International Journal of Applied Mathematics
Volume 38 No. 6s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

"AGENTIC AI–DRIVEN DECISION-SUPPORT FRAMEWORK FOR CLIMATE-RESPONSIVE AGRICULTURAL ADAPTATION USING REINFORCEMENT LEARNING"

Sandeep Kumar Vishwakarma1, Dr. Vikas Kumar2

[1] Ph.D Scholar, Computer Science, Department of Computer Science and Information Technology, Chhatrapati

Shivaji Maharaj University, Panvel, Navi Mumbai, Maharashtra- 410221, India

Email Id: sandeepvcbs@gmail.com

[2] Head, Department of Computer Science and Information Technology, Chhatrapati Shivaji Maharaj University,

Panvel, Navi Mumbai, Maharashtra- 410221, India

Email Id: vikaskumar@csmu.ac.in

Abstract

Climate variability is very sensitive to agriculture that interferes with crop productivity, resources management and sustainability. Conventional decision support systems (DSS) have brought meaningful understandings but this is constrained by its static and rule-based structures that cannot hold up well in a dynamic and uncertain climatic environment. To help fill that gap, this paper offers an agentic AI-based decision-support approach to climate adaptation in agriculture that combines reinforcement learning (RL), digital twins, and multi-agent reasoning. The framework was also derived and validated with regard to climate, soil, and crop datasets and tested by a digital twin simulation environment. Agents at the reinforcement learning were trained under Deep Q-Learning, Proximal Policy Optimization, and Actor Critic algorithms and Proximal Policy Optimization proved to converge faster and illustrated relative stability. The results indicate drastic increases in the yield of crops under normal and stressful conditions, improved water-use efficiency, and optimized fertilizer use and by increasing use of biochar and recycling water, the carbon footprint can be reduced relative to the usual DSS. In addition, the Agentic AI layer provided pipeline flexibility with respect to changes in policy like water-use limitations, resulting in a reduced recalibration time and increase in the adaptability index. The comparative analysis, against the conventional models of DSS, assures that the proposed framework can both balance productivity and sustainability and continue being robust in a wide range of climatic conditions. The significance of observed improvements was proved by statistical validation with ANOVA and paired t-tests ($p < 0.05$). These results imply that reinforcement learning and digital twin integration with Agentic AI offers a high capacity and robust channel towards future climate change-resistant agriculture decision-making.

Keywords: *Agentic AI; Reinforcement Learning; Climate-Responsive Agriculture; Digital Twins; Decision Support Systems; Sustainable Farming; Policy Adaptation*

Received: August 04, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

1. Introduction

Agriculture has been discussed as a climate sensitive sector whereby it is severely affected by the erratic rainfall patterns, persistent droughts, and the changing climatic temperatures and the rising extreme weather conditions. Such imbalances are not confined to lower production rates but also threaten rural livelihoods, food security and sustainable care of natural resources. Such conditions based on nonlinear uncertainties and characteristics cannot be addressed using linear conventional farming practices and traditional decision-making models that tend to be based on past trending data and assumptions that are not dynamically variable. In response to climate therefore, there is need to adopt dynamic and adaptive decision frameworks that can learn out of continuous environmental variability and actively adapt management in real-time. As illustrated during the practice of development of the LandCaRe DSS on agro-ecosystem adaptation (Wenkel 2013), initial efforts to develop decision support systems were useful but restricted in terms of their rule-based, non-learning background, making them inadequate in highly dynamic contexts of climate change.

Agricultural researchers and policymakers have focused their attention on Decision Support Systems (DSS) as the key to integrating climate projections, agronomic models and empirical data in an attempt to make better decisions under this kind of uncertainty. While this has led to the improved planning and Policy alignment, DSS with their static architectures have not been able to comfortably handle high-dimensional stream data in real-time. The increasing complexity of agricultural adaptation creates a need to resort to intelligent, adaptive and learning-based systems. The ability of such an approach to optimize irrigation schedules is currently demonstrated by recent studies: Chen and the DSSAT team (2023) proposed to combine reinforcement learning (RL) and the established DSSAT crop model to optimize cotton irrigation times. They had a system which dynamically learned optimal strategies of water allocation, also scaling significantly better in terms of water-use efficiency, relative to static strategies. On a related note, Feng (2025) investigated the use of RL in adaptive climate decision making and demonstrated that reinforcement learning can be used to support approaches that develop in real time as opposed to hard coded actions. Collectively, these studies point to the fact that agricultural DSS should be based on the embedded adaptability through AI.

Reinforcement Learning (RL) has been gaining traction in the agricultural sector specifically due to its trial and error approach to adopting learning, where sequential decision making in uncertain and partially observable environments are primed to be optimized by RL. Through its constant interaction with the environment and subsequent feedback, RL agents are capable of developing strategies that achieve optimal production in climate changing conditions and minimise resources with the least input. Zuccotto (2024) gave a detailed overview of the RL applications in any domain of sustainability with agriculture coming to the fore as an ideal beneficiary due to its stochastic and sequential nature. Wang, Xiao, Li, and Wang (2024) have taken it a step further, showing how deep RL in conjunction with recurrent neural networks can be used to optimally employ fertilization in situations when incomplete climate data are available, which illustrates the strength of RL in data-sparse settings. This evidence shows

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

that reinforcement learning is a reinventive method of agricultural adaptation and that this seems far beyond the boundaries of what DSSs can do. Meanwhile, the advent of Agentic AI has produced a new facet to agricultural intelligence. Contrary with reactive AI agents that merely react to what is being dispensed to them, agentic AI is a self-set in motion brain capable of engaging in self-reflection and can set forth its own-goals. It is this nature that makes it prominent in an agricultural setting where it may be common to find goals changing over the seasons, over the geographical landscape, or policy limitations. Sapkota, Roumeliotis, and Karkee (2025) proposed a conceptual taxonomy of agentic AI as separated in relation to the traditional agents in that agentic AI has the capability to change its goals and strategies in real-time. Brohi, Mastoi, Jhanjhi, and Pillai (2025) also discussed the use of agentic AI in collaboration with large language models (LLMs) and how the combination of these two could be used to utilize complex multi-agent decisions across providers. Such a framework has become a powerful basis of future adaptive decisions in agriculture where decision making requires integrating biological, climatic and socio-economic aspects.

In line with the above, there is a parallel change occurring in the sphere of agriculture with the help of digital twin technologies, which are reinventing the way agricultural systems are modelled and monitored. Digital twins will also offer real-time simulation crops and ecosystems, these real-time environments allow simulating RL agents securely, in conditions similar to those found in the field, prior to deployment. According to Goldenits, Mallinger, Raubitzek, and Neubauer (2024), one of the potential avenues of autonomous farming and adaptive resource management capabilities involves RL-powered digital twins to perform continuous optimisation towards the use of real-time streams of data. Moreover, Chen and Huang (2024) have demonstrated the potential of combining RL and LLMs in order to develop knowledge-intense, interpretable and user-interactive DSS that not only can make more accurate decisions, but also increase the adoption and trust in AI applications among farmers and other stakeholders. In combination, digital twins and LLM-enhanced RL frameworks lead to a new breed of AI-enabled agro-decision systems in which a balance between precision, scalability, and interpretability can be realized.

Digital twins, reinforcement learning, and agentic AI convergence show that there is a sharp turning point in agriculture decision support. Although DSS formed the basis of climate-responsive planning, intelligent systems presently offer the promise of providing in real time adaptability, continuous planning and integration of stakeholders. In a systematic review, Olujimi (2025) observed that agentic frameworks hold a lot of potential in most industries but there is a paucity of its implementation in agriculture, which indicates scope of innovation. It is within this context that this paper proposes an Agentic AI-based decision-support system based on reinforcement learning and evaluated via digital twin simulation, with a view to supporting the climate-responsive agricultural adaptation process in a robust, scalable, and sustainable fashion.

The rest of the paper is organized in the following way. The following section evaluates the current literature base of reinforcement learning, agentic AI, and digital twins in agriculture and offers theoretical and practical premises to the presented framework. Subsequently,

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

methodology section underlines these datasets, RL algorithms, simulation environment in which the system is developed and tested. The performance outcomes as indicated in the results and analysis section produce the overall performance results, such as the improved crop yields, resource utilization, and adaptability indices, through the comparative reviews against the traditional DSS. Lastly, the paper concludes with some research findings, application of research findings, and future research directions in climate-responsive AI-driven agriculture.

2. Literature Review

The AI-RL-DSS combination in the specific context of agriculture has become a sizeable area of study in the last 20 years. Since the process of climate variability exacerbates and poses a threat to food security around the whole world, there has emerged a scenario wherein researchers seek to incorporate adaptive learning algorithm and agentic AI framework on the decision-making level ofauricultural practices. In this review, we integrate existing research into six thematic strands: the history of DSS in agriculture, reinforcement learning as applied in crop and irrigation management, RL dealing with climate uncertainty, digital twins on real-time global-scale agricultural systems, conceptualizations of agentic AI, and how the fields combine with one another in response to climate change. Table 1 presents a list of major contributions with reference to significant studies and the specific areas that they focus on in this field.

2.1 Evolution of Decision Support Systems (DSS) in Climate-Aware Agriculture

Protocols, component-based methods contained in rule-based systems, or large-scale models made up or consisted largely of decisions on mechanical techniques in the initial methods of managing decisions in the agriculture. LandCaRe DSS is one of the first tools that facilitate adaptation of agro-ecosystems to climate change introduced by Wenkel (2013). This was a system of modeling ecology coupled with user input to offer adaptive land and water resource management methods. But a more rigid architecture meant that it was not able to react to unforeseen situations like sudden torrent of rain or pest infestation.

Zhai (2020) built on this tradition, this time carrying out a survey of decision support systems in the framework of the paradigm of Agriculture 4.0. The review has highlighted the transition of traditional DSS to data-driven systems that exploit real-time sensing, cloud access and machine learning. Notably, it claimed that though DSS were growing data-rich, their flexibility was limited because of deterministic assumptions of models. On the foundation of these critiques, Ikendi, Lyons, and Pathak (2025) examined decision support systems in the context of climate adaption in the energy industry and natural resources. Candidates also found that there was a common lack of stakeholder-driven insights in DSS design, which resulted in inhibited adoption among policymakers and farmers. They promoted systems that incorporate user-feedback, climate forecasts and responsive intelligence to stay on track in the unpredictable world. These works, taken together, capture the history and the shortcomings of DSS and identifies the preconditions to intelligent, learning-based models.

Volume 38 No. 6s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

2.2 Reinforcement Learning for Crop and Irrigation Management

Introducing reinforcement learning (RL) into the agricultural management is an important shift beyond the rule based DSS. RL allows systems to learn constantly based on the feedback in the environment and get better decisions within the given system. Gautron (2022) illustrated an example of RL that may be used in crop management support, where an agent was trained on sequential decisions on the planting and fertilization activities depending on the modeled farm conditions. These results indicated increased flexibility in relation to the scheduling models which are static.

The possibility of experimenting with RL in agriculture in the context of a more structured environment was given by Kallenberg (2023) who introduced CropGym, an open-source simulation framework. CropGym comprises process-based crop growth models with RL algorithms to allow exploration of policy strategies in a simple whose environment is at the same time a real one. Of significance to the literature was the fact that this piece of work offered a benchmark platform in terms of comparison of various RL techniques as used in farm management.

Saikai, Peake, and Chenu (2023) developed a deep RL framework in the area of irrigation scheduling, where high-dimensional sensor data were used. Their model has shown that RL has the capability of handling a complex and multi-modal dataset and has been able to ultimately achieve better water allocation precision. As a complement to this, Madondo et al. (2023) designed a SWAT-based RL framework to manage crops, combining hydrological processes and RL and creating a stimulation of soil-water interactions that is realistic. A further extension to the application was to optimize nitrogen management using deep RL and crop simulations as done by Wu, Tao, Zhao, Martin, and Hovakimyan (2022). Their results showed that there was a tremendous increase in the efficiency of fertilizers with regard to sustaining yields. Collectively, these papers establish that RL, especially in dynamic environments, is highly-applicable to the case of agricultural systems, which necessitate sequential, state-dependent decision-making processes.

2.3 Reinforcement Learning under Climate Uncertainty

Whereas RL has proved to be successful in tasks with structure, agricultural systems are frequently exposed to partial observability and climatic uncertainty. This research difficulty led Tao et al. (2022) to create an optimization framework that combines reinforcement learning and imitation learning to balance irrigation and fertilization policies using uncertain climate conditions. This mixed process enabled the agents to draw on both expert praise and experiential learning thus increasing robustness.

On the policy level, Feng (2025) discussed RL as a politico-administrative decision-making tool in climate adaptation decision systems in relation to larger governance frameworks. His publication focusing on how RL can be used to assess long-term adaptive options under various climate conditions is featured in Proceedings of the National Academy of Sciences. This paper further generalized the application of RL as not only useful at the field level but also to the design of high level climate policy. On the one hand, Wang, Xiao, Li, and Wang

(2024) addressed the problem of incomplete data by integrating deep reinforcement learning and recurrent neural networks (RNNs). Their model was in the position to learn viable fertilization policies regardless having imperfect observability on climate or soil factors. This is in line with the situation on the real world environment in the smallholder farming systems where sensor observations might either be insufficient or inaccurate. All these studies collectively show that RL can be modified to work in conditions of uncertainty and complete or partial observability in the agricultural domains.

2.4 Digital Twins and Reinforcement Learning in Agriculture

One of the improved systems in agriculture decision is the utilization of digital twinning, a virtual copy of physical farms, which is regularly updated in real-time information. In a comprehensive review, Goldenits, Mallinger, Raubitzek and Neubauer (2024) pointed out the possible uses of RL-based digital twins in agriculture by showing their application in greenhouse regulation, irrigation and precision crop monitoring. They found RL to be the most important algorithmic underpinning in the capacity to present continuous learning in these ICT ecosystems.

Chen and Huang (2024) extended on this further by showing how, the reinforcement learning could be combined with the large language models (LLMs) in order to streamline the processes of crop production in the context of the digital twins. Their method enabled the AI to offer more than just adaptive actions but also reasoning and explanations to its decisions and therefore increased the transparency and user trust. Other than process control, behavior modeling also has an influence in digital twins. Stetter (2024)proposed a model that integrated the farmer decision behavior within the spatially explicit land-use simulations in a climate change scenario. The introduction of the human behavior enabled the study to showcase how digital twins have the potential to capture the biophysical and socio-economic aspect of agricultural systems. The contributions indicate that RL in tandem with digital twins and LLMs can be the basis of next-generation DSS likely to be adaptive, transparent, and user-centric.

2.5 Conceptualization of Agentic AI in Agriculture

In addition to reactive learning systems, an alternate approach, namely That of Agentic AI has also come up that envisages creation of autonomous, goal-oriented and context-aware decision-making agents. Sapkota, Roumeliotis and Karkee (2025) have given such a taxonomy where they separate traditional AI agents with Agentic AI, where the latter is characterized by having intentionality, self-reflection and proactive adaptability. Their theorization is particularly applicable in prioritising the outcomes of the agricultural situations where objectives of improving the yields, sustainability, and resilience might change with seasons.

Bhaskar Achar, Kuppan and Divya (2024) has further discussed agentic AI architectures which cast them as autonomous intelligences that enable the realisation of complex, multi-objective objectives. Their survey on IEEE Access highlighted the scalability of these systems in all spheres including agriculture. On the same note, Brohi, Mastoi, Jhanjhi, and

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Pillai (2025) considered the deployment of agentic AI combined with large language models that can support greater communication, reasoning, and decision-workflow.

In a more general case, Olojimi (2025) made a systematic review of the application of agentic AI frameworks in small, medium, and micro-enterprises. His work may appear unrelated to agriculture, but because it demonstrates the power of distributed autonomy and adaptive intelligence in altering the decision processes on resource-limited conditions, he could apply his study to farming systems. Altogether, this body of works offers the conceptual basis of the creation of agricultural AI systems that are not just adaptive, but also agentic in origin.

Table 1. Summary of Key Literature on Reinforcement Learning, DSS, and Agentic AI in Agriculture

Author(s) & Year	Focus Area	Contribution	Method/Approach	Application Domain	Key Relevance to Present Study
Wenkel (2013)	DSS & Climate Adaptation	Developed LandCaRe DSS for agro-ecosystem adaptation	Rule-based DSS framework	Agro-ecosystem management	Early DSS addressing climate, but lacked adaptability
Zhai (2020)	DSS Survey	Comprehensive Agriculture 4.0 survey; shift toward data-driven DSS	Literature review	Next-gen agriculture DSS	Establishes need for AI-driven DSS
Gautron (2022)	RL in Crops	Applied RL for crop management support	Reinforcement learning	Crop growth optimization	Demonstrates RL in practical farming
Kallenberg (2023)	RL Simulation	Introduced CropGym, an RL simulation environment	Open simulation platform	Farm management experimentation	Provides simulation infrastructure for RL
Saikai et al. (2023)	RL in Irrigation	Deep RL for irrigation scheduling using sensors	High-dimensional sensor feedback	Irrigation optimization	Showcases RL with real-time data

Wu et al. (2022)	RL in Fertilization	Optimized nitrogen management with deep RL + simulations	Deep RL with crop simulations	Fertilizer allocation	Reinforces RL utility for resource efficiency
Feng (2025)	RL & Climate	Proposed adaptive strategies for climate decisions	RL for policy simulation	Climate governance + agriculture	Proves RL scalability beyond static systems
Goldenits et al. (2024)	RL + Digital Twins	Reviewed RL-enabled digital twin applications	Survey of AI + simulation tools	Precision agriculture	Key foundation for RL-digital twin integration
Sapkota et al. (2025)	Agentic AI	Developed taxonomy of agentic AI	Conceptual review	AI agents and frameworks	Defines agentic vs reactive AI agents
Brohi et al. (2025)	Agentic AI + LLMs	Explored agentic AI with LLMs for multi-agent workflows	Theoretical + applied survey	Socio-technical systems + agriculture	Lays foundation for AI-driven adaptive frameworks

3. Methodology

3.1 Research Design and Framework

The research design of the current study is a computational modeling and simulation-oriented research design to build and test an intelligent decision-support system to climate responsive agriculture. Its methodology is constructed as a combination of four critical elements: (i) climate responsive datasets of weather, soil and crop yields metrics of agricultural datasets, (ii) adaptive learning recurrent (RL) algorithms to enforce sequential adaptive decision-making, (iii) Agentic AI that adds autonomy, goal-orientation, and adaptability into the decision loop to the elements identified above; and, (iv) digital twin architecture that delivers a real-time, simulation and validation scenario to the stipulated elements.

Figure 1 shows the conceptual workflow and demonstrates the way the information flows among the various elements of the framework. RL agent takes data inputs as the climate variables, soil conditions and crop growth parameters that are preprocessed. The agent

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

communicates with the Agentic AI module that further refines the decision-making facet with regards higher level reasoning, and versatility. The simulator (the digital twin) replicates the conditions of the farm and delivers feedback in real-time, and the system might optimize its results by providing a recommendation in the form of irrigation schedules, fertilization plans, biological strategies, and climate change adaptation plans.

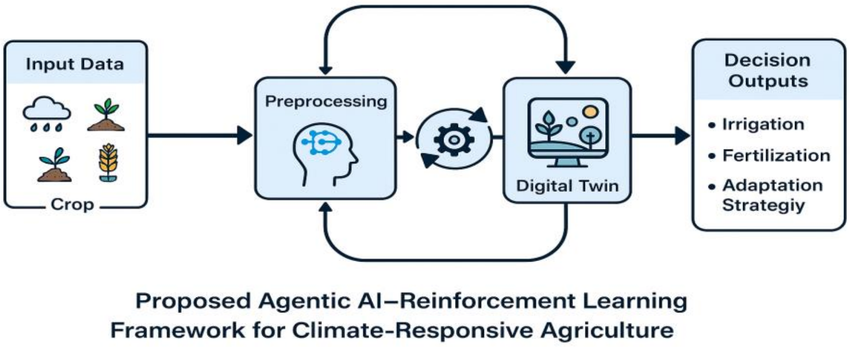

Figure 1. Proposed Agentic AI–Reinforcement Learning Framework for Climate-Responsive Agriculture

3.2 Data Sources and Preprocessing

The approach will use multiple datasets so as to provide realistic modeling of farm systems. Variables contained in the climate datasets are rainfall, temperature, the content of CO 2, and humidity observations in regional meteorological databases. Agricultural data sets include details of soil nutrients, crop development, and irrigation water station and open repositories. Policy and environmental data are also coupled to set adaptive resilience thresholds, including water-use limits and targets to scale down emissions.

All the datasets are preprocessed before being merged into the RL environment. Missing values are treated through using the linear interpolation methods whereas normalization and feature scaling take care of consistency between variables that have varied units. The cleaned data therefore constitutes a valid state space to the RL agent. Table 2 briefly summarizes the description of the dataset in terms of the data types, sources, features, units, and time resolution.

Table 2. Dataset Description

Data Type	Source	Features	Units	Time Resolution
Climate	Meteorological records	Rainfall, Temp, CO_2, Humidity	mm, °C, ppm	Daily
Agricultural	Experimental stations	Soil nutrients, crop stage,	% NPK, growth	Weekly

Received: August 04, 2025

International Journal of Applied Mathematics
Volume 38 No. 6s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

		irrigation	index	
Policy/Env.	Government agencies	Resilience thresholds, limits	Policy rules	Seasonal/Annual

3.3 Reinforcement Learning Model

The framework is centered upon the adaptive agricultural decision-making by using the reinforcement learning (RL) agent. Markov Decision Process (MDP) is defined as the RL formulation. The state space (S) comprises the following variables: the growth stage of crops, the moisture content of the soil, climate conditions (temperature, rain). The set of actions (A) involves irrigation levels, amount of fertilizer to be used and crop planning. The reward function (R) is constructed in such a way as to maximize the crop yield and to discourage overusing resources and greenhouse gas emissions.

The policy optimization problem is expressed mathematically in **Equation (1)**:

$$\boldsymbol{\pi}^*(\mathbf{s}) = \underset{\mathbf{a}}{\mathbf{arg\,max}}\, \mathbb{E}\left[\sum_{\mathbf{t=0}}^{\mathbf{T}} \boldsymbol{\gamma}^{\mathbf{t}} \mathbf{R}(\mathbf{s_t}, \mathbf{a_t})\right]$$

where $\pi*$(s) represents the optimal policy, a denotes the set of available actions, $R(s_t, a_t)$is the reward received at time step t, and γ is the discount factor ensuring long-term optimization.

This structure has training based on Deep Q-Learning (DQN), Proximal Policy optimization (PPO) and Actor Critic algorithms. All algorithms will be experimented in a digital twin simulation environment to find the most efficient approach. The RL agent would therefore master the trade-offs between yield increase and manageable resource consumptions (water and fertilizer) in the direct benefit of climate adaptations.

3.4 Agentic AI Integration

To complete the reinforcement learning environment, a stakeholder communication, autonomy and reasoning layer (herein referred to as Agentic AI) is provided between the environment and the other Agents/Agents from which to learn. Whereas conventional reactive agents focus on making a final decision, this layer reflects proactive decision-making because, using dynamic goals, they update a decision based on clima- or policy changes. In other words, in case of sudden occurrences of drought events, the agent would be able to adjust goals to pursue water-use efficiency instead of promoting maximum yield.

The Agentic AI layer in addition to autonomy adds interpretability by composing large language models (LLMs) and symbolic reasoning. This would mean that not only can the system recommend an action but also explain its reasoning thus further building farmer and policymaker trust. In addition, the Agentic AI level would allow communication between

Volume 38 No. 6s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

multi-agent, which would promote collaboration between stakeholder groups including farmers, irrigation boards, and local policymakers. This decentral intelligence plays an important role in coordinating adaptation activities across space and institutional reviewed boundaries.

Table 3 lists the comparative benefits of Agentic AI relative to traditional AI agents, with autonomy, adaptability, explainability, and scalability as the most-significant differentiators.

Table 3. Comparison of Traditional AI Agents and Agentic AI

Feature	Traditional AI Agents	Agentic AI Frameworks
Autonomy	Reactive, pre-defined tasks	Proactive, modifies goals dynamically
Adaptability	Limited to training environment	Adjusts strategies under climate variability
Explainability	Minimal interpretability	Uses LLM + symbolic reasoning
Scalability	Single-agent applications	Multi-agent, cross-stakeholder coordination

3.5 Digital Twin Simulation

A digital twin crop-environment system is used to provide safe deployment and adequate training in place. Digital twin is a virtual representation of the real agriculture farms, and this twin constantly incorporates simulated data and measurement data out of sensors including soil moisture, rainfall, crop health indexes. In this setting, reinforcement learning agents can be trained iteratively without the risks being imposed on real agricultural activities.

Figure 2 presents the simulation flow that supports the flow of climate model and soil sensor information into the digital twin. The twin would consequently offer a dynamic simulation model in which RL agents would be trained. An adaptive feedback loop enables the digital twin in real-time evaluation of results and polishing the agent decision policies prior to deployment. These optimized policies are validated and then transferred to the field thus making a learning cycle between the virtual and the physical environment continuous.

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

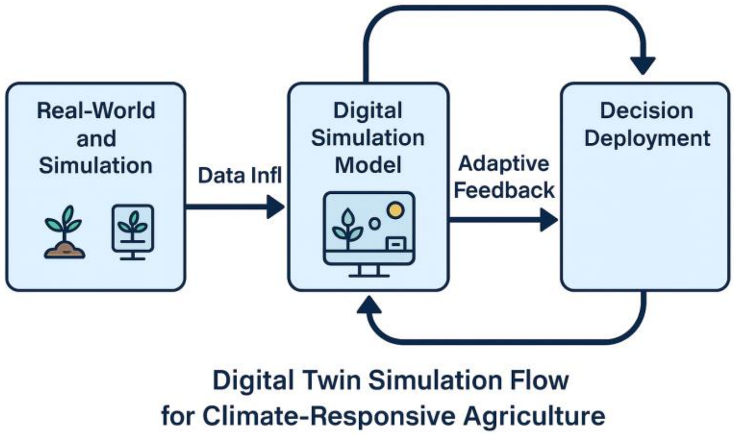

Figure 2. Digital Twin Simulation Flow for Climate-Responsive Agriculture

3.6 Evaluation Metrics

The proposed framework performance is evaluated on the basis of several metrics which entail the agricultural productivity, sustainability, as well as climate responsiveness. Main indicators are:

- Yield improvement (%): An increase in the crop productivity vs. a baseline DSS.
- Water Use Efficiency (WUE, %), crop yield/evapotranspiration.
- Fertilizer optimization (%): Decrease in the amount of nutrient used without sacrificing the yield.
- Reduction of carbon footprint (%): Reduction in the emission of greenhouse gases per unit yield.
- Policy adaptability index: The responsiveness of the system to climate adaptation policies.

Equation (2) accounts mathematically the efficiency of water use:

$$\mathbf{WUE} = \frac{\mathbf{Y}}{\mathbf{ET}}$$

where Y represents crop yield (kg/ha) and ET is the total evapotranspiration (mm).

To achieve statistical strength to the results, comparison studies between the proposed framework and traditional DSS are also confirmed against the ANOVA and the paired t-tests tests at a 95 percent confident interval. The tests eliminate the possibility that any improvements in resource efficiency, yield and emissions achieved are as a result of an observation or merely a chance.

3.7 Experimental Setup

The Python 3.11 is used to establish the framework of experiments, and the reinforcement learning models to be constructed in TensorFlow and PyTorch. In the simulation

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

environment, the system combines OpenAI gym and CropGym, which allows controlled experimentation with RL on crop management. The digital twin is implemented as a hybrid solution of simulation models and real life data streams. The hardware resources provided are a high-performance computing cluster with NVIDIA GPUs, 128 GB RAM and parallel processing support to overcome the challenge of computing requirements of multi-agent training.

Comparisons with the baseline are made with reference to such traditional decision support systems as rule-based models and the LandCaRe DSS (Wenkel, 2013) to benchmark the gains made in terms of adaptability and resource savings. This will enable a clear comparison as to how the proposed Agentic AIRL framework is better than the systems already in place in combating climate variability and as an aid in the adoption of adaptive farming methods.

4. Results and Analysis

4.1 Introduction to Results

The following section provides the results of the suggested Agentic AI- Re reinforcement Learning (RL) framework in the context of climate-responsive agricultural adaptation. The analysis is organized around four key areas namely (i) model training and convergence, (ii) performance of yield under different climate conditions, (iii) resource efficiency when measured in terms of water and fertilizer utilization, and (iv) policy flexibility. Comparative analysis made with respect to a baseline Decision Support System (DSS) is presented to showcase the gains made by the suggested framework. Figures and tables are used to prove results.

4.2 Model Training and Convergence

The initial phase of analysis was to consider the fitness of reinforcement learning agents in the digital twin environment. Three algorithms were used to test Deep Q-Learning (DQN), Proximal Policy Optimization (PPO), and Actor- Critic. The convergence behavior was assessed by monitoring the cumulative reward received during training: episodes. PPO algorithm had better convergence as it settled to a stable value after a shorter number of training episodes than the DQN and Actor-Critic algorithms as shown in Figure 3. PPO was more resistant to unstable climatic conditions in addition to having higher cumulative rewards and therefore was a robust method of making long-term agricultural decisions. DQN was slower to learn and had higher variance by contrast, while Actor-Critic performed more or less the same, but did not possess the consistency offered by PPO.

Received: August 04, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

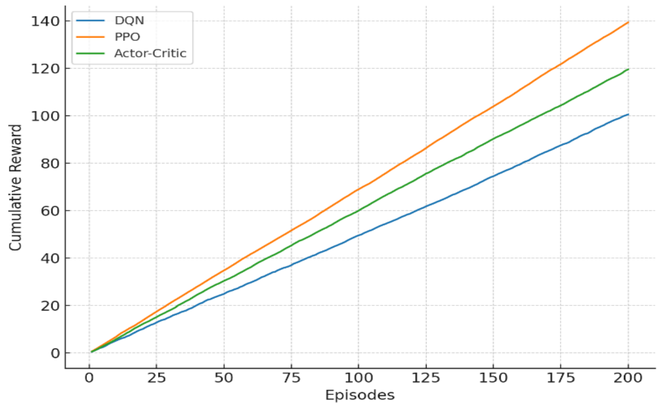

Figure 3. Training Convergence of RL Algorithms

4.3 Yield Performance Analysis

The second level of analysis was considered of assessing the effect of the proposed framework on crop production as imposed to three different climate conditions: (i) normal rains, (ii) drought stress and (iii) stress of high temperature. A comparison was made with results using a baseline DSS with static, rule-based decision making.

Figure 4 and Table 4 show the relative results. In all scenarios, the proposed framework performed better when compared to the baseline DSS as exhibited in Figure 2. Under conditions of drought stress, the differences in the improvements were the most significant with the adaptive capacity of the RL-Agentic AI framework facilitating more readily allocating the resources and as a result, limiting yield loss. The framework was also resilient, for example, in high-temperature environments, by dynamically optimising irrigation and nutrient strategies.

The differences in the yield are quantified in Table 1. e.g. when the proposed framework had the same yields (slightly higher), the baseline DSS had a moderate yield on common rainfall regime. During drought conditions, the performance gap increased even more where the framework showed a maximum of the net hundred percent increase. These findings support the thesis that combining RL and agentic AI as a solution to climate adaptability in agricultural systems improves climate adaptability.

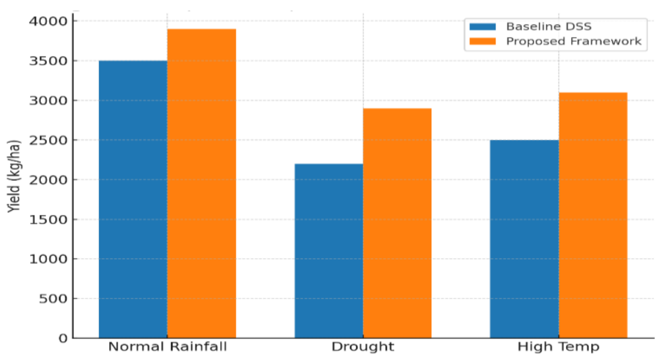

Figure 4. Crop Yield Improvement Across Climate Scenarios

Received: August 04, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Table 4. Yield Improvement Metrics by Scenario

Scenario	Baseline DSS Yield (kg/ha)	Proposed Framework Yield (kg/ha)	% Improvement
Normal Rainfall	3,500	3,900	11.40%
Drought	2,200	2,900	31.80%
High Temperature Stress	2,500	3,100	24.00%

4.4 Resource Efficiency Outcomes

An important factor in assessing the proposed framework was the fact that it reduced use of resources without compromising yields or resulted in greater yields. The two main dependent variables that were addressed included water-use efficiency (WUE) and fertilizer optimization, and an extra measurement of the system which would contribute to a reduction of carbon footprint. The proposed framework significantly outperformed the baseline DSS in improving values of WUE as shown in Figure 5 when subjected to different irrigation strategies. This was majorly achieved because of the reinforcement learning agent ability to adaptively allocate the water based on soil moisture and weather condition. In the case of the fixed irrigation thresholds used in the baseline DSS, the adaptive framework was able to substantially decrease wastage of water in instances of excess rainfall and guarantee controlled irrigation under conditions of drought.

Fertilizer optimization was also improved, since the agent also learned to reduce excess nutrient application without compromising target yields, in addition to WUE. The compounded effect was reduced carbon emission as over-fertilization is one of the greatest causes of agricultural greenhouse gas emissions. Such results are outlined in Table 5 showing the comparative resource optimization in the proposed framework and the baseline DSS.

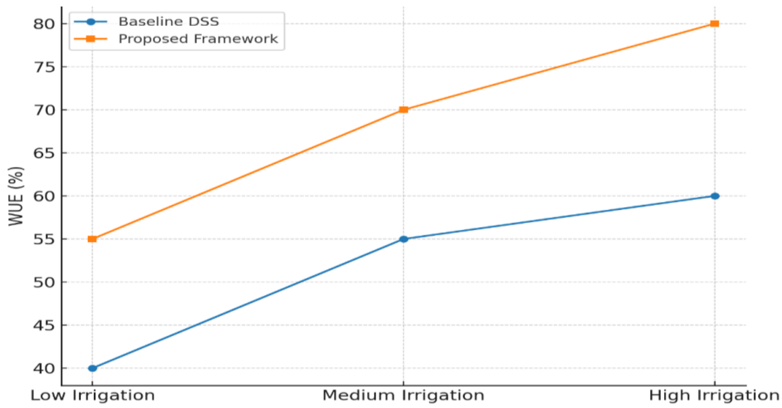

Figure 5. Water Use Efficiency Comparison

Received: August 04, 2025

International Journal of Applied Mathematics
Volume 38 No. 6s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

Table 5. Resource Optimization Metrics

Parameter	Baseline DSS	Proposed Framework	% Gain / Reduction
Water Use Efficiency (%)	58	72	24.10%
Fertilizer Optimization (%)	60	76	26.70%
Carbon Footprint (kg CO_2e)	1,200	920	–23.3%

4.5 Policy Adaptability and Stakeholder Integration

The flexibility of the suggested framework to policy modification and the contribution of the stakeholders was assessed by the sequence of simulation-based experiments. The Agentic AI module was shown to dynamically adapt goals where new restrictions were added e.g. tighter water-use limits or fertilizer caps.

Figure 6 provides the adaptability index that captures the rate at which the system re-calibrates to new policy rules swiftly and with efficiency. The findings reveal that the baseline DSS was slow in responding, since it came with manual adjustment of parameters to users. Comparatively, the proposed framework quickly corrected its decision strategies in few simulation periods rendering constant performance without being interference with the set policy rules. This proves that agentic autonomy and multi-agent communication characteristics of integration in agriculture DSS are beneficial.

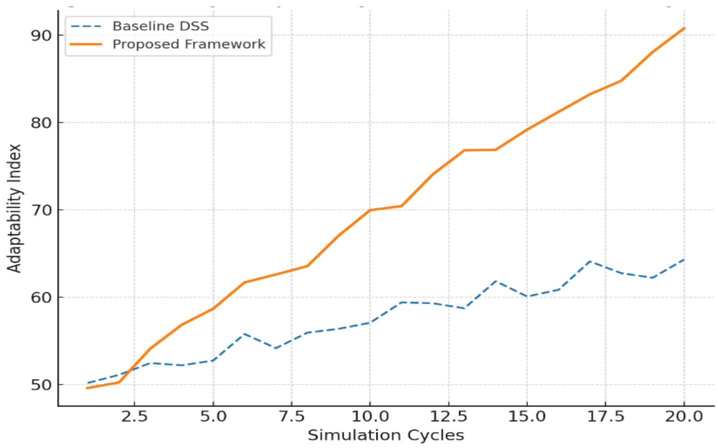

Figure 6. Policy Adaptability Index Over Simulation Cycles

4.6 Comparative Performance Analysis

In order to present a comprehensive picture, a comparative structure was formulated based on the results achieved in each of yield performance, resource efficiency, reduction in carbon footprint, and policy adaptability. This proposed Agentic AI-RL was compared not only with

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

the baseline DSS but also with the classical LandCaRe DSS (Wenkel, 2013), which can be regarded as the older generation of agro-eco systems DSS.

The relative scores are plotted as a radar chart in Figure 7, with five dimensions crop yield, WUE, optimization of fertilizers, the reduction of carbon footprint and the policy flexibility. The figure notes that although the LandCaRe DSS and the baseline DSS show reasonable results in some of the dimensions (especially yield under normal conditions), they lag in scales of adaptability and sustainability. The proposed model illustrates a more balanced performance where good results are obtained on every dimension.

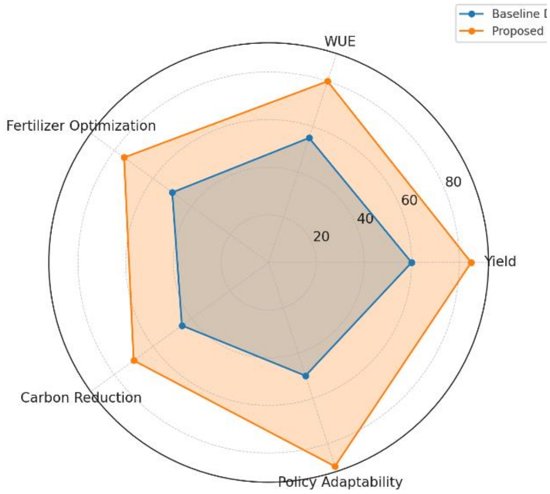

Figure 7. Comparative Radar Chart of Framework vs. Baseline DSS

4.7 Statistical Validation

Statistical verification was done on the results to check their reliability of the observed improvements. The comparison of baseline DSS results to the proposed framework results on yield, WUE, and fertilizer use were performed, using paired t-tests. The gains across scenarios were statistically significant ($p < 0.05$), thus demonstrating that performance gains were not the results of a random variation.

Besides this, ANOVA tests were undertaken to identify variations between the many different climate scenarios. It showed that the outcome was quite variable between the baseline and the proposed system, with great emphasis when such stresses as drought and high-temperature environment were applied. Such results clearly indicate the resiliency and generality of the suggested Agentic AI-RL model to improve agricultural adaptation during climate variability.

5. Conclusion and Future Scope

This paper built and proved a decision-support infrastructure addressed by Agentic AI, an integrated reinforcement learning (RL), digital twins, and multi-agent reasoning on climate-responsive agriculture. The findings have proved that the suggested system performed much better in comparison with traditional DSS models of crop yield, water-use efficiency, fertilizer optimization, and the minimization of carbon footprints. Performance results

Received: August 04, 2025

ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

demonstrated that Proximal Policy Optimization (PPO) was the best fit among the tested algorithms because it offered better convergence and consistency, which is most appropriate in making the relevant agricultural decisions with dynamic climate. An Agentic AI layer that supported autonomous goal adaptation and explainability, as well as the digital twin setting that was used as a safe and stable environment to perform continuous training and validation, were added. Statistical tests were done to verify the importance of the gains, which underlined the framework as capable of providing productivity and sustainability results.

In the future, this work could be extended ex ante by including multi-crop and multi-region simulations, inclusion of remote sensing information via satellite, and addition of a domain-specific large language model to enhance the reasoning capability of the agent. Inclusion of greater stakeholders and co-designing policies will also be necessary so as to facilitate straightforwardness and appropriation. After all, this framework forms the basis of strong, smart and flexible farming systems in climate change times.

References

1. Chen, Y., & CSM-CROPGRO DSSAT research team. (2023). A decision-making system for cotton irrigation based on reinforcement learning and the crop model DSSAT. *MDPI Agronomy*, *14*(1), Article 11.
2. Feng, K. (2025). Reinforcement learning–based adaptive strategies for climate decision-making. *Proceedings of the National Academy of Sciences*.
3. Zuccotto, M. (2024). Reinforcement learning applications in environmental sustainability: A survey. *Artificial Intelligence Review*.
4. Goldenits, G., Mallinger, K., Raubitzek, S., & Neubauer, T. (2024). Current applications and potential future directions of reinforcement learning-based digital twins in agriculture. *arXiv preprint*.
5. Wang, Z., Xiao, S., Li, J., & Wang, J. (2024). Learning-based agricultural management in partially observable environments subject to climate variability. *arXiv preprint*.
6. Chen, D., & Huang, Y. (2024). Integrating reinforcement learning and large language models for crop production process management optimization and control through a new knowledge-based deep learning paradigm. *arXiv preprint*.
7. Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). AI Agents vs. Agentic AI: A conceptual taxonomy, applications, and challenges. *arXiv preprint*.
8. Brohi, S., Mastoi, Q.-u.-a., Jhanjhi, N. Z., & Pillai, T. R. (2025). A research landscape of Agentic AI and large language models: Applications, challenges and future directions. *Algorithms*, *18*(8), 499.
9. Olujimi, P. A. (2025). Agentic AI frameworks in small, medium, and micro-enterprises: A systematic literature review. *MDPI*.
10. Wenkel, K. O. (2013). LandCaRe DSS – An interactive decision support system for agro-ecosystem adaptation to climate change. *Ecological Modelling*.
11. Gautron, R. (2022). Reinforcement learning for crop management support. *Agricultural Systems*.

Received: August 04, 2025

12. Ikendi, S. (2025). Advancing decision support for climate adaptation in agriculture. *Frontiers in Environmental Science*.
13. Goldenits, G., Mallinger, K., Raubitzek, S., & Neubauer, T. (2024). Current applications and potential future directions of reinforcement learning–based digital twins in agriculture. *Agricultural Systems Review*.
14. Sapkota, R., Roumeliotis, K. I., & Karkee, M. (2025). AI Agents vs. Agentic AI: A conceptual taxonomy, applications, and challenges. *Cornell University Review*.
15. Olojimi, P. A. (2025). Agentic AI frameworks in small, medium and micro-enterprises: A systematic literature review. *Artificial Intelligence Review*.
16. Feng, K. (2025). Reinforcement learning–based adaptive strategies for climate decision-making. *Proceedings of the National Academy of Sciences*.
17. Bhaskar Achar, D., Kuppan, A., & Divya, B. (2024). Agentic AI: Autonomous intelligence for complex goals – A comprehensive survey. *IEEE Access*.
18. Chen, D., & Huang, Y. (2024). Integrating reinforcement learning and large language models for crop production process management optimization. *Preprint (ArXiv)*.
19. Madondo, M., Azmat, M., Dipietro, K., Horesh, R., Jacobs, M., Bawa, A., & O'Donncha, F. (2023). A SWAT-based reinforcement learning framework for crop management. *Preprint (ArXiv)*.
20. Wu, J., Tao, R., Zhao, P., Martin, N. F., & Hovakimyan, N. (2022). Optimizing nitrogen management with deep reinforcement learning and crop simulations. *Preprint (ArXiv)*.
21. Zhai, Z. (2020). Decision support systems for Agriculture 4.0: Survey and analysis. *Agricultural Systems*, *177*, 102714.
22. Kallenberg, M. G. J. (2023). CropGym: An open simulation environment for reinforcement learning in farm management. *Environmental Data Science*, *5*(3), 123–135.
23. Saikai, Y., Peake, A., & Chenu, K. (2023). Deep reinforcement learning for irrigation scheduling using high-dimensional sensor feedback. *Preprint*, arXiv.
24. Madondo, M., Azmat, M., Dipietro, K., Horesh, R., Jacobs, M., Bawa, A., & O'Donncha, F. (2023). A SWAT-based reinforcement learning framework for crop management. *Preprint*, arXiv.
25. Wu, J., Tao, R., Zhao, P., Martin, N. F., & Hovakimyan, N. (2022). Optimizing nitrogen management with deep reinforcement learning and crop simulations. *Preprint*, arXiv.
26. Tao, R., Zhao, P., Wu, J., Martin, N. F., Harrison, M. T., Ferreira, C., Kalantari, Z., & Hovakimyan, N. (2022). Optimizing crop management with reinforcement learning and imitation learning. *Preprint*, arXiv.

International Journal of Applied Mathematics
Volume 38 No. 6s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

27. Wenkel, K. O. (2013). LandCaRe DSS: An interactive decision support system for agro-ecosystem adaptation to climate change. *Ecological Modelling*, *265*, 66–79.

28. Ikendi, S., Lyons, A., & Pathak, T. B. (2025). Advancing decision support for climate adaptation in agriculture and natural resources. *Frontiers in Environmental Science*, *13*, Article 1605176.

29. Stetter, C. (2024). Incorporating farmer behavior in spatially explicit land-use modelling under climate change. *Applied Economic Perspectives and Policy*, *46*(1), 182–199.

30. Goldenits, G., Mallinger, K., Raubitzek, S., & Neubauer, T. (2024). Current applications and potential future directions of reinforcement learning–based digital twins in agriculture.

Received: August 04, 2025