agronomy

MDPI

Article

Greenhouse Irrigation Control Based on Reinforcement Learning

Juan Pablo Padilla-Nates, Leonardo D. Garcia, Camilo Lozoya *, Luis Orona and Aldo Cortes-Perez

Tecnologico de Monterrey, School of Engineering and Science, Monterrey 64700, Nuevo Leon, Mexico; a01242223@tec.mx (L.D.G.)

* Correspondence: camilo.lozoya@tec.mx

Abstract

Precision irrigation provides a sustainable approach to enhancing water efficiency while maintaining crop productivity. This study evaluates a reinforcement learning approach, using the advantage actor–critic algorithm, for closed-loop irrigation control in a greenhouse environment. The reinforcement learning control is designed to regulate soil moisture near the maximum allowable depletion threshold, minimizing water use without compromising plant health. Its performance is compared against two common strategies: an on–off closed-loop controller and a time-based open-loop controller. The results show that the proposed controller consistently reduces irrigation water consumption relative to both benchmarks, while adapting effectively to environmental variability and the crop's increasing water demand during growth. These findings highlight the potential of reinforcement learning to achieve a more efficient balance between water conservation and crop health in controlled agricultural systems.

Keywords: smart farming; controlled environment agriculture; advantage actor–critic; sensor-based irrigation; water-use efficiency

check for updates

Academic Editor: Wei Wu

Received: 8 October 2025
Revised: 23 November 2025
Accepted: 28 November 2025
Published: 2 December 2025

Citation: Padilla-Nates, J.P.; Garcia, L.D.; Lozoya, C.; Orona, L.; Cortes-Perez, A. Greenhouse Irrigation Control Based on Reinforcement Learning. *Agronomy* **2025**, *15*, 2781. https://doi.org/10.3390/agronomy15122781

Copyright: © 2025 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

1. Introduction

According to the United Nations, the global population is projected to reach 9 billion by 2050, posing significant challenges for food production and sustainable resource management [1]. Agriculture must evolve to meet this demand while operating within the limits of increasingly scarce natural resources, particularly fresh water. In this context, greenhouse cultivation has gained prominence as a controlled-environment agriculture technique that enables higher crop yields with lower resource consumption compared to open-field farming [2].

One of the major resource-intensive processes in greenhouse farming is irrigation. Conventional irrigation strategies, such as time-based or manual scheduling, often result in water waste due to suboptimal soil-moisture levels, which can affect crop health and productivity. Precision irrigation is an approach that integrates sensors, data analysis, and control algorithms to enhance water-use efficiency without compromising plant growth and development. Improving water efficiency involves applying the liquid in real-time according to crop requirements. This can be achieved by controlling the root zone soil moisture between the specific thresholds of the field capacity and permanent wilting point [3].

Closed-loop irrigation, which utilizes feedback to provide crops with precise water amounts at optimal times, has been proposed as a solution for improving efficient crop water management [4]. Irrigation systems continuously monitor soil-moisture levels using sensors and trigger irrigation only when moisture drops below a predefined threshold. This ensures that crops receive water only when needed, while irrigation is halted before

overwatering occurs, avoiding water waste and maintaining moisture within optimal ranges for plant growth. In this context, different control strategies have been implemented, achieving water savings from 10% up to 40% [5]. However, typically, the predefined moisture threshold remained fixed during the complete irrigation period, without considering the different crop growth stages.

Recent advances in artificial intelligence, particularly reinforcement learning, offer new opportunities for developing adaptive irrigation systems that can learn optimal control policies through interaction with the environment [6]. Unlike traditional control methods, reinforcement learning controllers can dynamically adjust to changing environmental conditions and plant needs, making them ideal for complex, nonlinear systems such as crop irrigation.

The work presented in this paper investigates the application of a reinforcement learning approach, specifically the advantage actor–critic algorithm, for automated irrigation control in a greenhouse environment. The reinforcement learning control is evaluated in terms of its ability to maintain soil moisture near the maximum allowable depletion threshold while minimizing water usage. Its performance is compared with that of two commonly used approaches: an on–off closed-loop controller and a time-based open-loop controller. The findings suggest that reinforcement learning can significantly enhance water efficiency and moisture regulation in greenhouse operations, contributing to the broader goal of sustainable agriculture.

Crop irrigation has traditionally relied on predefined schedules or manual decision-making processes, often leading to inefficient water use and inconsistent soil-moisture levels. To address these limitations, several studies have explored the use of sensor-based systems and automatic controllers to enhance water-use efficiency and maintain optimal growing conditions.

One common approach is time-based irrigation, which applies water at fixed intervals regardless of actual soil moisture or plant needs. While easy to implement, this method frequently results in overwatering or underwatering, especially under dynamic environmental conditions [7]. Closed-loop control systems, such as on–off and proportional–integral–derivative (PID) controllers, have demonstrated improved performance by adjusting irrigation based on real-time sensor data [8]. Also, model predictive control (MPC) has demonstrated its applicability and efficiency in irrigation systems [9–11]. However, these classical control methods require manual tuning and often struggle with system nonlinearities and time delays inherent in plant–soil–water interactions.

Artificial intelligence techniques have increasingly been explored to address the limitations of conventional irrigation management in agriculture [12]. In recent years, advances in precision agriculture have promoted the integration of machine learning methods into irrigation decision-making. Supervised learning approaches—such as dynamic neural networks—have been applied to soil-moisture prediction and irrigation scheduling optimization [13–15]. While these models may perform well under controlled conditions, they require substantial labeled datasets and often exhibit limited adaptability to dynamic and heterogeneous field environments.

Reinforcement learning (RL) has emerged as a promising alternative because it enables agents to learn optimal control actions through trial-and-error interactions with the environment [16]. RL establishes a closed-loop decision-making framework in which cumulative rewards guide the system toward improved performance over time [17]. RL methods are commonly grouped into value-based, policy-based, and hybrid actor–critic approaches. Value-based algorithms (e.g., Q-learning) have been widely explored in irrigation applications due to their simplicity, model-free nature, and robustness when soil–plant–atmosphere interactions are difficult to characterize [18,19]. Actor–critic methods provide a balance

between value estimation and policy optimization and offer the important advantage of generating continuous actions, making them suitable for greenhouse systems where irrigation amounts can be precisely adjusted.

Despite these advantages, the application of RL in greenhouse irrigation remains limited. Existing work has primarily focused on greenhouse climate control [20,21] and crop management and decision optimization [22,23], mostly within simulated or digital-twin environments. Only a small number of studies have addressed RL-based closed-loop irrigation control, and very few include experimental validation under realistic greenhouse conditions or across a complete crop growth cycle.

Previous research has mostly concentrated on methodological or simulation-based contributions. For instance, [24,25] described general principles for implementing RL irrigation strategies but did not provide experimental evidence. In [26], irrigation decisions were evaluated exclusively through simulations supported by limited field data and without consideration of specific crop requirements. Similarly, [19] developed an RL approach for open-field rice irrigation guided by weather forecasts; however, the study was purely simulation-based.

Experimental studies are still scarce. In [27], an RL controller was tested on almond trees over a short 15-day period, achieving 9.52% water savings. More recently, [28] conducted a greenhouse experiment and reported a 14% reduction in water use for leafy vegetables using an RL-based irrigation algorithm. Nevertheless, the authors noted that computational complexity remains a barrier for practical deployment on low-cost hardware platforms.

The present work addresses these limitations by experimentally evaluating an actor–critic irrigation controller in a greenhouse environment and comparing its performance with established irrigation strategies over a complete crop production cycle. This study provides experimental evidence, practical insights, and performance benchmarks that support the realistic adoption of reinforcement learning for greenhouse irrigation, emphasizing the feasibility of implementation on affordable hardware using a lightweight algorithm.

2. Materials and Methods

2.1. Crop Irrigation Dynamics

The irrigation process in a greenhouse is governed by the interaction between soil moisture, crop, and environmental variables. Its dynamics follow the water balance principle, which is defined by the following discrete-time equation:

$$\theta_{k+1} = \theta_k + ir_k - ETc_k - dp_k \tag{1}$$

where θ_k is the volumetric soil moisture at time k, ir_k is the irrigation water applied, ETc_k is the crop evapotranspiration, and dp_k represents deep percolation or drainage losses.

Crop evapotranspiration is the combined water loss due to soil evaporation and plant transpiration. It depends on climatic variables (temperature, relative humidity, solar radiation, and wind speed) and crop characteristics (growth stage, leaf area, and canopy development). Actual crop evapotranspiration is difficult to measure, but it can be approximated by

$$ETc_k = Kc\, ETo_k, \tag{2}$$

where ETo_k is the reference evapotranspiration, which depends only on climatic variables according to the Penman–Monteith equation, as described by [29], and Kc is the crop coefficient that depends mainly on crop type and growth stage.

Soil moisture is the amount of water contained in the soil. It refers to the water held in the spaces between soil particles and is a critical factor for plant growth, as it determines

how much water is available for roots to absorb. In precision irrigation, soil moisture is typically measured as volumetric water content θ_k, which is defined as the ratio of the volume of water in the soil to the total volume of the soil sample at time k.

According to the volumetric water content (VWC), soil moisture in an irrigation process can be classified into three regions:

- **Gravitational:** Soil is oversaturated with water, waterlogging may be visibly present in the crop, and the water drains quickly due to gravity.
- **Available water:** In this region, water is available for crop roots. This region is delimited by the field-capacity (FC) and the permanent-wilting-point (PWP) thresholds.
- **Unavailable water:** In this region, water is not available to crops, and plants suffer from severe water stress.

FC is the amount of water content retained by the soil after excess gravitational water has drained away; it represents the upper limit of the available water region for the plants. PWP represents the lower limit of this region, below which plants can no longer extract water from the soil. The maximum allowable depletion (MAD) is the fraction of the available water in the root zone that can be depleted before irrigation is required to avoid any plant water stress. MAD defines how much of that available water can be used before irrigation should be applied. If the soil water content drops below the MAD level, plants may start to experience stress, which can reduce crop yield. Typically, the MAD value represents 50% of the available water.

Irrigation is inherently nonlinear, as the relationship between applied water volume and soil-moisture dynamics depends on factors such as soil texture, compaction, and crop water needs. Additional uncertainties arise from climatic variability (e.g., solar radiation and temperature), heterogeneous plant growth, and sensor measurement noise. Conventional irrigation schedules, such as fixed-time-based approaches, often overlook these dynamic interactions, resulting in inefficient water use or crop stress. Closed-loop strategies that rely on soil moisture feedback improve efficiency, yet their tuning remains challenging under nonlinear and stochastic conditions. Reinforcement learning provides a data-driven alternative, capable of adapting to complex dynamics and uncertainties without requiring explicit system modeling.

2.2. Reinforcement Learning Control

Reinforcement learning (RL) is a specific branch of machine learning, distinct from supervised and unsupervised learning approaches. In supervised learning, a dataset labeled by an expert supervisor is given to the algorithm. This label refers to a correct action that the system should perform in a specific situation. In unsupervised learning, this approach identifies hidden patterns or clusters within unlabeled data. Reinforcement learning aims to maximize the value of a reward; it neither identifies a pattern nor knows a priori which action is right or wrong.

One of the significant advantages of reinforcement learning algorithms is their ability to improve performance continually. Once exposed to an environment, the RL control, known as a policy, adapts to changes in its process according to predictions it makes about which action provides the largest reward. Precision irrigation can benefit from this approach. The lack of an accurate model and the unpredictable environmental changes in the soil can be addressed with the right tools for more robust irrigation control. One such algorithm that provides a solution to this problem is the actor–critic approach, which employs an agent called an actor to decide on actions and a critic to evaluate them. Specifically, the advantage actor–critic uses the relative performance of the action (advantage) as the learning signal, leading to faster and more stable learning.

As shown in Figure 1, the RL agent–environment relationship implements a closed-loop control system. However, unlike the traditional control approach, whose control law is deterministic, RL control is stochastic by nature. The environment is the irrigation process, while observations are conducted through soil moisture sensors. The policy implements the irrigation schedule, and the reinforcement learning algorithm continuously updates the policy according to the current state (s_k), action (a_k), and reward (r_k) conditions.

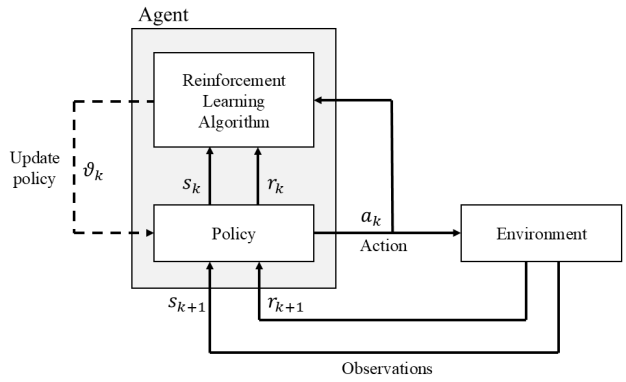

Figure 1. Reinforcement learning model architecture implementing the advantage actor–critic closed-loop control for the greenhouse irrigation process.

For the implementation of the RL control, a Markov decision process (MDP) is considered as the mathematical framework. MDP describes decision-making in the environment under the control of an agent. MDP is formally defined as

$$MDP = (\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R}), \tag{3}$$

where $\mathcal{S}$ is the set of all possible states that the process can be in. For the irrigation process, the system states represent the soil moisture regions:

$$\mathcal{S} = \{Gravitational, Available, Unavailable\}. \tag{4}$$

$\mathcal{A}$ is the space of all actions that are permitted in the dynamic process. Two possible actions are available for the system: irrigate or do not irrigate.

$$\mathcal{A} = \{ir_k, \bar{ir}_k\}. \tag{5}$$

$\mathcal{P}$ is the probability function that determines the numerical degree of certainty that a state s_k will transition to a state s_{k+1} after taking action a_k. The probability dynamics that describe how the state evolves after an action is taken can be defined as

$$p(s_{k+1} \mid s_k, a_k). \tag{6}$$

$\mathcal{R}$ is the space defining the immediate reward that the agent receives after transitioning from state s_k to state s_{k+1} when taking action a_k. The reward is designed to encourage water savings while keeping soil moisture within the allowable range defined by FC and PWP. Therefore, a Gaussian curve can be used to represent the reward as

$$r_k = \frac{1}{\sigma\sqrt{2\pi}} \exp\left[-\frac{1}{2}\left(\frac{\theta_k - \mu}{\sigma}\right)^2\right], \tag{7}$$

where σ is the standard deviation, μ is the mean value, and θ_k is the current soil-moisture value. The graphical relation between the soil-moisture regions and the reward is depicted in Figure 2.

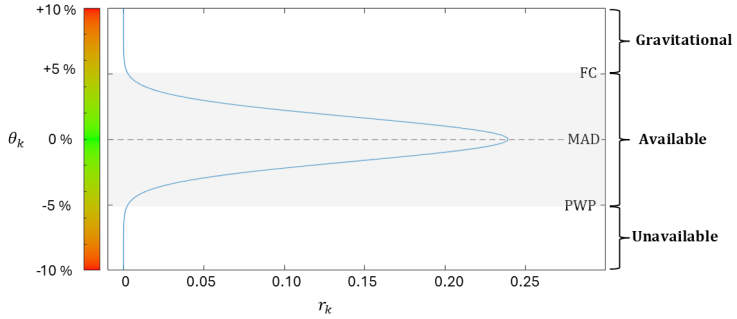

Figure 2. Reward values r_k according to the different soil-moisture regions (gravitational, available water, and unavailable water), where θ_k indicates the actual soil-moisture value. Notice that the maximum reward is reached at the maximum allowable depletion (MAD), while at the field capacity (FC) and at the permanent wilting point (PWP), the reward is practically zero.

2.3. Advantage Actor–Critic Algorithm

The advantage actor–critic (A2C) algorithm is a reinforcement learning algorithm that combines the strengths of policy-based and value-based approaches. It is particularly useful for continuous control problems, such as irrigation scheduling, where the decision on how much water to apply must adapt to nonlinear and uncertain dynamics.

The actor represents the policy $\pi(a|s, \vartheta)$, which maps the current state s_k (current soil-moisture region) to a probability distribution over irrigation actions a_k. The actor decides what action to take, then its policy parameters ϑ are updated using feedback from the critic. The critic approximates the value function $\hat{v}(s, \mathbf{w})$, which estimates the expected long-term rewards from state s_k. The critic evaluates how good the action was, given the policy. Its weight parameters $\mathbf{w}$ are updated to minimize the temporal difference error.

The temporal difference (TD) error indicates how much the current estimate needs to be adjusted based on the latest information and is defined as

$$\delta = r_{k+1} - \bar{r} + \gamma\, \hat{v}(s_{k+1}, \mathbf{w}) - \hat{v}(s_k, \mathbf{w}), \tag{8}$$

where δ is the TD error, $\gamma \in [0, 1]$ is the discount factor that determines how much future rewards matter compared to immediate rewards, and $\bar{r}$ is the reward baseline. Then, the baseline is updated with the TD error as

$$\bar{r} = \bar{r} + \alpha_{\bar{r}}\delta, \tag{9}$$

where $\alpha_{\bar{r}}$ is the learning rate parameter. Updating $\bar{r}$ helps stabilize the learning process by normalizing rewards and reducing variance in policy updates.

The A2C algorithm utilizes eligibility traces for both the critic and the actor. An eligibility trace is a short-term memory vector that records which states and actions contributed to the current estimated reward and is updated after each time step. The use of traces contributes to improving efficiency when the system operates in a continuously changing environment.

The critic's eligibility trace $\mathbf{z_w}$ keeps track of the history of state visits, helping propagate the TD error back to states that contributed to it, and is obtained as

$$\mathbf{z_w} = \gamma\lambda_{\mathbf{w}}\mathbf{z_w} + \nabla_{\mathbf{w}}\hat{v}(s_k, \mathbf{w}), \tag{10}$$

where $\lambda_{\mathbf{w}} \in [0, 1]$ controls how long past states remain eligible and $\nabla_{\mathbf{w}}$ is the gradient of the value function used by the critic to improve its estimation accuracy. The TD error δ is used to update the value function weights to improve the critic's estimate of the state values as

$$\mathbf{w} = \mathbf{w} + \alpha_{\mathbf{w}}\delta\mathbf{z_w}. \tag{11}$$

The actor's eligibility trace $\mathbf{z}_\vartheta$ accumulates the gradients of the policy, allowing for more effective updates that consider the influence of past actions. The actor's eligibility trace describes how historical corrections to the policy π by means of ϑ have been affecting its performance. Hence,

$$\mathbf{z}_\vartheta = \gamma\lambda_\vartheta \mathbf{z}_\vartheta + \nabla_\vartheta \ln \pi(a_k \mid s_k, \vartheta), \tag{12}$$

where $\lambda_\vartheta \in [0, 1]$ is the decay factor (same as in the critic trace, balances short- vs. long-term credit) and ∇_ϑ is the policy gradient. Then, the TD error δ is used to update the policy parameters to favor actions that lead to higher returns as

$$\vartheta \leftarrow \vartheta + \alpha_\vartheta \delta \mathbf{z}_\vartheta. \tag{13}$$

The A2C with eligibility traces is particularly effective in environments where the agent needs to adapt continuously and improve its decision-making process based on a comprehensive understanding of past and present states. The pseudocode for the algorithm, incorporating all the elements previously described, is presented in Algorithm 1. Its inputs are the function approximation parametrization and the hyperparameters. At each step, it interacts with the environment and updates the weights and eligibility trace vectors.

Algorithm 1: Advantage actor–critic with eligibility traces algorithm

Input: $\hat{v}(s, \mathbf{w})$
Input: $\pi(a \mid s, \vartheta)$

`// Declare hyperparameters.`
Data: $\alpha_\mathbf{w} > 0, \alpha_\vartheta > 0, \alpha_{\bar{r}} > 0$
Data: $\lambda_\mathbf{w} \in [0, 1], \lambda_\vartheta \in [0, 1]$
Data: $\gamma \in [0, 1]$

`// Initialize the vectors and the baseline to initial values.`
$\bar{r} \leftarrow r_0$;
$\mathbf{z}_\mathbf{w} \leftarrow \mathbf{0}$;
$\mathbf{z}_\vartheta \leftarrow \mathbf{0}$;
$\mathbf{w} \leftarrow \mathbf{w}_0$;
$\vartheta \leftarrow \vartheta_0$;

while *True* **do**
  `/* Select an action from the policy and observe the next state and reward. */`
  $a_k \sim \pi(\mathcal{A} \mid s_k, \vartheta)$;
  $s_{k+1}, r_{k+1} \leftarrow \text{Environment}(s_k, a_k)$;
  `/* Compute the TD error and update the reward baseline. */`
  $\delta \leftarrow r_{k+1} - \bar{r} + \gamma\hat{v}(s_{k+1}, \mathbf{w}) - \hat{v}(s_k, \mathbf{w})$;
  $\bar{r} \leftarrow \bar{r} + \alpha_{\bar{r}}\delta$;
  `/* Update the eligibility trace vectors. */`
  $\mathbf{z}_\mathbf{w} \leftarrow \gamma\lambda_\mathbf{w}\mathbf{z}_\mathbf{w} + \nabla_\mathbf{w}\hat{v}(s_k, \mathbf{w})$;
  $\mathbf{z}_\vartheta \leftarrow \gamma\lambda_\vartheta\mathbf{z}_\vartheta + \nabla_\vartheta \ln \pi(a_k \mid s_k, \vartheta)$;
  `/* Update the weights of the function approximations. */`
  $\mathbf{w} \leftarrow \mathbf{w} + \alpha_\mathbf{w}\Delta\mathbf{z}_\mathbf{w}$;
  $\vartheta \leftarrow \vartheta + \alpha_\vartheta\Delta\mathbf{z}_\vartheta$;
end

3. Results

3.1. Experimental Setup

The experimental greenhouse was equipped with integrated climate control, irrigation control, and soil-moisture monitoring systems, as illustrated in Figure 3. Climate regulation was achieved through an evaporative cooling system consisting of a wet pad and extraction fans, a standard approach for hot and arid regions such as northern Mexico. For irrigation, three independently controlled drip lines were installed, each capable of supplying water to up to 14 green pepper plants. The following elements composed the automation system:

Soil-moisture module: Nine soil-moisture sensors ECH2O EC-5 (METER Group Inc., Pullman, WA, USA) were deployed to measure the volumetric water content, with three sensors assigned to representative plants on each irrigation line. Data acquisition was carried out using an ESP32 microcontroller (Espressif Systems, Shanghai, China), which enabled real-time monitoring and integration with the irrigation control system.

Irrigation module: This module manages the actuation of the three irrigation electrovalves and measures the corresponding water flow using YF-S201 sensors (Digiten, Shenzhen, China). The system is implemented on an ESP32 microcontroller, which enables precise valve control, flow-data acquisition, and integration with the overall irrigation management framework.

Climate module: The climate monitoring module integrates three sensors: a solar-radiation PYR sensor (Apogee Instruments, Logan, UT, USA), a Davis Cup wind-speed sensor (Davis Instruments Corp., Hayward, CA, USA), and a combined Atmos 14 air temperature–relative humidity sensor (METER Group Inc., Pullman, WA, USA). In addition to data acquisition, the module controls the wet-pad pump and the extraction fans to regulate greenhouse climate conditions. The system is implemented on an ESP32 microcontroller, providing both sensor integration and actuator management within a unified platform.

Controller module: The module executes the three evaluated irrigation control algorithms: (1) time-based control, (2) on–off control, and (3) reinforcement learning-based control. It is implemented on a Raspberry Pi 4 single-board computer (Raspberry Pi Ltd., Cambridge, UK), running Python 3.12 (Python Software Foundation, Wilmington, DE, USA), providing sufficient computational resources for real-time algorithm execution, data processing, and integration with the greenhouse monitoring and actuation systems.

Wireless sensor network: A WiFi communication network is established using a publish–subscribe messaging model implemented via the MQTT protocol (OASIS Open, Burlington, MA, USA), a standard widely used in IoT (Internet-of-Things) applications. This lightweight and flexible protocol enables real-time, asynchronous data exchange between all modules, allowing any device to transmit or receive information at any time with minimal overhead.

Web server: A computer hosts Node-RED v3.1.6 services (OpenJS Foundation, San Francisco, CA, USA), providing a web-based interface that allows users to monitor and interact with the greenhouse systems in real time. This platform enables intuitive visualization, data logging, and system setup through a graphical, browser-accessible dashboard.

The jalapeño green pepper (*Capsicum annuum* L.) was used to evaluate the three irrigation control algorithms. This crop requires precise water management, as excessive irrigation increases its susceptibility to diseases, while insufficient irrigation limits plant growth and development. The jalapeño pepper is among the most extensively cultivated and economically relevant chili cultivars in Mexico. In greenhouse conditions with controlled temperature, humidity, and irrigation, the first harvest can be expected around 65–70 days after transplanting. Greenhouse production systems offer several advantages over open-field cultivation, including the potential for continuous year-round production,

increased yields, enhanced fruit quality, and a reduced incidence of pests and diseases. Crop development begins with seed germination, followed by the seedling stage, which takes place in nursery trays. Once seedlings reach adequate physiological maturity, they are transplanted by relocating the young plants to their final growing containers, where the choice of substrate is crucial for optimal establishment. Substrates employed in greenhouse horticulture must ensure appropriate aeration, water-holding capacity, drainage, and nutrient availability. Table 1 specifies the physical properties of the substrate used during the evaluation to provide a balanced medium for plant growth.

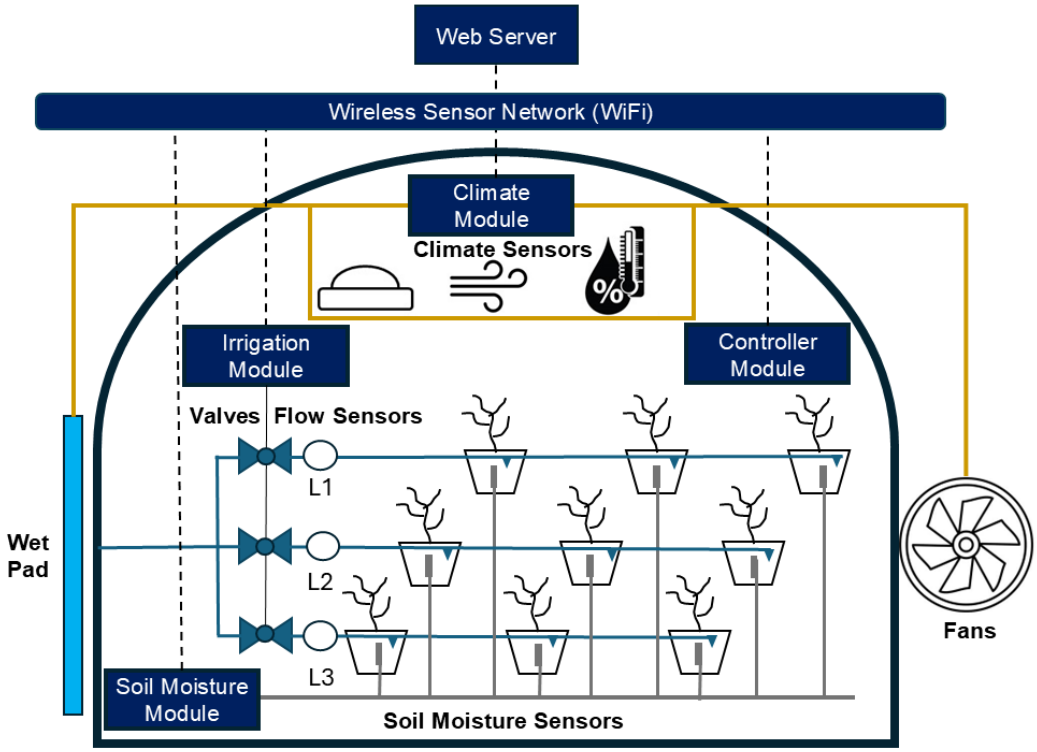

Figure 3. Automated greenhouse system composed of four modules: soil-moisture measurement, irrigation control, climate control, and a controller module that simultaneously executes the three irrigation control algorithms under evaluation. A WiFi network implements the communication for the modules and the web server.

Table 1. Physical properties of the substrate mixture and container specifications for jalapeño pepper cultivation.

Parameter	Value/Range	Description
Substrate composition	1:1:1 ($v/v/v$)	Loamy soil: peat moss: perlite
Total porosity	65–75%	Typical range for greenhouse horticultural media
Water-holding capacity	45–55%	Volume of water retained at container capacity
Air-filled porosity	15–25%	Air volume after drainage
Container volume per plant	8 L	Common volume for pepper in greenhouse cultivation

The experiment lasted 10 weeks, including a 2-week adaptation phase, followed by 8 weeks of evaluation. It began immediately after transplantation and concluded with the first harvest. The adaptation phase was devoted to plant establishment, fertilization, and acclimation to greenhouse conditions. During the subsequent 8-week evaluation period, the three irrigation control algorithms were applied. This interval encompassed the

crop's principal developmental stages—vegetative growth, flowering, and fruit development—during which water demand increases markedly due to canopy expansion and fruit set. Insufficient irrigation in these stages has been associated with reduced biomass accumulation and significant yield losses. The crop growth timeline, along with the evaluation period, is illustrated in Figure 4.

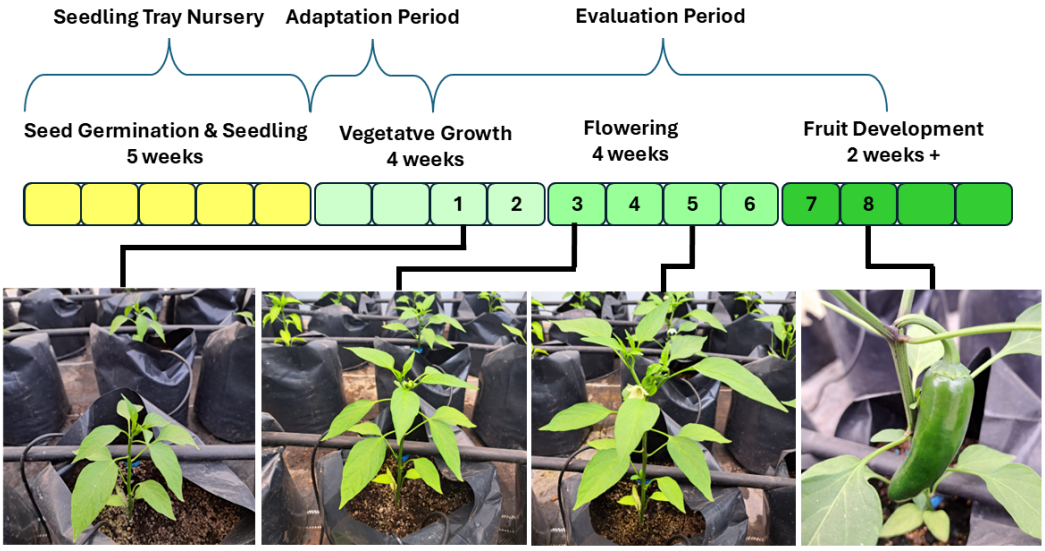

Figure 4. Jalapeño green pepper growth timeline. Images illustrate the plant's growth at specific timestamps: week 1 (late vegetative growth stage), week 3 (initial flowering stage), week 5 (late flowering stage), and week 8 (initial fruit development stage).

The assessment considered both water consumption and crop development, ensuring healthy plant growth. Each control strategy was implemented in a separate drip irrigation line equipped with emitters delivering 8 L per hour per dripper. Environmental conditions inside the greenhouse were maintained at a set-point temperature of 28 °C, and an average relative humidity of 40%.

Soil moisture was monitored using nine sensors, which had been previously calibrated against gravimetric measurements of the experimental substrate to ensure accuracy. For each irrigation line, three sensors were installed in independent pots and placed at the root zone of representative plants, providing aggregated values for soil-moisture estimation [30]. The field-capacity and permanent-wilting-point thresholds were determined during the week prior to the evaluation period. To this end, the substrate was saturated with water and allowed to drain under gravity until the volumetric water content stabilized, approximately 48 h after saturation, following the procedure described by [31].

The maximum allowable depletion was set at 50% [32]. Plants used during this calibration phase were not included in the subsequent experimental evaluation. Volumetric water content measurements were normalized with respect to the MAD threshold. In this scheme, positive values indicated VWC above MAD, whereas negative values indicated VWC below MAD.

The time-based scheduling control was programmed to irrigate twice daily for five minutes per event, supplying a total of 1.33 L per plant per day. This irrigation volume was determined based on local growers' practices, who consider it adequate for jalapeño pepper cultivation under greenhouse conditions. The on–off control strategy was configured to maintain VWC within a range of 0–5% above MAD to avoid any type of crop water stress [33]. For the reinforcement learning control, hyperparameters were tuned through preliminary simulations, where multiple scenarios were evaluated with respect to learning rate, settling time, reward function performance, oscillatory behavior, and convergence

performance. Simulations were conducted in MATLAB using the crop irrigation model defined by [9] to obtain adequate values for the learning rates α_*, discount factor γ, and decay factors λ_*:

$$\alpha_{\mathbf{w}} = 0.0075,\ \alpha_{\vartheta} = 0.0750,\ \alpha_{\bar{r}} = 0.0250, \tag{14}$$

$$\gamma = 1.0000,\ \lambda_{\mathbf{w}} = 0.4000,\ \lambda_{\vartheta} = 0.4000. \tag{15}$$

The discount factor was set to its maximum value, $\gamma = 1$, to fully account for future observations of the process. Stability analysis showed that oscillations are avoided when the λ_* parameter remains below 0.7500. The learning rate for the reward baseline, $\alpha_{\bar{r}}$, provided stable performance around a value of 0.0250. The weight-update learning rate, $\alpha_{\mathbf{w}}$, was effective within the range of 0.0040–0.0080. Finally, setting the policy learning rate, α_{ϑ}, to a value approximately ten times greater than $\alpha_{\mathbf{w}}$ enabled the policy to adapt faster to environmental dynamics.

3.2. Experimental Results

Figure 5 presents the soil-moisture dynamics at the root zone, expressed as VWC, for each irrigation control strategy. The results show that both the on–off control and the reinforcement learning algorithm maintained soil moisture below the FC threshold, i.e., within +5% of MAD. In contrast, the time-based control consistently moved soil moisture into the gravitational water region, leading to excessive drainage and water loss. In all cases, soil moisture remained above the PWP threshold at –5%, thereby preventing water stress. The RL control exhibited moisture peaks during the first two weeks, reflecting the initial learning phase; however, after this period, the algorithm successfully stabilized irrigation events and avoided oversaturation in the gravitational water region. Overall, the RL approach demonstrated a superior balance between water-use efficiency and plant safety, outperforming the time-based strategy and matching or exceeding the stability of the on–off control.

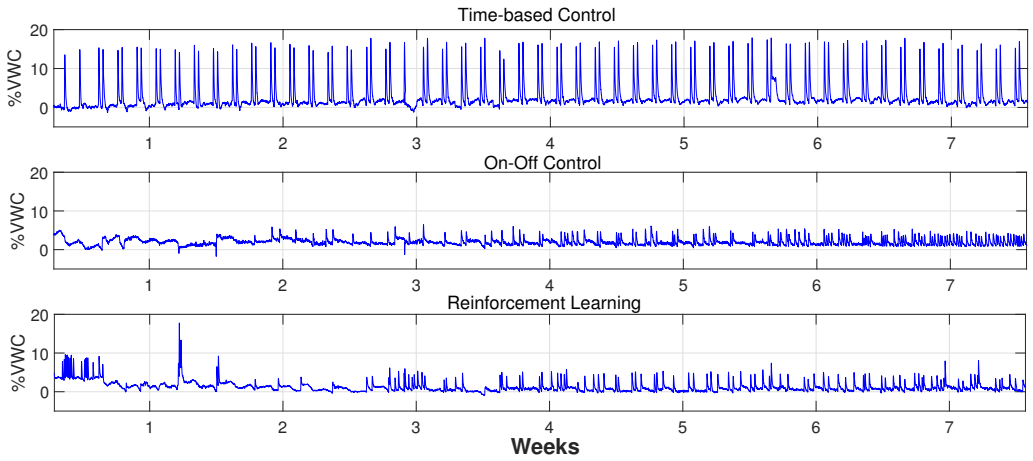

Figure 5. Comparison of soil-moisture dynamics, measured as a percentage of volumetric water content (%VWC) for the three evaluated irrigation control methods.

To assess irrigation accuracy, the root mean square error (RMSE) was calculated to quantify how closely soil moisture was maintained relative to the MAD threshold. Deviations above this level indicate water loss due to over-irrigation, while deviations below it signal the onset of crop water stress. Figure 6 illustrates the irrigation accuracy achieved by the three evaluated methods. The time-based control exhibited the highest error, stabilizing around 5% after two weeks. The on–off control showed an initial peak in error but subsequently maintained a steady deviation of approximately 3%. The reinforcement

learning control demonstrated the highest accuracy, with a progressive reduction in RMSE throughout the evaluation period, ultimately converging to values close to 2%. These results indicate that RL not only ensures plant safety but also improves water-use efficiency by minimizing both excess application and the risk of stress.

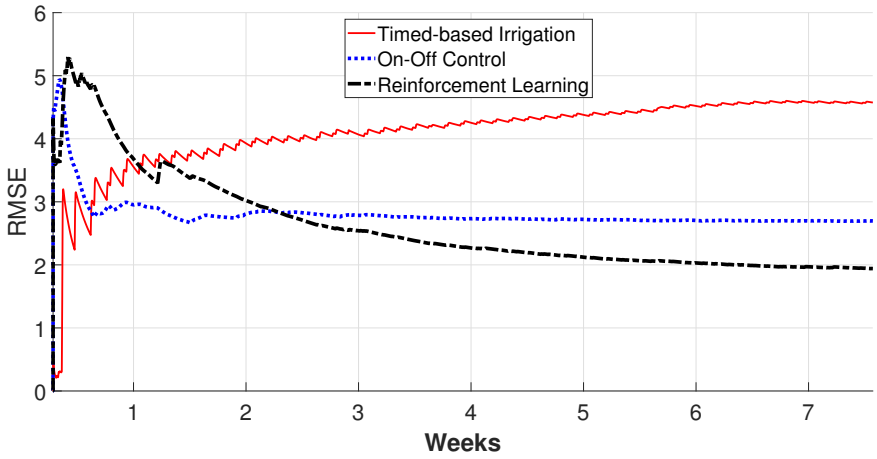

Figure 6. Crop irrigation accuracy measured as the root mean square error (RMSE) from the soil-moisture reference defined by the maximum allowable depletion (MAD).

Figure 7 presents the irrigation events recorded during the evaluation period, along with the reference evapotranspiration dynamics, to examine how ambient conditions influenced the response of the three control algorithms. The time-based control applied two irrigation events per day, independent of ETo fluctuations or crop growth stage. The on–off control exhibited some sensitivity to ambient conditions, increasing irrigation frequency when ETo exceeded 5 mm/day, particularly evident during weeks 3 and 5. However, the most pronounced adjustments were associated with crop development; during weeks 6–8, corresponding to late flowering and fruit set, irrigation events increased substantially. The reinforcement learning (RL) strategy also adjusted irrigation in response to crop growth, as well as environmental conditions, to a lesser extent. Unlike the on–off control, the RL control produced more gradual changes in irrigation frequency, thereby avoiding abrupt increases in water application. Overall, the RL control achieved a more adaptive and balanced response, outperforming the rigid behavior of the time-based control and the reactive nature of the on–off strategy.

Water consumption was adopted as the primary performance metric, since the main objective of the evaluation was to minimize irrigation while maintaining comparable crop yield. Figure 8 presents the accumulated water consumption over the evaluation period. The time-based control served as the reference, as its water use could be predetermined a priori from the fixed schedule. The on–off controller initially achieved water savings, remaining consistently below the reference for the first five to six weeks. However, its accumulated consumption increased during the final two weeks, nearly converging with the time-based strategy, and resulted in negligible net savings by the end of the trial. By contrast, the RL control exhibited higher consumption in the first two weeks, reflecting its exploration and adaptation phase. From week three onward, the RL control maintained a consistently lower rate of water use, ultimately achieving substantial reductions relative to both benchmarks: 36.9% compared to the time-based control and 31.5% compared to the on–off control. These results demonstrate the superior potential of RL for enhancing water-use efficiency and advancing the sustainability of greenhouse crop production.

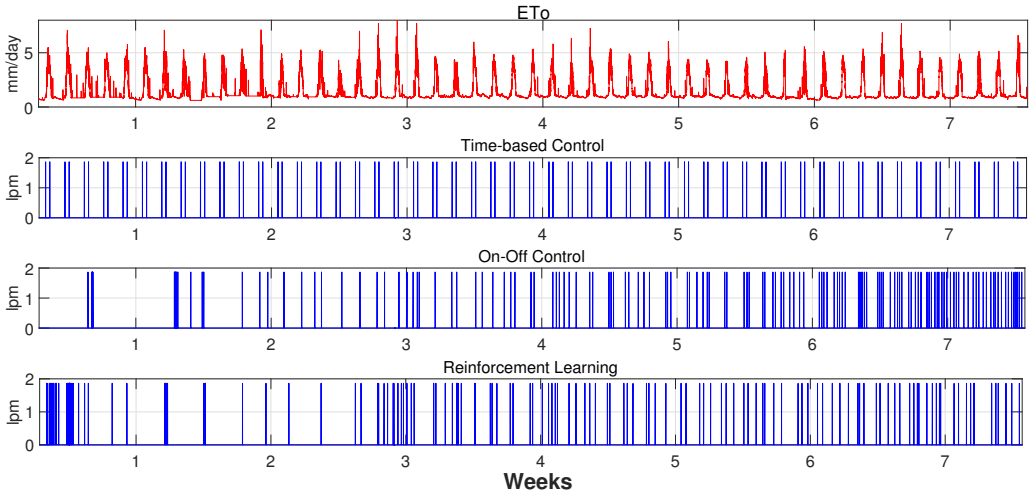

Figure 7. Irrigation events, expressed in liters per minute, generated by applying the three control algorithms. In addition, the reference evapotranspiration (ETo) was included to establish the relationship between the prevailing climatic conditions and the corresponding irrigation patterns.

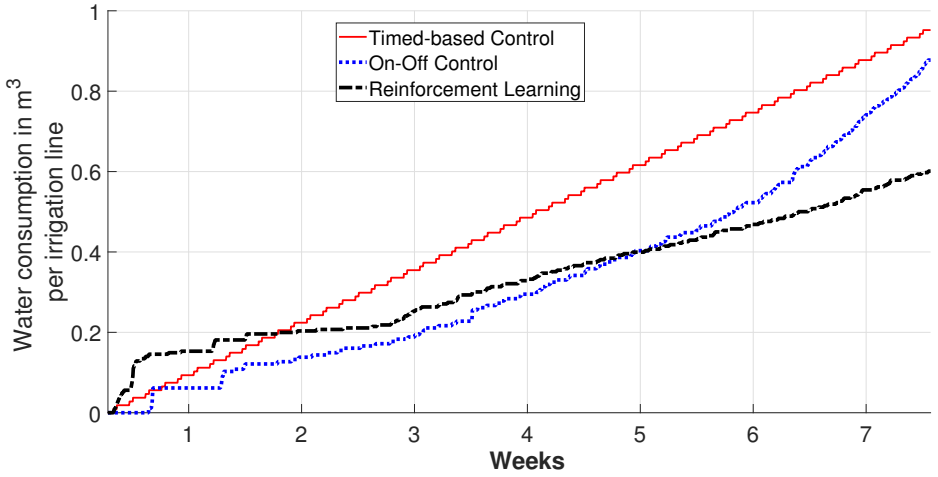

Figure 8. Accumulated water consumption (m^3) for each irrigation line during the evaluation period of the three control algorithms. Each algorithm was tested on a dedicated irrigation line containing 14 plants.

In terms of crop development, all three irrigation strategies supported healthy and vigorous plant growth. Across treatments, plants reached an average height of approximately 80 cm—within the expected range for mature jalapeño peppers—and produced a vegetative dry weight of roughly 120 g per plant (excluding fruits). These consistent values indicate that the irrigation approaches were similarly effective in sustaining normal growth and biomass accumulation, with no observable reduction in plant vigor. Consequently, crop yield remained comparable across treatments, demonstrating that the reduced-water RL strategy did not negatively affect productivity.

4. Discussion

The results demonstrated clear differences among the three irrigation strategies in their ability to regulate soil water content. The time-based control produced soil moisture that frequently entered the gravitational water region above the FC threshold, promoting deep percolation losses and inefficient water use. By contrast, both the on–off and reinforcement learning strategies maintained soil moisture below FC and within the efficient soil-moisture range, bounded by MAD, avoiding any water stress. Notably, the RL con-

trol exhibited initial fluctuations during the learning phase but progressively stabilized, ultimately achieving a consistent balance between water conservation and crop health.

A quantitative assessment, as measured by the RMSE, highlighted the limitations of time-based scheduling, which resulted in persistent deviations from the target range, averaging 5%. The on–off strategy performed better, maintaining a steady error of 3% after initial adjustment. However, the RL strategy progressively refined its control, reducing error over time and converging near 2%. This improvement indicates that RL not only adapts irrigation timing but also learns to minimize both over-irrigation and under-irrigation, keeping soil moisture more consistently within an adequate range.

The irrigation event analysis provided further insight into the responsiveness of control. The time-based method applied water rigidly, regardless of evapotranspiration or growth stage, confirming its lack of adaptability. The on–off control showed partial responsiveness, increasing irrigation during periods of high ETo and at later growth stages, though often through abrupt shifts in frequency. In contrast, the RL algorithm exhibited a smoother adjustment process, aligning irrigation frequency with both crop development and environmental demand. This indicates that RL integrates feedback signals more effectively, avoiding sudden water-application spikes while keeping moisture within the safe soil-moisture range.

The cumulative water consumption underscores the efficiency gains achievable with adaptive control. The time-based strategy consistently utilized the largest volume of water, whereas the on–off strategy conserved water during early growth but lost this advantage at later stages. The RL algorithm initially consumed more water during its exploratory phase but rapidly adapted to reduce consumption, ultimately achieving the lowest total use by the end of the trial. These savings were obtained without driving soil moisture below the PWP, confirming that efficiency improvements were compatible with maintaining the crop in the safe soil-moisture range.

Collectively, the findings highlight the potential of reinforcement learning as a viable alternative for irrigation control strategies. Unlike fixed scheduling or threshold-based logic, RL demonstrated the ability to dynamically balance water-use efficiency and soil-moisture regulation, keeping the root zone consistently within the safe range. This adaptability positions RL as a promising tool for advancing sustainable greenhouse production, where optimizing water-use efficiency while safeguarding crop growth is a critical priority.

The experimental results demonstrate that RL can outperform both conventional time-based scheduling and simple threshold-based approaches by dynamically adapting to crop demand and environmental conditions while optimizing water use. Because the proposed method is model-free, it does not rely on an explicit representation of system dynamics and, therefore, does not require retraining, provided that operating conditions remain within the range of previously experienced states. Instead, it primarily requires a careful selection of initial parameters. In this study, the initial values defined in (14) and (15) were chosen based on preliminary simulation analyses, as inappropriate parameter choices may degrade overall system performance. Nonetheless, the underlying RL framework—the learned policy structure, state representation, and reward formulation—can be reused across different irrigation conditions. Although further validation across crop types, environmental settings, and complete production cycles is needed, the findings indicate that RL holds strong potential as a practical tool for precision irrigation, contributing to more sustainable and resilient greenhouse production systems.

Comparing the water savings obtained in this study with previously published work on greenhouse irrigation shows that closed-loop strategies based on evapotranspiration typically reduce water consumption by approximately 25–35% [5,34,35], while soil-moisture-driven approaches achieve similar savings in the range of 25–40% [36]. In this study,

the proposed method achieved a 36.9% reduction in water use relative to empirical time-based irrigation, placing its performance at the upper end of values reported for established greenhouse irrigation strategies. Reported water savings from RL-based irrigation systems in both open-field and controlled environments generally fall within 5–20% [4,19,26,28], suggesting that the efficiency gains observed here are comparable to, and in some cases exceed, those documented in recent reinforcement learning studies. An additional advantage of the proposed approach is that, unlike many RL implementations, it was validated across a complete plant growth cycle and is capable of running on low-end computing devices, facilitating deployment in resource-constrained greenhouse settings.

In the context of greenhouse jalapeño production in Mexico, implementing water-saving irrigation strategies can significantly boost farm profitability. By reducing water consumption, producers can lower the costs associated with water extraction, pumping, and any related energy use. Over an entire crop cycle, even modest water savings can add up to substantial operational savings. These efficiency gains not only reduce variable costs but also increase the economic resilience of growers, especially under growing water scarcity and pressure on water resources in northern Mexico. Furthermore, by maintaining high yields while using less water, farmers can improve their net margins without sacrificing productivity. This makes advanced irrigation control techniques—such as reinforcement learning-based strategies—not just a tool for sustainability but also a practical lever for enhancing the financial viability of greenhouse jalapeño operations in Mexico.

5. Conclusions

The reinforcement learning control strategy consistently maintained soil moisture within the efficient soil-moisture range, between the maximum allowable depletion and the field capacity, after its initial learning phase, avoiding both water stress and oversaturation. In contrast, the time-based control led to recurrent over-irrigation, while the on–off control showed acceptable regulation but with less stability. Root mean square error analysis confirmed that reinforcement learning achieved the highest precision in matching irrigation to the maximum allowable depletion threshold, converging to 2%, compared with 3% for the on–off control and 5% for the time-based control.

Reinforcement learning demonstrated the ability to gradually adjust irrigation frequency in response to both crop growth stage and environmental demand, outperforming the rigid time-based method and the reactive on–off control, which tended to produce abrupt changes. Over the 8-week evaluation period, the proposed approach achieved the greatest net reduction in cumulative water consumption, resulting in savings of more than 30%, despite an initial learning cost. The time-based strategy used the most water, whereas the on–off control only achieved temporary savings. Regarding overall performance, reinforcement learning emerged as the most effective approach, providing a dynamic balance between water savings and maintaining soil moisture within safe limits. These results highlight its potential to support sustainable irrigation management in greenhouse systems, where both efficiency and crop health are critical.

As part of future research, we plan to investigate the cross-site and cross-crop transferability of the proposed method, assessing whether the learned policy structure and state representations can generalize effectively to other greenhouse environments and diverse horticultural crops.

Author Contributions: Conceptualization, J.P.P.-N., L.D.G., and C.L.; methodology, C.L., L.O., and A.C.-P.; software, J.P.P.-N. and L.D.G.; validation, C.L., L.O., and A.C.-P.; formal analysis, C.L., L.O., and A.C.-P.; investigation, J.P.P.-N., L.D.G., and C.L.; resources, C.L.; data curation, J.P.P.-N. and L.D.G.; writing—original draft preparation, J.P.P.-N., L.D.G., and C.L.; writing—review and editing, C.L., L.O., and A.C.-P.; visualization, C.L.; supervision, C.L., L.O., and A.C.-P.; project

administration, C.L.; funding acquisition, C.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Chihuahua State Government (Mexico) through the "Fondo Estatal de Ciencia, Innovación y Tecnología", grant number FECTI/2024/CV-CDF/024.

Data Availability Statement: The data presented in this study are available on request from the corresponding author due to privacy agreements with agricultural partners that restrict public sharing of field-level sensor data.

Acknowledgments: The authors would like to acknowledge the financial and technical support of Tecnológico de Monterrey, Mexico, in the production of this work, as well as Mexico's SECIHTI (Secretaría de Ciencia, Humanidades, Tecnología e Innovación) for its Master's scholarship grant.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

A2C	Advantage Actor–Critic
FC	Field Capacity
IoT	Internet of Things
MAD	Maximum Allowable Depletion
MDP	Markov Decision Process
MQTT	Message-Queuing Telemetry Transport
MPC	Model Predictive Control
PID	Proportional–Integral–Derivative
PWP	Permanent Wilting Point
RL	Reinforcement Learning
RMSE	Root Mean Square Error
TD	Temporal Difference
VWC	Volumetric Water Content

References

1. FAO. *The State of Food and Agriculture 2022. Leveraging Automation in Agriculture for Transforming Agrifood Systems*; Technical Report; Food and Agriculture Organization of the United Nations: Rome, Italy, 2022. [CrossRef]
2. González-Amarillo, C.A.; Corrales-Muñoz, J.C.; Mendoza-Moreno, M.A.; Amarillo, A.M.G.; Hussein, A.F.; Arunkumar, N.; Ramirez-González, G. An IoT-Based Traceability System for Greenhouse Seedling Crops. *IEEE Access* **2018**, *6*, 67528–67535. [CrossRef]
3. Bwambale, E.; Abagale, F.K.; Anornu, G.K. Data-Driven Modelling of Soil Moisture Dynamics for Smart Irrigation Scheduling. *Smart Agric. Technol.* **2023**, *5*, 100251. [CrossRef]
4. Agyeman, B.T.; Naouri, M.; Appels, W.M.; Liu, J.; Shah, S.L. Learning-based multi-agent MPC for irrigation scheduling. *Control Eng. Pract.* **2024**, *147*, 105908. [CrossRef]
5. Abioye, E.A.; Abidin, M.S.Z.; Mahmud, M.S.A.; Buyamin, S.; Ishak, M.H.I.; Rahman, M.K.I.A.; Otuoze, A.O.; Onotu, P.; Ramli, M.S.A. A review on monitoring and advanced control strategies for precision irrigation. *Comput. Electron. Agric.* **2020**, *173*, 105441. [CrossRef]
6. Kelly, T.; Foster, T.; Schultz, D.M. Assessing the value of deep reinforcement learning for irrigation scheduling. *Smart Agric. Technol.* **2024**, *7*, 100403. [CrossRef]
7. Klein, L.J.; Hamann, H.F.; Hinds, N.; Guha, S.; Sanchez, L.; Sams, B.; Dokoozlian, N. Closed Loop Controlled Precision Irrigation Sensor Network. *IEEE Internet Things J.* **2018**, *5*, 4580–4588. [CrossRef]
8. Romero, R.; Muriel, J.; García, I.; Muñoz de la Peña, D. Research on automatic irrigation control: State of the art and recent results. *Agric. Water Manag.* **2012**, *114*, 59–66. [CrossRef]
9. Lozoya, C.; Mendoza, C.; Aguilar, A.; Román, A.; Castelló, R. Sensor-based model driven control strategy for precision irrigation. *J. Sens.* **2016**, *2016*, 9784071. [CrossRef]

10. Agyeman, B.T.; Sahoo, S.R.; Liu, J.; Shah, S.L. LSTM-based model predictive control with discrete actuators for irrigation scheduling. *IFAC-PapersOnLine* **2022**, *55*, 334–339. [CrossRef]
11. Cáceres, G.B.; Ferramosca, A.; Gata, P.M.; Martín, M.P. Model Predictive Control Structures for Periodic ON–OFF Irrigation. *IEEE Access* **2023**, *11*, 51985–51996. [CrossRef]
12. Bwambale, E.; Abagale, F.K.; Anornu, G.K. Smart irrigation monitoring and control strategies for improving water use efficiency in precision agriculture: A review. *Agric. Water Manag.* **2022**, *260*, 107324. [CrossRef]
13. Adeyemi, O.; Grove, I.; Peets, S.; Domun, Y.; Norton, T. Dynamic Neural Network Modelling of Soil Moisture Content for Predictive Irrigation Scheduling. *Sensors* **2018**, *18*, 3408. [CrossRef] [PubMed]
14. Gu, Z.; Zhu, T.; Jiao, X.; Xu, J.; Qi, Z. Neural network soil moisture model for irrigation scheduling. *Comput. Electron. Agric.* **2021**, *180*, 105801. [CrossRef]
15. Custódio, G.; Prati, R.C. Comparing modern and traditional modeling methods for predicting soil moisture in IoT-based irrigation systems. *Smart Agric. Technol.* **2024**, *7*, 100397. [CrossRef]
16. Sharma, A.; Jain, A.; Gupta, P.; Chowdary, V. Machine Learning Applications for Precision Agriculture: A Comprehensive Review. *IEEE Access* **2021**, *9*, 4843–4873. [CrossRef]
17. Umutoni, L.; Samadi, V. Application of machine learning approaches in supporting irrigation decision making: A review. *Agric. Water Manag.* **2024**, *294*, 108710. [CrossRef]
18. Bu, F.; Wang, X. A smart agriculture IoT system based on deep reinforcement learning. *Future Gener. Comput. Syst.* **2019**, *99*, 500–507. [CrossRef]
19. Chen, M.; Cui, Y.; Wang, X.; Xie, H.; Liu, F.; Luo, T.; Zheng, S.; Luo, Y. A reinforcement learning approach to irrigation decision-making for rice using weather forecasts. *Agric. Water Manag.* **2021**, *250*, 106838. [CrossRef]
20. Ajagekar, A.; You, F. Deep Reinforcement Learning Based Automatic Control in Semi-Closed Greenhouse Systems. *IFAC-PapersOnLine* **2022**, *55*, 406–411. [CrossRef]
21. Wang, L.; He, X.; Luo, D. Deep Reinforcement Learning for Greenhouse Climate Control. In Proceedings of the 2020 IEEE International Conference on Knowledge Graph (ICKG), Nanjing, China, 9–11 August 2020; pp. 474–480. [CrossRef]
22. Zhang, W.; Cao, X.; Yao, Y.; An, Z.; Xiao, X.; Luo, D. Robust model-based reinforcement learning for autonomous greenhouse control. In Proceedings of the Asian Conference on Machine Learning, Virtual, 17–19 November 2021; pp. 1208–1223.
23. Goldenits, G.; Mallinger, K.; Raubitzek, S.; Neubauer, T. Current applications and potential future directions of reinforcement learning-based Digital Twins in agriculture. *Smart Agric. Technol.* **2024**, *8*, 100512. [CrossRef]
24. Ding, X.; Du, W. Poster Abstract: Smart Irrigation Control Using Deep Reinforcement Learning. In Proceedings of the 2022 21st ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN), Milano, Italy, 4–6 May 2022; pp. 539–540. [CrossRef]
25. Zhou, N. Intelligent control of agricultural irrigation based on reinforcement learning. *J. Phys. Conf. Ser.* **2020**, *1601*, 052031. [CrossRef]
26. Sun, L.; Yang, Y.; Hu, J.; Porter, D.; Marek, T.; Hillyer, C. Reinforcement Learning Control for Water-Efficient Agricultural Irrigation. In Proceedings of the 2017 IEEE International Symposium on Parallel and Distributed Processing with Applications and 2017 IEEE International Conference on Ubiquitous Computing and Communications (ISPA/IUCC), Guangzhou, China, 12–15 December 2017; pp. 1334–1341. [CrossRef]
27. Ding, X.; Du, W. Optimizing irrigation efficiency using deep reinforcement learning in the field. *ACM Trans. Sens. Netw.* **2024**, *20*, 99. [CrossRef]
28. Tang, R.; Tang, J.; Abu Talip, M.S.; Aridas, N.K.; Guan, B. Reinforcement learning control method for greenhouse vegetable irrigation driven by dynamic clipping and negative incentive mechanism. *Front. Plant Sci.* **2025**, *16*, 1632431. [CrossRef]
29. Zotarelli, L.; Dukes, M.D.; Romero, C.C.; Migliaccio, K.W.; Morgan, K.T. Step by Step Calculation of the Penman-Monteith Evapotranspiration (FAO-56 Method): AE459, 2/2010. EDIS. Available online: https://edis.ifas.ufl.edu/pdffiles/AE/AE45900.pdf (accessed on 27 November 2025).
30. Lozoya, C.; Mendoza, G.; Mendoza, C.; Torres, V.; Grado, M. Experimental evaluation of data aggregation methods applied to soil moisture measurements. In Proceedings of the SENSORS, 2014 IEEE, Valencia, Spain, 2–5 November 2014; pp. 134–137. [CrossRef]
31. Garcia, L.D.; Lozoya, C.; Castañeda, H.; Favela-Contreras, A. A discrete sliding mode control strategy for precision agriculture irrigation management. *Agric. Water Manag.* **2025**, *309*, 109315. [CrossRef]
32. Allen, R.G.; Pereira, L.S.; Raes, D.; Smith, M. *FAO Irrigation and Drainage Paper No. 56*; Food and Agriculture Organization of the United Nations: Rome, Italy, 1998; M-56.
33. Lozoya, C.; Favela-Contreras, A.; Aguilar-Gonzalez, A.; Orona, L. A Precision Irrigation Model Using Hybrid Automata. *Trans. Asabe* **2019**, *62*, 1639–1650. [CrossRef]
34. Yang, H.; Du, T.; Qiu, R.; Chen, J.; Wang, F.; Li, Y.; Wang, C.; Gao, L.; Kang, S. Improved water use efficiency and fruit quality of greenhouse crops under regulated deficit irrigation in northwest China. *Agric. Water Manag.* **2017**, *179*, 193–204. [CrossRef]

35. Incrocci, L.; Thompson, R.B.; Fernandez-Fernandez, M.D.; De Pascale, S.; Pardossi, A.; Stanghellini, C.; Rouphael, Y.; Gallardo, M. Irrigation management of European greenhouse vegetable crops. *Agric. Water Manag.* **2020**, *242*, 106393. [CrossRef]
36. Nikolaou, G.; Neocleous, D.; Katsoulas, N.; Kittas, C. Irrigation of Greenhouse Crops. *Horticulturae* **2019**, *5*, 7. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.