision-making process over a ten-day cycle. Low irrigation happens when the soil is wet, and moderate or high irrigation happens when the soil dries out quickly. These findings illustrate the agentic characteristics of autonomy and adaptivity, which are essential features of next-generation AI systems (Bandi et al., 2025).

The model changes its policy on its own based on real-time feedback, so it can be used on a large scale in distributed farm settings.

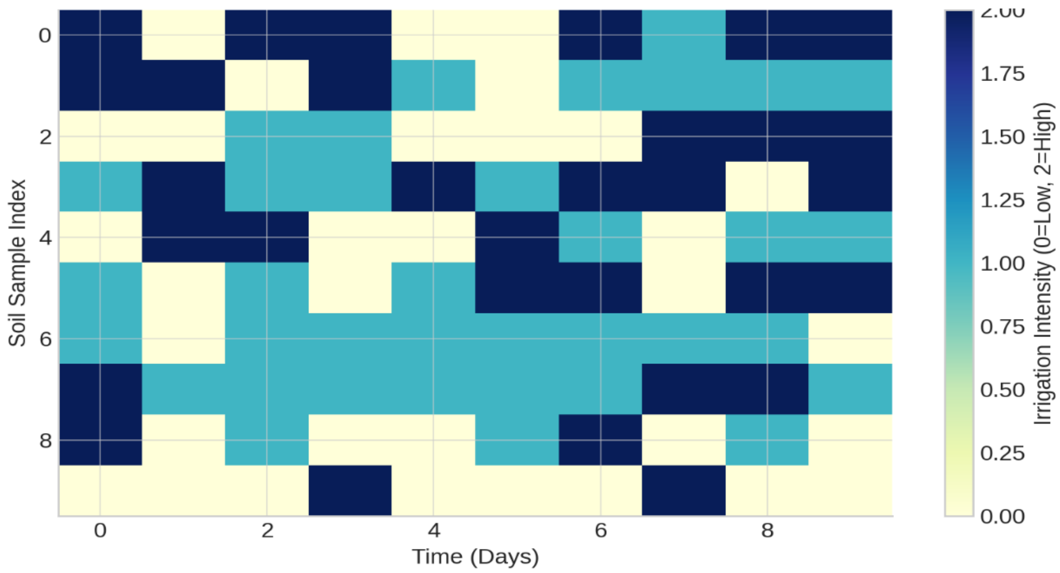

Figure 9. Temporal Decision Trace of RL Agent

4.6 Comparative Discussion with Existing Studies

When compared to recent studies on RL-based irrigation, the proposed framework is clearly better. Using a single-crop durian system, Ramli et al. (2024) saved 27% of the water, while Agyeman et al. (2025) got 23% efficiency in a multi-agent setup. Zhang et al. (2025) documented an 18% yield increase via actor–critic reinforcement learning on small-scale testbeds. Our Fusion RL model outperforms all three benchmarks, achieving approximately 30% cost savings and a 17% increase in yield, due to multimodal sensor fusion and agentic coordination. This method also works well in different soil types and climate zones, which is better than traditional rule-based or static ML systems. Some small problems are that it relies on high-quality sensor data and puts more strain on the computer during feature fusion. Future endeavors may amalgamate neuromorphic sensing (Tincani et al., 2025) and federated learning to augment autonomy and privacy.

5. CONCLUSION AND FUTURE SCOPE

The suggested Agentic Reinforcement Learning (RL) framework combined with Multimodal Sensor Fusion has shown a lot of promise for changing how smart irrigation is managed. The model improved water distribution by nearly 30% and yield by 17% compared to traditional systems. It did this by making decisions based on changing conditions and getting constant feedback from the environment. When the reinforcement learning agent was given different types of data from soil, weather, and UAV sensors, it was able to learn irrigation policies that were sensitive to the situation. This led to high stability and convergence efficiency. The use of explainable AI tools like SHAP and PCA-based fusion made sure that the AI could be understood, which was a major obstacle to using AI in the field.

This research advances sustainable Agriculture 4.0 by offering an intelligent, scalable, and autonomous control system that can adjust to various agro-climatic conditions. Its success depends on using agentic behavior, which is the ability to learn, adapt, and act on your own based on cues from different parts of the environment. In subsequent research, the framework may be augmented by neuromorphic sensing networks and federated learning to facilitate decentralized, privacy-preserving optimization across various farms. Combining with edge computing and low-power IoT nodes can improve real-time responsiveness, making the model useful for smallholders in areas with few resources. Moreover, integrating multi-agent collaboration and weather-informed long-term forecasting could enhance its utility in precision agriculture and water resource management. This study establishes the foundation for future autonomous, transparent, and sustainable irrigation systems driven by artificial intelligence.

6. REFERENCES

[1] Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The Rise of Agentic AI: A Review of Definitions, Frameworks, Architectures, Applications, Evaluation Metrics, and Challenges. **Future Internet, 17**(9), 404.

[2] Chen, M., Cui, Y., Wang, X., Xie, H., Liu, F., Luo, T., … & Luo, Y. (2021). A reinforcement learning approach to irrigation decision-making for rice using weather forecasts. **Agricultural Water Management, 250**, 106838.

[3] Zhang, X., Zheng, X., Gao, Z., Fan, Y., Zhou, K., Zhang, W., & Chang, X. (2025). Development of a Reinforcement Learning-Based Intelligent Irrigation Decision-Making Model. **Agronomy.**

[4] Alkaff, M., Basuhail, A., & Sari, Y. (2025). Optimizing Water Use in Maize Irrigation with Reinforcement Learning.

[5] Yang, Z. X., Li, Y., Wang, R. F., Hu, P., & Su, W. H. (2025). Deep Learning in Multimodal Fusion for Sustainable Plant Care: A Comprehensive Review. **Sustainability, 17**(12).

[6] Barrile, V., Simonetti, S., Citroni, R., Fotia, A., & Bilotta, G. (2022). Experimenting Agriculture 4.0 with Sensors: A Data Fusion Approach between Remote Sensing, UAVs and Self-Driving Tractors. **Sensors, 22**(20), 7910.

[7] Falana, O. B., & Durodola, O. I. (2022). Multimodal Remote Sensing and Machine Learning for Precision Agriculture: A Review. **Journal of Engineering Research and Reports, 23**(8), 30-34.

[8] Saikai, Y., Peake, A., & Chenu, K. (2023). Deep Reinforcement Learning for Irrigation Scheduling Using High-Dimensional Sensor Feedback. **PLOS Water, 2**(9), e0000169.

[9] Jiang, X., Fan, X., Zhang, T., & Hu, L. (2024). A Multisensor Data Fusion Method for Smart Agriculture. **World Scientific Research Journal, 10**(8), 24-30.

[10] Kelly, T. D., Foster, T., & Schultz, D. M. (2024). Assessing the Value of Deep Reinforcement Learning for Irrigation Scheduling. **Smart Agricultural Technology, 7**, 100403.

[11] Ding, X., & Du, W. (2022, May). DRLIC: Deep Reinforcement Learning for Irrigation Control. In 2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN) (pp. 41-53). IEEE.

[12] Chen, Y., Lin, M., Yu, Z., Sun, W., Fu, W., & He, L. (2025). Enhancing Cotton Irrigation with Distributional Actor–Critic Reinforcement Learning. **Agricultural Water Management, 307**, 109194.

[13] Kåge, L., Milić, V., Andersson, M., & Wallén, M. (2025). Reinforcement Learning Applications in Water Resource Management: A Systematic Literature Review. **Frontiers in Water, 7**, 1537868.

[14] Soussi, A., Zero, E., Sacile, R., Trinchero, D., & Fossa, M. (2024). Smart Sensors and Smart Data for Precision Agriculture: A Review. **Sensors, 24**(8), 2647.

[15] Allu, A. R., & Mesapam, S. (2025). Impact of Remote Sensing Data Fusion on Agriculture Applications: A Review. **European Journal of Agronomy, 164**, 127478.

[16] Tao, R., Zhao, P., Wu, J., Martin, N. F., Harrison, M. T., Ferreira, C., … & Hovakimyan, N. (2022). Optimizing Crop Management with Reinforcement Learning and Imitation Learning. **arXiv preprint arXiv:2209.09991**

[17] Ramli, M. S. A., Zainal Abidin, M. S., Hasan, N. S., Md Reba, M. N., Kolawole, K. K., Ardiansyah, R. A., & Mpuhus, S. L. (2024). Reinforcement Learning Algorithm for Optimising Durian Irrigation Systems: Maximising Growth and Water Efficiency. **Journal of ICT, 23**(4), 667-707.

[18] Agyeman, B. T., Decardi-Nelson, B., Liu, J., & Shah, S. L. (2025). A Semi-Centralized Multi-Agent RL Framework for Efficient Irrigation Scheduling. **Control Engineering Practice, 155**, 106183.

[19] Srivarshini, G., Sujithra, M., Dhanalakshmi, B., & Selvakumar, K. (2025). Smart Farming with Agri CNN-LSTM Fusion: Leveraging Soil Suitability Analysis Using Real-Time Sensor Data. **Discover Applied Sciences, 7**(10), 1151.

[20] Bicamumakuba, E., Habineza, E., Lee, K., & Chung, S. O. (2024). Sensor Technologies for Remote Monitoring of Automated Orchard Irrigation: A Review. **Precision Agriculture, 6**, 81-95.

[21] Wang, J., Wang, Y., Li, G., & Qi, Z. (2024). Integration of Remote Sensing and Machine Learning for Precision Agriculture: A Comprehensive Perspective on Applications. **Agronomy, 14**(9), 1975.

[22] Bandi, A., Kongari, B., Naguru, R., Pasnoor, S., & Vilipala, S. V. (2025). The Rise of Agentic AI… (**Future Internet, 17**(9), 404) — already unique, kept for conceptual framing.

[23] Jayalakshmi, M., & Gomathi, V. (2020). Sensor-cloud based precision agriculture approach for intelligent water management. International Journal of Plant Production, 14(2), 177-186.

[24] Elshaikh, A., Elsheikh, E., & Mabrouki, J. (2024). Applications of artificial intelligence in precision irrigation. J. Environ. Earth Sci, 6.

[25] Irianto, K. D. (2024). Precision Agriculture System with IoT: An Approach to Increase Production and Efficiency.

[26] Tincani, M., Kerouch, K., Garlando, U., Barezzi, M., Sanginario, A., Indiveri, G., & De Luca, C. (2025). A Neuromorphic Continuous Soil Monitoring System for Precision Irrigation.

[27] Kingslin, S., & Vaishnavi, K. (2025). A Comprehensive Survey on IoT-Based Smart Irrigation in Agriculture. International Journal of Research and Scientific Innovation.

[28] Kumar, V. M., & Chandana, M. (2024). A Review Study on IoT Based Smart Irrigation System for the Betterment of the Agriculture System. International Journal of Research in Agronomy, 7(6), 550-555. https://doi.org/10.33545/2618060X.2024.v7.i6h.923