iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 Issue: IX Month of publication: September 2026

DOI: https://doi.org/10.22214/ijraset.2026.84820

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

An Oracle-Distilled Neural Network for Hardware-Free Precision Water Management across Indian Agroclimatic Zones

Shiv Mandlik[1], Aarush Khilosia[2]
[1]Adani International School, Ahmedabad, Gujarat, India
[2]Shanti Asiatic School, Ahmedabad, Gujarat, India
Both authors contributed equally to this work.

Abstract: Agriculture accounts for roughly 90% of India's freshwater withdrawals and is driving unsustainable groundwater depletion. Sensor-based precision irrigation, which could reduce this demand, costs USD 500–2,000 per installation and is therefore out of reach of the smallholders who cultivate about 86% of Indian farmland. This paper presents an irrigation policy learned by behavioral cloning from a dynamic programming oracle that has perfect foresight of an entire season's weather. The resulting neural network (approximately 105,000 parameters; 420 KB in ONNX format) requires only a three-day public weather forecast, runs on a mid-range smartphone in under 5 ms, and uses no soil sensors or field hardware of any kind. The policy was evaluated over 15 locations spanning 12 states and four crops (cotton, maize, groundnut and soybean), including three locations and three seasons withheld entirely from training. On the held-out data it attained 97–99% of maximum attainable yield while reducing irrigation by 28–44% relative to a physics-based threshold controller functionally equivalent to sensor scheduling, and by 59–67% at crop level relative to prevailing high-yield farmer practice (43–92% depending on location; 61% on average). Behavioral analysis shows that the learned policy applies frequent, small irrigations timed to the growth stages of highest yield sensitivity, defers irrigation when rainfall is forecast, and tolerates controlled stress during stage-insensitive periods. Because it generalises to unseen locations and requires no capital expenditure, the approach offers a practical route to precision water management for more than 100 million Indian smallholders.
Keywords: Smart Irrigation; Behavioral Cloning; Water Conservation; Precision Agriculture; Smartphone-Based AI; Smallholder Farming; Groundwater Sustainability

Nomenclature

AI	Artificial Intelligence	Kc	Crop Coefficient
BC	Behavioral Cloning	Ky	Yield Response Factor
DAS	Days After Sowing	MDP	Markov Decision Process
DP	Dynamic Programming	ONNX	Open Neural Network Exchange
ET_0	Reference Evapotranspiration	RAW	Readily Available Water
ETc	Crop Evapotranspiration	TAW	Total Available Water
FAO	Food and Agriculture Organization	θFC	Field Capacity
ICAR	Indian Council of Agricultural Research	θPWP	Permanent Wilting Point

I. INTRODUCTION

India supports about 18% of the world's population with roughly 4% of its freshwater resources. Agriculture is the backbone of the national economy, supporting over 45% of the workforce [3] and consuming more than 90% of the country's freshwater withdrawals [1]. This dependence has produced groundwater extraction that is widely regarded as unsustainable: India is the world's largest user of groundwater, and over 60% of its irrigated agriculture depends on aquifers [13], whose levels are declining across the Indo-Gangetic plain and the peninsular region [2]. Climate change compounds the problem, rendering the monsoon, historically the lifeline of Indian agriculture, increasingly unpredictable.

The consequences fall hardest on smallholder farmers, who cultivate about 86% of India's landholdings on plots of less than two hectares each [3]. They face a persistent dilemma: irrigate generously to protect yield, at the cost of exhausting their water supply, or follow conventional schedules that conserve water but sacrifice productivity. Traditional irrigation calendars were developed over generations of predictable monsoons and become unreliable as rainfall variability increases.

Precision irrigation techniques exist, but they remain inaccessible to those who need them most. Sensor-based systems require USD 500–2,000 per installation in hardware alone, together with technical maintenance and periodic sensor replacement. These costs lie far beyond the reach of a smallholder. Drip and micro-irrigation, though subsidised and widely promoted, still demand substantial capital investment and, more fundamentally, do not answer the questions of when and how much to irrigate. Calendar-based recommendations issued by agricultural extension services [4] do not respond to weather variability, wasting water during wet periods and causing yield loss during dry spells. The result is a technological divide: precision agriculture is adopted by commercial farms while smallholders persist with flood irrigation, yield stagnation and falling water tables.

This paper presents a fundamentally different approach. We engineered an artificial intelligence system that uses only openly available weather forecasts and yet exceeds the performance of current sensor-based precision irrigation. The system uses no soil sensors, no moisture probes and no field hardware whatsoever. Its only input from the farmer is how much irrigation was applied, and when. The policy is trained by behavioral cloning [5], a supervised learning technique in which a neural network learns to imitate expert decisions. In our framework the expert is a dynamic programming (DP) oracle that computes the theoretically optimal irrigation schedule using perfect foresight of the entire season's weather. The oracle represents an absolute ceiling that no deployable system can exceed. By capturing its core decision principles, the trained network approaches that ceiling using only a three-day forecast obtained free of charge through a public API on any smartphone.

Our research objectives were as follows:

1) To develop an AI-driven irrigation system that matches or surpasses sensor-based performance without requiring capital-intensive hardware.
2) To validate the system across four major crops (cotton, groundnut, maize and soybean) at 15 locations spanning the major agroclimatic zones of India.
3) To quantify water savings relative to existing baselines without compromising food security.
4) To evaluate generalisation to locations and seasons entirely unseen during training, using both spatially and temporally held-out data.

II. MATERIALS AND METHODS

A. Study Area and Data Sources

1) Selection of Locations

The study covers 15 locations across 12 Indian states (Table I), spanning latitudes from 10.8° N (Coimbatore, Tamil Nadu) to 30.9° N (Ludhiana, Punjab) and representing arid, semi-arid and sub-humid climates. The selection criteria were: (i) representation of the major agroclimatic zones of India; (ii) significance for the four target crops; (iii) availability of long-term weather records; and (iv) diversity of rainfall regimes and soil types.

TABLE I
STUDY LOCATIONS AND CHARACTERISTICS

Location	State	Climate Zone	MAP (mm)	Soil Type	Dataset
Jodhpur	Rajasthan	Arid	363	Sandy loam	Train
Ahmedabad	Gujarat	Semi-arid	782	Clay loam	Train
Ludhiana	Punjab	Semi-arid	734	Loam	Train
Hisar	Haryana	Semi-arid	429	Sandy loam	Train
Meerut	Uttar Pradesh	Sub-humid	845	Silt loam	Train
Varanasi	Uttar Pradesh	Sub-humid	1,020	Clay	Train
Nagpur	Maharashtra	Semi-arid	1,082	Clay	Train
Bellary	Karnataka	Semi-arid	534	Red loam	Train
Coimbatore	Tamil Nadu	Semi-arid	657	Red loam	Train
Jabalpur	Madhya Pradesh	Sub-humid	1,386	Clay	Train

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

Location	State	Climate Zone	MAP (mm)	Soil Type	Dataset
Patna	Bihar	Sub-humid	1,099	Alluvial	Train
Indore	Madhya Pradesh	Semi-arid	946	Black clay	Train
Dharwad	Karnataka	Semi-arid	751	Red loam	Test
Hyderabad	Telangana	Semi-arid	812	Red sandy	Test
Guntur	Andhra Pradesh	Sub-humid	856	Alluvial	Test

MAP = Mean Annual Precipitation. Train: 2010–2022 (13 years); Test: 2023–2025.

To ensure a rigorous evaluation, the dataset was partitioned both spatially and temporally. The training set comprised 12 locations over 2010–2022 (13 years). Three locations (Dharwad, Guntur and Hyderabad) were withheld entirely and evaluated on 2023–2025 data. This spatial–temporal split tests whether the policy generalises to environments and time periods never encountered during training.

2) *Weather and Soil Data*

Daily historical weather was obtained from the Open-Meteo Historical Weather API [6] at 0.1° spatial resolution, comprising maximum and minimum temperature (°C), relative humidity (%), precipitation (mm), wind speed (m s^{-1}) and solar radiation (MJ m^{-2} d^{-1}). Reference evapotranspiration (ET_0) was computed using the FAO-56 Penman–Monteith equation [7]:

$$ET_0 = \frac{0.408\,\Delta(R_n - G) + \gamma \frac{900}{T+273} u_2 (e_s - e_a)}{\Delta + \gamma(1 + 0.34\,u_2)}$$

where Δ is the slope of the saturation vapour pressure curve, R_n is net radiation, G is soil heat flux, γ is the psychrometric constant, T is mean temperature, u_2 is wind speed at 2 m height, e_s is saturation vapour pressure and e_a is actual vapour pressure.

Soil hydraulic properties, namely field capacity (θFC), permanent wilting point (θPWP) and saturated hydraulic conductivity (Ksat), were derived from ISRIC SoilGrids 2.0 at 250 m resolution [8]; where texture-based cross-checks were required, the pedotransfer relationships of Saxton and Rawls [9] were applied. Total available water was computed as:

$$TAW = 1000 \times (\theta_{FC} - \theta_{PWP}) \times Z_r$$

where Z_r is the effective root depth, which varies by crop and growth stage.

For operational deployment the system uses a three-day forecast from the Open-Meteo Forecast API. All data services used are open access and require neither API keys nor subscription fees.

B. *Crop Simulation Model*

1) *Water Balance Simulation*

The simulation engine implements the FAO-56 single crop coefficient method [7], tracking daily root zone depletion (D_r) as:

$$D_{r,t} = D_{r,t-1} - P_t - I_t + ET_{c,t} + D_{p,t}$$

where P is precipitation, I is irrigation, ETc is crop evapotranspiration and Dp is deep percolation (all in millimetres). Deep percolation occurs when water input exceeds field capacity. Crop evapotranspiration is computed as ETc = Kc × ET_0, with Kc varying across the four phenological stages in accordance with FAO-56 guidelines. Crop parameters are given in Table II.

TABLE II
CROP PARAMETERS USED IN THE SIMULATION MODEL

Parameter	Cotton	Groundnut	Maize	Soybean
Growing season (days)	165	130	120	110
Kc initial	0.35	0.40	0.30	0.40
Kc mid-season	1.20	1.15	1.20	1.15
Kc late	0.50	0.60	0.35	0.50
Ky (critical stage)	1.15	0.70	1.25	0.80
Max root depth (m)	1.30	0.70	1.00	0.80

Kc values from FAO-56 [7]; Ky values from FAO-33 [10].

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

2) Yield Model

Yield response to water stress follows the Stewart multiplicative model [11]:

$$\frac{Y_a}{Y_m} = \prod_{i=1}^{n}\left[1 - k_{y,i}\left(1 - \frac{ET_{a,i}}{ET_{c,i}}\right)\right]$$

where Y_a/Y_m is relative yield, $k_{y,i}$ is the yield response factor for growth stage i, and $ET_{a,i}/ETc_{,i}$ is relative evapotranspiration during that stage. Higher k_y values indicate greater yield sensitivity to water deficit. Cotton, for example, has $k_y \approx 1.15$ during flowering, so a 10% evapotranspiration deficit at that stage reduces yield by approximately 11.5%.

C. Baseline Irrigation Methods

Three baselines were implemented, spanning the spectrum of current practice.

1) ICAR Government Schedule: calendar-based irrigation at fixed intervals and depths, following the production guidelines of the Indian Council of Agricultural Research [4]. This method is widely practised and actively promoted by agricultural extension services. It accounts for neither soil moisture conditions nor actual weather.
2) High-Yield Farmer Practice: irrigation is applied whenever soil moisture falls below 70% of field capacity, eliminating water stress entirely. This risk-minimising strategy is widespread where electricity for pumping is subsidised or effectively free. It attains maximum yield, but at the cost of heavy deep percolation losses and excessive water extraction.
3) Physics-Based Threshold: irrigation is triggered automatically when root zone depletion exceeds a crop-specific threshold (typically 50–60% of TAW). This represents the best available reactive scheduling and is functionally equivalent to a soil moisture sensor installation, but it cannot optimise across growth stages or anticipate forecast rainfall.

D. AI Training Methodology

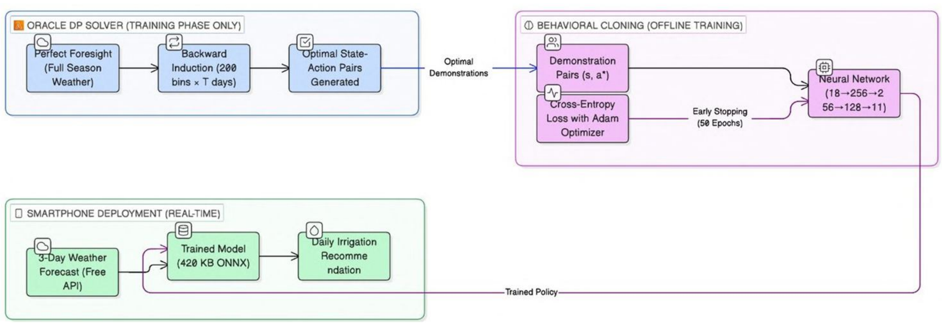

Fig. 1 System architecture. The DP oracle is used only during training; deployment requires nothing beyond a three-day public weather forecast and a smartphone.

1) Problem Formulation

Irrigation was formulated as a finite-horizon Markov decision process (MDP) defined by the tuple (S, A, R, T) [12].

State space (18 dimensions): normalised root zone depletion ratio, days after sowing, current Kc, effective root depth, current-day precipitation and ET_0, a three-day forecast of precipitation and ET_0 (six values), crop type (four one-hot values), and soil parameters (θFC, θPWP, Ksat).

Action space (11 levels): irrigation amounts ranging from 0 mm to 50 mm in 5 mm increments: {0, 5, 10, 15, 20, 25, 30, 35, 40, 45, 50}.

Reward function:

$$R_t = -c_w \cdot I_t - c_s \cdot \max(0,\ D_{r,t} - RAW) - c_d \cdot D_{p,t} + Y_{final}$$

where c_w penalises water use, c_s penalises stress when depletion rises above readily available water (RAW ≈ 0.5 × TAW), c^d penalises percolation loss, and Y^{final} is the normalised yield reward. This structure encodes the central trade-off between water conservation and yield protection.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

2) Oracle Dynamic Programming Solver

The novelty of our approach lies in the source of the training data. Rather than trial-and-error reinforcement learning, which typically requires millions of environment interactions, or imitation of historical farmer decisions, which are known to be suboptimal, we learn from a theoretically optimal oracle.

A backward induction DP solver was implemented with perfect foresight of the full growing season's weather. The continuous depletion state was discretised into 200 bins, and the optimal action for each state–time pair was computed recursively as:

$$V_t(s) = \max_{a \in A}[R(s,a) + \gamma \cdot V_{t+1}(s')]$$

with $\gamma = 0.99$, and s' determined by the water balance dynamics. The solver generated over 50,000 optimal state–action pairs by processing 12 training locations × 13 years × 4 crops.

Because no system can predict an entire season's weather in advance, this oracle cannot itself be deployed. It does, however, provide the best possible teacher. The question this study addresses is therefore whether a neural network can learn the principles underlying these optimal decisions using only information available in real time.

3) Behavioral Cloning

The policy network was trained by supervised learning to imitate the oracle's actions given the same state observations. The architecture is as follows:

- Input layer: 18 neurons (state dimensions)
- Hidden layer 1: 256 neurons, ReLU activation
- Hidden layer 2: 256 neurons, ReLU activation
- Hidden layer 3: 128 neurons, ReLU activation
- Output layer: 11 neurons, softmax (action probabilities)

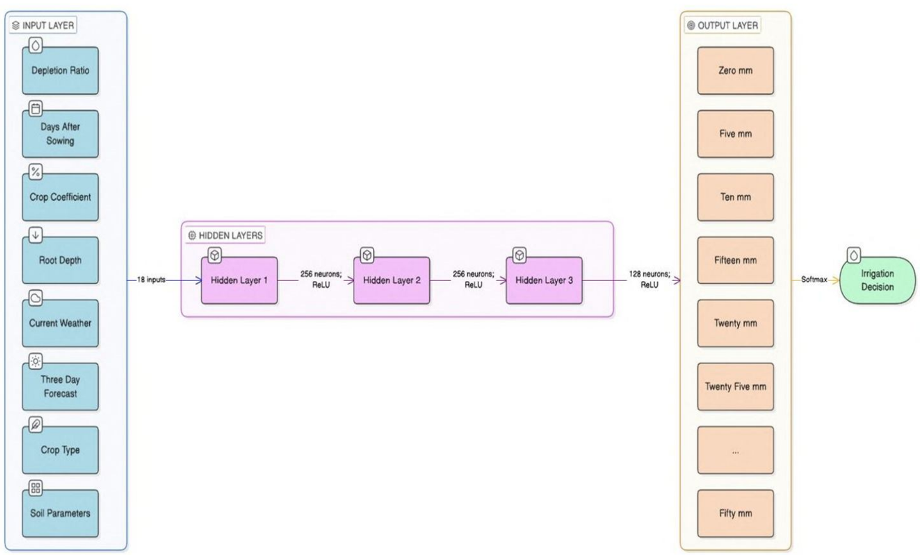

Fig. 2 Policy network architecture. Eighteen state features are mapped through three fully connected ReLU layers to a softmax distribution over eleven discrete irrigation depths.

The network has approximately 105,000 trainable parameters. Training used the Adam optimiser (learning rate 0.001), a batch size of 256 and cross-entropy loss, with early stopping on a 10% validation holdout over 50 epochs. Convergence was reached in approximately 45 minutes on a standard laptop CPU.

Critically, the network is trained on exactly the same 18-dimensional state vector that it will receive at test time. This vector includes the three-day weather forecast but not the oracle's full-season foresight, which compels the network to infer the oracle's decision from limited information. This is the mechanism by which perfect-foresight knowledge is distilled into a real-time, deployable policy.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

E. Evaluation Protocol

Each irrigation strategy was evaluated through complete growing season simulations (110–165 days, varying by crop) across all 15 study locations and across multiple years, so as to capture inter-annual variability. The performance metrics were:

- Relative yield: $Y_a/Y_m \times 100$ (%)
- Total irrigation applied (mm season^{-1})
- Water savings: percentage reduction relative to each baseline
- Water efficiency: millimetres of irrigation per 10% of maximum yield achieved
- Stress days: number of days on which D_r exceeded RAW

All simulations and analyses were conducted in Python 3.11, principally using SciPy and NumPy.

F. Implementation and Deployment

The trained model occupies 420 KB in ONNX format and executes inference in under 5 ms on a mid-range smartphone. Internet connectivity is required only to retrieve weather forecasts from the free and open-access Open-Meteo API. No soil sensors, subscription or specialised equipment are needed. The system issues daily irrigation recommendations, specifying whether to irrigate and, if so, by how much, on the basis of the farmer's report of the most recent irrigation amount and how long ago it was applied, from which it internally reconstructs the current soil moisture state. Farmers may optionally supply direct soil moisture readings to refine the recommendations, but these are not required.

III. RESULTS

A. Yield and Water Performance

Performance summarised across all four crops and all 15 locations is presented in Table III and Figs. 3 and 4.

TABLE III

PERFORMANCE COMPARISON ACROSS IRRIGATION METHODS (MEAN VALUES ACROSS ALL LOCATIONS)

Method	Yield (% of max)	Irrigation (mm)	Water savings vs. high-yield	Efficiency (mm per 10% yield)
BC (Ours)	98.5	147	61%	14.9
Physics-Based	99.7	215	43%	21.6
ICAR Schedule	88.5	208	45%	23.5
High-Yield Farmer	99.9	377	0%	37.7

Efficiency is computed as total irrigation divided by (relative yield / 10). Lower values indicate greater water productivity.

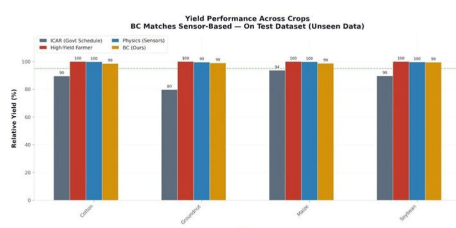

Fig. 3 Relative yield by crop on the held-out test dataset. The behavioral cloning policy matches sensor-based scheduling to within roughly one percentage point across all four crops.

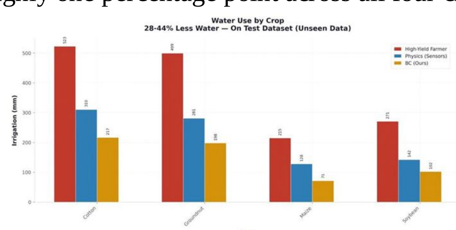

Fig. 4 Seasonal irrigation applied by crop on the held-out test dataset. Water savings relative to the physics-based controller range from 28% to 44% depending on crop.

On the held-out test data, per-crop results were as follows. Cotton reached 99% relative yield with 217 mm of irrigation, compared with 310 mm for the physics-based controller (30% saving) and 523 mm for high-yield practice (59% saving). Groundnut reached 99% with 198 mm, against 281 mm (30%) and 499 mm (60%). Maize reached 99% with 71 mm, against 128 mm (44%) and 215 mm (67%). Soybean reached 99% with 102 mm, against 142 mm (28%) and 271 mm (62%).

The behavioral cloning policy consistently outperformed the ICAR calendar schedule, which achieved only 80–94% relative yield across the four crops because its fixed schedule is unresponsive to actual weather. The resulting yield penalty of 6–20% represents a direct and recurring income loss for smallholder growers.

B. Geographic Generalization

The policy remained effective at the three held-out locations (Dharwad, Guntur and Hyderabad), which present combinations of soil type, climate and weather never encountered during training (Table IV).

TABLE IV
PERFORMANCE AT HELD-OUT TEST LOCATIONS (BC SYSTEM, 2023–2025)

Location	State	Yield (%)	Water savings vs. high-yield (%)
Dharwad	Karnataka	98	66
Hyderabad	Telangana	99	75
Guntur	Andhra Pradesh	98	84

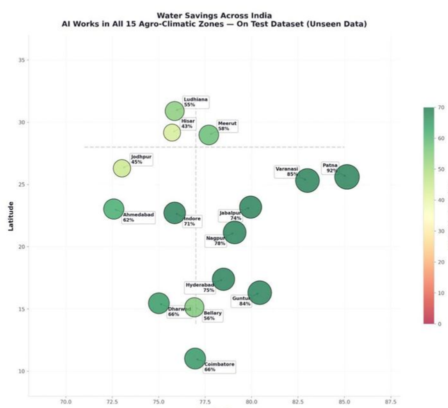

Fig. 5 Water savings relative to high-yield farmer practice at all 15 study locations. Savings are positive across every agroclimatic zone represented.

Across all 15 locations, water savings relative to high-yield practice ranged from 43% at Hisar to 92% at Patna, with a mean of 61% (Fig. 5). Relative yield at the three held-out locations (98–99%) was indistinguishable from that at the 12 training locations, indicating that performance does not degrade on unseen sites. This confirms that the network learned generalisable principles of crop and water management rather than memorising location-specific patterns.

The policy also adapts to inter-annual climate variability. It maintained consistent yield across the full range of monsoon years, whether drought or wet, by reducing irrigation in wet years and increasing it in dry years, delivering a mean 61% water saving relative to high-yield practice across all levels of rainfall variability (Fig. 6).

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IX Sep 2026- Available at www.ijraset.com

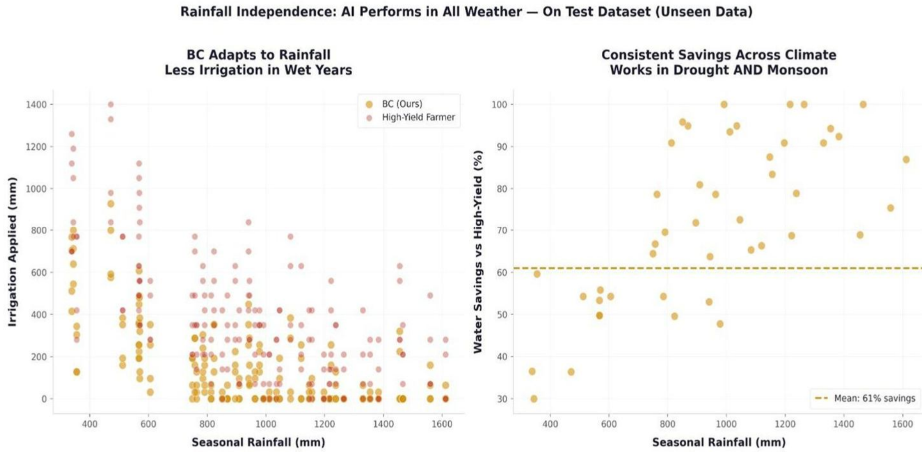

Fig. 6 Rainfall independence. Left: the policy reduces irrigation as seasonal rainfall increases. Right: water savings relative to high-yield practice remain consistent across drought and monsoon years.

C. *Behavioral Analysis and Efficiency*

Examining day-by-day irrigation schedules reveals why the behavioral cloning system is more efficient (Fig. 7).

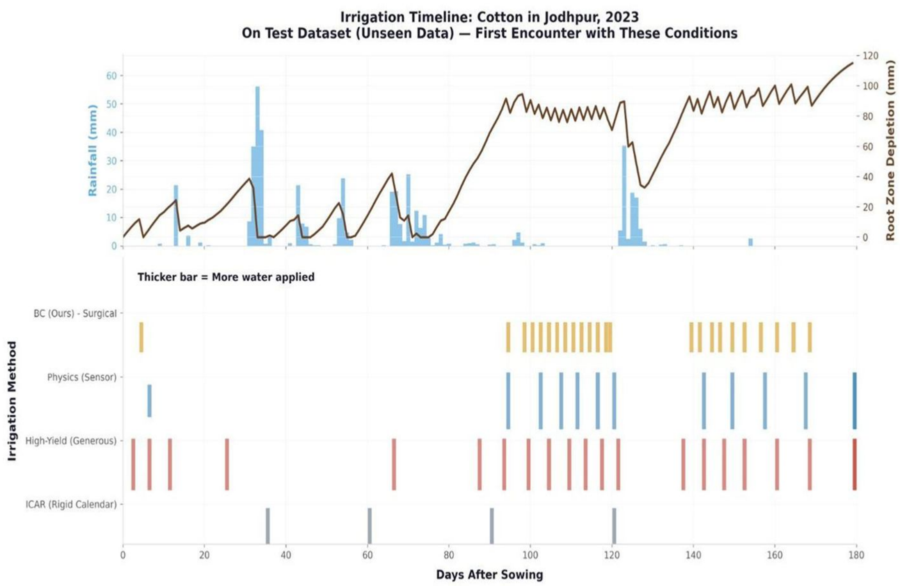

Fig. 7 Irrigation timeline for cotton at Jodhpur, 2023, a location-year never seen during training. Bar thickness is proportional to the depth applied.

For cotton at Jodhpur in 2023, the high-yield farmer baseline applied 17 large irrigations of 70–80 mm each, including applications immediately before and during rainfall events, totalling approximately 1,310 mm, much of which was lost to percolation. The ICAR schedule applied only three irrigations of 60 mm, incurring a yield penalty of about 10%. The physics-based threshold applied nine irrigations of 65–80 mm, which was effective but purely reactive.

The behavioral cloning system displayed a distinctly different signature: applications of approximately 25–30 mm each, concentrated when cotton's yield sensitivity to water stress is highest, during flowering and boll development ($k_\gamma = 1.15$). It strategically accepted moderate stress (5–15 stress days) during the vegetative stage, where k_γ is lower, thereby conserving water for the periods that matter most for yield. In this location-year it achieved 98% relative yield with 490 mm of water, 63% less than the high-yield approach at the same site (Fig. 8).

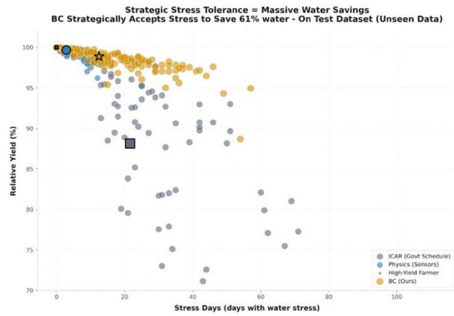

Fig. 8 Strategic stress tolerance. The behavioral cloning policy accepts moderate stress days at stage-insensitive periods, sustaining high relative yield at far lower water use.

This advantage is quantified by the water efficiency ladder (Fig. 9), expressed as millimetres of irrigation required per 10% of maximum yield:

- Behavioral cloning: 14.9 mm, which is 31% more efficient than the physics-based controller and 60% more efficient than high-yield practice
- Physics-based threshold: 21.6 mm
- ICAR schedule: 23.5 mm, despite achieving a lower yield
- High-yield farmer practice: 37.7 mm

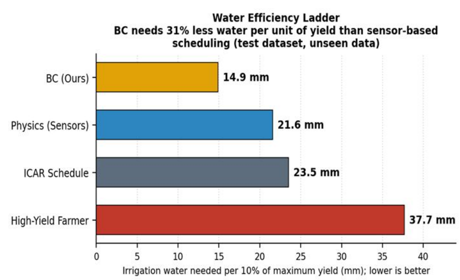

Fig. 9 Water efficiency ladder. Irrigation required per 10% of maximum yield achieved; lower is better. The ICAR schedule is inefficient despite its low total water use because its yield is also low.

The efficiency advantage derives from three learned behaviours: (a) minimising deep percolation by applying water at depths matched to soil intake rates; (b) timing applications to the growth stages with the greatest yield response to water; and (c) exploiting the rainfall forecast to defer irrigation when rain is expected. Each of these behaviours was acquired solely from oracle demonstration, and together they represent a level of agronomic reasoning that a purely sensor-based irrigation system cannot reach.

IV. DISCUSSION

A. *Mechanism of Water Saving*

Despite having only a three-day forecast, against the oracle's full-season foresight, the behavioral cloning system achieves near-identical performance. The key insight this reveals is that optimal irrigation decisions are governed largely by crop state and near-term weather rather than by long-range prediction. The 18-dimensional state vector captures the essential physics (current soil moisture, crop developmental stage, short-term forecast and atmospheric demand), and the network learns a mapping from these features to actions that closely reproduces perfect-foresight decisions.

There is a sharp contrast between the frequent, small, precisely timed behavioural signature of the learned policy and that of all three baselines. The analogy is intravenous hydration versus drinking large glasses of water at fixed intervals: matching supply to uptake rate minimises waste. The system's strategic acceptance of controlled stress during stress-tolerant vegetative stages, while protecting flowering and grain filling, demonstrates an operational understanding of crop physiology as encoded in the stage-specific k_γ values of the Stewart yield model. Agronomists have long recognised this differential sensitivity; the behavioral cloning policy learns to exploit it systematically, which is rarely achieved in practice.

B. *Implications for Water Security*

India's groundwater levels are declining at a rate that demands scalable intervention. Roughly 90% of the electricity supplied to Indian agriculture is consumed by groundwater pumping [14], and aquifer levels are falling by 0.3–1.0 m annually across the major agricultural states [2]. A system delivering an average 60% improvement in water-use efficiency could materially slow aquifer depletion even at partial adoption.

The social equity dimension is of equal importance. The smallholders who work 86% of India's landholdings [3] are currently excluded from precision agriculture purely by the cost of the required tools. A smartphone-based system with zero hardware dependence removes that barrier. India had 1.31 billion telephone subscribers at the end of 2025, of which 544 million were rural [15], so precision irrigation intelligence can be delivered through a device that many farmers already possess.

The system also offers climate resilience. As monsoon variability increases, rigid calendar schedules become progressively less reliable. The behavioral cloning system self-adjusts to seasonal conditions, increasing irrigation in drought years and reducing it in wet years. This is precisely the adaptive capacity that the Indian agricultural sector requires.

C. *Limitations of the Study and Future Directions*

Several limitations must be acknowledged. First, and most significantly, the study was conducted entirely in simulation. Although the FAO-56 water balance and the Stewart yield model are internationally validated, both make simplifying assumptions that require validation in the field. Factors that were not modelled, such as soil spatial heterogeneity, non-uniformity of irrigation application, and pest and disease pressure, may affect field performance. We are currently developing partnerships with farmer cooperatives in Gujarat and Maharashtra in order to conduct real-time trials.

Second, the simulation assumes uniform soil properties within each field, whereas real fields vary spatially. Incorporating field-scale soil mapping and uncertainty quantification would benefit future versions of the system.

Third, weather forecast accuracy degrades beyond three days. Although the study incorporated realistic forecast error, systematic errors differ geographically and may therefore affect accuracy elsewhere. Probabilistic decision making using ensemble forecasts is a natural extension of this work.

Fourth, validation covered four crops; extension to the rice–wheat system is a priority, as it dominates the Indo-Gangetic Plain and is among the most water-intensive systems in Indian agriculture. This will require modelling flood irrigation conditions. Post-deployment considerations include farmer trust, an interface with voice support for users with limited literacy, and integration with existing agricultural extension programmes. Over 250 million Soil Health Cards had been distributed to Indian farmers as of July 2025 [16], which provides a ready distribution channel.

Notwithstanding these limitations, the core finding remains robust: irrigation policies distilled from optimal control demonstrations generalise across crops, climates and geographies while delivering substantial water savings. Moreover, the methodology of using an oracle to create a deployable policy may prove applicable well beyond irrigation, to resource management problems more generally.

V. CONCLUSIONS

India must feed a growing population of 1.4 billion while its groundwater reserves deplete and the monsoon grows more erratic. Smallholder farmers, who work over 86% of India's landholdings, are compelled to choose between conserving water and protecting yield. Accurate precision agriculture exists but remains capital-intensive, while traditional calendar-based scheduling irrigates on predetermined dates regardless of actual conditions and wastes large volumes of water.

This paper has demonstrated how artificial intelligence can resolve that trade-off. By training a neural network to imitate a theoretically optimal irrigation controller, we developed a system that achieves 97–99% of the maximum attainable yield while reducing water consumption by 28–44% relative to sensor-based scheduling and by 59–67% relative to prevailing farmer practice. The only requirements for a farmer are an ordinary smartphone and access to free public weather forecasts.

Three properties establish the scalability of the approach. First, performance generalises to entirely unseen locations, indicating that the network learned universal irrigation principles rather than localised patterns. Second, the system self-adjusts to seasonal variation, conferring substantial resilience to climate change. Third, the model is small (approximately 420 KB) and runs instantly on basic smartphones, allowing immediate deployment with no additional infrastructure cost.

This approach could give approximately 100 million smallholders access to the benefits of precision irrigation without subscription, hardware or technical prerequisites. As a growing number of regions in India face uncertainty over water availability, an AI-based irrigation system offers a tangible route for small farmers to use their land productively and secure abundant harvests through crop intelligence rather than uncontrolled water use.

VI. ACKNOWLEDGMENT

We thank our mentors and advisors, whose guidance in both artificial intelligence and agricultural systems enabled us to combine these fields in this research. We thank Open-Meteo for the provision of comprehensive historical weather data, and ISRIC (World Soil Information) for access to detailed soil datasets, both of which were foundational to our analysis. We are grateful to Adani International School and Shanti Asiatic School for fostering environments that encourage innovation in sustainable technologies, and to our families for their unfailing support throughout this work.

REFERENCES

[1] Food and Agriculture Organization, “AQUASTAT database: Water use in agriculture,” FAO, Rome, 2023. [Online]. Available: https://www.fao.org/aquastat/en/

[2] Central Ground Water Board, “National compilation on dynamic ground water resources of India 2022,” Ministry of Jal Shakti, Government of India, New Delhi, 2022. [Online]. Available: https://cgwb.gov.in/

[3] Ministry of Agriculture and Farmers Welfare, “Agricultural statistics at a glance 2021,” Government of India, New Delhi, 2021. [Online]. Available: https://agricoop.gov.in/

[4] Indian Council of Agricultural Research, “Crop production guidelines for major crops in India,” ICAR, Ministry of Agriculture and Farmers Welfare, Government of India, New Delhi, 2022.

[5] D. A. Pomerleau, “Efficient training of artificial neural networks for autonomous navigation,” Neural Computation, vol. 3, no. 1, pp. 88–97, 1991, doi: 10.1162/neco.1991.3.1.88.

[6] Open-Meteo, “Open-Meteo API documentation: Historical weather and forecast API,” 2024. [Online]. Available: https://open-meteo.com/

[7] R. G. Allen, L. S. Pereira, D. Raes, and M. Smith, Crop Evapotranspiration: Guidelines for Computing Crop Water Requirements, FAO Irrigation and Drainage Paper 56. Rome, Italy: Food and Agriculture Organization of the United Nations, 1998. [Online]. Available: https://www.fao.org/3/x0490e/x0490e00.htm

[8] L. Poggio, L. M. de Sousa, N. H. Batjes, G. B. M. Heuvelink, B. Kempen, E. Ribeiro, and D. Rossiter, “SoilGrids 2.0: Producing soil information for the globe with quantified spatial uncertainty,” SOIL, vol. 7, no. 1, pp. 217–240, 2021, doi: 10.5194/soil-7-217-2021.

[9] K. E. Saxton and W. J. Rawls, “Soil water characteristic estimates by texture and organic matter for hydrologic solutions,” Soil Science Society of America Journal, vol. 70, no. 5, pp. 1569–1578, 2006, doi: 10.2136/sssaj2005.0117.

[10] J. Doorenbos and A. H. Kassam, Yield Response to Water, FAO Irrigation and Drainage Paper 33. Rome, Italy: Food and Agriculture Organization of the United Nations, 1979.

[11] J. I. Stewart, R. M. Hagan, W. O. Pruitt, R. E. Danielson, W. T. Franklin, R. J. Hanks, J. P. Riley, and E. B. Jackson, “Optimizing crop production through control of water and salinity levels in the soil,” Utah Water Research Laboratory, Utah State Univ., Logan, UT, USA, Reports Paper 67, 1977. [Online]. Available: https://digitalcommons.usu.edu/water_rep/67

[12] R. S. Sutton and A. G. Barto, Reinforcement Learning: An Introduction, 2nd ed. Cambridge, MA, USA: MIT Press, 2018. [Online]. Available: http://incompleteideas.net/book/the-book-2nd.html

[13] Independent Evaluation Group, “Addressing groundwater depletion: Lessons from India, the world’s largest user of groundwater,” World Bank Group, Washington, DC, USA, 2021. [Online]. Available: https://ieg.worldbankgroup.org/blog/addressing-groundwater-depletion-lessons-india-worlds-largest-user-groundwater

[14] A. Rajan and K. Ghosh, “Energy productivity of Indian agriculture: Are energy guzzling districts generating higher agricultural value?” in Proc. 3rd World Irrigation Forum (WIF3), Bali, Indonesia, Sep. 2019, Paper ST-1.1. [Online]. Available: https://www.icid.org/wif3_bali_2019/wif3_1-1_7-min.pdf

[15] Telecom Regulatory Authority of India, “The Indian telecom services performance indicators, October–December 2025,” TRAI, New Delhi, India, Mar. 2026. [Online]. Available: https://trai.gov.in/

[16] Ministry of Agriculture and Farmers Welfare, “Soil Health Card: Swasth Dharaa, Khet Haraa,” Press Information Bureau, Government of India, New Delhi, Aug. 2025. [Online]. Available: https://www.pib.gov.in/