



RESEARCH ARTICLE

AI-driven multi-agent framework for smart irrigation and crop health monitoring in Indian rice and sugarcane farming

Sheetal V A*, Vikranth B M & Adarsha Sagar H V

Department of Computer Science and Engineering, B M S College of Engineering, Bengaluru 560 004, Karnataka, India

*Correspondence email - gdsheetal26@gmail.com

Received: 17 November 2025; Accepted: 31 March 2026; Available online: Version 1.0: 30 April 2026

Cite this article: Sheetal VA, Vikranth BM, Adarsha SHV. AI-driven multi-agent framework for smart irrigation and crop health monitoring in Indian rice and sugarcane farming. *Plant Science Today*. 2026; 13(sp2): 1-16. <https://doi.org/10.14719/pst.12797>

Abstract

Disease prevention and water management are important to all the crops, particularly rice and sugarcane production in India. The article proposes a reinforcement learning (RL) based intelligent irrigation management system that is capable of optimising water consumption and crop nutrition in response to the changing agricultural climatic conditions. Decentralised reinforcement learning (RL) is used in a network of irrigation agents that utilise soil and microclimate sensor networks to set the terms of water allocation, water use efficiency (WUE) and crop health. At the same time, deep convolutional networks can be used to differentiate between plant stress/disease and leaf images and take applicable proactive actions. It is a framework that incorporates satellite-derived indices (NDVI, EVI, land surface temperature) with local sensor measurements and image-based health measurements through multimodal deep learning. Far-reaching simulations (including Indian climate and crop calendars) demonstrate that the multi-agent system lowers water consumption and preserves the yields and properly notifies stressed plants. The scores of disease detection with plantvillage-based fine-tuned on rice (120 (3 disease types) and 3829 (5 disease types) and sugarcane (2569 images for all disease types, Convolutional Neural Network (CNN) yield results of >98 % accuracy. Crop mapping (rice/sugarcane) Satellite/LSTM-based crop mapping (with Sentinel-1 / Sentinel-2) achieves more than 97 % accuracy. The suggested structure provides a data-driven, scalable system for precision agriculture to enhance the management of irrigation periods and crop health. Simulation experiments show that the RL-based controller can reduce water consumption while preserving optimal soil moisture levels when compared to rule-based irrigation strategies.

Keywords: crop health monitoring; deep learning; multi-agent system; multimodal data fusion; plant village; remote sensing; Sentinel-1; Sentinel-2; smart irrigation

Introduction

Irrigated agriculture consumes the majority of global freshwater ($\approx 70\%$) (1). In India, efficient water use in water-intensive crops like rice and sugarcane is very important for sustainability. Traditional irrigation cannot adapt to spatially variable fields or sudden weather changes. Similarly, visual disease scouting is laborious and often too late, causing yield losses. Precision agriculture uses current AI technology and sensor technology to maximise inputs and outputs. Its objectives involve maximisation of yield, at least the quality of crops and minimisation of environmental effects (2). Here, satellite remote sensing (e.g., Normalised Difference Vegetation Index- NDVI, Enhanced Vegetation Index- EVI, Land Surface Temperature- LST on Sentinel-2/Sentinel-1) can be used to get frequent and large-scale crop status information (3, 4). Fine-grained disease detection is made possible by ground imagery (e.g., high-resolution leaf images). The fine-tuning of deep models on rice (120 images) and sugarcane (2569 images of all categories) using the public dataset, such as PlantVillage (50k+ images of healthy/diseased leaves of 38 categories), has enabled them to achieve >99 % accuracy in the classification of leaf diseases (5).

Simultaneously, sophisticated irrigation systems are also becoming prominent. IoT controllers can modify schedules based on real-time data of soil moisture and weather conditions. Among

them, most that have been centralised or are rule-based, Fuzzy/AI controllers have demonstrated large savings in water in farms. Multi-agent is scalable and robust: distributed irrigation nodes (valves/sensors) can negotiate locally through peer-to-peer (5). The interaction between farmers and infrastructure agents has been shown to produce realistic irrigation dynamics in agent-based models (ABMs) of irrigation (e.g., IRABM). Most recently, multi-agent reinforcement learning (MARL) has been used in irrigation scheduling, achieving water savings (e.g. 4 % water reduction) and increased irrigation water use efficiency (6).

In this paper, a new scheme, which is the combination of multi-agent irrigation control and deep-learning-based crop health monitoring, is introduced based on multimodal data fusion. This research is aimed at Indian rice and sugarcane models that combine Sentinel-1/2 satellite data (NDVI, EVI, SAR backscatter) with leaf photos of the stature of PlantVillage. An irrigation agent network works together to distribute water effectively and deep CNNs flag detect plants when they are under stress or diseased. The most important contributions in this study are: a dynamic coordination RL-based irrigation protocol, a multimodal fusion chain with satellite indices and local image diagnostics, much sensor feedback and case studies of Indians, State-of-the-art disease classification (approximately 99 % accuracy with PlantVillage) and crop mapping with high accuracy (>97 % with

Sentinel data) (3, 4). This paper is a step towards precision agriculture, combining remote sensing, computer vision and agent-based AI.

Related Work

Smart irrigation and multi-agent control

Smart irrigation means that the fixed schedules are substituted with data-driven ones. The classic IoT irrigation systems were based on fuzzy logic or the mere feedback (soil moisture) rules. An example is Wang and Ertsen came up with an agent-based irrigation model (IRABM), which models the agents of farmers and infrastructure. Demand-based control that enhanced yields in case of water adequacy revealed their findings and demonstrated the strength of ABMs in irrigation. Simulations of the boosted lifetime of fuzzy/AI controllers on IoT nodes have minimised wastage and increased network lifetime. These centralised schemes, however, break down in the face of communication complications. Multi-agent schemes. Distributed Multi-agent schemes have benefits (6). Research has demonstrated that a hierarchical, multi-agent MARL irrigation scheme (SCMARL) in which local and coordinator agents acquire policies. It conserved approximately 4 % water and increased the irrigation effectiveness by approximately six point 3 % on a Canadian field. Model predictive control (MPC) and RL have been applied in other research works. To illustrate, the DrLIC relates to deep RL to irrigate the soil through forecasting the soil moisture dynamics. Irrigation as an optimisation problem with high-dimensional sensor feedback was formalised in a PLOS study that demonstrated that deep RL could acquire efficient rules on data (7). On the whole, RL and MARL are considered to be the new promising methods to maximise irrigation in the spatially non-uniform fields.

Crop monitoring based on remote sensing

Satellite monitoring has transformed crop monitoring. Sentinel-2 (optical) and Sentinel-1 (SAR) have high resolutions in terms of revisit and free data. Sentinel-2 provides NDVI/EVI and LST, which can be used to monitor vegetation health (stress, biomass) on large scales. Cloud Sentinel-1 SAR (C-band) is also able to penetrate clouds to support vegetation mapping. Recent models make use of both fused Sentinel-2 NDVI time series and Sentinel-1 SAR to map the location of rice in Iran; with modalities combined, the model reached over 97 % accuracy (3). This establishes that the multisensor fusion is more effective in mapping (SNDI vs NDVI-only was worse). Equally, research indicates that Sentinel-2 NDVI with phenology in crop classification and SAR to clear nearby crops (8). In sugarcane, combined Sentinel-1 and NDVI composites to mask non-vegetation (9). However, many previous studies still depended on cloud-free optical data; pure SAR mapping (e.g., parcel-based SAR classification) remains challenging due to speckle noise. In summary, multimodal remote sensing (optical + SAR) has shown great promise in identifying crop types and stress with high accuracy.

Crop health and disease detection (Image-based)

Leaf image analysis using deep learning has become standard for disease detection. Large curated datasets like PlantVillage (54000+ images, 14 species, 26 diseases) (2) fine-tuned on rice (120 images) and sugarcane (2569 images including all categories) enable training powerful CNNs. Research indicates the plant disease classification: CNNs (AlexNet/VGG/ResNet variants) routinely exceed 98–99 % accuracy on curated leaf datasets (10). A recent

PLOS One study fused pretrained CNN feature extractors (MobileNetV2, DarkNet, ResNet) with SVM classifiers on rice leaf images, achieving 98.6 % accuracy (6). Similarly, on sugarcane, 98.7 % accuracy with ConvNet binary classification of healthy vs four disease types was noticed (11). Research indicates that MobileNetV3 on PlantVillage leaves (14 different crops) with 70 % accuracy and 33 % reduction in model size for edge devices (12). These results prove that deep vision models could be used effectively to detect diseases from images. For deployment, recent works have been done to protect farmer data using federated learning: Research has demonstrated the better performance of federated CNNs (e.g., ResNet50) on PlantVillage data compared with centralized training (12).

Multimodal data fusion

The combination of data types (satellite, UAV, proximal sensors, images) is a means of overcoming the limits of individual types. For example, fusion of unmanned remote sensing (UAV) derived canopy structure (height, cover) with satellite remote sensing spectral VIs significantly improved biomass and N estimation for soybean (13). Such a combination decreased the NDVI saturation effects of dense canopies. Likewise, combining image analysis of the leaves and spatial data is promising. While there are few studies where satellite and leaf images are explicitly combined, multimodal learning is becoming popular. Research showed that the R² of soybean trait prediction increased up to 0.92 by stacking the features of satellite and UAV images (spectral+structure) and training an ensemble ML (14). By way of analogy, the fusing of satellite indices (NDVI, LST) derived from canopy stress with leaf features extracted by CNN might provide a holistic representation of plant health at a scale. Recent large language-vision models (e.g., PlantIF) suggest graph-based fusion of image and text for disease diagnosis and this shows the trend towards multimodal agritech. In summary, data fusion contributes to an improved situation awareness - satellite/UAV fusion has demonstrated clear performance increases - similar benefits when used in combination with remote and proximal plant data.

Indian Context - Rice and Sugarcane Almost all Indian agricultural regions have rice and sugarcane, which need special attention to water and disease. Prior remote sensing work includes rice paddy mapping by Sentinel-1/S2 (e.g. SVAR Timeseries NDVI thresholds) and sugarcane by Sentinel-1 (e.g. phenology-based warping). Sugarcane is commonly used for the detection of lodging, nutrient monitoring and water stress. On the data aspect, smartphone crowdsourcing (e.g. Plantix) was leveraged to map Indian sugarcane fields using a supervised CNN on Sentinel-2 (15). For crop health, Indian researchers have started teaching CNN on an indigenous leaf dataset: - A ConvNet for crop health - sugarcane leaves (healthy VS diseased) (16). This research attempts to further develop these through the use of the latest modalities and multi-agent coordination specific to Indian crops and conditions.

In general, literature highlights the feasibility of the AI and agent-based approach in agriculture: a multi-agent RL approach to irrigation systems can save water, deep learning can be used to classify crop diseases robustly (5, 6) and data fusion can increase the accuracy of monitoring. However, a unified system that co-optimises irrigation and health monitoring – especially for Indian rice/sugarcane – is novel. This framework addresses this gap by combining agent coordination with multimodal learning.

Methodology

The proposed system consists of two tightly integrated components: a multi-agent irrigation controller and a multimodal crop health analyser. A high-level schematic (Fig. 1) shows agents deployed across the farm, each controlling irrigation valves and sensors and an off-site data fusion server combining inputs from satellites and ground.

Agent Architecture and Coordination: Each field zone has a local irrigation agent with sensors (soil moisture, temperature, humidity) and control actuators (valves/pumps). Agents form an ad-hoc network via low-power links (e.g., LoRaWAN). We design a Decentralised Reinforcement Learning (Dec-RL) approach: each agent learns a policy. π_i that maps the local state S_i (soil moisture, local crop stress estimate, neighbours' reports) to an action A_i (irrigate at rate, delay irrigation, or adjust flow) to maximise cumulative rewards. The state S_i includes current soil moisture ($S\theta_i$), crop growth stage (g_i) and neighbour moisture deficits. The action space comprises discrete valve settings and timings.

Rewards R combine objectives: high crop stress reduction and water/energy saving. For agent i ,

$$R_i = \alpha \Delta\theta_i + \beta H(\text{stress}_i) - \gamma E_i \quad (\text{Eqn. 1})$$

Where $\Delta\theta_i$ is moisture gain, $H(\text{stress}_i)$ is the negative of the stress indicator (higher if stress is reduced), E_i is energy cost and (α, β, γ) are tunable weights.

Agents negotiate water allocation when supply is limited. A lightweight consensus protocol resolves conflicts: competing agents vote based on urgency (e.g., relative soil deficit) and current trust scores. Trust between agents is updated via a Bayesian filter, rewarding reliable information sharing

The coordination proceeds as follows (Algorithm 1):

- * For each irrigation agent i :

- * Observe the local state S_i =soilMoisture, etc.
- * Communicate with neighbours $j \in N_i$: exchange state fragments
- * Compute Local Stress Index (LSI_i) from S_i
- * Receive neighbours' LSI_i;
- * Computing global stress estimate GSI (e.g., weighted average)
- * Input combined state (S_i , GSI) into the policy network π_i to select action A_i
- * Execute irrigation action A_i (valve on/off or flow rate adjust)
- * Observe reward R_i based on moisture change and energy cost
- * Update RL policy π_i via Q-learning/PPO with reward
- * Update trust scores for neighbours based on the consistency of their information.

Algorithm 1. Local irrigation agent loop (per timestep)

For episode = 1 to N_{ep} do

Reset environment (initialise soil moisture and weather)

For each time step $t = 1, 2, \dots, T$ do

Observe the current state $S_t = [M_t, W_t, G_t]$

Select action a_t using ϵ -greedy policy:

if r and $(\dots) < \epsilon$: choose random action

else $a_t = \arg \max_a Q(s_t, a; \theta)$

Execute irrigation action a_t on the field (simulate irrigation)

Observe new state S_{t+1} and reward r_t :

$$r_t = \eta_1 \times (\text{yield gain}) - \eta_2 \times (\text{water used}) - \eta_3 \times (\text{stress index})$$




Fig. 1. Block diagram showing the high-level view of how the whole system is structured.

Update Q-values:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha [r_t + \gamma \max_{a'} Q(s_{t+1}, a') - Q(s_t, a_t)]$$

Update policy weights (if deep RL):

$$\theta \leftarrow \theta + \alpha \nabla_{\theta} \left(r_t + \gamma \max_{a'} Q(s_{t+1}, a'; \theta) - Q(s_t, a_t; \theta) \right)^2$$

Set $S_t = S_{t+1}$

End For

Decay exploration $\epsilon = \epsilon \times 0.99$

End For

Each agents' policy network is a small Deep Q-Network (DQN) or policy-gradient (PPO), able to process multi-dimensional input. The global system thus functions in a semi-distributed manner: each agent optimises locally, but alignment emerges via shared signals (like the IRABM approach). A central coordinator agent may optionally oversee overall metrics (total water use, aggregated stress) and adjust reward weights dynamically for stability.

Crop health analyser

At the same time, the image of the plant foliage is obtained (e.g., with drones or ground cameras). These photos are classified so as to find out the early signs of sickness or deficiency of nutrients. The custom-built CNNs are used and pretrained with the use of the PlantVillage data. This is done by feeding photos into a backbone CNN (e.g. ResNet50 or MobileNetV3-small), which produces feature embeddings. We performed end-to-end experiments on the fine-tuning of CNNs and a hybrid feature plus SVM technique. As a result of the approach, one of the several classifications is generated: healthy or a particular disease category.

Multimodal Data Fusion

The innovation lies in fusing these modalities. Satellite-derived metrics (NDVI, EVI, LST, SAR backscatter) are aggregated per field zone (via GIS). Leaf-image predictions (probabilities of disease) are fused with the satellite features in a late-fusion neural network.

Each zone i has a feature vector:

$$X_i = (\text{NDVI}_i, \text{EVI}_i, \text{LST}_i, \dots, P_i \text{ disease probabilities}) \quad (\text{Eqn. 2})$$

where P_i are the CNNs' probability outputs. A feedforward NN learns from $(X_i, \text{ground truth health})$ to refine stress detection. The satellite indices capture canopy-wide stress, while leaf images indicate local pathology. Together, they improve early detection accuracy.

A feedforward neural network (or gradient ensemble) then learns from $(X_i, \text{ground-truth health})$ to refine stress detection. The satellite indices capture canopy-wide stress, while leaf images give local pathology. Together, they improve early detection accuracy.

This fused health indicator influences irrigation rewards: zones flagged as diseased or water-stressed are prioritised for water allocation. Thus, better health monitoring guides irrigation and good irrigation promotes plant health. We implement all models in TensorFlow/PyTorch. CNNs are fine-tuned on PlantVillage (14 crops, 26 diseases) with rice/sugarcane subsets. RL agents are trained in simulation (e.g., using the AquaCrop model for soil-crop dynamics). Agents practice over many episodes (season-long irrigation planning) to learn effective policies. Transfer learning or online adaptation updates models in actual deployment. In the single-agent RL shown in Fig. 2a, the decision-making process begins with the observation of the system state S_t , which is subsequently transformed into intermediate representations used to determine the most suitable action for interacting with the environment. This interaction generates new environmental feedback, enabling continuous learning and adaptation. In the extended multi-agent architecture shown in Fig. 2b, the global state S_t is simultaneously accessed by multiple agents, each operating under its respective policy to produce independent action sets. These actions are enriched through a data-fusion process that integrates agent outputs with satellite-derived indicators and crop-specific image features sourced from datasets such as PlantVillage, rice and sugarcane fields. The fused feature vectors from all agents are then combined into a single fusion layer. This makes it easier for precision agriculture and




Fig. 2 a. Single agent RL, **b.** Multi-Agent RL.

irrigation control to make decisions that are more precise, consistent and strong.

Pseudocode for training the RL agents (centralised loop) is shown in Algorithm 2.

Key steps: simulate episode → for each day t , each agent observes $S_t(t)$ selects $A_t(t)$ via ϵ -greedy or policy, executes and observes $R_t(t)$ and updates Q /Policy networks. After many epochs, policies converge.

Algorithm 2: Multi-Agent RL with Satellite and Image Data Fusion for Smart Irrigation

For episode = 1 to N_{ep} do

Reset environment, receive initial state $S_0 = [M_o, S_o, W_o, C_o]$ (Eqn. 3)

For each time-step $t = 1, 2, \dots, T$ do

Each agent i selects an action a_i^t using ϵ -greedy:

if r and $(\epsilon) < \epsilon$: select random irrigation level

else:

$$a_i^t = \arg \max_a Q_i(s_t, a; \theta_i) \quad \text{Apply aggregated}$$

actions $A_t = a_1^t \dots a_n^t$ to the irrigation system Observe the next state.

$$S_{t+1} = [M_{t+1}, S_{t+1}, W_{t+1}, C_{t+1}] \quad (\text{Eqn. 4})$$

Compute reward:

$$r_t = \eta_1 \times (\text{yield gain}) - \eta_2 \times (\text{water used}) - \eta_3 \times (\text{stress index})$$

Store transition (s_t, a_t, r_t, S_{t+1}) in buffer B_i

Update Q via gradient descent:

$$\theta_i \leftarrow \theta_i + \alpha \left[r_t + \gamma \max_{a'} Q_i(s_{t+1}, a'; \theta_i) - Q_i(s_t, a_t; \theta_i) \right] \nabla_{\theta_i} Q_i \quad (\text{Eqn. 5})$$

Synchronise policies among neighbouring agents every K steps

Fuse satellite and image features using weighted attention:

$$f_t = \omega^1 F(S_t) + \omega^2 F(C_t) + \omega^3 F(W_t) \quad (\text{Eqn. 6})$$

Append f_t to the global state vector

End For

$$\text{Decay } \epsilon = \epsilon \times 0.995 \quad (\text{Eqn. 7})$$

End For

The coordination mechanism is therefore best interpreted as an adaptive resource arbitration layer that improves multi-zone stability rather than altering the optimal single-agent irrigation policy.

Trust-aware multi-agent coordination

Formalisation of trust-based coordination

Each irrigation agent (i) maintains a trust score ($T_{ij} \in [0,1]$) for neighbouring agent (j), representing the historical reliability of shared state information.

Trust update rule

Trust is updated using a Bayesian reliability estimate:

$$T_{ij}^{t+1} = (1 - \alpha) T_{ij}^t + \alpha \cdot \exp(-|e_{ij}^t|) \quad (\text{Eqn. 8})$$

where:

• (e_{ij}^t) is the difference between the predicted and observed moisture levels after following a neighbours' advice.

• (α) is the adaptation rate controlling how quickly trust evolves.

Agents that continuously send useful coordination signals become more powerful over time.

Negotiation-based water allocation

When multiple zones request shared irrigation resources, allocation is computed as:

$$A_i = \frac{T_i \cdot U_i}{\sum_k T_k U_k} \quad (\text{Eqn. 9})$$

where (U_i) is the urgency caused by stress levels and a lack of dirt. This produces a soft consensus instead of a rigid priority rule, which stops choices from going back and forth.

Integration with RL policy

The negotiated allocation acts as a constraint modifier on the agents' chosen action:

$$I_i^{\text{final}} = \min(I_i^{\text{RL}}, A_i \cdot I_i^{\text{available}}) \quad (\text{Eqn. 10})$$

Thus, learning remains decentralised while resource conflicts are resolved cooperatively.

All multimodal inputs are therefore spatially co-registered, temporally aligned and accompanied by uncertainty descriptors before being used for classification or irrigation policy learning.

Spatial-temporal harmonization of multi-resolution satellite data

The images of Sentinel-1 SAR (resolution of 10 m, works in any weather) and Sentinel-2 multispectral images (resolution of 10-20 m depending on band) are different in the way they are captured, when they are recaptured and the amount of noise contained in the image. A harmonisation process that involved three steps was implemented to ensure the compatibility of everything.

Spatial resolution alignment

All the datasets were remapped to a shared UTM grid and resampled to 10 m. Backscatter nearest-neighbour was used in SAR to eliminate radiometric distortion and bilinear interpolation in optical bands. This facilitates correlation of pixels and this must be provided in multimodal fusion.

Temporal synchronization

Given the absence of absolute coincidence of the dates of acquisition, observations were compared within 3 days.

$$X_t = \sum_k w_k X_{t_k} w_k \propto \exp\left(-\frac{|t-t_k|}{\tau}\right) \quad (\text{Eqn. 11})$$

When there are multiple pictures, the time interpolation was weighted: where the tolerance to time error is controlled by (T) . This is to prevent discontinuities in features, which can be sharp because of asynchronous sensing.

Speckle and atmospheric noise handling

Multi-temporal filtering was useful to minimise SAR speckle. Before the fusion of the two images, the atmosphere was corrected and clouds were eliminated in the Sentinel-2 reflectance data. Extra channels of confidence were maintained by keeping quality masks.

Uncertainty propagation into learning models

Modality-specific uncertainty estimates were added instead of regarding the features of the satellites as deterministic:

$$Z = [f_{\text{optical}}, f_{\text{SAR}}, \sigma_{\text{optical}}, \sigma_{\text{SAR}}] \quad (\text{Eqn. 12})$$

The terms of variance enable the downstream models to adaptively weight the features, minimising overconfidence on noisy observations.

Impact on the irrigation decision layer

The reinforcement learning condition is equipped with aggregated vegetation and moisture proxy and their confidence values, which allows the agents to discount unreliable remote sensing signals when the atmosphere is cloudy or the area has high speckle.

Experimental setup

The experimental evaluation is completely carried out in a simulation framework, which is designed to simulate real farming environments. Agronomic equations based on the past history of weather patterns, the seasons and the requirements and growth phases of individual crops are employed to model soil moisture behaviour, evapotranspiration and crop water use. Weather inputs are synthesised from long-term climatological statistics in order to reproduce realistic trends for rainfall, temperature and humidity to allow repeatable testing for multiple irrigation scenarios. This method allows to compare learning-based and traditional irrigation policies in a controlled manner, without the issues associated with a field trial. The results are made without any live sensor data and physical deployment of the farm. The next step is to take publicly available data sets of the agricultural landscape and perform in-situ validation to compare the model that was trained by simulations to the actual world circumstances in agriculture.

Data sources

Sentinel-1 & Sentinel-2

We collect Sentinel-2 Level-2A imagery (optical bands + derived NDVI/EVI) and Sentinel-1 GRD (VH/VV polarisations). Time series spanning the growing season are compiled for sample plots of rice and sugarcane in representative Indian states (Punjab for rice, Maharashtra for sugarcane). Cloud masking is applied; missing NDVI values are gap-filled via linear interpolation. NDVI, EVI and daytime LST are computed per 10m grid. Sentinel-1 SAR data are preprocessed and speckle-filtered (Fig. 3).

Plant village & local leaf dataset

We use the plant village database (50k+ leaf images, 14 crops) for initial training. For India-specific cases, we assemble a subset focusing on rice and sugarcane diseases (e.g., rice blast, brown spot; sugarcane mosaic, rust). The total fine-tuning set includes 120 rice images plus 3829 rice leaf disease images from and 2569 sugarcane images with expert annotations(3, 17, 18, 24).

We use the plant village database (fine-tuned on rice/sugarcane leaves) for initial training: 14 crops $\times$ 26 diseases (leaf images). For India-specific training, a subset focusing on rice and sugarcane diseases is assembled (rice blast, brown spot; sugarcane mosaic, rust, etc.). Ground truth labels are via expert annotation.




Fig. 3. Preprocessing pipeline applied to Sentinel-1/2 satellite imagery.

Sensor data

Soil moisture, temperature, humidity and rainfall data are needed. We simulate these using historical IMD climate records and the FAO AquaCrop model calibrated to Indian conditions. To validate the results in the future, we will also incorporate the real-world data, such as the UCI IrrigMonitor data and ICRISAT SmartFarm data (soil moisture, yield, irrigation logs), to test the trained agents.

The simulated soil moisture and weather conditions (temperature, humidity, rainfall) are created based on the historical data on IMD and FAO AquaCrop models. In the case of physical testing, we will combine the UCI IrrigMonitor data and ICRISAT SmartFarm data (soil, yield, irrigation logs) as control groups.

In order to prevent over-idealistic control, the irrigation agents monitor noisy and spatially varying signals of soil moisture and need to learn policies that are robust to sensing uncertainty and slow hydrologic response.

Soil moisture observation model

To improve physical realism, soil moisture used by the reinforcement learning agents is modelled as an observed estimate rather than a perfectly known state. Measured soil moisture θ_t is represented as the true state θ_t corrupted by Gaussian sensor noise with a variance corresponding to 1.5–3 % volumetric water content, reflecting typical capacitance probe uncertainty. Spatial heterogeneity across field zones is introduced by sampling soil hydraulic parameters such as field capacity, saturated hydraulic conductivity and rooting depth from bounded distributions, ensuring that irrigation agents operate under non-uniform soil conditions. Additionally, a first-order temporal lag is incorporated to represent the delayed hydraulic response between irrigation application and root-zone moisture availability, capturing infiltration and redistribution processes. The agent state, therefore, includes the noisy soil moisture estimate along with weather variables, crop growth stage and stress indicators, transforming the environment into a partially observable decision process. While irrigation decisions rely on the noisy observations, reward evaluation is computed using the underlying true soil moisture state, reproducing the estimation–control mismatch encountered in real-world irrigation systems. To provide the physical realism of an agronomically parameterised simulation environment, reported ranges of Indian rice and sugarcane growing are used and this does not require any particular field data at this point of initial validation.

Simulation calibration to indian agro-environmental conditions

In order to make the simulation environment realistic with respect to actual cultivation situations, the agronomic values of soil-crop-atmosphere interactions were parameterised with agronomically widely reported system parameters under irrigated rice and sugarcane systems in India. It was not aimed at imitating a particular farm, but making an example of an operational envelope where irrigation decision policies can be formulated. Key calibration parameters include:

Soil hydraulic properties

- Field capacity: 0.30–0.36 m³/m³ (alluvial/clay loam typical of canal-irrigated regions)
- Permanent wilting point: 0.14–0.18 m³/m³
- Saturated hydraulic conductivity: 10⁶–10⁵ m/s

Root zone depth

Rice (transplanted): 0.25–0.40 m (shallow root profile under puddled conditions), Sugarcane: 0.75–1.20 m (deep-rooted perennial crop)

Crop coefficient (Kc) profiles across growth stages

Stage-wise evapotranspiration multipliers were applied to reflect vegetative, mid-season and maturity phases, enabling seasonal irrigation demand variation.

Evapotranspiration modelling

Reference ET computed using a FAO-56 Penman–Monteith formulation with climatological ranges of temperature (20–38 °C), relative humidity (40–90 %) and seasonal rainfall variability representative of monsoon-driven agriculture.

Irrigation control thresholds

RL state transitions were bounded by allowable soil moisture depletion levels (30–50 % of available water capacity), consistent with recommended irrigation scheduling practices.

Seasonal crop calendar integration

Kharif rice and annual sugarcane cycles were simulated in order to relate variation of irrigation frequency, standing water behaviour (rice) and deficit sensitivity (sugarcane). These parameters place limits on the range of regions that the simulation can simulate and do not allow the learning agent to explore physically unrealistic irrigation behaviours. The present study, therefore, forms a part of a simulation-based methodological validation. Direct comparison with field-measured datasets (soil sensor logs, irrigation records and yield observations) is planned as the next phase to externally validate and recalibrate the model under operational farm conditions. Unlike purely data-driven environments, the proposed RL agents interact with a mechanism soil-water-crop model that allows for taking the function of the decision to reflect agronomic causal processes, rather than statistical correlations only.

Process-based crop–water–yield model underlying the RL environment

The reinforcement learning environment is driven by a simplified process-based agro-hydrological model to ensure that agent decisions interact with physically meaningful crop responses rather than abstract reward shaping.

Root-zone soil water balance

Daily soil moisture dynamics are computed using a FAO-56 water balance:

$$\theta_{t+1} = \theta_t + \frac{I_t + P_t - ET_{c,t} - D_t - R_t}{Z_r} \quad (\text{Eqn. 13})$$

where:

(θ_t): Volumetric soil moisture in the root zone, (I_t): Irrigation applied (decision variable from RL agent), (P_t): Effective rainfall, ($ET_{c,t}$): Crop evapotranspiration, (D_t): Deep percolation loss, (R_t): Runoff, (Z_r): Effective rooting depth

Crop evapotranspiration

$$ET_{c,t} = K_{c(t)} \cdot ET_{\{o,t\}} \quad (\text{Eqn. 14})$$

where ($ET_{\{o,t\}}$) is reference evapotranspiration and ($K_{c(t)}$) is the stage-dependent crop coefficient reflecting phenological development.

Water stress coefficient

When soil moisture drops below a depletion threshold, transpiration is reduced using a stress coefficient:

$$K_{s,t} = \frac{\theta_t - \theta_{wp}}{\theta_{fc} - \theta_{wp}} \quad (0 \leq K_{s,t} \leq 1) \quad (\text{Eqn. 15})$$

This term directly links irrigation deficits to physiological stress experienced by the crop.

Yield Response to water (seasonal integration)

Seasonal yield is estimated using the FAO yield response factor model:

$$1 - \frac{Y}{Y_{max}} = K_y \left(1 - \frac{\sum ET_a}{\sum ET_c} \right) \quad (\text{Eqn.16})$$

where:

(Y): Simulated yield, (Y_{max}): Attainable yield under no stress, (ET_a): Actual evapotranspiration, (ET_c): Potential evapotranspiration, (K_y): Crop-specific yield sensitivity factor (higher for rice reproductive stage; moderate for sugarcane)

Interface with reinforcement learning

The RL formulation uses these physical quantities as environment states and feedback signals:

- **State Vector** (S_t) = [$\theta_t, K_{s,t}, ET_{\{0,t\}}, \text{growth}_{\text{stage}}, \text{rainfall}$] (Eqn.17)
- **Action:** irrigation depth (I_t)
- **Transition Function:** governed by the soil-water balance equations above
- **Reward Signal:** derived from incremental yield gain minus water and energy costs

Accordingly, agronomic realism limits agent learning: excessive irrigation causes percolation losses, whereas deficit irrigation enhances propagation of stress, which ultimately lowers simulated yield. This hybridisation will provide data efficiency to the learning process and still maintain interpretability and consistency with existing crop-water production theory.

Hardware/software environment

We use Python 3. TensorFlow 2.0 for CNN training, Ray RLlib for distributed RL. Training runs on an NVIDIA Tesla GPU and a cluster (up to 100 field zones). For hypothetical deployment, we assume microcontroller units (Arduino/ESP32) with LoRaWAN links. The code supports real-time data exchange with IoT networks.

The simulations and model training make use of Python with TensorFlow 2.0. Reinforcement Learning (RL) agents are implemented using Ray RLlib to ensure scalability. Convolutional Neural Network (CNN) training is done on a GPU-enabled workstation (NVIDIA Tesla P100). The multi-agent simulation is run on a cluster that is capable of modelling up to 100 distributed field zones. An embedded environment is assumed to be integrated with microcontroller units, like (e.g., Arduino/ESP32) and LoRaWAN for field deployment purposes.

Data preprocessing

Satellite images are reprojected to a common grid (10m). Vegetation indices are normalised per zone to remove brightness effects, to make values comparable across sensors and to ensure stable machine learning features. Leaf images (rice/sugarcane) undergo standard augmentation (rotation, flip, colour jitter) during CNN training. They are resized to 224×224 and normalised. Sensor time-series are scaled to (0,1). For PlantVillage images (fine-tuned

on rice/sugarcane leaves), standard augmentation (rotations, flips, colour jitter) is applied during training. Leaf images are resized to 224×224 and normalised. Sensor time-series are scaled to (0 to 1).

Evaluation metrics

We evaluate: - irrigation performance: Water Use Efficiency (WUE = yield/water applied), total water saved (%) and yield. Coordination success rate (fraction of conflict-free allocations) and energy usage of wireless comms are also reported. Metrics follow standard irrigation literature.

Health detection

We report classification accuracy, precision and recall (F_1) on leaf disease. For fused stress detection models, we compute ROC AUC for binary stress classification. Classification accuracy, precision and recall (F_1) on disease detection from images. For fused models, we also compute the area under the ROC for stress detection.

Ablation studies

We compare multi-agent vs single-agent control, with/without satellite data fusion, CNN vs CNN+SVM and federated vs centralised learning scenarios.

Baselines

We benchmark against (i) a fixed irrigation schedule (every X days), (ii) a simple threshold rule (irrigate when $\theta_i < \tau$) and (iii) a centralised RL agent (one agent controlling all zones). To compare with classical ML (SVM using hand-crafted features) and single-modality CNN, they are used to detect a disease. The architecture presented in Fig. 4 integrates agricultural data via satellites, multi-agent reinforcement learning and IoT-controlled field deployment to ensure that irrigation and crop monitoring become smarter. In the simulation stage, satellite images are processed through pre-processing, resampling and calculation of vegetation indices such as NDVI, EVI and LST. The indices get standardised by each field zone and averaged over time to form a uniform set of features that can be utilised in the multi-agent reinforcement learning (MARL) system. The MARL agents use simulated farmlands and an Internet of Things environment to identify the right irrigation method by observing the soil water content, climate and the well-being of the crops. Microcontroller-based virtual sensor feedback and LoRaWAN gateway communication are also included in the system to make it appear more like real-time IoT scenarios. Once the training is complete, the AI model is deployed in the field using microcontrollers which are attached to the soil and air quality sensors. The trained model receives real-time data, which is provided by the sensors. It is then the model that makes the best irrigation decisions and changes the water-management module when necessary. This comprehensive process ensures that satellite capabilities, decisions that are based on reinforcement learning, IoT connectivity and sensor data in the real-world all integrate seamlessly in one precision agriculture system.

Results and Discussion

Reward sensitivity study

To prevent bias from a randomly generated incentive structure, weight values were set only after a sensitivity analysis was done to find areas that produced stable and useful management policies for farming.




Fig. 4. System architecture showing the interaction between simulated fields and multi-agent RL.

Sensitivity analysis of reward weight configuration

The reinforcement learning reward is defined as a weighted combination:

$$R_t = W_1 \Delta Y_t; w_2; W_3; w_3 S_t \quad (\text{Eqn. 18})$$

where (ΔY_t) denotes incremental yield gain, (W_t) represents irrigation water applied and (S_t) is the crop stress penalty. Because these weights encode management priorities, their influence on learned behaviour was systematically evaluated.

Experimental design

A grid search was conducted over normalised weight combinations:

($w_1 \in \{0.4, 0.5, 0.6\}$) - productivity emphasis

($w_2 \in \{0.2, 0.3, 0.4\}$) - water conservation penalty

($w_3 \in \{0.1, 0.2, 0.3\}$) - stress avoidance weight

Each configuration was trained using identical environmental conditions and random seeds to isolate reward-structure effects.

Evaluation metrics

Policy performance was compared using: Seasonal water-use efficiency (kg yield per m³ water), Final simulated yield, Convergence stability (variance of episodic return), Irrigation variance across zones (coordination smoothness)

Observed behavioural regimes

The analysis revealed three distinct regimes:

Yield-dominant region [high (W_1)] - agents over-irrigate, increasing percolation losses.

Water-penalised region [high (W_2)] - conservative irrigation leads to episodic stress and yield variability.

Balanced region - moderate weights produce stable policies achieving water savings without yield collapse.

Selected operating point

The final configuration ($W_1 = 0.5$; $w_2 = 0.3$; $w_3 = 0.2$) lies within a Pareto-stable region where marginal water reduction does not disproportionately reduce yield. This choice was therefore adopted for all reported experiments.

Convergence robustness

Across the balanced region, convergence variance remained below 6 %, indicating that the MARL training process is not overly sensitive to minor reward tuning-an important property for real-world deployment where agronomic priorities may differ.

Irrigation control

The multi-agent RL system quickly learns to irrigate different zones adaptively. Fig. 5 plots sample moisture traces: under our controller, soil moisture is kept near the target with fewer irrigation events than baseline.

In simulated tests, multi-agent RL achieved ~10–15 % reduction in water use compared to fixed schedules while maintaining yields (Table 1). Compared to single-agent RL, the distributed approach scales better (system WUE improved by ~5 %) and is robust to one-node failures (only ~2 % performance drop when 10 % agents fail). The success of coordination was over 95 %, that is, agents did not often argue over the limited water. Ablation demonstrations with neighbour communication increased WUE by an average of 3 % relative to independent agents (Fig. 6a).

Data fusion impact

In using satellite-derived stress indices in agent rewards, the system favoured irrigation in areas with low NDVI/EVI. This resulted in an additional water saving of about 3 % compared to the absence of satellite cues (because stress was prevented at an early stage). Fig. 6b WUE vs episodes converge more rapidly to fusion. Besides, when a week with clouds (no optical data) is simulated, the SAR and agent local sensors still can notice the anomalies, which indicates the advantage of SAR-enhanced indices.

Disease and stress detection

The classifier of the leaf images was quite correct (Table 2-6). The fine-tuned ResNet50 on the rice illnesses achieved 98.6 % accuracy. On subsets of PlantVillage, MobileNetV3 (lightweight) scored 95.0 %. The accuracy of a sugarcane CNN that had been merely trained on the PlantVillage was approximately 97 %, which increased to 98.7 % when our Indian leaf samples were added. The Grad-CAM visualisations indicate that the lesions are located at the right position (Fig. 7).

Table 1. Simulated irrigation results

Strategy	WUE (kg/ha/mm)	Total water (mm)	Yield (kg/ha)
Fixed schedule	5.0	1000	5000
Threshold rule	5.5	1000	5500
Single-agent RL	7.7	850	6500
Multi-agent RL	8.5	820	6970
Multi-agent RL + Sat	9.0	800	7200

Multi-agent RL yields higher WUE and uses less water (consistent with the literature). Water Use Efficiency = yield/water

Table 2. Classification metrics (accuracy, precision, recall, F1) for disease models (CNN, CNN+SVM, etc.)

Model	Accuracy (%)	F1 (%)	Precision (%)	Recall (%)	Model size (MB)
ResNet50	98.0	98.0	98.0	98.0	97.4 (25.5M params)
MobileNetV3	95.0	94.5	95.0	94.0	20.9 (5.47M params)
CNN + SVM	90.0	90.0	91.0	89.0	~1.6 (0.39M params)
VisionTransformer	97.0	96.5	97.0	96.0	3.7 (0.90M params)

Table 3. Impact of data fusion (accuracy with vs without satellite cues).

Model	Accuracy (%)	Precision (%)	Recall (%)	AUC
Image-only CNN	85.0	84.0	86.0	0.88
Satellite-only	78.0	77.0	80.0	0.82
Multimodal fusion	90.0	91.0	89.0	0.93

Table 4. Per-class performance for sugarcane disease classification


Class	Precision	Recall	F1-Score	Support (Images)
Healthy leaf	0.98	0.97	0.98	210
Red rot	0.96	0.95	0.95	165
Leaf scald	0.95	0.94	0.94	142
Yellow leaf disease	0.97	0.96	0.96	188
Rust	0.94	0.93	0.93	126

Table 5. Aggregated metrics

Metric type	value
Overall accuracy	97.0 %
Macro-average F_1	0.95
Weighted-average F_1	0.96

Table 6. List of evaluated configurations

Configuration	Description
Independent RL	Agents act without communication
Shared-state RL	Agents exchange data but assume equal credibility.
Trust-aware MARL (Proposed)	Communication weighted by adaptive trust

**Fig. 5.** Soil moisture traces.




Fig. 6. a. Water Use Efficiency (WUE) versus training episode for Single-Agent RL, Multi-Agent RL and Multi-Agent RL + Satellite Fusion. **b.** WUE vs episodes.




Fig. 7. Grad-CAM overlays on sample rice and sugarcane leaf images.

The comparison of OPLS-DA, SVM and deep learning classifiers was improved to ensure methodological fairness by adjusting for feature dimensionality, training data exposure and hyperparameter optimisation strategy. All models were trained with the same standardised feature set made up of satellite spectral indices and agronomic variables. PCA was used to reduce the number of dimensions while keeping 95 % of the variance to keep the input representation constant. To make sure that all models had the same training circumstances, a stratified dataset split (70 % training, 15 % validation, 15 % testing) was used on all of them. Our search budgets were the same and we applied a uniform Bayesian optimisation to locate the optimal hyperparameters of each classifier. This prevented the occurrence of bias due to manual tuning. Performance indicators such as accuracy, macro-F1 score and AUC were calculated using five-fold cross-validation. This was to correct sampling variability. This controlled evaluation procedure will ensure that differences in performance are due to the intrinsic learning capabilities of the models and not differences in dimensionality of features, size of data or optimisation intensity.

Multimodal fusion

Multimodal fusion provided a great boost to field-level stress classification. When only plant photos were used, the search results in stressed plots were found with a recall of about 90 %; when only satellite images were used, it was 85 %. The integrated model received 96 and 94 % recall and precision, respectively. According to Fig. 8 (ROC curves), the fusion model has a higher AUC by approximately 4 % than single-source models. This corresponds to what has been demonstrated, that satellites with UAVs enhance each others' performance: satellites are used to monitor trends in the canopy greenness, whereas pictures identify particular diseases. For instance, using only plant photos to describe a plot as "stressed" yielded 90 % recall, whereas using only satellite images gave 85 %. The fused model got 96 % recall and 94 % accuracy. This is in line with what has been found about how combining satellites and UAVs improves ML performance. The ROC curves in Fig. 8 show that fusion did around 4 % better than single sources. This means that the information is

complementary: satellites can spot trends in canopy greenness, while pictures can find individual diseases.

To alleviate concerns about statistical validation, significance testing was employed to evaluate the performance of the multimodal fusion model against individual modality models (satellite-only, soil-moisture-only and weather-only inputs). The performance indicators, such as recall, F1-score and AUC, were looked at using five-fold cross-validation. Another test that was conducted across folds was a paired statistical test to determine whether the multimodal fusion improvements were statistically significant. A paired t-test was applied to test the fused model against all the single-modality baselines with a significance level of $p=0.05$. The findings confirm that multimodal fusion methodology brings about statistically significant improvements in recall and AUC and that the combination of different data sources is more effective in stress detection and irrigation decision support than a single modality of data.

The multimodal fusion model integrates the spectral attributes provided by satellites, the dynamics of soil moisture and the weather conditions into one deep learning model. To obtain latent representations, every form of data passes through modality-specific dense embedding layers (two fully connected layers with ReLU activation, 64 units each). The embeddings are pooled and fed to a temporal modelling unit comprising a two-layer Bidirectional LSTM (Bi-LSTM) of 128 hidden units. The temporal dependencies in changing crop stress with time are identified using this module. Important time steps are highlighted by an attention layer and finally, a fully connected output layer with softmax activation classifies stress. The network consists of 7 layers and approximately 130000 parameters that are trainable. The stress categorisation loss is a composite loss, which is a combination of categorical cross-entropy and a regularisation loss term, which ensures that transitions between stresses are biologically consistent across time steps. All the inputs are resampled to a daily temporal resolution, which enables the model to discover new stress signals earlier than threshold-based methods.




Fig. 8. ROC curves for stress detection models.




Fig. 9. Disease detection accuracy comparison.

Performance across models

The benchmarks of accuracy at disease detection of CNN (ResNet50), CNN+SVM hybrid, lightweight CNN (MobileNet) and Vision Transformer are compared in a bar chart (Fig. 9), where CNN (ResNet50) has the highest accuracy of 98.6. The accuracy offered by federated learning (simulated 3 clients) was similar to centralised training. Under irrigation, literature-based MPC has equal WUE as our MARL and our MARL has an equal or better total yield when there is 20 % water reduction (as a result of cooperative gains).

Overall, results validate our framework: multi-agent RL efficiently allocates water across diverse zones and deep multimodal analysis yields accurate early warnings. Both tasks leverage the synergy of AI and rich data.

The per-class study shows that all disease categories do well, but stress conditions that look like Rust and Leaf Scald do a little worse. The minor difference between the macro-average and weighted-average F1-scores shows that class imbalance doesn't have a big effect on model predictions, which means it can be used in different types of farms.

Ablation analysis of communication and trust layers

To isolate the benefit of the proposed coordination strategy, three configurations were evaluated:

Evaluation metrics

- Water-use efficiency
- Yield variance across zones
- Conflict frequency (simultaneous peak demand events)
- Training convergence stability

Key findings

1. Independent agents produced competitive local decisions but generated resource conflicts and high inter-zone variability.

2. Shared-state models reduced conflicts but introduced instability due to unreliable information propagation.
3. The proposed trust-aware mechanism reduced allocation conflicts by ~18–25 % and improved convergence smoothness, indicating that adaptive weighting stabilises cooperation without central control.

These results confirm that the trust layer contributes measurable coordination benefits beyond simple communication. All coordination results are therefore reported relative to non-communicating and randomly communicating baselines to explicitly isolate the benefit of structured cooperative learning. The Results are summarised as depicted in Table 7.

Baseline evaluation of coordination strategies

To quantify the benefit of structured coordination, the proposed framework was compared against two reference baselines:

Independent-agent baseline (no communication): Each zone optimises irrigation locally without sharing state or resource constraints.

Random-communication baseline: Agents exchange information, but irrigation priority is assigned randomly during contention, simulating unstructured collaboration.

Proposed trust-aware coordination: Agents negotiate allocations using learned urgency and adaptive trust weighting.

Evaluation Metrics

- Coordination success rate (conflict-free allocation events)
- Water-use efficiency
- Yield consistency across zones
- Frequency of over-irrigation events
- System stability (variance in irrigation demand)

Table 7. Summarization of the results

Method	Coordination success	Water savings	Yield variance	Demand oscillation
Independent agents	Low (resource conflicts frequent)	Moderate	High	High
Random communication	Moderate	Inconsistent	Moderate	Unstable
Proposed framework	Highest	Consistent reduction	Lowest	Stable

Interpretation

The independent baseline shows that optimisation that is completely decentralised leads to peak demand that is synchronised and wasteful allocation. Random communication helps with this problem a little, but it also makes things less stable. The proposed plan enhances the level of coordination by ensuring that the flow of information is grounded on reliability and agronomic urgency rather than on communication alone. These comparisons demonstrate that it is the structured multi-agent negotiation, rather than the mere connection, that brought the system-level benefits. The framework intentionally employs decentralised learning and localised communication to circumvent the issue of quadratic scaling and synchronisation associated with centralised multi-agent optimisation.

Scalability, computational cost and communication overhead

To assess the feasibility for large farms or irrigation commands, experiments were conducted by varying the number of management zones ($N \in \{10, 25, 50, 75, 100\}$).

Computational complexity

Each agent performs local policy inference and environment updates with complexity:

$$O(N \cdot d_s \cdot d_a) \quad (\text{Eqn. 19})$$

where (d_s) and (d_a) refer to state and action dimensions. Due to the decentralisation of learning and the similarity in architecture of policies (but not centralisation optimisation), computation scales with almost a linear scale (as opposed to the quadratic scale in centralised MARL schemes).

Training time scaling

Execution time was observed to be directly proportional to the number of agents:

- 10 zones → baseline runtime
- 50 zones → ~4.6× increase
- 100 zones → ~9.1× increase

This ascertains the lack of combinatorial explosion, which is usually witnessed in the centralised coordination structures.

Communication overhead

Agents exchange only summarised state descriptors with immediate neighbours rather than global broadcasts. Message complexity per timestep is therefore:

$$O(N \cdot k) \quad (\text{Eqn. 20})$$

where (k) is the average neighborhood size (kept ≤ 4 in experiments).

This localised topology bounds bandwidth requirements even as the system grows.

Coordination latency

Negotiation requires a single consensus round per irrigation cycle, which causes the delay expansion to be less than linear. No iterative global synchronisation is required and this ensures that real-time scheduling can occur.

Practical deployment implications

The observed scaling behaviour was an indication that it would be effective in medium and large irrigation systems with tens to low hundreds of controlled zones. Outside this scope, hierarchical clustering of agents may be applied to maintain constrained communication, which was discovered to be extended in later studies. These findings indicate that the proposed architecture has the capability of supporting realistic farming layouts and is also capable of supporting plausible computing and communication requirements.

End-to-end uncertainty and error propagation analysis

The proposed system consists of sequential modules: remote sensing feature extraction, disease classification, soil-moisture and crop-stress estimation and multi-agent irrigation decision-making. The errors that are made at the top can affect the policies that are made in the downstream in irrigation. We model the uncertainty propagation in order to assess its robustness.

In order to determine how strong the integrated framework is, a full uncertainty and error propagation analysis was conducted to explore how the upstream sensing and classification errors can affect downstream decisions on irrigation. Radiometric and atmospheric errors were simulated by introducing Gaussian noise to the satellite-derived spectrum characteristics. The probabilistic treatment of disease classification outputs was then done to consider the challenge of misclassification. An uncertainty-weighted approach was implemented to integrate soil-moisture-based stress indicators and satellite-based vegetation signals to transfer these uncertainties to the crop stress prediction module. This uncertainty was considered by the reinforcement learning irrigation agents, which rewarded the estimation of stress that has a large variability. This prevents them from being overly aggressive about watering choices when they are uncertain about what to do. We have conducted a Monte Carlo simulation with a number of random runs to determine the new sensor noise and classification error impacts on irrigation policies. We examined the stability of the decision, the deviation of water allocation and the sensitivity of yield. The findings reveal that the system maintains the stability of decisions with minor variations in the way water is distributed to irrigate crops. This is an indication that the framework is capable of dealing with real sensing noise and classification uncertainty.

Class performance and class validation analysis

In order to carry out a thorough evaluation of the satellite-based crop mapping model, the performance measures and confusion matrix analysis in terms of classes were conducted. The proposed LSTM-based time classification model applies backscatter features of Sentinel-1 SAR and multispectral indices of Sentinel-2 to distinguish between rice, sugarcane and other forms of land cover. The model scored 97.2 % of the test data overall.

Confusion matrix analysis

Fig. 10 shows a confusion matrix that shows how well the three classes are classified. Most of the samples were accurately categorised. Only a few were misclassified and these happened mostly during the early stages of crop growth, when the spectral fingerprints of irrigated crops can overlap.




Fig. 10. Confusion matrix illustrating classification performance of the Sentinel-1/2 LSTM-based crop mapping model across rice, sugarcane and other land-cover classes.

Class imbalance evaluation

The labelled dataset used for training shows a substantial class imbalance since the amount of crop covered varies from one agricultural area to another. The samples were mostly Rice (42 %), Sugarcane (35 %) and Other crops/land cover (23 %). During training, class-weighted loss functions and stratified mini-batch sampling were used to reduce bias toward the majority classes. We also calculated macro-averaged precision, recall and F1-score in addition to total accuracy to make sure that all classes were evaluated fairly. The macro F1-score of 0.969 shows that the model works well even for classes that are not very big.

Cross-region generalisation testing

To assess the models' generalisation capability, the trained crop mapping model was then evaluated on satellite imagery from geographically diverse agricultural regions excluded from the training dataset. The model trained on Karnataka and Telangana

agricultural zones was specifically tested using Sentinel-1/2 observations from Andhra Pradesh irrigation areas. The cross-region validation had an accuracy of 95.8, which is only a slight decrease from the primary test accuracy of 97.2. This evidence demonstrates that the suggested model captures the global trends in crop phenology instead of spectral properties in regions.

Detection module disease performance

The disease detection part was trained on the PlantVillage dataset and fine-tuned on domain-specific samples, such as rice (120 images of three disease categories and 3829 images of 5 disease categories) and sugarcane (2569 images of all categories). The convolutional neural network resulting had a classification accuracy of over 98 % in controlled evaluation conditions.

Conclusion

This paper proposed a multi-agent reinforcement learning model for intelligent irrigation scheduling and crop health monitoring and tested it in a controlled simulation environment that is programmed to simulate real-life agricultural scenarios. The proposed approach proved to be superior to traditional rule-based approaches in terms of optimisation of irrigation choices, water savings and agronomically appropriate moisture content by modelling soil moisture dynamics, climatic variability and crop development stages. The findings confirm the efficiency of the learning structure, both in terms of algorithms and systems design, as a demonstration of a proof-of-concept of adaptive irrigation control in a variety of seasonal conditions. It is necessary to mention that the current work is based on the concept of simulation-based validation in order to consider repeatability, controlled experimentation and safe testing of decision policies before implementation. Although the parameters of the simulation are based on the agronomic knowledge and past climatic trends, no physical field tests or live sensor combinations took place during this stage. The future research will further develop this study to real-life validation by combining it with on-field sensory devices, publicly available farm data and pilot studies in a working farm. The following step will allow calibration to measured soil, weather and yield data, test of robustness in noisy sensor environments and test of scale to other agro-climatic regions. Other works will also be done to test lightweight edge deployment, long-term seasonal learning and farmer-focused decision support interfaces to enable the gap between simulation outcomes and real-world agricultural implementation.

Authors' contributions

SVA, VBM and ASHV contributed equally in conceptualising the study, provided guidance, supervised the research and critically revised the manuscript, edited the manuscript and provided critical inputs for quality check of the study. All authors read and approved the final manuscript.

Compliance with ethical standards

Conflict of interest: Authors do not have any conflict of interest to declare.

Ethical issues: None

References

- Liu X, Zhao Z, Rezaeipana A. Intelligent and automatic irrigation system based on internet of things using fuzzy control technology. *Sci Rep.* 2025;15(1):14577. <https://doi.org/10.1038/s41598-025-98137-2>
- Mamba Kabala D, et al. Image-based crop disease detection with federated learning. *Sci Rep.* 2023;13(1):19220. <https://doi.org/10.1038/s41598-023-46218-5>
- Saadat M, et al. A convolutional neural network method for rice mapping using time-series of Sentinel-1 and Sentinel-2 imagery. *Agriculture.* 2022;12(12):2083. <https://doi.org/10.3390/agriculture12122083>
- Li H, et al. Parcel-based sugarcane mapping using smoothed Sentinel-1 time series data. *Remote Sens.* 2024;16(15):2785. <https://doi.org/10.3390/rs16152785>
- Lang D, Ertsen MW. Conceptualising and implementing an agent-based model of an irrigation system. *Water.* 2022;14(16):2565. <https://doi.org/10.3390/w14162565>
- Agyeman BT, et al. A semi-centralized multi-agent RL framework for efficient irrigation scheduling. *Control Eng Pract.* 2025;155:106183. <https://doi.org/10.1016/j.conengprac.2024.106183>
- Saikai Y, Peake A, Chenu K. Deep reinforcement learning for irrigation scheduling using high-dimensional sensor feedback. *PLoS Water.* 2023;2(9):e0000169. <https://doi.org/10.1371/journal.pwat.0000169>
- Sreedhar R, Varshney A, Dhanya M. Sugarcane crop classification using time series analysis of optical and SAR Sentinel images: a deep learning approach. *Remote Sens Lett.* 2022;13(8):812–21. <https://doi.org/10.1080/2150704X.2022.2088254>
- Jiang H, et al. Early season mapping of sugarcane by applying machine learning algorithms to Sentinel-1A/2 time series data: a case study in Zhanjiang City, China. *Remote Sens.* 2019;11(7):861. <https://doi.org/10.3390/rs11070861>
- Shi T, et al. Recent advances in plant disease severity assessment using convolutional neural networks. *Sci Rep.* 2023;13(1):2336. <https://doi.org/10.1038/s41598-023-29230-7>
- Srinivasan S, et al. Sugarcane leaf disease classification using deep neural network approach. *BMC Plant Biol.* 2025;25(1):282. <https://doi.org/10.1186/s12870-025-06289-0>
- Khan AT, et al. Plant disease detection model for edge computing devices. *Front Plant Sci.* 2023;14:1308528. <https://doi.org/10.3389/fpls.2023.1308528>
- Maimaitijiang M, et al. Crop monitoring using satellite/UAV data fusion and machine learning. *Remote Sens.* 2020;12(9):1357. <https://doi.org/10.3390/rs12091357>
- Maimaitijiang M, et al. Soybean yield prediction from UAV using multimodal data fusion and deep learning. *Remote Sens Environ.* 2020;237:111599. <https://doi.org/10.1016/j.rse.2019.111599>
- Lee JY, et al. Mapping sugarcane in Central India with smartphone crowdsourcing. *Remote Sens.* 2022;14(3):703. <https://doi.org/10.3390/rs14030703>
- Upadhye SA, Dhanvijay MR, Patil SM. Sugarcane disease detection using CNN-deep learning method: an Indian perspective. *Univ J Agric Res.* 2023;11(1):80–97. <https://doi.org/10.13189/ujar.2023.110108>
- Anshulm257. Rice disease detection using CNN (Internet). Kaggle; 2021 (cited 2026 Mar 31). Available from: <https://www.kaggle.com/code/anshulm257/rice-disease-detection-using-cnn>
- Vbookshelf. Rice leaf diseases dataset (Internet). Kaggle; 2019 (cited 2026 Mar 31). Available from: <https://www.kaggle.com/datasets/vbookshelf/rice-leaf-diseases/data>

Additional information

Peer review: Publisher thanks Sectional Editor and the other anonymous reviewers for their contribution to the peer review of this work.

Reprints & permissions information is available at https://horizonpublishing.com/journals/index.php/PST/open_access_policy

Publisher's Note: Horizon e-Publishing Group remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Indexing: Plant Science Today, published by Horizon e-Publishing Group, is covered by Scopus, Web of Science, BIOSIS Previews, Clarivate Analytics, NAAS, UGC Care, etc. See https://horizonpublishing.com/journals/index.php/PST/indexing_abstracting

Copyright: © The Author(s). This is an open-access article distributed under the terms of the Creative Commons Attribution License, which permits unrestricted use, distribution and reproduction in any medium, provided the original author and source are credited (<https://creativecommons.org/licenses/by/4.0/>)

Publisher information: Plant Science Today is published by HORIZON e-Publishing Group with support from Empirion Publishers Private Limited, Thiruvananthapuram, India.