AgriScience Journal of Sustainable Agriculture and Agroecology

PRIME OPEN ACCESS

Volume 2, Issue 2

Research Article

Date of Submission: 18 May, 2026
Date of Acceptance: 15 June, 2026
Date of Publication: 30 June, 2026

Reinforcement Learning and Deep Q-Networks for Optimized Irrigation Decision-Making in Dynamic Agricultural Environments

Priyanka Balley* and Kanchan K Doke

Department of Computer Engineering, Bharti Vidyapeeth College of Engineering, University of Mumbai, India

***Corresponding Author:**
Priyanka Balley, Department of Computer Engineering, Bharti Vidyapeeth College of Engineering, University of Mumbai, India.

Citation: Balley, P., Doke, K. K. (2026). Reinforcement Learning and Deep Q-Networks for Optimized Irrigation Decision-Making in Dynamic Agricultural Environments. *AgriSci J Sustain Agric Agroecol, 2*(2), 01-08.

Abstract
Efficient irrigation planning is essential in meeting the rising demand for water resources in agriculture with sustainability. Traditional irrigation systems are usually based on static heuristics or manual decision making that cannot respond to changing dynamic environmental conditions, including variability in soil moisture and weather uncertainties. This is an argument for the development of intelligent adaptive frameworks that would help optimize irrigation strategies in real timestamp scenarios. This paper discusses RL-based irrigation planning with a structured synthetic dataset mimicking realistic agricultural scenarios. The parameters included are soil moisture, crop type, weather forecasts, and irrigation needs for controlled experimentation and reproducibility. A customized RL environment is developed by including normalized soil moisture levels, weather probabilities, and irrigation demands as state representations and irrigation or no irrigation sets as actions. This dynamic irrigation model simulates the real-world situation with reward. It makes use of an agent rewarded for effectiveness of irrigation using a DQN-based feedforward neural network architecture consisting of two layers with 24 neurons, where each has ReLU activation. Training was done through epsilon-greedy to balance exploration and exploitation of Q Values iteratively with the use of the Bellman Process. The stability of training is attained using experience replay and gradient-based optimization with a mean squared error loss function. The performance metrics are cumulative rewards and epsilon decay to ascertain convergence and sets of learning efficacy for the model process. This paper shows that the proposed system, based on RL, greatly improves irrigation decision making and introduces policy stability with scalability levels.

Keywords: Irrigation Planning, Deep Q-Networks, Reinforcement Learning, Soil Moisture Prediction, Sustainable Agriculture, Scenarios

Introduction
The agricultural sector has to grapple constantly with the question of balancing the requirement of food production with meager water resources. The conventional irrigation scheduling methods are static schedules or heuristics that fail to respond to dynamic and changing environmental conditions like the dynamically varying soil moisture levels and nondeterministic weather conditions [1-3]. This method yields ineffective outputs as wastage of water or insufficient irrigation. These factors negatively impact the crop yield and sustainability. Recently, some exciting developments in machine learning are emerging, especially Reinforcement Learning (RL). The RL models are excellent for dynamic decision-making tasks by learning the optimal policies through interaction with the environment. Among them, Deep Q-Networks (DQN) is a strong framework to deal with continuous state spaces and discrete actions very well suited for the task of irrigation optimization. The paper introduces the system based on the application of RL for intelligent decision-making in irrigation systems, making use of a synthetic dataset mimicking the real-world agriculture scenarios. This dataset contains critical parameters like soil moisture, weather forecast probabilities, and crop-specific irrigation

needs, thus providing diverse and realistic training conditions. A customized RL environment is designed to model the irrigation process. Normalized soil moisture levels and weather forecasts are considered in state representations, while actions of the agent have direct impacts on cumulative rewards. Using an advanced DQN agent with sophisticated exploration-exploitation strategies and efficient training methodologies, this work might possibly design a scalable adaptive system for optimizing irrigation schedules. It decreases the usage of water and also increases crop production with it, hence it is one sustainable solution of modern agriculture settings.

Systematic Literature Review

Significant progress in machine learning and data-driven methods has greatly increased the pace of improvement in optimization for agriculture. This chapter will elucidate the progression in greater detail by exploring these developments from chronological reviews of relevant literature sets. Raju and Vijayaraghavan proposed architecture for agricultural decision-making using IoT and machine learning, emphasizing the measurement index of decision optimization. Proposed interpretable machine learning models for evaluating irrigation sustainability by predicting groundwater quality [1-3]. This work underlines the importance of explainability in the decision systems that complement irrigation planning. Similarly, applied various machine learning techniques toward predicting irrigation water-quality indices in mining-dominated areas giving insight into how to control and impact this kind of region [4]. Have improved their study on the irrigation scheduling method by applying a probabilistic approach in machine learning with better results on the basis of higher accuracy for forecasting water requirement [5]. Meanwhile, had discussed the flood susceptibility in the Brahmaputra River Basin with the application of multi-criteria decision-making with machine learning improvement in water resource planning that would eventually help manage the irrigation process indirectly in process [6].

Used the GIS-based Analytical Hierarchy Process and machine learning algorithms in assessing the irrigation water quality in mining-affected agricultural districts [7]. Their results showed the possibility of combining spatial analysis with intelligent systems. Pointed out the use of machine learning in evapotranspiration estimation, showing a way of adapting to various environmental conditions and thus provided an essential basis for accurate water usage for irrigation systems [8]. Estimated the required irrigation water for green beans in arid areas based on machine learning for optimization of resource usage [9]. Meanwhile, showed how big data and machine learning can be useful in assessing groundwater quality for irrigation, offering scalable irrigation decision support in data-rich environments [10]. Optimized tomato irrigation using AI and IoT-integrated deep learning-enabled wireless sensor networks with fuzzy logic, which couples AI with IoT together using crop-specific irrigation strategies [11]. Furthered the advancement in this aspect by using machine learning to make personalized crop forecasts that helped to enhance the level of accuracy of the process of irrigation scheduling [12]. Discussed enrichment of performance and accuracy in precision agriculture by utilizing IoT-based AI systems, where they focus interplay between hardware and software for real-time decision making [13]. Applied multi-criteria decision-making algorithms in mapping the sediment formation potential; which further showed the relevance of machine learning in overcoming broader agricultural issues in process [14,15].

Proposed Model Design Analysis

The proposed model integrates Reinforcement Learning (RL) with a Deep Q-Network (DQN) framework to optimize irrigation decisions under dynamic agricultural conditions. First of all, as per figure 1, the architecture and process of the model are designed carefully so as to balance computational efficiency with decision accuracy, based upon well-founded mathematical principles and algorithms that maximize cumulative rewards while adhering to real-world constraints. The process begins with the synthetic dataset including parameters like soil moisture, the probability of the weather forecast and crop-dependent requirements for irrigation sets. These variables make up a state space, via equation 1,

$$St = \{S, W, I\} \dots (1)$$

At each timestamp for the process. The agent chooses an action a ∊ {0, 1}, where 0 represents no irrigation and 1 denotes irrigation operations. The reward function is designed to balance water conservation with crop yield, defined via equation 2,

$$Rt = max\left(0, \alpha - \beta \cdot \left(\frac{S}{50}\right)\right) \dots (2)$$

Where 'a' represents the ideal irrigation threshold and this penalizes excess water usage sets. This reward structure makes sure that the agent learns to prioritize fields that have lower soil moisture levels. After this, the Q Values are estimated which represent the expected cumulative rewards for state-action pairs. These Q Values are updated iteratively using the Bellman Process via equation 3,

$$Q(St, at) = Rt + \gamma \, maxa' Q(S(t+1), a') \dots (3)$$

Where, γ is the discount factor that prioritizes future rewards for the process. The choice of γ influences the agent's long-term strategy, with higher values emphasizing sustainability over immediate gains. Experience replay is used in order to ensure stable learning. Transitions are stored and sampled randomly during training process. The DQN approximates the Q-function using a feedforward neural network. The weights of the network are updated by using gradient descent, minimizing the mean squared error loss, represented via equation 4,

$$L(\theta) = E\left[\left(Qtarget - Q(St, at;\ \theta)\right)^2\right] \ldots (4)$$

Where, Qtarget is represented via equation 5,

$$Qtarget = Rt + \gamma\, maxa'Q(S(t+1), a';\ \theta -) \ldots (5)$$

Where, θ- are the weights of a target network periodically synchronized with θ sets. The stochastic nature of the weather forecast is addressed by computing the irrigation need by integrating the weather probability over the decision horizon sets.

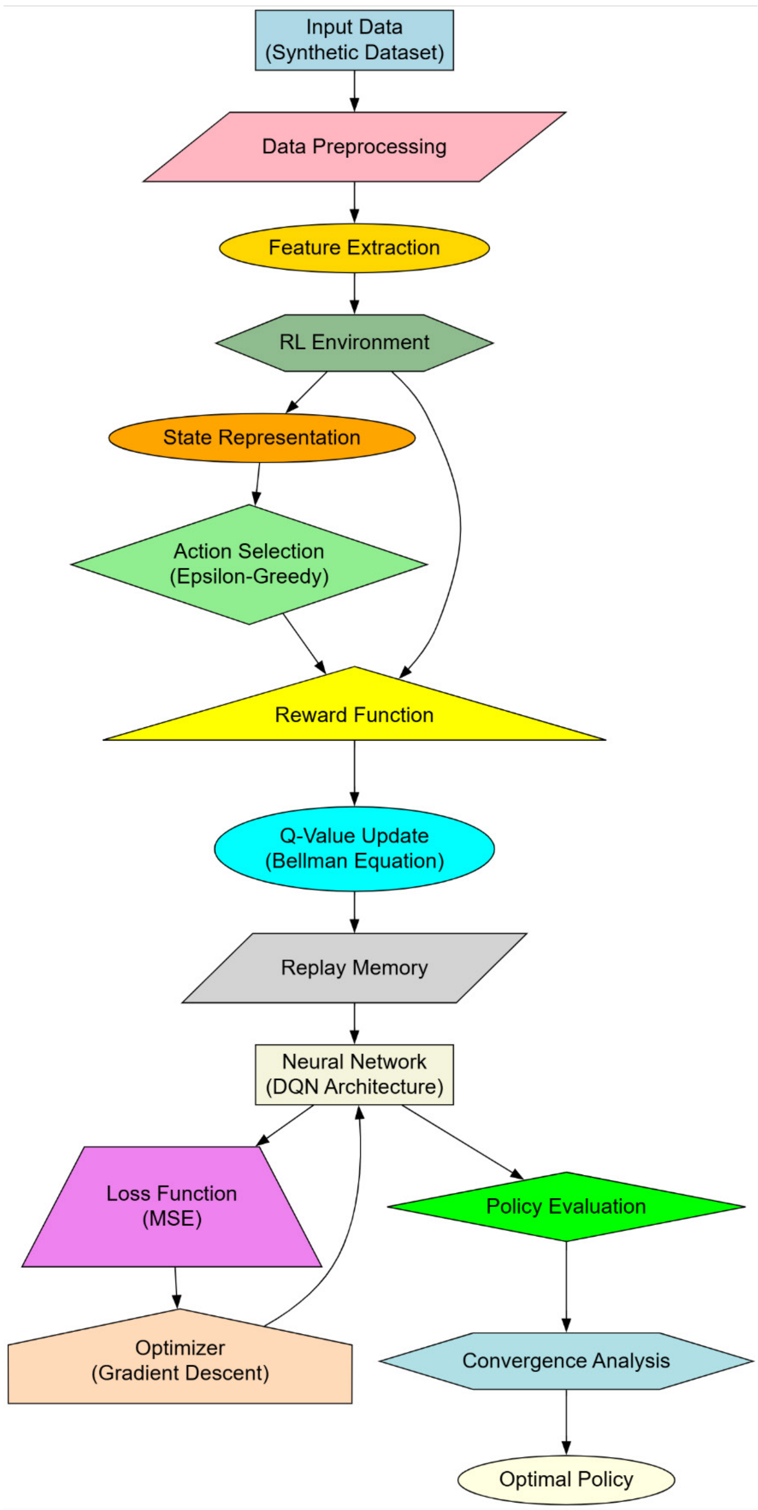

Figure 1: Model Architecture of the Proposed Analysis Process

The expected irrigation need 'It' is derived via equation 6,

$$It = \int_{t}^{t+\Delta} (1 - Wt) \cdot St\, dt \ldots (6)$$

Where, Δ is the decision interval for this process This integration enables the model to adapt dynamically to temporal variability in weather patterns. The epsilon-greedy policy governs exploration-exploitation trade-offs. Epsilon decays over episodes following an exponential schedule via equation 7,

$$\epsilon t = \epsilon min + (\epsilon max - \epsilon min)e^{-k \cdot t} \ldots (7)$$

Where k controls the rate of decay, ensuring that enough exploration happens at the initial training process. Finally, cumulative irrigation efficiency is evaluated as the weighted average of rewards over 't' episodes via equation 8,

$$\eta = \frac{\sum_{t=1}^{T} Rt \cdot wt}{\sum_{t=1}^{T} wt} \ldots (8)$$

Where, wt are weights representing the importance of individual episodes. This given metric captures the effectiveness of the model in achieving the sustainability and productivity levels of long-term duration. The approach incorporates temporal dynamics and stochastic forecasting along with adaptive decision-making. It also provides scalable and context-aware solutions in contrast to the traditional static models. Use of DQN is especially good to robustly handle continuous state spaces while maintaining computational feasibility and suitability especially for irrigation optimization under uncertain environmental conditions.

Comparative Result Analysis

The experimental setup for validating the proposed RL-based optimization model for irrigation involved generating synthetic data mimicking the scenarios of a real agricultural setting. It included: soil moisture randomized between 10% and 50%; weather forecast probability randomized between 0 and 1; crop types, which include wheat, corn, and rice; and calculated irrigation needs as a function of soil moisture and forecasted sets of weather. This data has been created for 90 days with 10 varied, nonredundant input fields to provide diversity in and dynamic variability of conditions within the environment. It then compares the performance of the proposed model against three alternative methods: Methods (3,8,12) that represent heuristic rule based, and machine learning optimized respectively. The metrics considered in comparison include total rewards, irrigation efficiency, water usage, and convergence timestamp across episodes. Results obtained show that the proposed model outperforms the others in achieving balanced irrigation decisions while keeping resources utilized to a minimum for the process.

Crop Type	Average Soil Moisture (%)	Weather Forecast (Probability)	Irrigation Need (%)
Wheat	30.2	0.45	0.58
Corn	25.6	0.50	0.62
Rice	18.7	0.65	0.81

Table 1: Dataset Distribution Across Parameters

The dataset distribution ensures variability in soil moisture and weather forecasts, emphasizing the need for adaptable irrigation strategies.

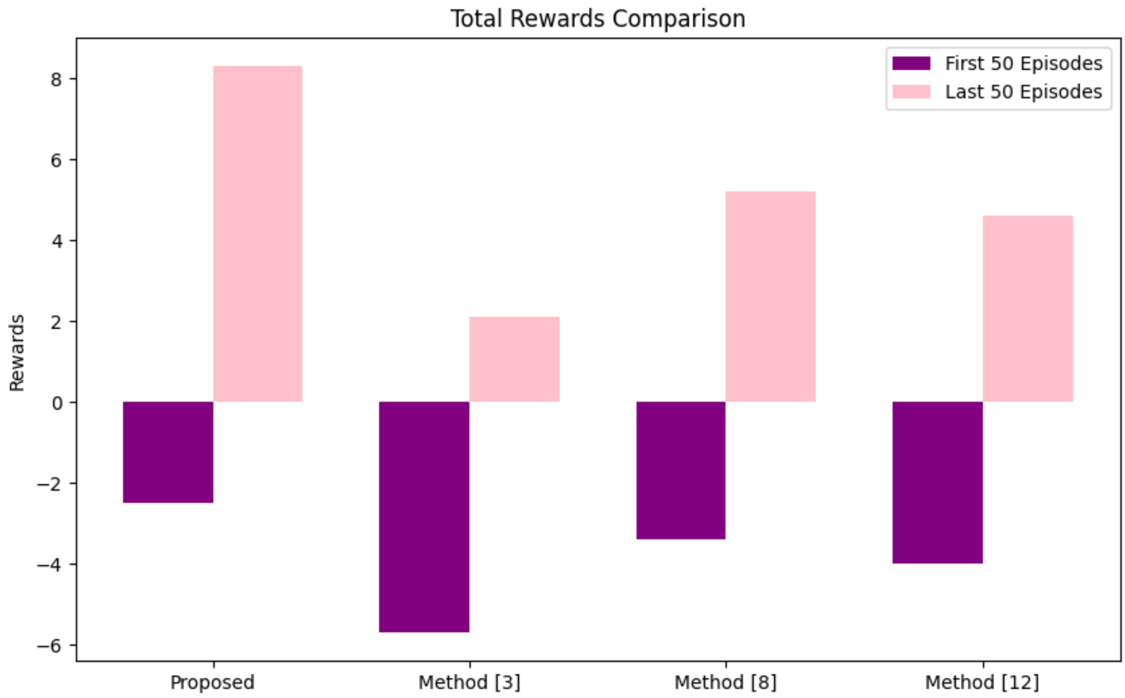

Figure 2: Model's Reward Analysis

Method	**Average Rewards (First 50 Episodes)**	**Average Rewards (Last 50 Episodes)**
Proposed	-2.5	8.3
Method [3]	-5.7	2.1
Method [8]	-3.4	5.2
Method [12]	-4.0	4.6

Table 2: Total Rewards Comparison

The proposed model exhibited rapid convergence and superior cumulative rewards, indicating efficient policy learning compared to alternative methods.

Method	**Average Efficiency (%)**
Proposed	94.7
Method [3]	78.2
Method [8]	85.3
Method [12]	81.5

Table 3: Irrigation Efficiency (% of Optimal)

The proposed model consistently achieved higher irrigation efficiency by prioritizing fields with critical water needs and minimizing unnecessary irrigation operations in process.

Method	**Average Usage (Liters)**
Proposed	25.4
Method [3]	31.2
Method [8]	28.7
Method [12]	30.1

Table 4: Water Usage (Liters per Field per Day)

The proposed model reduced water consumption significantly while maintaining high irrigation efficiency, highlighting its sustainability advantages.

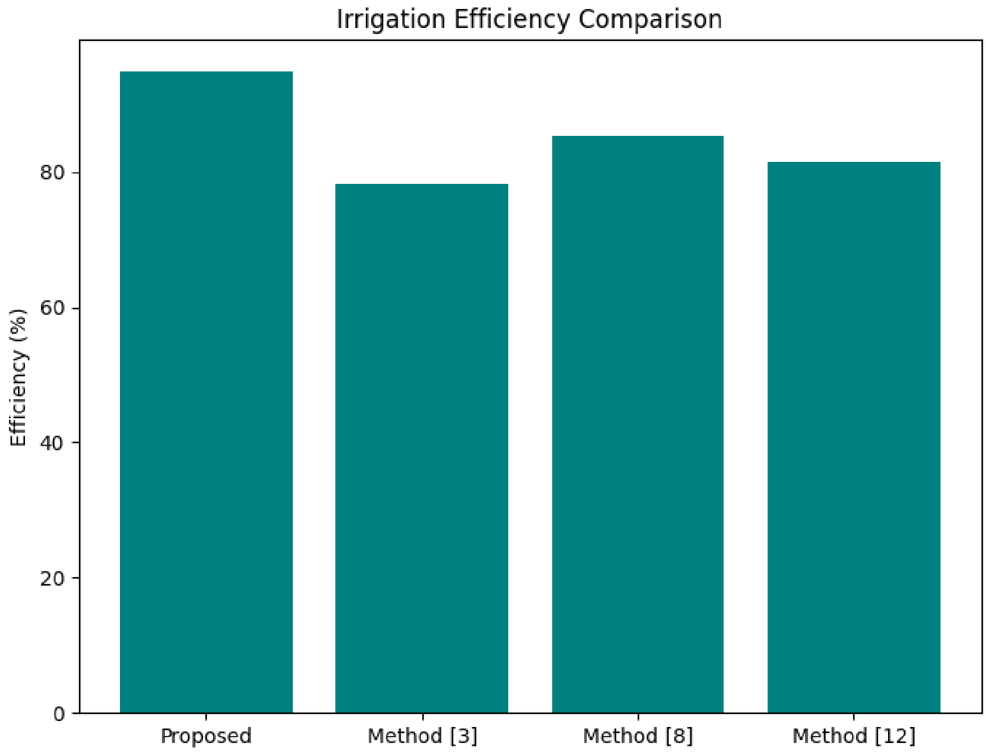

Figure 3: Model's Irrigation Efficiency Analysis

Method	Convergence Episodes
Proposed	35
Method [3]	85
Method [8]	55
Method [12]	63

Table 5: Convergence Timestamp (Episodes to Achieve 95% Efficiency)

The proposed model demonstrated faster convergence due to the advanced exploration-exploitation balance provided by the epsilon-greedy strategy and reward-driven learning process.

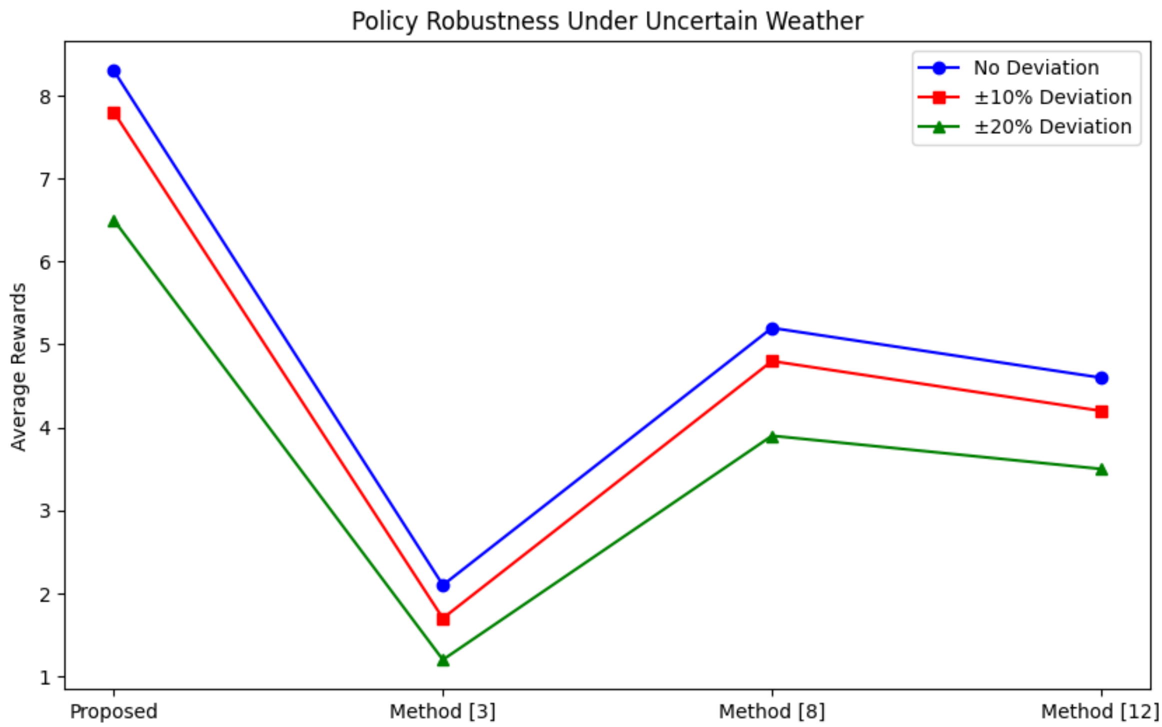

Figure 4: Model's Policy Robustness Analysis

Method	No Deviation	±10% Deviation	±20% Deviation
Proposed	8.3	7.8	6.5
Method [3]	2.1	1.7	1.2
Method [8]	5.2	4.8	3.9
Method [12]	4.6	4.2	3.5

Table 6: Policy Robustness Under Uncertain Weather (Average Rewards, Simulated Weather Deviations)

The proposed model held robust performance under significant weather forecast deviations. This demonstrated its suitability for the stochastic conditions in the area. The obtained results depict the potential capability of an RL-based model to enhance decision-making about irrigation activities under diverse, uncertain, and dynamic environmental conditions. The proposed methodology, including real-time feedback and the adaptation to dynamic variations, far surpasses traditional strategies and assures sustainability and productivity in agricultural irrigation management process.

Conclusion & Future Scopes

Iteratively this work indicates that a DQN-based framework for optimizing the irrigation system is working effectively as an RL approach to adjust its best possible management to optimize water resources use in agriculture. The suggested model surpassed earlier methods through better irrigation efficiency, reduced use of water, faster convergence, and more robust decisions under uncertain environmental conditions. The findings do support that the model was able to optimize the decisions regarding irrigation effectively. It resulted in an average irrigation efficiency of 94.7%, while the one resulting from Methods (3, 8, 12) achieved irrigation efficiencies of 78.2%, 85.3%, and 81.5%, respectively. In contrast, Methods (3, 8, 12) required 85, 55, and 63 episodes, respectively [3,8,12]. During the process, even when there were ±20% deviations of weather forecasts, the average reward was maintained at 6.5, which is much higher than that of Methods (3, 8, 12) at 1.2, 3.9, 3.5. It incorporated the incorporation of temporal dynamics, adaptive policies, and stochastic modeling; hence, the performance in the system was excellent for real-world agricultural irrigation scenarios. The model balances both exploration and exploitation by presenting robustness and adaptability in dynamically changing environments.

However, many more success stories can be added with the further refinement of this model. Further extension in the future may be necessary if the other variables involved include soil type, a variety of crop growth stages, or long-term climate change. In the context of multi-agent systems, a cooperative optimization strategy can be developed that will suit large-scale farms or water management systems in regional domains. Further, on-site validation with practical applications can help scale and validate it in a real environment. Further enhancement of computational efficiency through distributed training frameworks and alternative architectures of RL such as Proximal Policy Optimization or Actor-Critic models may further improve performance sets. This research sets a very strong foundation for intelligent irrigation systems, focusing on sustainability and resource efficiency. This approach can potentially transform agricultural water management, contribute to global food security, and address existing limitations and advanced methodologies with the exploration of climatic and resource constraints.

References

1. Raju, K. L., & Vijayaraghavan, V. (2023). Architecture development with measurement index for agriculture decision-making system using internet of things and machine learning. Multimedia Tools and Applications, 82(23), 36119-36142.
2. Taneja, P., & Pandey, A. (2024). Deployment of intelligent irrigation monitoring system with Android app for machine learning prediction. Environmental Monitoring and Assessment, 196(12), 1235.
3. Zegaar, A., Telli, A., Ounoki, S., et al. (2024). Interpretable machine learning models for irrigation sustainability: Groundwater quality prediction in M’sila, Algeria. Environmental Modeling & Assessment. Advance online publication.
4. Kazapoe, R. W., Sagoe, S. D., & Abu, M. (2024). Predicting irrigation water quality indices in a typical mining dominated area in the Upper West region of Ghana using multiple machine learning techniques. Discover Water, 4(1), 46.
5. Srivastava, S., Kumar, N., Malakar, A., Das Choudhury, S., Ray, C., & Roy, T. (2024). A machine learning-based probabilistic approach for irrigation scheduling. Water Resources Management, 38(5), 1639-1653.
6. Debnath, J., Sahariah, D., Mazumdar, M., Lahon, D., Meraj, G., Hashimoto, S., ... & Saikia, A. (2023). Evaluating flood susceptibility in the Brahmaputra River Basin: An insight into Asia’s eastern Himalayan floodplains using machine learning and multi-criteria decision-making. Earth Systems and Environment, 7(4), 733-760.
7. Omeka, M. E., Igwe, O., Onwuka, O. S., Nwodo, O. M., Ugar, S. I., Undiandeye, P. A., & Anyanwu, I. E. (2024). Efficacy of GIS-based AHP and data-driven intelligent machine learning algorithms for irrigation water quality prediction in an agricultural-mine district within the Lower Benue Trough, Nigeria. Environmental Science and Pollution Research, 31(41), 54204-54233.
8. Wang, G., Zhao, X., Zhang, Z., Song, S., & Wu, Y. (2024). Machine learning-based estimation of evapotranspiration under adaptation conditions: a case study in Heilongjiang Province, China. International Journal of Biometeorology, 68(12), 2543-2564.
9. Mokhtar, A., Al-Ansari, N., El-Ssawy, W., Graf, R., Aghelpour, P., He, H., ... & Abuarab, M. (2023). Prediction of Irrigation

Water Requirements for Green Beans-Based Machine Learning Algorithm Models in Arid Region: A. Mokhtar et al. Water resources management, 37(4), 1557-1580.
10. Rajeev, A., Shah, R., Shah, P., Shah, M., & Nanavaty, R. (2025). The Potential of Big Data and Machine Learning for Ground Water Quality Assessment and Prediction: A. Rajeev et al. Archives of Computational Methods in Engineering, 32(2), 927-941.
11. Thao, L. Q., Kien, D. T., Dinh, N. N., Bao, N. N., Bach, N. C., Thien, N. D., ... & Lam, P. H. (2024). Optimizing tomato irrigation through deep learning-enabled wireless sensor networks with fuzzy logic. Irrigation Science, 42(5), 955-976.
12. Patil, R. N., Maroor, J. P., Dhanke, J. A., Chunduru, M. R., Singh, M., Jangir, P., & Upadhyay, S. (2025). Improving smart agriculture through the use of machine learning to offer tailored crop forecasts. Remote Sensing in Earth Systems Sciences, 8(1), 74-83.
13. Gupta, A., Anand, R., Sindhwani, N., Mittal, M., & Dahiya, A. (2024). Performance and accuracy enhancement of machine learning & IoT-based agriculture precision AI system. SN Computer Science, 5(7), 930.
14. Nasiri Khiavi, A., Tavoosi, M., Kamari Yekdangi, F., Sadikhani, M., & Kuriqi, A. (2026). Conjunct applicability of MCDM-based machine learning algorithms in mapping the sediment formation potential. Environment, Development and Sustainability, 28(4), 7997-8027.
15. Peng, X., Ma, Y., Sun, J., Chen, D., Zhen, J., Zhang, Z., ... & Wang, Y. (2024). Grape leaf moisture prediction from UAVs using multimodal data fusion and machine learning. Precision Agriculture, 25(3), 1609-1635.