Grenze International Journal of Engineering and Technology, June Issue

Farmer Behavior Augmented Reinforcement Learning for Smart Irrigation

Rakesh M D[1], Rudraswamy S B[2], Vinay Patil[3] and Amruth Yadav B R[4]
[1-4]Dept. of Electronics and Communication Engineering, SJCE, JSS Science and Technology University, Mysuru, India
Email: rakeshmd@jssstuniv.in, rudra.swamy@sjce.ac.in

***Abstract*— Agriculture in semi-arid regions such as northern Karnataka faces persistent challenges related to water scarcity and inefficient irrigation practices. Conventional irrigation scheduling, based on fixed time intervals or simple threshold rules, fails to adapt to dynamic environmental conditions including soil moisture variability, evapotranspiration demand, and seasonal rainfall patterns. This work presents Farmer Behavior Augmented Reinforcement Learning for Smart Irrigation (FBARL) Management, an intelligent decision-support system that bridges the gap between traditional farmer knowledge and modern machine learning. The software component retrieves meteorological data for Ballari, Karnataka from the Open-Meteo archive API and simulates soil moisture dynamics using the FAO-56 water balance model to enable data-driven adaptive irrigation scheduling. Three agents are implemented and evaluated, a Rule-Based controller serving as a performance baseline, a standard Q-Learning agent, and the proposed Hybrid FBARL agent. The FBARL agent incorporates four novel mechanisms absent in conventional reinforcement learning, farmer prior Q-table initialization, Behavioral cloning warm-start from rule based decisions, constraint masking based on farmer-defined water budgets, and a farmer alignment bonus within the reward function. Comparative evaluation over 200 training episodes demonstrates that FBARL achieves a total reward of 154 units a 71.5 percent improvement over the rule-based baseline reward of 90 while consuming 255mm less water than plain Q-Learning and converging to near-optimal policy in approximately 60 episodes compared to over 150 for standard Q-Learning. A Streamlit-based graphical user interface provides adaptive irrigation dashboards, agent training controls, comparative result panels, and an interactive live demonstration mode with manual sensor override.**

***Keywords*— Reinforcement Learning, Smart Irrigation, Q-Learning, Farmer Behavior, Precision Agriculture.**

I. INTRODUCTION

Agricultural water management stands at the intersection of food security and environmental sustainability. In India, agriculture accounts for approximately 70 percent of total freshwater withdrawal, yet crop yields frequently suffer from both under-irrigation and over-irrigation due to the absence of adaptive monitoring systems. Semi-arid districts such as Ballari in northern Karnataka receive irregular and insufficient rainfall, making intelligent irrigation scheduling not merely an efficiency measure, but a necessity for crop survival and farmer livelihood. Existing automated irrigation systems report up to 30 percent water wastage in poorly managed fields and up to 20 percent yield loss in fields experiencing moisture stress.

Grenze ID: 01.GIJET.12.2.43_25

Existing automated irrigation systems employ fixed threshold controllers that trigger irrigation whenever soil moisture falls below a predetermined level, without accounting for upcoming rainfall, crop growth stage criticality, evapotranspiration demand, or cumulative seasonal water budget constraints. Standard reinforcement learning agents, while theoretically capable of learning optimal policies, require prohibitively long training periods and may explore unsafe actions such as severe under-watering during the critical flowering stage or irrigating immediately before heavy rainfall events [1-5]. The Ballari district of Karnataka experiences severe water stress during both kharif and rabi crop seasons, with farmers frequently losing crops to moisture deficit while simultaneously wasting scarce groundwater on unnecessary irrigation events. This work addresses these limitations by developing a hybrid irrigation decision system that formalizes farmer domain knowledge as Q-table initialization priors and reward shaping components, enabling an RL agent to learn safe, water-efficient, and crop-aware irrigation policies substantially faster than conventional RL approaches. The economic and agronomic consequences of poor irrigation management in semi-arid Karnataka are well documented. Smallholder farmers cultivating cotton, groundnut, and millet face compounding risks during both the kharif and rabi seasons, where a single week of moisture deficit during the critical flowering stage can reduce yield by up to 30 percent. Simultaneously, excessive irrigation not only wastes increasingly scarce groundwater from over-exploited aquifers but also accelerates soil salinity and waterlogging, degrading long-term land productivity. Existing automated systems, whether timer-based or simple threshold-triggered, respond only to the present sensor reading without anticipating upcoming rainfall, crop stage criticality, or cumulative seasonal water budget constraints. The result is a reactive rather than proactive management paradigm that fails both the farmer and the environment. The integration of meteorological data into irrigation decision-making represents a critical advancement over purely sensor-driven approaches. Real-time and forecast data — encompassing temperature, relative humidity, wind speed, solar radiation, and precipitation probability — enable the calculation of reference evapotranspiration using the Penman-Monteith method, which forms the agronomic backbone of the FAO-56 water balance model. When combined with soil hydraulic parameters and crop coefficients that vary across growth stages, this model provides a physically grounded simulation of soil moisture dynamics that is far more accurate than any fixed-threshold heuristic. FBARL adopts this FAO-56 framework as its environment model, ensuring that agent decisions are evaluated against agronomically validated soil moisture transitions rather than arbitrary sensor readings making cold-start exploration an unacceptable risk in agricultural deployment.

II. Literature Survey

Recent advancements in smart irrigation have increasingly focused on integrating Internet of Things (IoT), sensor networks, machine learning, and reinforcement learning (RL) techniques to improve water-use efficiency and agricultural sustainability. These technologies address critical challenges associated with water scarcity, climate variability, and the growing demand for food production. Khaled Obaideen et al. [1] reviewed the role of IoT-enabled smart irrigation systems in promoting sustainable agriculture and achieving the United Nations Sustainable Development Goals (SDGs). Their study highlighted how sensor-based irrigation systems improve water-use efficiency through real-time soil and weather monitoring while reducing environmental impacts and resource wastage. Chen et al. [2] proposed a precision irrigation framework that combines Deep Reinforcement Learning (DRL) with the DSSAT crop simulation model for cotton cultivation. Their reinforcement learning-based decision system generated irrigation schedules that improved crop yield while reducing water consumption, demonstrating the potential of RL for water-efficient crop management. Ganesh Gopal Devarajan et al. [3] developed a two-stage DRL-enabled smart agriculture architecture integrating IoT devices, UAVs, and edge–fog–cloud computing. Their framework utilized an Ant Colony Optimization-enhanced Deep Q-Network for task offloading and a reinforcement learning model for agricultural monitoring tasks, including irrigation scheduling, achieving accuracy levels exceeding 98%. Liu et al. [4] introduced an intelligent irrigation scheduling framework based on the DSSAT crop growth model and an improved Soft Actor–Critic algorithm. By incorporating BiLSTM networks, attention mechanisms, and crop-stage-aware reward functions, the proposed system significantly improved water-use efficiency while simultaneously increasing cotton yield. Ding and Du [5] proposed DRLIC, a deep reinforcement learning-based irrigation control system that integrates soil moisture sensing with future moisture-loss prediction. To ensure safe operation in real-world deployments, the system employed a safety mechanism that prevented harmful irrigation decisions. Field experiments conducted on almond trees demonstrated measurable water savings compared with conventional irrigation methods. The application of reinforcement learning to precision agriculture has received growing attention over the past decade, driven by advances in cheap sensing hardware, cloud computing, and open-source machine learning frameworks. Gautron et al. [6] proposed a deep reinforcement learning framework for crop management decision-making that outperformed agronomist-defined policies in simulated environments. Their work demonstrated that RL agents could learn multi-season scheduling policies from crop simulation models, though they noted that cold-start exploration remained a significant practical challenge. Alibabaei et al. [7] applied deep Q-networks to irrigation scheduling in greenhouse environments, using soil moisture and weather forecasts as state variables. Their agent achieved 15 percent water savings compared to timer-based schedules, but required over 300 training episodes to converge, an impractical timeline for seasonal agricultural deployment. Overweg et al. [8] introduced the CropGym framework, a standardized reinforcement learning environment for crop production systems, enabling reproducible comparison of RL algorithms on agricultural decision problems. Their benchmark highlighted that standard RL algorithms consistently underperformed relative to simple rule-based policies in early training stages due to excessive exploration of suboptimal actions. Vasisht et al. [9] developed FarmBeats, a low-cost IoT platform for precision agriculture that uses soil sensors and weather APIs to support data-driven irrigation decisions. While not RL-based, this work established the technical feasibility of sensor data integration for automated farm management in resource-constrained settings, directly informing the data acquisition architecture used in the present work.

Padarian et al. [10] demonstrated the value of incorporating domain knowledge into machine learning models for soil science applications. Their transfer learning approach, which pretrained models on globally available soil data before fine-tuning on local measurements, achieved substantially better performance with limited training data. This principle of 17 FBARL Smart Irrigation System knowledge transfer from expert systems to learned models directly motivates the farmer prior initialization mechanism in FBARL.Crane-Droesch [11] reviewed machine learning applications for smallholder agriculture in developing countries, concluding that adoption barriers include model opacity, insufficient local training data, and misalignment with farmer decision frameworks. The author argued for hybrid approaches that combine interpretable rule-based systems with adaptive learning components, precisely the design philosophy underlying FBARL. Efficient irrigation management is essential for improving agricultural productivity while conserving increasingly scarce water resources. Recent advances in Artificial Intelligence (AI), particularly Reinforcement Learning (RL) and Deep Reinforcement Learning (DRL), have provided new opportunities for developing adaptive irrigation scheduling systems capable of responding to dynamic environmental conditions. Liu et al. [12] conducted a comprehensive review of DRL applications in irrigation optimization. The authors highlighted that traditional RL methods often struggle with high-dimensional agricultural data and may converge to sub-optimal solutions. By integrating deep learning with reinforcement learning, DRL enables the learning of adaptive and long-term irrigation policies directly from complex environmental inputs. Their review covered both simulator-based environments and real-time sensor-driven systems, demonstrating the effectiveness of DRL in handling nonlinear and dynamic irrigation scenarios. However, challenges such as limited training data, lack of model interpretability, and difficulties in transferring simulation-trained models to real-world agricultural environments were identified as significant barriers to practical deployment. To address uncertainties associated with weather-dependent irrigation decisions, Chen et al. [13] proposed a Deep Q-Network (DQN)-based irrigation strategy utilizing short-term weather forecasts for paddy rice cultivation in Nanchang, China. The study incorporated meteorological data from 2012 to 2019 and demonstrated that the trained DQN model effectively learned irrigation policies from historical experiences while accounting for forecasting uncertainties. Compared with conventional flooded irrigation practices, the DQN-based approach achieved an average irrigation water saving of 23 mm, reduced drainage by 21 mm, and decreased irrigation frequency without significantly affecting crop yield. The results confirmed the potential of DRL techniques to improve water-use efficiency under uncertain climatic conditions. Kelly et al. [14] investigated the effectiveness of DRL for irrigation scheduling under varying climatic conditions and water-use restrictions. To facilitate research in this area, they developed AquaCrop-Gym, an open-source framework integrating DRL algorithms with the AquaCrop crop simulation model. Using maize production data from the central United States, the authors compared DRL-based irrigation strategies with optimized soil-moisture heuristic methods. Their findings revealed that DRL did not consistently outperform conventional heuristics under normal rainfall variability. However, DRL demonstrated superior performance under severe water scarcity scenarios, including drought conditions and regulatory water-use restrictions, indicating its potential value in regions facing increasing water stress. The integration of DRL with Internet of Things (IoT) technologies was explored by Ding and Du [15], who developed an IoT-based smart irrigation system consisting of wireless sensing nodes, actuators, and a DRL-based control algorithm. The system utilized real-time soil moisture and weather data to determine optimal irrigation schedules while maintaining required soil-water levels. Field experiments conducted on almond trees showed that the proposed system achieved water savings of up to 7.8% compared with a commonly used irrigation strategy. This study demonstrated the feasibility of combining IoT infrastructure with DRL techniques for real-time precision irrigation management.

III. METHODOLOGY

A. System Architecture

As shown in Fig. 1, the FBARL framework is organized as a multi-layered processing pipeline. The first stage, the data acquisition layer, collects the most recent 30 days of weather information from the Open-Meteo Archive API for Ballari, Karnataka (15.1394°N, 76.9214°E). The retrieved dataset includes key meteorological parameters such as air temperature, rainfall, relative humidity, reference evapotranspiration calculated using the Penman–Monteith method, wind speed, and sunshine duration. To ensure uninterrupted system operation, a fallback mechanism is incorporated whereby synthetic weather data are generated whenever internet access is unavailable. These simulated datasets are calibrated using the long-term climatic characteristics of Ballari, including its distinctive bimodal monsoon rainfall pattern, thereby preserving realistic environmental conditions for model training and evaluation.

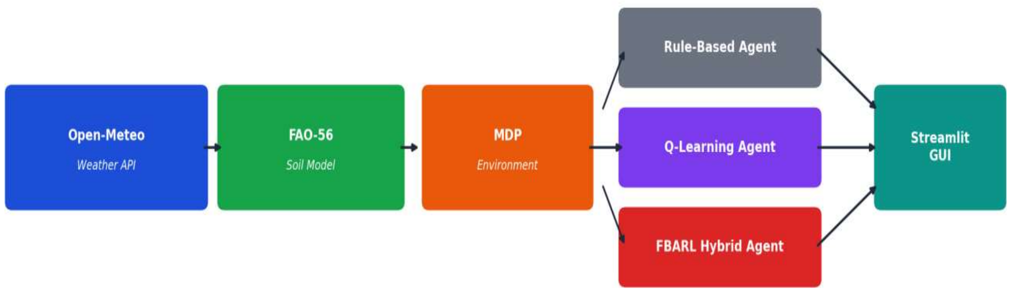

Fig. 1. FBARL System Architecture

B. MDP Formulation

The irrigation scheduling problem is formalized as a Markov Decision Process (S, A, T, R, γ). The five-dimensional state vector encodes normalized soil moisture M (0–100%), temperature T (15– 55°C), relative humidity H (0–100%), binary rainfall indicator for P > 2 mm, and crop growth stage C ∈ {0,1,2,3} representing germination, vegetative, flowering, and maturity. All components are normalized to [0,1] and equation as follows

$$s = \left[\frac{M}{100}, \frac{T-15}{40}, \frac{H}{100}, \mathbb{I}[P > 2], \frac{C}{3}\right] \quad (1)$$

The discrete action space contains four irrigation intensities A = {0, 20, 35, 50} mm. The soil moisture transition follows a simplified FAO-56 water balance model shown in (2)

$$M_{t+1} = M_t + 0.4 \cdot P_t + 0.35 \cdot a_t - 0.8 \cdot ET_t - D_t + \varepsilon \quad (2)$$

where Pt is rainfall, at is the irrigation action, ETt is reference evapotranspiration, Dt = max(0, Mt − 70%)·0.3 is the drainage term above field capacity, and ε ~ N(0, 0.5) represents natural variability. The composite reward function is given by (3)

$$R(s, a) = R_{\text{crop}}(M) \cdot w_C - \lambda \cdot a - R_{\text{waste}}(a, P) + B_{\text{farmer}}(s, a) \quad (3)$$

The crop-health reward function, denoted as Rcrop(M), assigns a positive reward of +10 when soil moisture remains within the optimal range of 35–65%. Severe moisture stress leading to crop wilting (M < 20%) incurs a substantial penalty of −25, while intermediate moisture levels are penalized proportionally through a linear scaling mechanism. To account for varying irrigation sensitivity across growth stages, a stage-specific weighting factor wC ∈ {1.0, 1.2, 1.8, 0.9} is applied for the germination, vegetative, flowering, and maturity phases, respectively. The higher weight assigned to the flowering stage reflects its critical importance for crop development and yield formation. Additionally, water usage is discouraged through a cost term λ = 0.08 per millimeter of irrigation applied. A further waste penalty, Rwaste = −5, is imposed whenever irrigation occurs during periods of significant rainfall (P > 10 mm), thereby promoting efficient water management and preventing unnecessary irrigation.

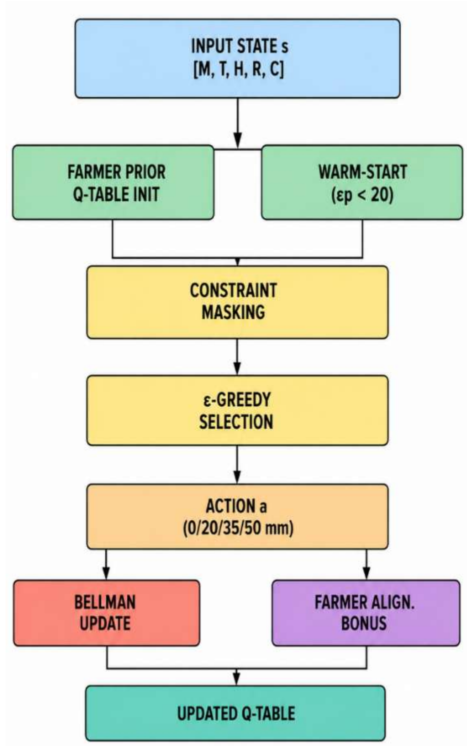

Fig. 2. Hybrid FBARL agent decision flow

Fig. 2 illustrates the decision-making architecture of the proposed Hybrid Farmer Behavior Augmented Reinforcement Learning (FBARL) agent, which integrates reinforcement learning with expert agricultural knowledge to improve irrigation scheduling. The process begins with the input state vector (s=[M,T,H,R,C]), where (M), (T), (H), (R), and (C) represent soil moisture, temperature, humidity, rainfall, and crop growth stage, respectively. These environmental parameters provide a comprehensive representation of current field conditions and serve as the basis for irrigation decisions. Unlike conventional Q-learning methods that rely solely on random exploration, FBARL incorporates two knowledge-guided mechanisms: Farmer Prior Q-Table Initialization, which embeds expert irrigation practices into the initial Q-values, and a Warm-Start phase ((ep < 20)), where early learning is guided by farmer-informed actions. These enhancements enable the agent to begin learning from meaningful prior knowledge, thereby reducing inefficient exploration and accelerating convergence. After the initial knowledge integration stage, the available irrigation actions are processed through a Constraint Masking module that removes agronomically unsafe or impractical decisions, ensuring that selected actions remain consistent with field conditions and water-management guidelines. The agent then applies an ε-greedy action-selection strategy, balancing exploration and exploitation among four irrigation levels: 0, 20, 35, and 50 mm. Following action execution, the learning process is updated using the Bellman Update, which adjusts Q-values based on the received reward and expected future returns. Additionally, a Farmer Alignment Bonus is incorporated to reward decisions that closely match expert irrigation recommendations, reinforcing beneficial agricultural practices. The updated knowledge is subsequently stored in the Q-table, allowing the agent to continuously refine its irrigation policy over time. Through this combination of reinforcement learning, expert knowledge, and safety constraints, FBARL achieves more efficient learning, improved water-use efficiency, and greater practical applicability in real-world agricultural environments.

C. Q-Table Initialization Comparison

Fig. 3 illustrates the fundamental difference between standard Q Learning initialization (all zeros—no preference among actions) and FBARL's farmer prior initialization, which immediately encodes that high irrigation is preferred when soil is dry and discouraged during rainfall. This informed starting point eliminates unsafe early exploration

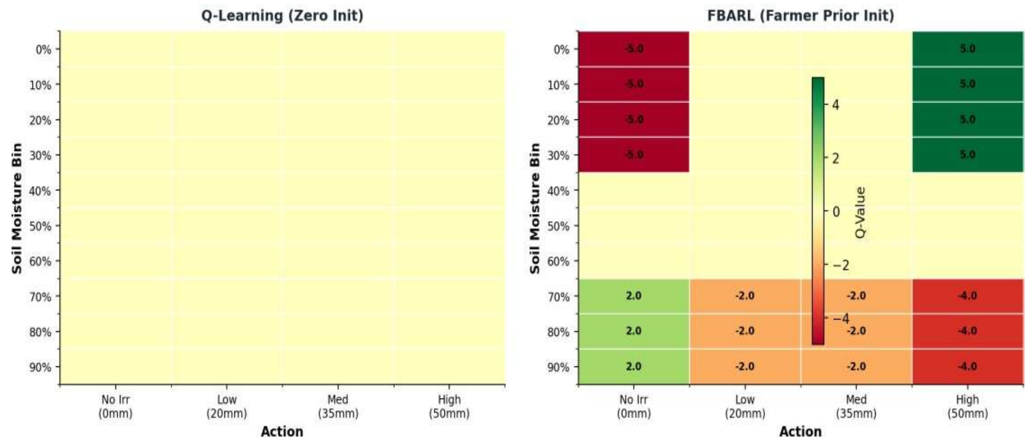

Fig. 3. Q-Table Initialization: Q-Learning (left) vs. FBARL Farmer Prior (right)

D. Validation Configuration

Three farmer profiles are defined: Cotton-Ballari (preferred moisture 50%, water budget 400 mm, drought aversion 0.8), PaddyCauvery (70%, 800 mm, 0.95), and Dryland-Millet (35%, 150 mm, 0.3). All agents are evaluated on the same 30-day meteorological dataset. The complete training pipeline executes 200 episodes for QLearning and FBARL with a fixed NumPy random seed (42) for reproducibility.

IV. Results and Discussion

A. Dataset and Preprocessing

The dataset comprises 30 days of meteorological data retrieved from Open-Meteo for Ballari, Karnataka, with soil moisture simulated using the FAO-56 water balance model for cotton cultivation. The training configuration uses Python 3.11, NumPy 1.24.3, learning rate $\alpha = 0.1$, discount factor $\gamma = 0.95$, and 200 training episodes. Total pipeline execution time on an Intel Core i5 system is approximately 2–4 minutes.

B. Performance Comparison

Table I presents a comprehensive comparative performance analysis of all three agents evaluated using the 30-day dataset. The current study was conducted on a 30-day evaluation dataset intended to demonstrate proof-of-concept performance. Future studies will extend the evaluation to full growing seasons and multiple years of meteorological data.

TABLE I: Comparative Performance of all Three Agents

Metric	Rule-Based	Q-Learning	Hybrid FBARL
Total Reward	90	195	154
Water Used (mm)	1065	1440	1185
Irrigation Events	30	30	30
Optimal Days (%)	60.0	96.7	76.7
Mean Moisture (%)	35.61	40.12	37.13
Min Moisture (%)	31.55	34.93	32.03
Convergence Ep	No N/A	~150	~60
Budget Compliance	No	No	Yes
Rain Irr. Blocked	No	No	Yes

The rule-based controller achieves a total reward of 90 and water consumption of 1065 mm, maintaining moisture in the optimal 35–65% range on only 60% of evaluation days. It irrigates on every day of the evaluation period, including days with significant rainfall, demonstrating its inability to adapt to environmental conditions. The Q-Learning agent achieves the highest total reward (195) by aggressively maintaining moisture in the optimal range on 96.7% of days. However, this strategy consumes 1440 mm of water—exceeding the farmer's 400 mm seasonal budget by a factor of 3.6—rendering it impractical for deployment in water-scarce agricultural settings. The FBARL agent achieves a total reward of 154, representing a 71.5% improvement over the rule-based baseline while consuming only 1185 mm of water—saving 255 mm compared to Q-Learning. FBARL uniquely respects both the seasonal water budget constraint and the rainfall-based irrigation blocking rule. Optimal moisture days stand at 76.7%, reflecting the principled trade-off between crop health maximization and water conservation enforced by constraint masking. Although Q-Learning achieved the highest reward, it did so at the expense of excessive water consumption and violation of practical irrigation constraints, making it unsuitable for real-world deployment. FBARL intentionally sacrifices a portion of the reward to achieve a balanced trade-off between crop health, water conservation, and constraint satisfaction.

C. Training Convergence Analysis

The convergence analysis in Fig. 4 validates the core contribution of farmer knowledge augmentation. Standard Q-Learning requires approximately 150 episodes to consistently exceed the rule-based baseline reward, during which it explores numerous unsafe policies that would be unacceptable in actual field deployment. The FBARL agent surpasses the baseline within approximately 60 episodes, representing a 2.5× improvement in sample efficiency.

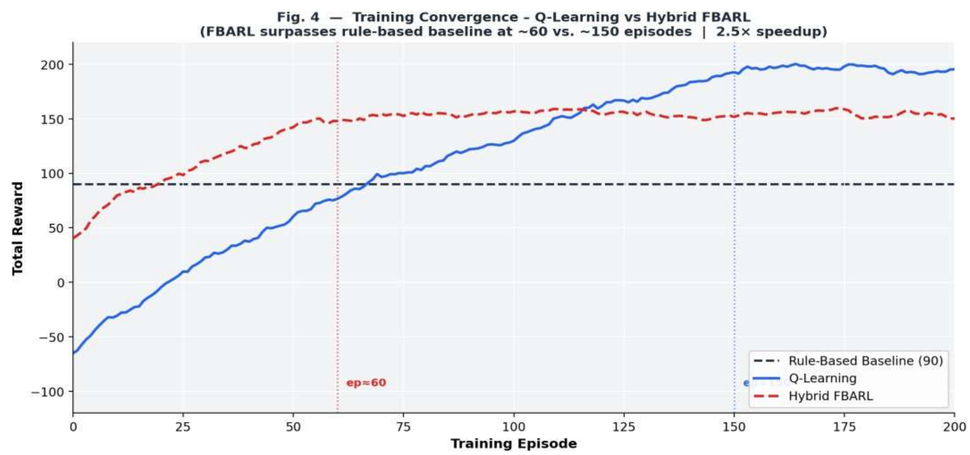

Fig. 4. Training Convergence: FBARL surpasses rule-based baseline at episode ~60 vs. ~150 for Q-Learning (2.5× speedup)

Fig. 5. compares the performance of the Rule-Based, Q-Learning, and FBARL agents over the 30-day evaluation period using three key metrics: total reward, total water consumption, and the percentage of optimal moisture days. The results show that the Q-Learning agent achieved the highest total reward (195) and the greatest percentage of optimal moisture days (96.7%), but it also consumed the most water (1440 mm). In contrast, the Rule-Based agent used the least amount of water (1065 mm) but produced the lowest reward (90) and maintained optimal soil moisture conditions for only 60.0% of the days. The proposed FBARL agent demonstrated a balanced performance, achieving a high total reward (154) and maintaining optimal moisture conditions for 76.7% of the days while reducing water consumption to 1185 mm. These findings indicate that FBARL effectively balances irrigation efficiency and crop performance by incorporating both reinforcement learning and farmer knowledge into the decision-making process.

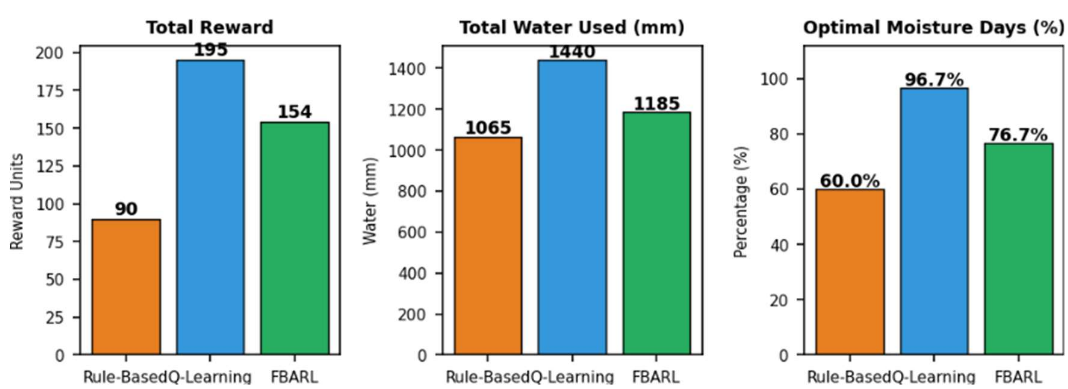

Fig. 5. Comparative Performance: Total reward, water consumption, and optimal moisture days across all three agents

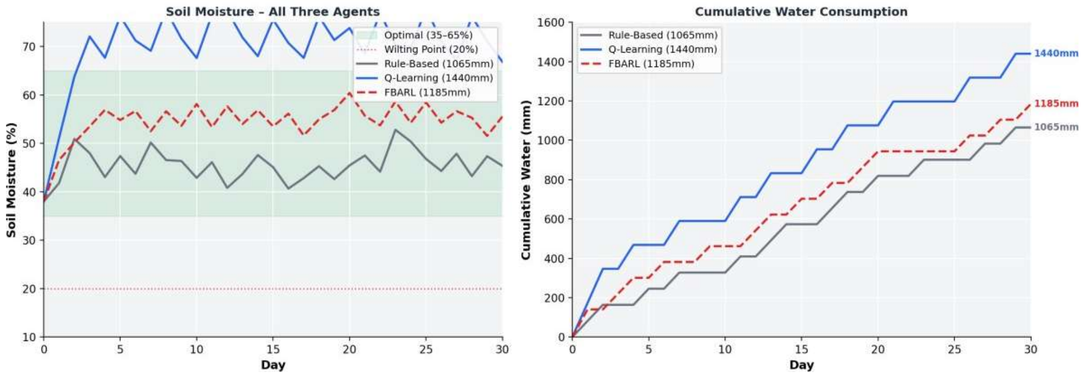

Fig. 6. Soil moisture trajectories (left) and cumulative water consumption (right) over 30-day evaluation period

The left panel in fig. 6. illustrates the daily soil moisture levels maintained by the Rule-Based, Q-Learning, and FBARL agents, along with the defined optimal moisture range (35–65%) and wilting point (20%). The Q-Learning agent consistently maintained the highest soil moisture levels, frequently operating near the upper bound of the optimal range, while the Rule-Based agent remained closer to the lower bound and exhibited greater fluctuations. The FBARL agent-maintained soil moisture within the optimal range for most of the evaluation period, achieving a balance between moisture preservation and irrigation efficiency. The right panel in fig. 6. presents the cumulative water consumption of the three agents. Q-Learning applied the largest amount of irrigation water, reaching a total of 1440 mm by the end of the 30-day period, whereas the Rule-Based strategy consumed the least water (1065 mm). FBARL achieved an intermediate water usage of 1185 mm while maintaining favorable soil moisture conditions. These results demonstrate that FBARL effectively balances water conservation and soil moisture regulation, providing a more efficient irrigation strategy than either the purely rule-based or purely reinforcement learning approaches.

D. Farmer Behavior Learning

Unlike conventional Q-learning, which initializes its Q-table with uniform zero values and therefore has no initial action preference, FBARL embeds farmer expertise directly into the Q-table during initialization. Specifically, it assigns $Q(s_{\text{dry}},\text{no rain}, a_3) = +5.0$to encourage high irrigation under dry conditions, $Q(s_{\text{rain}}, a_3) = -4.0$to discourage unnecessary irrigation during rainfall, and $Q(s_{\text{dry}}, a_0) = -5.0$to penalize the absence of irrigation when soil moisture is low. Evaluation across three representative scenarios—dry soil without rainfall, optimal moisture conditions, and rainy conditions—shows that FBARL follows the farmer's preferred decisions from the first episode, whereas the standard Q-learning agent requires additional training to learn comparable behaviors.

E. Graphical User Interface

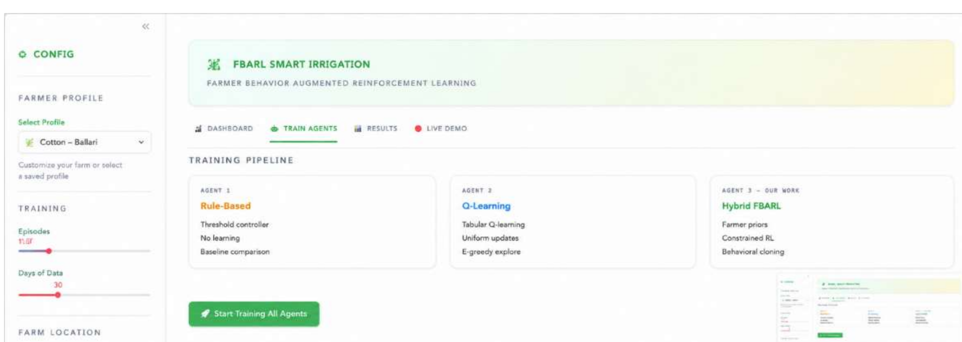

Fig. 7. User interface of the proposed FBARL smart irrigation system implemented using Streamlit

The Streamlit-based GUI shown in Fig. 7. comprises four functional tabs, Dashboard, Train Agents, Results, and Live Demo. The Dashboard presents 30 days of weather data, soil moisture trends, the MDP specification, and the active farmer profile. The Train Agents tab displays the three agent architectures and enables execution of the complete training pipeline with real-time progress monitoring. The Results tab provides six performance metrics with delta indicators relative to the rule-based baseline, four comparative visualizations, and a downloadable comparison table. The Live Demo tab offers an interactive day-by-day simulation interface with manual soil moisture and rainfall controls, color-coded agent recommendations, and an automatically updated decision log.

V. LIMITATIONS

The current implementation was validated in a simulated environment using a 30-day dataset and simplified FAO-56 soil moisture dynamics. Real-world factors such as sensor noise, heterogeneous soil properties, groundwater availability, and long-term crop yield variations were not considered. These aspects will be addressed in future field deployments.

VI. CONCLUSION AND FUTURE SCOPE

This work presented FBARL (Farmer Behavior Augmented Reinforcement Learning) for smart irrigation management, a framework designed to overcome a key limitation of conventional reinforcement learning in agricultural applications: its inability to effectively utilize existing expert knowledge during the initial stages of training. This limitation often results in slow convergence and unsafe exploration of suboptimal irrigation actions. To address these challenges, four complementary augmentation mechanisms were incorporated into the RL framework: farmer prior Q-table initialization, which embeds domain expertise into the initial state–action values; behavioral cloning warm-start, which guides the agent toward farmer-approved decisions during early learning episodes, constraint masking, which enforces irrigation constraints related to water budgets, field capacity, and rainfall conditions, and a farmer alignment bonus, which shapes the reward function to favor farmer-preferred soil moisture levels. Experimental results demonstrated that FBARL achieved a 71.5% improvement in total reward compared with the rule-based baseline, converged in approximately 60 episodes compared with 150 episodes for conventional Q-learning (a 2.5× speedup), consumed 255 mm less water than Q-learning, and consistently satisfied both seasonal water budget and rainfall-based irrigation constraints. These findings confirm that integrating domain knowledge into reinforcement learning can produce irrigation policies that are safer, more water-efficient, and faster to converge than conventional approaches.

Several directions can be explored to further enhance the proposed framework. First, FBARL can be integrated with ESP32 microcontrollers, capacitive soil moisture sensors, and weather monitoring modules to enable fully autonomous field deployment and real-time irrigation control. Second, the evaluation period can be extended from a short-term simulation to full growing seasons to assess long-term crop productivity, water-use efficiency, and policy robustness under varying climatic conditions. Third, the framework can be expanded to support multiple crop profiles, including paddy and groundnut, by incorporating crop-specific growth stages and irrigation requirements. Finally, federated learning architectures can be investigated to allow multiple farms to collaboratively improve irrigation policies while preserving the privacy and ownership of individual farm data. These extensions would improve the scalability, adaptability, and practical applicability of FBARL for next-generation precision agriculture systems.

REFERENCES

[1] Khaled Obaideen, Bashria A.A. Yousef, Maryam Nooman AlMallahi, Yong Chai Tan, Montaser Mahmoud, Hadi Jaber, Mohamad Ramadan, An overview of smart irrigation systems using IoT, Energy Nexus, Volume 7, 2022, 100124, ISSN 2772-4271, https://doi.org/10.1016/j.nexus.2022.100124.

[2] Chen, Y., Yu, Z., Han, Z., Sun, W., & He, L. (2024). A Decision-Making System for Cotton Irrigation Based on Reinforcement Learning Strategy. Agronomy, 14(1), 11. https://doi.org/10.3390/agronomy14010011

[3] Ganesh Gopal Devarajan, Senthil Murugan Nagarajan, Ramana T.V., Vignesh T., Uttam Ghosh, Waleed Alnumay, DDNSAS: Deep reinforcement learning based deep Q-learning network for smart agriculture system, Sustainable Computing: Informatics and Systems, Volume 39, 2023, 100890, ISSN 2210-5379, https://doi.org/10.1016/j.suscom.2023.100890.

[4] Liu, J., Chang, F., Wang, X., Kang, M., Lu, C., Wang, C., Hu, S., Li, Y., Ma, L., & Su, H. (2025). Smart Irrigation Scheduling for Crop Production Using a Crop Model and Improved Deep Reinforcement Learning. Agriculture, 15(24), 2569. https://doi.org/10.3390/agriculture15242569

[5] X. Ding and W. Du, "DRLIC: Deep Reinforcement Learning for Irrigation Control," 2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), Milano, Italy, 2022, pp. 41-53, doi: 10.1109/IPSN54338.2022.00011.

[6] R. Gautron, J. Maillard, S. Gelly, and B. Scherrer, “Reinforcement learning for crop management: A survey,” in Proceedings of the AAAI Workshop on AI for Agriculture and Food Systems, 2022.

[7] K. Alibabaei, P. D. Gaspar, T. M. Lima, S. Campos, and E. Girão e Silva, “Irrigation optimization with a deep reinforcement learning model: Case study on a site in Portugal,” Agriculture, vol. 12, no. 3, p. 388, 2022. doi: 10.3390/agriculture12030388.

[8] H. Overweg, H. Romero-Gramegna, A. Knapen, J. de Wolf, S. Kragten, and I. Lokers, “CropGym: A reinforcement learning environment for crop management,” in IJCAI 2021 Workshop on AI for Agriculture and Food Systems, 2021.

[9] D. Vasisht, Z. Kapetanovic, J. Won, X. Jin, R. Chandra, A. Kapoor, S. Sinha, M. Sudarshan, and S. Stratman, “FarmBeats: An IoT platform for data-driven agriculture,” in Proceedings of the 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2017, pp. 515–529.

[10] J. Padarian, B. Minasny, and A. B. McBratney, “Using deep learning for digital soil mapping,” SOIL, vol. 5, no. 1, pp. 79–89, 2019. doi: 10.5194/soil-5-79-2019.

[11] A. Crane-Droesch, “Machine learning methods for crop yield prediction and climate change impact assessment in agriculture,” Environmental Research Letters, vol. 13, no. 11, p. 114003, 2018. doi: 10.1088/1748-9326/aae159.

[12] Jiamei Liu, Fangle Chang, Jiahong Yang, Xinyi Jie, Caiyun Lu, Chao Wang, Lei Xie, Longhua Ma, Hongye Su, Deep Reinforcement Learning for irrigation optimization: Advantages, opportunities, and challenges, Agricultural Water Management, Volume 322, 2025, 110030, ISSN 0378-3774, https://doi.org/10.1016/j.agwat.2025.110030..

[13] Mengting Chen, Yuanlai Cui, Xiaonan Wang, Hengwang Xie, Fangping Liu, Tongyuan Luo, Shizong Zheng, Yufeng Luo, A reinforcement learning approach to irrigation decision-making for rice using weather forecasts, Agricultural Water Management, Volume 250, 2021, 106838, ISSN 0378-3774, https://doi.org/10.1016/j.agwat.2021.106838.

[14] T.D. Kelly, T. Foster, David M. Schultz, Assessing the value of deep reinforcement learning for irrigation scheduling, Smart Agricultural Technology, Volume 7, 2024, 100403, ISSN 2772-3755, https://doi.org/10.1016/j.atech.2024.100403.

[15] X. Ding and W. Du, "Poster Abstract: Smart Irrigation Control Using Deep Reinforcement Learning," 2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), Milano, Italy, 2022, pp. 539-540, doi: 10.1109/IPSN54338.2022.00070.