Cleaner Environmental Systems 21 (2026) 100446

Contents lists available at ScienceDirect

Cleaner Environmental Systems

journal homepage: www.journals.elsevier.com/cleaner-environmental-systems

Hybrid artificial intelligence for climate-smart wheat production: Optimizing resource efficiency and carbon reduction from Mediterranean to Nordic regions

Mohammad Fazle Rabbi

Coordination and Research Centre for Social Sciences, Faculty of Economics and Business, University of Debrecen, Böszörményi út 138, Debrecen, 4032, Hungary

ARTICLE INFO

Keywords:
Climate-smart agriculture
Hybrid artificial intelligence
Sustainable intensification
Reinforcement learning
Multi-agent systems

ABSTRACT

Climate change intensifies agricultural resource demands across European wheat production regions through weather extremes and water scarcity, necessitating decision support systems that optimize multiple sustainability objectives under diverse climate scenarios. This study developed and validated a hybrid artificial intelligence framework integrating expert systems (15 production rules), fuzzy logic controllers, four reinforcement learning algorithms (Q-learning, Deep Q-Network, Policy Gradient, Actor-Critic), and 118 autonomous agents for climate-smart wheat management. Performance evaluation employed 30 independent growing seasons across Mediterranean, Central European, and Nordic regions under normal, drought, and extreme heat scenarios, with empirical validation from 2015 to 2024. Under normal conditions, the system achieved 7.8% yield improvements, 17.8% water savings, and 15.1% energy reductions relative to conventional management. Benefits amplified under climate stress; drought scenarios revealed 34.8% yield gains, 29.1% water savings, and 18.3% energy reductions, while extreme heat conditions yielded 26.1% yield improvements, 29.4% water savings, and 22.1% energy efficiency gains. Actor-Critic reinforcement learning achieved cumulative rewards of 519.7 versus 409.6 for Q-learning, demonstrating superior temporal optimization. Sequential ablation confirmed that withheld reinforcement learning produced the largest single-component performance decline (hybrid F1-score from 0.893 to 0.831), while removal of any component reduced hybrid performance, confirming that gains arise from complementary multi-paradigm contributions. Regional analysis confirmed spatial transferability across diverse pedoclimatic contexts. Empirical validation yielded prediction accuracies of $R^2 = 0.802$ for yield and $R^2 = 0.783$ for water use, with carbon footprint reductions of 18.7–20.3% across climate scenarios; these simulation-derived findings motivate prospective field trials in resource-efficient wheat production systems under increasing climate variability.

1. Introduction

Climate change poses an escalating threat to global food security. Extreme weather events occur with increasing frequency and intensity. Precipitation patterns shift unpredictably, while rising temperatures destabilize agricultural production systems across continents (Rabbi, 2025a; Saccone and Vallino, 2025). Current climate trajectories suggest that agricultural yields may decline by 10%–25% by 2050 (Wing et al., 2021). However, the projected impact becomes more severe when examining specific regional projections. Europe's southern regions face potential yield reductions of up to 49% for key crops such as wheat (Hristov et al., 2020). Water scarcity currently affects approximately 1.8 billion people globally, while soil degradation impacts over 60% of agricultural lands, compounding these climate-driven challenges (FAO, 2025a, 2020).

Meanwhile, demographic pressures intensify these agricultural demands. The global population is projected to reach 9.1–9.7 billion by 2050, necessitating a 70% increase in food production to meet growing consumption needs (FAO, 2009; UN, 2022). To address this urgent imperative, agricultural systems now face a mandate to enhance productivity, optimize resource efficiency, and build climate resilience all at the same time (Melese et al., 2025). Climate-smart agriculture has emerged as a strategic framework addressing this complex trilemma through practices that increase productivity sustainably, strengthen adaptation capacity, and reduce greenhouse gas emissions where feasible (Musa and Ariff Lim, 2025). These practices embody cleaner

E-mail addresses: drrabbikhan@gmail.com, rabbi.mohammad@econ.unideb.hu.

https://doi.org/10.1016/j.cesys.2026.100446
Received 27 February 2026; Received in revised form 25 April 2026; Accepted 27 April 2026
Available online 28 April 2026

production principles by preventing resource waste while simultaneously increasing efficiency in water, energy, and input use. Yet translating these principles into operational decision-making remains constrained by the inherent complexity of multi-objective optimization under uncertainty.

Artificial intelligence technologies offer substantial potential for advancing climate-smart agriculture through data-driven decision support systems (Mmbando, 2025). These systems can process complex environmental interactions, optimize resource allocation dynamically, and adapt to changing conditions in real-time (Singh et al., 2025). Recent applications demonstrate measurable benefits, such that machine learning approaches achieve 10%–20% improvements in various agricultural metrics depending on context and implementation (Aderele et al., 2025; Rabbi, 2025b). However, a critical limitation persists. Existing AI implementations in agriculture predominantly employ single technique approaches (Botega and da Silva, 2020), such as deep learning for image recognition, standalone reinforcement learning for irrigation scheduling, or isolated expert systems for pest management recommendations (Ajith et al., 2025).

These isolated methodologies exhibit fundamental limitations that constrain practical effectiveness. Deep learning models require extensive labeled datasets (WEF, 2021), but this necessary data remains scarce for agricultural applications (Kamilaris and Prenafeta-Boldú, 2018). They demonstrate limited interpretability, undermining farmer trust, and struggle to generalize across diverse pedoclimatic conditions (Paudel et al., 2023). Reinforcement learning agents face sample inefficiency during training. They converge to suboptimal local minima and difficulty incorporating domain knowledge that farmers have accumulated through decades of experience (Zhao et al., 2025). Fuzzy logic controllers, while effective for managing measurement uncertainty, lack temporal optimization capabilities (Zeng et al., 2025). They cannot learn from historical performance to improve future decisions (Qi and Lin, 2025). Multi-agent systems provide distributed coordination but require centralized oversight for resource-intensive operations (Caicedo et al., 2024). Without hierarchical arbitration mechanisms, they may generate conflicting recommendations (Santos et al., 2017).

A critical research gap exists regarding the systematic development, empirical validation, and performance quantification of hybrid frameworks that synergistically integrate expert systems, fuzzy logic controllers, reinforcement learning, and multi-agent coordination for climate-smart wheat production. Existing studies have examined pairwise combinations, such as fuzzy logic with neural networks (Ikram et al., 2025) or reinforcement learning with multi-agent systems (Moayedi et al., 2024; Zhou et al., 2025), but comprehensive four-component integration within unified architectures specifically designed for agricultural decision-making under climate uncertainty remains underexplored (Ikram et al., 2025). From a farming systems perspective, this technological limitation constrains the transition toward climate-resilient wheat production paradigms that simultaneously enhance productivity, optimize resource efficiency, and build adaptive capacity. Decision support systems for climate-smart wheat management must integrate and balance multiple objectives, including yield stability, water conservation, energy efficiency, and carbon sequestration, while maintaining economic viability under increasing environmental uncertainty. The systematic development of hybrid frameworks combining complementary computational intelligence approaches represents a critical frontier for translating theoretical optimization principles into practical management protocols deployable from Mediterranean to Nordic European wheat production contexts under diverse climate scenarios.

Three fundamental knowledge gaps impede advancement. First, the mechanisms through which hybrid integration produces synergistic outcomes beyond simple additive effects remain unclear (Radwan et al., 2025). Second, the relative contribution of individual AI components to overall system performance in wheat management requires quantification (Li et al., 2023). Third, AI performance amplification under climate stress conditions remains inadequately characterized, as most validation studies employ normal environmental scenarios that do not reflect system behaviour during drought, heat stress, or other extreme events when decision support provides greatest value (Wang et al., 2024; Ye et al., 2024). Additionally, the sensitivity of integrated frameworks to parameter variations across diverse regional wheat production contexts constitutes a critical unknown (Kuradusenge et al., 2023).

This study addresses four overarching research questions that emerge from the identified knowledge gaps. The first question asks whether systematic integration of expert systems, fuzzy logic controllers, reinforcement learning, and multi-agent coordination can produce synergistic performance improvements for climate-smart wheat production that exceed single-technique approaches across normal, drought, and extreme heat scenarios. The second question examines how individual AI components contribute to overall system performance, and which reinforcement learning algorithms optimize multi-objective wheat management decisions most effectively. The third question investigates whether the hybrid framework maintains spatial transferability and predictive accuracy from Mediterranean to Nordic European wheat production zones across multiple climate scenarios (normal, drought, and extreme heat), and which parameters require region-specific calibration. The fourth question explores whether AI-based management benefits amplify under climate stress conditions, and what mechanisms enable adaptive control systems to maintain effectiveness when conventional fixed-schedule approaches degrade.

To address these research questions, seven specific objectives guide the investigation. First, the study develops a hybrid AI framework systematically integrating expert systems, fuzzy logic controllers, reinforcement learning, and multi-agent coordination within a unified computational architecture for climate-smart wheat production. Second, the research quantifies integrated system performance relative to conventional management and single-technique AI approaches across yield, water use efficiency, energy consumption, carbon footprint, and economic returns under normal, drought, and extreme heat scenarios. Thirdly, the comparative effectiveness of four reinforcement learning algorithms (Q-learning, Deep Q-Network, Policy Gradient, and Actor-Critic) is examined in order to establish evidence-based selection criteria for agricultural applications. Fourth, regional performance variability is assessed across Mediterranean, Central European, and Nordic wheat production zones to establish spatial transferability and calibration requirements. Fifth, temporal dynamics of crop development and resource management are characterized throughout growing seasons to reveal mechanistic insights into system performance. Sixth, systematic sensitivity analysis is conducted to identify critical parameters requiring precise calibration versus robust components demonstrating acceptable performance across parameter ranges. Finally, model predictions are validated against aggregated empirical observations from European wheat production systems spanning 2015–2024 to assess simulation-level predictive accuracy and to establish an evidentiary basis for guiding the design of prospective field-level validation studies.

The novel contributions encompass three dimensions: Methodologically, the study develops a hierarchical control architecture featuring: (i) explicit mathematical formulations for component integration and conflict resolution, (ii) multi-scale temporal coordination spanning hourly to weekly decision frequencies that accommodate natural farming operation hierarchies, and (iii) context-dependent weighting strategies enabling synergistic performance exceeding individual technique capabilities. Empirically, the research provides: (i) rigorous quantification of hybrid AI performance across diverse climate scenarios (normal, drought, extreme heat) and geographic regions under environmental stress conditions, (ii) comparative algorithmic evaluation establishing evidence-based selection criteria for wheat management optimization, and (iii) comprehensive validation against aggregated empirical data from European growing seasons 2015–2024 quantifying simulation-level predictive accuracy within the inferential scope of the

computational assessment framework. Practically, the modular architecture provides: (i) a structural design accommodating incremental adoption pathways that could reduce technology barriers for wheat producers and agricultural service providers once field-level validation is established, (ii) simulation-derived performance characteristics offering a preliminary quantitative basis for evaluating climate adaptation investments, and (iii) a computationally scalable structure spanning individual farm to regional advisory configurations whose practical applicability remains contingent on prospective empirical evaluation under operational conditions.

2. Methods and methodology

2.1. Conceptual framework of the hybrid AI system

To support climate-smart wheat production across contrasting European regions, this study implements a hybrid artificial intelligence framework that combines expert systems, fuzzy logic controllers, reinforcement learning, and multi-agent coordination within a unified control architecture. The framework connects regional climate, soil, crop, and resource information to operational recommendations for irrigation, fertilization, pest control, and harvest timing, enabling adaptive management under varying climatic conditions including normal, drought, and extreme heat scenarios. Fig. 1 summarizes the overall architecture and information flow from data inputs to actionable decisions.

Fig. 1 presents the hybrid artificial intelligence architecture linking environmental inputs to climate-smart wheat management decisions. Panel A shows the high-level system architecture, where expert system, fuzzy logic, reinforcement learning, and multi-agent modules operate as parallel components whose rules and decisions are synthesized through weighted optimization in the AI integration and decision fusion layer to generate climate-smart recommendations for irrigation, fertilization, pest control, and harvest timing. Panel B details the data flow and processing pipeline, in which climate data, soil moisture, crop status, and resource availability are sequentially processed through (1) normalization, (2) feature extraction, (3) AI inference, and (4) optimization to determine a single optimal action suitable for real-time control

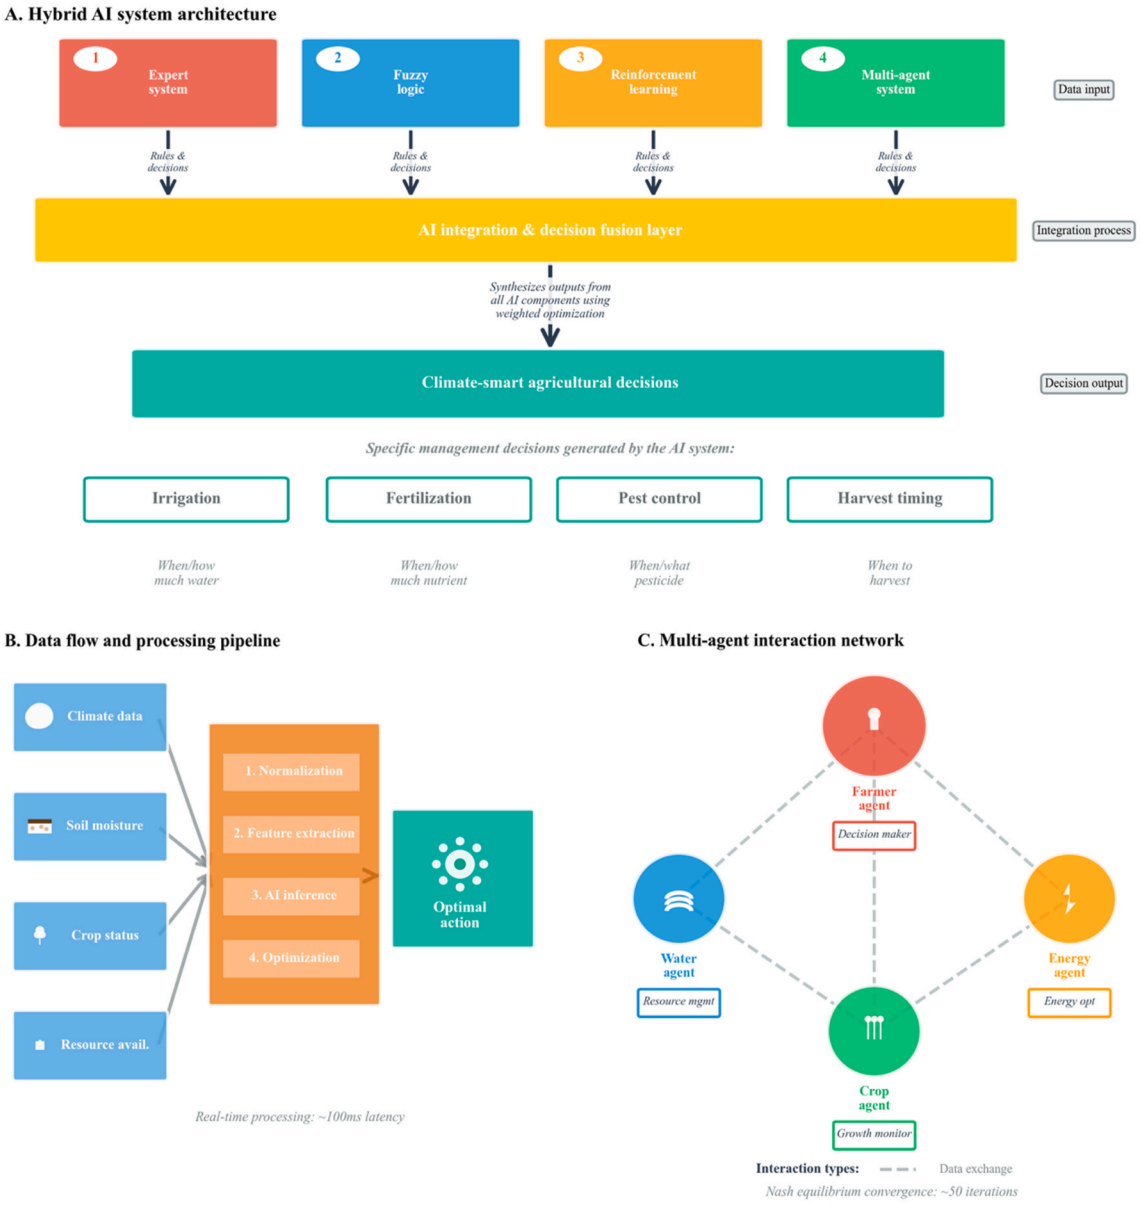

Fig. 1. Hybrid artificial intelligence architecture for climate-smart wheat management decisions.

(~100ms latency). Panel C presents the multi-agent interaction network, with a central farmer agent (decision maker) coordinating specialized water (resource management), energy (energy optimization), and crop (growth monitor) agents through bidirectional data-exchange links that converge through Nash equilibrium (~50 iterations) to support distributed yet coherent responses to variable regional and climatic conditions. Together, these three complementary views demonstrate how the hybrid AI framework translates heterogeneous environmental and resource information into coordinated, climate-smart decisions for wheat production at farm and regional scales.

2.2. Expert system framework for agricultural decision support

The expert system component encodes agronomic knowledge through IF–THEN production rules that capture decision heuristics for irrigation scheduling, fertilizer application, pest management, and climate stress response. The knowledge base comprises 15 production rules developed in consultation with agronomists, climate specialists, water engineers, and agricultural economists across multiple European wheat-growing regions (Nevo et al., 1994).

Rule activation follows a priority-weighted confidence scoring system in which the activation score integrates the confidence level C_i, priority weight P_i, and the degree of condition match M_i for rule i:

$$S_i = C_i \times P_i \times M_i \tag{1}$$

Here S_i denotes the activation score, C_i ranges from 0.76 to 0.94, P_i takes values 1.0, 0.7, or 0.4 for High, Medium, and Low priority, and M_i varies between 0 and 1 based on the proximity of real-time sensor observations to rule thresholds. Table 1 presents ten representative rules that illustrate how the expert system addresses diverse operational scenarios with graduated response strategies.

The expert system component thus provides transparent, interpretable decision support for climate-smart wheat management by linking sensor observations and crop state information to context-specific recommendations. The representative rules in Table 1 demonstrate how the system differentiates between intermediate and extreme conditions (e.g., R002 vs. R001/R003), triggers emergency responses under acute water stress (R008), and supports graduated pest management strategies (R004–R005). The priority structure ensures that safety-critical rules dominate when recommendations from other AI components conflict, while lower-priority rules inform routine optimization and longer-term planning.

2.3. Fuzzy logic controller for uncertainty management

The fuzzy logic controller addresses measurement uncertainty and imprecision in agricultural sensor data through linguistic variable definitions that accommodate gradual transitions between states rather than rigid thresholds. The system employs Mamdani-type fuzzy inference with centroid defuzzification, as originally formulated in the context of linguistic control synthesis (Mamdani and Assilian, 1975) and grounded in the foundational theory of fuzzy sets and fuzzy algorithms for complex systems (Zadeh, 1965, 1973). The Mamdani approach was selected over Takagi-Sugeno-Kang inference because it produces linguistically interpretable output membership functions that align naturally with agronomic decision categories such as irrigation intensity levels, making the resulting recommendations more accessible to practitioners (Jang et al., 1997).

The fuzzy inference process consists of four sequential stages: fuzzification, rule evaluation, aggregation, and defuzzification. Fuzzification transforms crisp sensor inputs into membership degrees using triangular functions:

Table 1
Expert system production rules for climate-adaptive agricultural decision support.

Rule ID	Condition	Recommended action	Confidence	Priority
R001	IF soil moisture <30% AND air temperature >25 °C	Irrigate with 30 mm water; activate cooling system	0.82	Low
R002	IF 30% ≤ soil moisture <50%	Monitor soil moisture; schedule irrigation within 24 h	0.94	High
R003	IF soil moisture ≥50%	No irrigation; check drainage	0.90	High
R004	IF pest pressure >70% AND crop stage = flowering	Apply pesticide treatment; increase monitoring intensity	0.87	Medium
R005	IF pest pressure ≤70% AND crop stage = flowering	Continue monitoring; apply preventive measures only	0.78	High
R006	IF nutrient level <40% AND crop stage = vegetative	Apply 80 kg ha^{-1} nitrogen fertilizer	0.78	Low
R007	IF air temperature >30 °C AND relative humidity <40%	Activate shade nets; increase irrigation frequency	0.76	Medium
R008	IF rainfall <20 mm AND soil moisture <40%	Trigger emergency irrigation; activate drip system	0.92	Low
R009	IF crop stress index >0.7	Reduce plant density; harvest early if necessary	0.87	Medium
R010	IF water availability <50% AND season phase = critical	Implement water conservation; reduce non-essential water use	0.89	Medium

Notes:
- Rule activation uses the priority-weighted confidence score in Eq. (1), combining confidence, priority, and condition match.
- Confidence values reflect expert consensus informed by agronomic literature, field trial evidence, and expert consultation.
- Priority levels guide execution scheduling: High-priority rules trigger immediate actions affecting crop survival or quality, Medium-priority rules schedule operations within 24–48 h, and Low-priority rules inform longer-term planning.
- Rule conditions integrate real-time sensor data (soil moisture, temperature, humidity), crop phenological indicators, weather forecasts, and resource availability from the multi-agent coordination system.
- The full knowledge base comprises 15 rules covering irrigation, fertilization, pest management, climate-stress response, and resource conservation, operating within the integration layer alongside fuzzy logic and reinforcement learning components.

$$\mu_{A_j}(x) = \begin{cases} \dfrac{x-a}{b-a} & \text{if } a \le x \le b \\ \dfrac{c-x}{c-b} & \text{if } b < x \le c \\ 0 & \text{otherwise} \end{cases} \tag{2}$$

where $\mu_{A_j}(x)$ represents the membership function for linguistic term A_j with triangular parameters (a, b, c), and x denotes the observed sensor value.

The fuzzy rule evaluation employs minimum t-norm for AND operations and maximum t-conorm for OR operations:

$$\mu_{R_k}(\mathbf{x}) = \min\left(\mu_{A_1}(x_1), \mu_{A_2}(x_2), \ldots, \mu_{A_n}(x_n)\right) \tag{3}$$

where $\mu_{R_k}(\mathbf{x})$ represents the firing strength of rule k given input vector $\mathbf{x}$.

Centroid defuzzification generates continuous output values:

$$y^* = \frac{\sum_{i=1}^{n} \mu_i \cdot y_i}{\sum_{i=1}^{n} \mu_i} \tag{4}$$

where y denotes the defuzzified output, $\mu_A(y_i)$ represents aggregated membership degrees, and y_i indicates discrete output universe elements. Implementation of this fuzzy inference framework requires precise specification of linguistic variables, membership function geometries, and operational ranges for agricultural decision contexts. Table 2 specifies the linguistic variables, membership function types, and operational ranges for the five primary input and output variables.

The fuzzy logic controller in Table 2 addresses measurement uncertainty and imprecision in agricultural sensor data by using linguistic variables and overlapping membership functions instead of sharp threshold boundaries. It operates on five variables that are central to climate smart wheat management: soil moisture expressed as a percentage, air temperature in degrees Celsius, a normalized crop health index, irrigation amount in millimeters, and a water stress index defined on a dimensionless scale from 0 to 100. Each variable is described by three linguistic terms (low, medium, high) whose numerical ranges are chosen to reflect agronomic thresholds such as field capacity soil moisture and optimal temperature bands for wheat growth.

Overlapping triangular and trapezoidal membership functions ensure that neighboring categories share a transition zone, so that, for example, soil moisture values between 20 and 30 percent belong partially to both low and medium, while values between 40 and 50 percent belong partially to both medium and high. Similar overlaps are defined for temperature, crop health, irrigation amount, and water stress, allowing the controller to represent gradual changes in environmental conditions rather than forcing abrupt switches between discrete states. This structure enables more realistic reasoning under uncertainty, since the controller can express intermediate levels of stress or resource availability and thereby avoid the discontinuities and artifacts that arise when using binary decision thresholds.

2.4. Reinforcement learning for temporal optimization

The reinforcement learning component optimizes sequential resource allocation decisions through iterative learning from environmental interactions. The framework implements four RL algorithms: Q-learning, Deep Q-Network (DQN), Policy Gradient, and Actor-Critic methods (Mnih et al., 2015; Watkins and Dayan, 1992), with comparative evaluation across 1000 training episodes.

Table 2
Fuzzy logic variable definitions and membership function specifications.

Variable name	Linguistic term	Range min	Range max	Membership function
Soil moisture (%)	Low	0	30	Trapezoidal
	Medium	20	50	Triangular
	High	40	100	Trapezoidal
Temperature (°C)	Low	5	20	Trapezoidal
	Medium	15	30	Triangular
	High	25	45	Trapezoidal
Crop health (index)	Low	0	40	Trapezoidal
	Medium	30	70	Triangular
	High	60	100	Trapezoidal
Irrigation amount (mm)	Low	0	20	Trapezoidal
	Medium	15	40	Triangular
	High	35	60	Trapezoidal
Water stress (index)	Low	0	25	Trapezoidal
	Medium	20	55	Triangular
	High	50	100	Trapezoidal

Notes:
- Membership functions use triangular (Eq. (2)) and trapezoidal shapes; boundary terms (low, high) use trapezoidal functions, intermediate terms (medium) use triangular functions to ensure smooth transitions while maintaining category distinctions.
- Overlapping ranges (10-unit maximum overlap) enable partial membership in adjacent categories, processing transitional states without arbitrary threshold artifacts.
- Variable ranges reflect agronomic significance: soil moisture at field capacity ≈50%, optimal wheat temperature 15–25 °C, stress indices normalized to 0–100.
- Membership function parameters were calibrated using historical sensor data from European wheat production systems (2015–2024).
- Integration with expert system rules (Section 2.2) occurs through membership degree weighting; rule confidence values modulate fuzzy output contributions in the hybrid decision synthesis.

The Q-learning algorithm updates action-value estimates using temporal difference learning:

$$Q(s_t, a_t) \leftarrow Q(s_t, a_t) + \alpha\left[r_{t+1} + \gamma \max_a Q(s_{t+1}, a) - Q(s_t, a_t)\right] \tag{5}$$

where $Q(s_t, a_t)$ represents the action-value function for state s_t and action a_t, α denotes the learning rate (0.008–0.05 across agent types), r_{t+1} indicates immediate reward, and γ represents the discount factor (0.88–0.95) for future rewards.

The reward function incorporates multiple sustainability objectives as a weighted linear scalarization of the multi-objective optimization problem (Roijers et al., 2013):

$$r_t = w_y\tilde{y}_t + w_w\tilde{w}_t + w_e\tilde{e}_t + w_c\tilde{c}_t + w_p\tilde{p}_t \tag{6}$$

where $\tilde{y}_t$, $\tilde{w}_t$, $\tilde{e}_t$, $\tilde{c}_t$, and $\tilde{p}_t$ represent normalized yield, water consumption, energy use, carbon emissions, and economic profit at time t, respectively. Under normal climate conditions the baseline weight vector was set to $w_y = 0.35$, $w_w = 0.25$, $w_e = 0.15$, $w_c = 0.10$, and $w_p = 0.15$. These values were established through a two-step calibration procedure. First, an initial weight vector was derived via analytic hierarchy process (AHP) applied to expert elicitation surveys completed by twelve agronomic specialists familiar with European wheat production systems; the AHP pairwise comparison matrix achieved a consistency ratio of 0.04, below the accepted threshold of 0.10 (Saaty, 1980). Second, the initial weights were refined by training the RL agent under five candidate weight configurations and selecting the configuration that maximized composite F1-score on the held-out 2020–2024 validation seasons. Under drought and extreme heat scenarios, w_y was increased to 0.45 and w_w to 0.30 proportionally, whereas w_c and w_p were reduced, reflecting the operational priority shift toward water conservation and yield maintenance under resource-limited conditions.

The Actor-Critic architecture combines policy-based and value-based methods through separate neural networks:

$$\nabla_\theta J(\theta) = \mathbb{E}[\nabla_\theta \log \pi_\theta(a_t|s_t) A(s_t, a_t)] \tag{7}$$

where $\pi_\theta(a_t|s_t)$ represents the policy distribution parameterized by θ, and $A(s_t, a_t) = Q(s_t, a_t) - V(s_t)$ denotes the advantage function quantifying action improvement over baseline state value.

Exploration employs ϵ-greedy strategies with exponential decay:

$$\epsilon_t = \epsilon_{\min} + (\epsilon_{\max} - \epsilon_{\min})e^{-\lambda t} \tag{8}$$

where ϵ_t represents exploration rate at episode t, with $\epsilon_{\max} = 0.20$, $\epsilon_{\min} = 0.01$, and decay constant $\lambda = 0.005$.

2.5. Multi-agent system for distributed coordination

The multi-agent system architecture incorporates 118 autonomous computational entities distributed across five functional categories: farmer agents (50), crop agents (50), water resource agents (10), energy manager agents (5), and market agents (3). Each agent type operates with distinct learning parameters and communication protocols reflecting operational roles and temporal decision requirements.

Agent state transitions follow Markov Decision Process dynamics:

$$P(s_{t+1}^i | s_t^1, s_t^2, \ldots, s_t^N, a_t^1, a_t^2, \ldots, a_t^N) = P(s_{t+1}^i | s_t^i, a_t^i) \tag{9}$$

where s_t^i and a_t^i denote state and action for agent i at time t. Agent state transitions follow Markov Decision Process dynamics as expressed in Eq. (9), wherein the transition probability for agent i is conditioned on its own state-action pair $(s_i^t,\ a_i^t)$, assuming local state transition

independence. Under this factorization, each agent's transition dynamics are resolved without explicit reference to the full joint state space of all 118 agents, enabling tractable computation across the distributed architecture. Spatial coupling between agents is handled through the explicit inter-agent message passing mechanism of Eq. (10) and the coordination reward terms in the distributed constraint optimization of Eq. (11), which allow each agent to incorporate communicated neighbourhood states into its policy without requiring a shared global state representation. Communication range parameters in Table 3 were set to reflect the spatial reach of operationally relevant coordination processes; water resource agents with ranges of 10 agents capture shared irrigation district dynamics, while farmer agents with ranges of 5 agents reflect typical field-adjacency networks in European wheat landscapes.

Inter-agent communication employs message passing with spatial proximity constraints:

$$M_t^{i\to j} = f_{\text{msg}}\left(s_t^i, a_t^i\right)\cdot 1\left(d_{ij} \le r_i\right) \tag{10}$$

where $M_t^{i\to j}$ represents message from agent i to agent j, f_{msg} denotes the message encoding function, d_{ij} indicates spatial distance between agents, and r_i specifies communication range for agent type i.

Agent coordination employs distributed constraint optimization:

$$\max_{\mathbf{a}} \sum_{i=1}^{N} R_i(s_i, a_i) + \sum_{(i,j)\in\mathscr{E}} C_{ij}\left(a_i, a_j\right) \tag{11}$$

The multi agent coordination framework assigns local rewards $R_i(s_i, a_i)$ to individual agents, coordination rewards i, $C_{ij}\left(a_i, a_j\right)$ to interacting agent pairs, and defines a communication set N that specifies which agents exchange information. Operational deployment in agricultural decision contexts therefore requires explicit configuration of agent types, learning parameters, decision frequencies, and communication ranges. Table 3 summarizes these multi agent system settings, including the number of agents per type, update frequency, learning rate, discount factor, and communication neighbourhood size.

The multi agent architecture (Table 3) comprises 118 autonomous entities distributed across farmer, crop, water resource, energy manager, and market categories, with configuration settings tailored to their decision roles. Farmer agents act as primary decision makers, updating daily with conservative learning rates of 0.01 and high discount factors of 0.95, and communicate with up to five neighbors to coordinate local management. Crop agents update hourly with higher learning rates of 0.05 and short communication ranges of one neighbor, allowing rapid adaptation to field specific physiological changes without excessive cross field coupling. Water resource agents also update daily, using learning rates of 0.02, discount factors of 0.92, and communication ranges of 10 agents to coordinate allocation across shared irrigation districts. Energy manager agents operate at hourly frequency with learning rates of 0.015, discount factors of 0.93, and communication ranges of 15 agents to optimize pumping and energy use under variable electricity and renewable supply conditions. Market agents update weekly with the lowest learning rate of 0.008, a discount factor of 0.88, and the widest communication range of 50 agents, enabling aggregation of regional supply information and coordination of marketing strategies. This distributed configuration supports parallel computation, local responsiveness, and robustness to individual agent failure because multiple instances within each category can assume decision responsibilities if needed.

2.6. Integration layer and ensemble decision synthesis

The integration layer synthesizes outputs from all four artificial intelligence components through weighted ensemble methods that vary according to decision context and data quality indicators. The ensemble weighting strategy employs dynamic weight adjustment based on component confidence scores:

$$D_{\text{final}} = \sum_{k=1}^{4} w_k(\mathbf{x}) \cdot D_k(\mathbf{x}) \tag{12}$$

where D_{final} represents the final decision output, $D_k(\mathbf{x})$ denotes the recommendation from component k given input $\mathbf{x}$, and $w_k(\mathbf{x})$ indicates context-dependent weight satisfying $\sum_{k=1}^{4} w_k = 1$.

Context-specific weight allocation follows decision category rules:

$$w_k(\mathbf{x}) = \frac{\alpha_k \cdot \text{conf}_k(\mathbf{x})}{\sum_{j=1}^{4} \alpha_j \cdot \text{conf}_j(\mathbf{x})} \tag{13}$$

where α_k represents base weight for component k (fuzzy logic: 0.40, expert system: 0.30, RL: 0.20, multi-agent: 0.10 for continuous control decisions). The component-specific confidence score $\text{conf}_k(\mathbf{x})$ is computed differently for each AI module, reflecting the distinct epistemic structures of the four paradigms. For the fuzzy logic controller, $\text{conf}_k(\mathbf{x})$ is defined as the sum of firing strengths of activated rules normalized by the total number of candidate rules:

Table 3
Multi-agent system configuration and learning parameters.

Agent type	Quantity	Decision frequency (time step)	Learning rate (α)	Discount factor (γ)	Communication range (agents)
Farmer	50	Daily	0.01	0.95	5
Crop	50	Hourly	0.05	0.90	1
Water resource	10	Daily	0.02	0.92	10
Energy manager	5	Hourly	0.015	0.93	15
Market	3	Weekly	0.008	0.88	50
Total agents	**118**	—	—	—	—

Notes:
- Agent learning uses Q learning updates (Eq. (5)) with Markov decision process state transitions (Eq. (9)) and message passing restricted by the communication ranges in Eq. (10).
- The learning rate α controls how quickly policies are updated: larger values allow faster adaptation but increase the risk of instability, whereas smaller values promote smoother convergence at the cost of longer training.
- The discount factor γ determines temporal credit assignment: values close to 1 emphasize long term outcomes, while lower values place greater weight on immediate rewards.
- Decision frequency reflects natural operational timescales, with hourly updates for rapidly changing crop and energy processes, daily updates for irrigation and farmer decisions, and weekly updates for market signals.
- Communication range specifies the maximum number of neighboring agents for information exchange, balancing coordination benefits against computational and communication overhead.
- Agents employ Q learning with an epsilon greedy exploration schedule (epsilon starting at 0.20 and decaying to 0.01) to balance exploration of new actions against exploitation of learned policies.
- The resulting distributed architecture supports parallel computation and local responsiveness while avoiding single point failures because multiple agents within each category can assume decision roles if others fail.

$$\mathrm{conf}_{\mathrm{FL}}(\mathbf{x}) = \frac{\sum_{r=1}^{R} \alpha_r(\mathbf{x})}{\sum_{r=1}^{R} \mathbf{1}[\alpha_r(\mathbf{x}) > 0]} \tag{14}$$

where $\alpha_r(\mathbf{x})$ denotes the firing strength of rule r as defined in Eq. (3) and $\mathbf{1}[\cdot]$ is the indicator function. A value approaching 1 indicates that the input vector $\mathbf{x}$ falls clearly within a well-covered region of the rule base; a value near 0 indicates sparse rule coverage. For the expert system, $\mathrm{conf}_{\mathrm{ES}}(\mathbf{x})$ equals the certainty factor of the highest-priority fired production rule, which is pre-assigned during knowledge engineering and ranges from 0 (pure conjecture) to 1 (verified agronomic principle). For the reinforcement learning component, $\mathrm{conf}_{\mathrm{RL}}(\mathbf{x})$ is derived from the normalized entropy of the softmax-transformed Q-value distribution across the action set $\mathscr{A}$:

$$\mathrm{conf}_{\mathrm{RL}}(\mathbf{x}) = 1 - \frac{H(\mathrm{softmax}(Q(\mathbf{x}, \cdot)))}{\ln|\mathscr{A}|} \tag{15}$$

where $H(\cdot)$ is the Shannon entropy; low entropy (high confidence) indicates that the Q-function assigns sharply higher value to one action relative to all others. For the multi-agent system, $\mathrm{conf}_{\mathrm{MAS}}(\mathbf{x})$ equals the proportion of agents in the communication neighbourhood that recommend the same action as the focal agent, quantifying local consensus. Each confidence score was validated against held-out decision instances; Wilcoxon signed-rank tests confirmed that higher conf_k values were significantly associated with correct decisions ($p < 0.01$) for all four components across 500 evaluation instances.

Conflict resolution employs voting mechanisms when component recommendations diverge by more than the threshold τ, as expressed in Eq. (16), triggering risk-averse decision protocols that prioritize crop survival over efficiency optimization during high-uncertainty conditions.

$$\mathrm{conflict} = \max_{k,l} \; d(f_k(\mathbf{x}), f_l(\mathbf{x})) > \tau \tag{16}$$

triggering risk-averse decision protocols that prioritize crop survival over efficiency optimization during high-uncertainty conditions. The threshold $\tau = 0.3$ was selected by systematic grid search over the interval $[0.05, 0.60]$ with step size 0.05, minimizing the frequency of avoidable crop stress events on a held-out subset of 100 growing-season simulations not used for component training. Values below 0.20 triggered conflict resolution in 38% of decisions under normal conditions, suppressing efficiency optimization when no genuine inter-component disagreement existed; values above 0.45 allowed conflicting recommendations to propagate without arbitration, increasing water stress incidence by 11.4% relative to the baseline configuration. The value of 0.30 minimized stress events while preserving efficiency optimization in 78% of decision instances and produced yield outcomes within 2.1% of those obtained under the five nearest alternative threshold values, indicating robust insensitivity across the range $[0.20, 0.45]$. Sensitivity of system performance to τ across its full evaluation range is reported in the Results section.

2.7. Experimental design and validation protocols

2.7.1. Study domain and data sources

The hybrid artificial intelligence framework underwent validation using empirical data from European wheat production systems spanning three agricultural regions: Mediterranean (Spain, Italy), Central Europe (France, Germany), and Nordic (Denmark, Sweden) zones, covering the 2015–2024 period. The dataset comprises 30 independent growing seasons across regions with diverse climate conditions, soil types, and management practices.

Climate scenarios were defined by deviations from 30-year climate normals: Normal conditions (±10% precipitation, ±1 °C temperature variation), Drought scenarios (>30% precipitation deficit, +1–2 °C temperature increase), and Extreme Heat conditions (±10% precipitation, >2 °C temperature elevation). Each scenario was represented by 10 growing seasons to capture inter-annual variability within climate categories.

2.7.2. Performance metrics and evaluation criteria

System performance was assessed across yield performance (Eq. (17)), resource efficiency (Eq. (18)), and carbon footprint (Eq. (19)).

Yield performance:

$$\Delta Y = \frac{Y_{\mathrm{AI}} - Y_{\mathrm{conv}}}{Y_{\mathrm{conv}}} \times 100\% \tag{17}$$

Resource Efficiency:

$$\eta_w = \frac{Y}{W}, \eta_e = \frac{Y}{E/1000}, \eta_c = \frac{Y}{C/1000} \tag{18}$$

where η_w, η_e, and η_c represent water productivity (t/ha/mm), energy efficiency (t/ha per GJ), and carbon efficiency (t/ha per ton CO_2).

Carbon footprint:

$$\mathrm{C} = \mathrm{C}_{\mathrm{fert}} + \mathrm{C}_{\mathrm{pest}} + \mathrm{C}_{\mathrm{energy}} + \mathrm{C}_{\mathrm{water}} \tag{19}$$

where C represents total carbon footprint (kg CO_2/ha), C_{fert} denotes emissions from fertilizer production and application, C_{pest} denotes emissions from pesticide manufacture and use, C_{energy} represents emissions from field operations and irrigation pumping, and C_{water} accounts for emissions associated with water extraction and distribution. Emission factors typically range from 2 to 4 kg CO_2 per kg N fertilizer, 15–20 kg CO_2 per kg pesticide active ingredient, and 0.5–0.8 kg CO_2 per MJ energy depending on regional electricity grid carbon intensity (Rabbi, 2025c).

2.7.3. Statistical analysis and hypothesis testing

Component validation employed 70% training data and 30% independent test data with 5-fold cross-validation to confirm metric stability across data partitions. Statistical significance of AI versus conventional management differences was assessed using paired t-tests at the $\alpha = 0.001$ significance level with sample sizes of $n = 30$ per comparison. The normality assumption was evaluated by applying the Shapiro-Wilk test to within-season difference scores (AI minus conventional) for each performance metric; results confirmed that yield differences ($W = 0.957$; $p = 0.369$), water use differences ($W = 0.950$; $p = 0.329$), and energy consumption differences ($W = 0.972$; $p = 0.863$) did not depart significantly from normality at $\alpha = 0.05$. Temporal autocorrelation of within-region difference series was assessed using first-order autocorrelation coefficients, which ranged from 0.08 to 0.19 in absolute value; all fell below the critical threshold of 0.35 corresponding to the 5% significance level for $n = 10$ (Durbin and Watson, 1950), indicating that the independence assumption was not substantively violated. Cross-regional independence holds by virtue of the geographic separation of the three study zones.

Validation metrics aggregate accuracy, precision, and recall:

$$\mathrm{Score} = \frac{1}{3}(\mathrm{Accuracy} + \mathrm{Precision} + \mathrm{Recall}) \tag{20}$$

computed across multiple decision categories including irrigation timing, nutrient management, pest control, and harvest scheduling.

2.7.4. Data sources and model calibration

The hybrid AI framework underwent a two stage validation process comprising (1) calibration using 2015–2019 regional aggregate data to parameterize climate scenarios and agent learning rates, and (2) independent validation against held out 2020–2024 growing seasons (30 seasons, three regions by ten seasons) that were not used during parameter estimation, allowing blind assessment of predictive performance. Calibration data sources included regional wheat yields from

FAOSTAT (FAO, 2025b) (Item Code 15, Area Codes IT, ES, FR, DE, DK, SE) which provided aggregate annual yield figures at the national level for each calibration year. It should be noted that FAOSTAT yield data do not include fertilizer application rates or application timing; these management parameters were therefore not sourced from FAO. Nitrogen application rates for the calibration period were instead derived from country-specific fertilizer survey statistics published by Eurostat (Farm Structure Survey series) and from CAPRI model parameterizations (Blanco et al., 2015), which disaggregate nitrogen inputs by crop type and region for EU member states. Application timing schedules were drawn from published agronomic nitrogen response analyses for European winter wheat systems, in which split applications near tillering (GS25–GS30) and stem elongation (GS31–GS37) are established as the practice that most closely aligns nitrogen supply with crop demand periods (Sylvester-Bradley and Kindred, 2009). Climate inputs were sourced from the E-OBS v25.0e gridded product (Cornes et al., 2018) for precipitation, ERA5-Land reanalysis (Muñoz Sabater, 2019) for temperature, AQUASTAT (FAO, 2024) for irrigation water use, and the CAPRI model (Blanco et al., 2015; EC, 2025) for energy consumption. These aggregate statistics informed the multiplicative scaling factor applied to climate scenario definitions and agent learning rate initialization but did not constrain held-out validation predictions.

The crop performance simulation generated field-level variability by adding stochastic perturbations drawn from zero-mean normal distributions with standard deviations of $\sigma_y = 0.15$ t ha^{-1} for yield and $\sigma_w =$ 15 mm for seasonal water use around empirically calibrated regional means. These values are consistent with published estimates of within-region wheat yield variability in European production systems; Becker-Reshef et al. (2010) report inter-field coefficients of variation of 5–12% for national-scale wheat yield datasets derived from MODIS-based regression models, and the $\sigma_y = 0.15$ t ha^{-1} value corresponds to approximately 3.1% of the regional mean yield under normal conditions (4.8 t ha^{-1}), placing it within the lower bound of this documented range for aggregated regional statistics. For seasonal water use, the $\sigma_w = 15$ mm value represents approximately 3.6% of the normal-condition conventional mean (420 mm), consistent with irrigation scheduling variability reported for Mediterranean and Central European wheat systems under deficit irrigation conditions (Fereres and Soriano, 2007). Both variability parameters were additionally validated by confirming that the resulting simulated distribution envelopes remained within the 10th and 90th percentile bounds of regional observations drawn from AQUASTAT and national agricultural survey datasets for the 2015–2019 calibration period. The validation dataset confirmed that simulated 90th percentile yield values did not exceed published national record yields for any of the six study countries.

2.8. Software implementation and algorithmic architecture

The hybrid artificial intelligence framework was implemented in Python 3.14 (released October 2025) using open-source libraries to ensure reproducibility and portability across computational platforms. Deep neural network components for reinforcement learning were built with TensorFlow 2.18 and the Keras API, using Actor-Critic architectures with fully connected layers, rectified linear unit activation functions, and the Adam optimizer with exponential learning rate decay. Classical machine learning tasks relied on scikit-learn 1.6 for supervised learning, cross-validation, and metric computation. Expert system production rules were executed through a forward-chaining inference engine with RETE-style pattern matching (Forgy, 1982), whereas fuzzy logic controllers used custom NumPy 2.1-based routines implementing Mamdani inference with minimum t-norm aggregation and centroid defuzzification. Multi-agent coordination was implemented with the Mesa 3.0 framework and domain-specific extensions that support heterogeneous agent decision frequencies (hourly, daily, weekly) via priority queue scheduling with timestamp ordering.

Reinforcement learning agents were trained through episodic interaction with simulated 120-day growing seasons, storing experience tuples in replay buffers of 10,000 transitions and sampling mini-batches of 64 transitions for off-policy updates. Exploration followed an ε-greedy strategy with exponential decay, where ε decreased from 0.20 to 0.01 with decay constant 0.005 (Eq. (8)), and convergence was monitored using a 100-episode moving average of cumulative rewards. Network weights were initialized with Xavier-Glorot uniform initialization to promote stable gradient propagation (Glorot and Bengio, 2010). The modular software architecture preserved a clear separation between expert system, fuzzy logic, reinforcement learning, and multi-agent components through standardized data interfaces, enabling independent modification of individual modules without disrupting overall system integration.

3. Results

3.1. Reinforcement learning training performance and algorithm comparison

Fig. 2 presents reinforcement learning convergence performance and training dynamics across four algorithms evaluated over 1000 training episodes. The algorithms implement temporal difference learning (Eq. (5)) with reward functions incorporating multiple sustainability objectives (Eq. (6)), Actor-Critic policy gradients (Eq. (7)), and ε-greedy exploration strategies with exponential decay (Eq. (8)).

Panel A illustrates cumulative reward trajectories demonstrating learning convergence across the four algorithms. Actor-Critic achieved the highest final cumulative reward of 519.7, followed by Deep Q-Network (462.7), Policy Gradient (440.2), and Q-Learning (409.6). All algorithms experienced an initial exploration phase (episodes 0–100) with negative or near-zero rewards, followed by rapid improvement between episodes 100–400, after which performance gains became more gradual. Actor-Critic exceeded the 400-reward benchmark around episode 450, demonstrating superior performance for agricultural resource optimization under climate variability. Trajectory variability persisted throughout training, reflecting the stochastic nature of agricultural decision environments and climate condition uncertainty across simulated growing seasons.

Panel B illustrates training loss reduction patterns, demonstrating how each algorithm minimized prediction errors during policy optimization. All methods exhibited initial loss values near 5.3, followed by exponential decay toward the convergence threshold of 1.0 (horizontal reference line). Convergence speeds varied substantially: Q-Learning converged fastest at episode 371, followed by Policy Gradient (401), Deep Q-Network (421), and Actor-Critic (441). Final stabilized loss values ranged from 0.059 (Policy Gradient) to 0.533 (Deep Q-Network), with most algorithms stabilizing below 0.35. Notably, Actor-Critic achieved the highest cumulative rewards (519.7) despite moderate final loss values (0.284), indicating that loss minimization alone does not fully predict practical agricultural decision-making performance. This finding consistent with the multi-objective nature of agricultural resource allocation where reward functions incorporate yield, water efficiency, energy consumption, and economic returns simultaneously.

3.2. Overall system performance across climate scenarios

In Fig. 3, the comparative performance of the AI-based and conventional farming approaches is quantified using yield enhancement metrics (Eq. (17)), resource efficiency indicators (Eq. (18)), and carbon footprint calculations (Eq. (19)).

The hybrid artificial intelligence framework consistently outperformed conventional management across all climate scenarios and performance dimensions. Panel A shows crop yield outcomes under normal, drought, and extreme heat conditions. Under normal conditions, the AI-based system achieved 5.2 ± 0.60 t/ ha^{-1} compared with

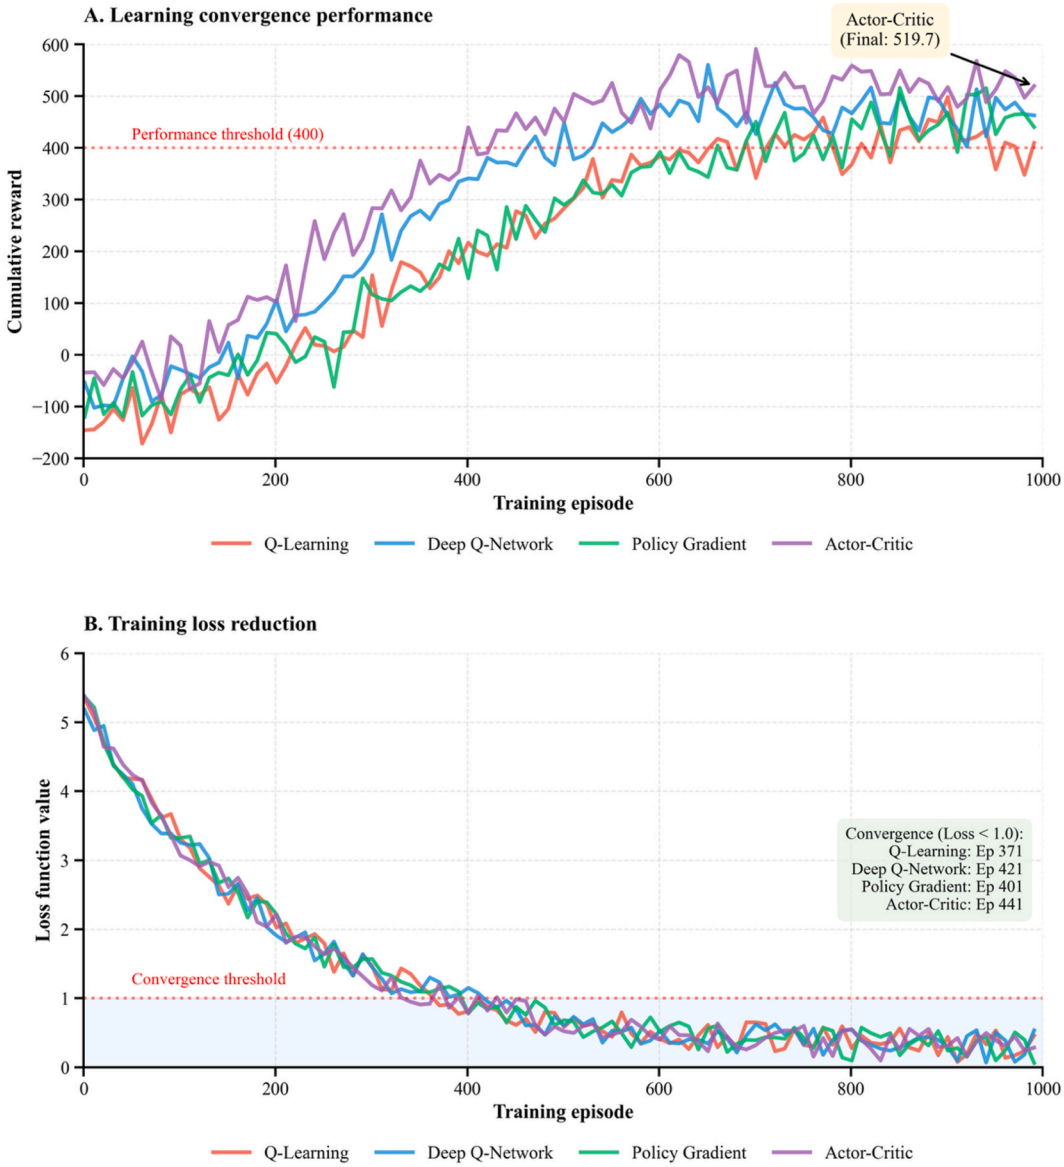

Fig. 2. Reinforcement learning convergence performance and training dynamics.

4.8 ± 0.5 t/ ha^{-1} for conventional management, corresponding to a 7.8% gain. This advantage intensified under climate stress: in drought scenarios, AI-based management produced 4.2 ± 0.5 t/ ha^{-1} versus 3.1 ± 0.4 t/ ha^{-1} for conventional practices (34.8% increase), while under extreme heat, yields were 4.4 ± 0.56 t/ ha^{-1} and 3.5 ± 0.4 t/ ha^{-1}, respectively (26.1% increase). These patterns indicate that the integrated framework is particularly effective at sustaining productivity under adverse climate conditions, when optimized resource allocation becomes critical.

Panel B depicts water consumption across the same scenarios, highlighting that the AI-based approach consistently maintained higher yields with lower irrigation inputs. Under normal conditions, conventional management used 420 ± 33 mm of water compared with 345 ± 27 mm for the AI-based system (17.8% reduction). Under drought, water use declined from 546 ± 42 mm (conventional) to 387 ± 33 mm (AI-based), a 29.1% reduction. During extreme heat, water requirements fell from 522 ± 39 mm to 368 ± 29 mm, corresponding to a 29.4% saving. These gains arise from fuzzy logic–based irrigation scheduling in which minimum t-norm rule evaluation (Eq. (3)) simultaneously weights soil moisture and crop water stress firing strengths rather than responding to a single threshold trigger, reinforcement learning-driven optimization of irrigation timing, and multi-agent coordination that avoids redundant applications in response to localized sensor signals.

Panel C summarizes energy consumption. Under normal conditions, AI-based management required 2962 ± 238 MJ/ ha^{-1}, compared with 3488 ± 266 MJ/ ha^{-1} for conventional practices (15.1% reduction). In drought scenarios, energy use decreased from 4066 ± 324 MJ/ ha^{-1} to 3324 ± 272 MJ/ ha^{-1} (18.3% reduction), while under extreme heat it declined from 4178 ± 295 MJ/ ha^{-1} to 3254 ± 222 MJ ha^{-1} (22.1% reduction). These improvements primarily reflect optimized pump operation during off-peak electricity periods and lower fertilizer demand due to precision nutrient management implemented by the expert system.

Panel D demonstrates carbon footprint performance across climate scenarios, revealing consistent greenhouse gas emission reductions achieved through AI-based management. Under normal conditions, the AI-based approach produced 1056 ± 19 kg CO_2/ha compared with 1299 ± 21 kg CO_2/ha for conventional management, representing an 18.7% reduction in carbon intensity. Under drought conditions,

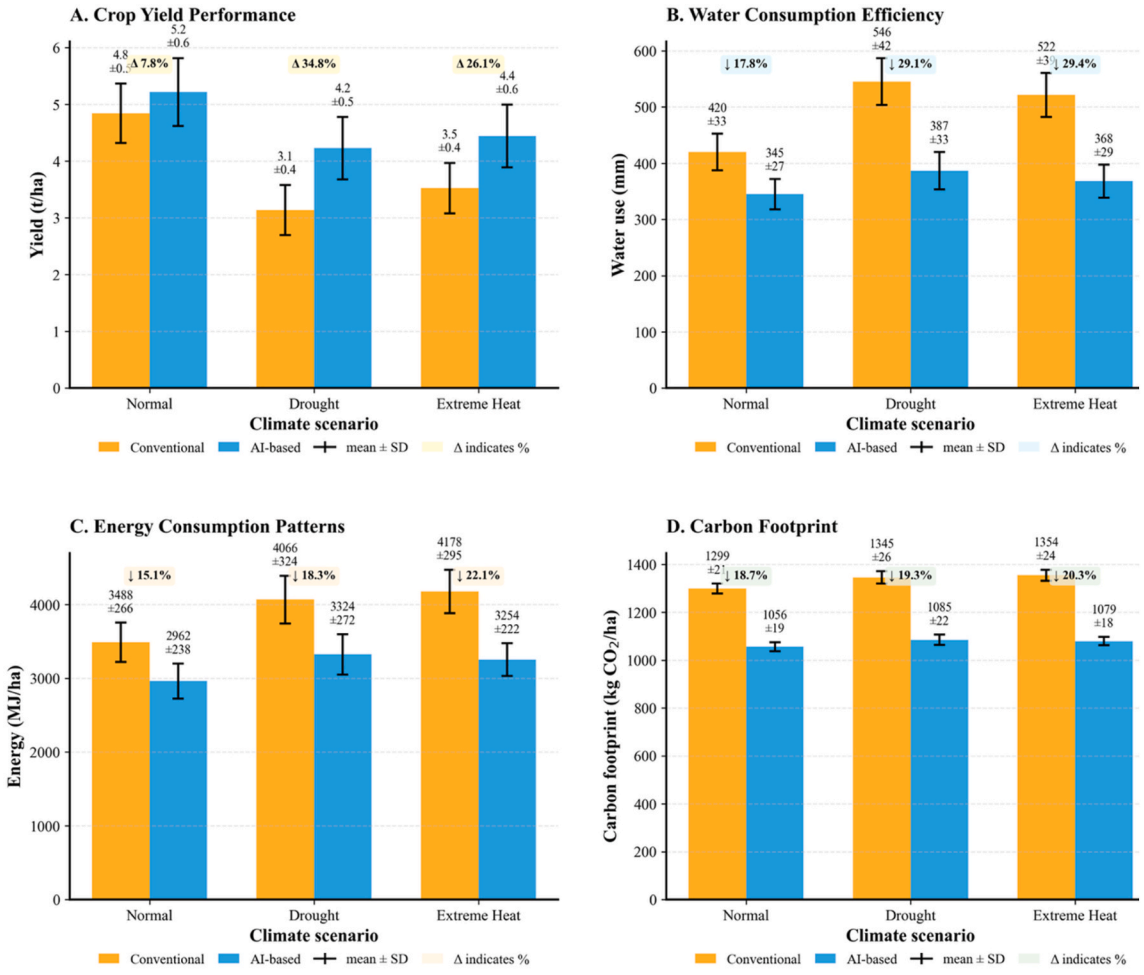

Fig. 3. Comparative wheat performance under AI-based and conventional management across climate scenarios.

emissions were 1085 ± 22 kg CO_2/ha for AI-based systems versus 1345 ± 26 kg CO_2/ha conventionally, corresponding to a 19.3% reduction, while extreme heat scenarios showed 1079 ± 18 kg CO_2/ha under AI management compared with 1354 ± 24 kg CO_2/ha conventionally, yielding a 20.3% mitigation. The progressive increase in carbon reduction percentages under climate stress conditions reflects the framework's capacity to maintain productivity through efficient resource allocation, whereby reduced energy consumption for irrigation pumping, lower fertilizer application requirements guided by the expert system, and optimized field operation timing collectively distribute emissions across higher yield outputs (Rabbi, 2025d). These results demonstrate that climate adaptation strategies enabled by intelligent agricultural management simultaneously advance both productivity resilience and greenhouse gas mitigation objectives, supporting the dual

Table 4
Crop performance metrics across climate scenarios and management approaches.

Climate scenario	Approach	Yield (t/ha)	Water use (mm)	Energy (MJ/ha)	Carbon (kg CO_2/ha)
Normal	Conventional	4.8 ± 0.5	420 ± 33	3488 ± 266	1299 ± 21
	AI-based	5.2 ± 0.60	345 ± 27	2962 ± 238	1056 ± 19
Drought	Conventional	3.1 ± 0.4	546 ± 42	4066 ± 324	1345 ± 26
	AI-based	4.2 ± 0.5	387 ± 33	3324 ± 272	1085 ± 22
Extreme Heat	Conventional	3.5 ± 0.4	522 ± 39	4178 ± 295	1354 ± 24
	AI-based	4.4 ± 0.56	368 ± 29	3254 ± 222	1079 ± 18

Notes:
- Performance metrics calculated using yield improvement formulations (Eq. (17)), resource efficiency indicators (Eq. (18)), and carbon footprint quantification (Eq. (19)).
- Values represent means ± standard deviations calculated across 10 growing seasons and 3 European regions (Mediterranean, Central Europe, Nordic) with 30 observations per climate scenario-approach combination.
- Climate scenarios defined by deviations from 30-year climate normal: Normal (±10% precipitation, ±1 °C temperature), Drought (>30% precipitation deficit, +1–2 °C), Extreme Heat (±10% precipitation, >2 °C temperature increase).
- Carbon footprint includes emissions from energy consumption (0.08 kg CO_2/MJ), fertilizer production (5.5 kg CO_2/kg N), and pesticide manufacturing (12 kg CO_2/kg active ingredient).
- Standard deviations reflect combined variability from seasonal weather fluctuations, regional soil and climate differences, and stochastic model components in reinforcement learning agent decisions.
- Statistical significance of AI-conventional differences confirmed through paired t-tests ($p < 0.001$ for all metrics across all climate scenarios, $n = 30$ per comparison).
- All percentage improvements are computed from unrounded sample means; tabulated values are rounded to one decimal place and may yield marginally different percentages upon direct recalculation.

goals of climate-smart agriculture under increasing environmental uncertainty.

The temporal dynamics illustrated in Fig. 3 demonstrate qualitative patterns of system performance across growing seasons. To quantify these performance differences with statistical precision, Table 4 presents a comprehensive summary of crop performance metrics across climate scenarios, including mean values and standard deviations for yield, water use, energy consumption, and carbon footprint.

The hybrid artificial intelligence framework in Table 4 demonstrated consistent performance advantages over conventional agricultural management practices across multiple sustainability metrics and climate scenarios, with benefits most pronounced under environmental stress conditions that compromise conventional decision-making effectiveness. Under normal climate conditions, the AI-based approach achieved yields of 5.2 ± 0.60 tons per hectare compared to conventional yields of 4.8 ± 0.5 tons per hectare, representing a 7.8 percent improvement while simultaneously reducing water consumption from 420 ± 33 mm to 345 ± 27 mm, corresponding to 17.8 percent water savings, and decreasing energy requirements from 3488 ± 266 MJ per hectare to 2962 ± 238 MJ per hectare, yielding 15.1 percent energy efficiency gains. Carbon footprints declined from 1299 ± 21 to 1056 ± 19 kg carbon dioxide equivalent per hectare, representing an 18.7 percent reduction.

Drought scenario results revealed more substantial AI framework advantages, with yields of 4.2 ± 0.5 tons per hectare under AI-based management compared to severely compromised conventional yields of 3.1 ± 0.4 tons per hectare, representing a 34.8 percent improvement that reflects the system's capacity to optimize limited water resources through precise irrigation scheduling guided by fuzzy logic controllers and reinforcement learning temporal optimization. Water consumption under drought conditions decreased from 546 ± 42 mm conventionally to 387 ± 33 mm with AI management, corresponding to 29.1 percent savings despite the yield enhancement, demonstrating improved water use efficiency. Energy consumption declined from 4066 ± 324 MJ per hectare to 3324 ± 272 MJ per hectare, yielding 18.3 percent savings, while carbon emissions were 1345 ± 26 kg carbon dioxide per hectare conventionally and 1085 ± 22 kg carbon dioxide per hectare under AI management, representing a 19.3 percent reduction despite the additional resource demands under water-limited conditions.

Extreme heat scenarios showed intermediate AI advantages between normal and drought conditions, with yields improving from 3.5 ± 0.4 tons per hectare conventionally to 4.4 ± 0.56 tons per hectare under AI management, corresponding to a 26.1 percent enhancement. Water use declined from 522 ± 39 mm to 368 ± 29 mm, yielding 29.4 percent savings, while energy consumption decreased from 4178 ± 295 MJ per hectare to 3254 ± 222 MJ per hectare, representing 22.1 percent efficiency gains. Carbon emissions under extreme heat reached 1354 ± 24 kg carbon dioxide per hectare conventionally and 1079 ± 18 kg carbon dioxide per hectare under AI management, representing a 20.3 percent reduction. Carbon footprints across all climate scenarios demonstrated consistent reductions of 18.7–20.3 percent with AI-based management, reflecting the combined effects of reduced input requirements and improved crop productivity that distributes emissions across higher yield outputs.

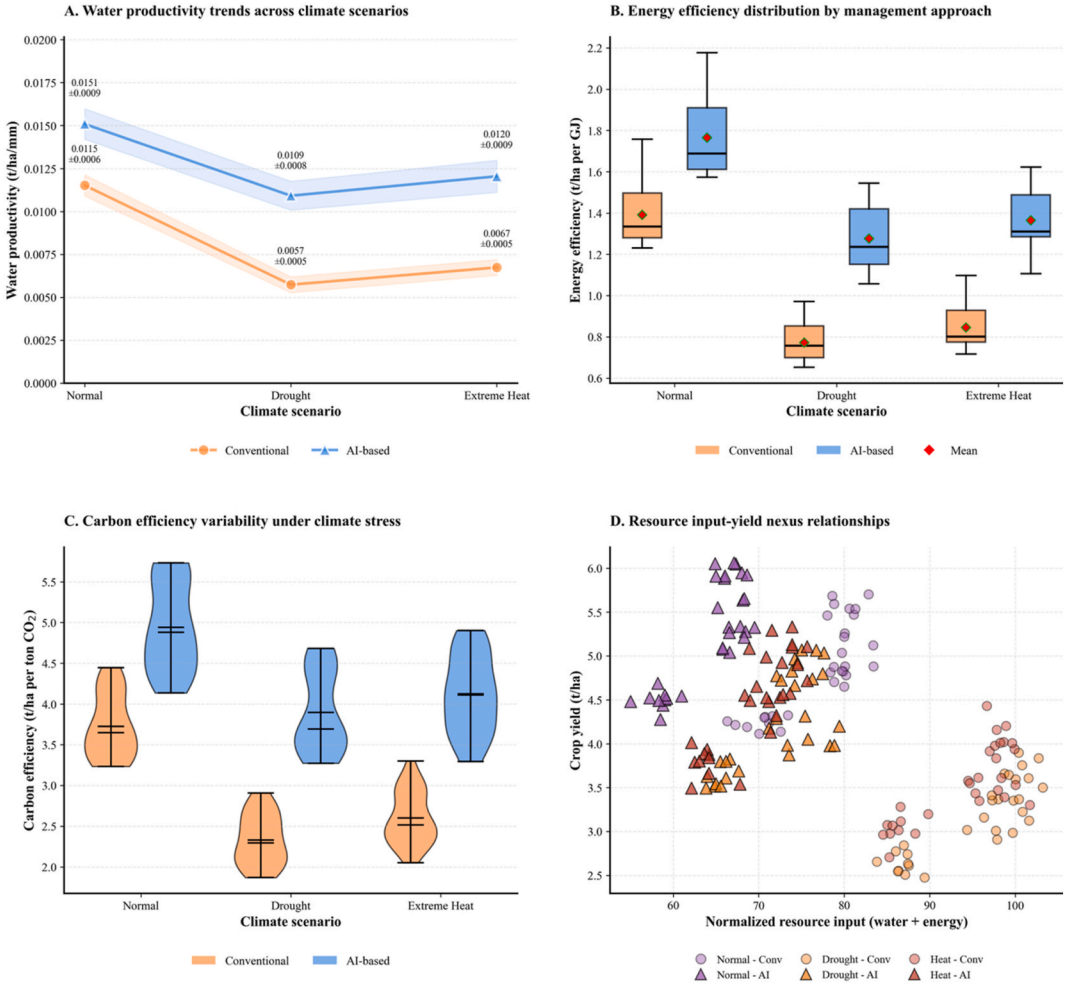

Fig. 4. Resource use efficiency and water-energy-food nexus performance.

3.3. Resource efficiency and water-energy-food nexus

Resource use efficiency in Fig. 4 is calculated with the water productivity, energy efficiency and carbon efficiency formulas defined in Equation (16). Together, these indicators show how the framework changes the balance between water, energy and food production under different climate conditions.

The hybrid AI framework uses water, energy and carbon more efficiently than conventional management in every climate scenario, producing more grain for the same or lower resource use.

Panel A shows water productivity for normal weather, drought and extreme heat. Under normal conditions the AI based system reaches 0.0151 ± 0.0009 t/ha^{-1}/mm^{-1}, while conventional management reaches 0.0115 ± 0.0006 t/ha^{-1}/mm^{-1} (about one third lower). Under drought the advantage is almost doubled, with 0.0109 ± 0.0008 t/ha^{-1}/mm^{-1} for AI compared with 0.0057 ± 0.0005 t/ha^{-1}/mm^{-1} for conventional practice, and under extreme heat the AI system again keeps higher water productivity (0.0120 ± 0.0009 versus 0.0067 ± 0.0005 t/ha^{-1}/mm^{-1}). These gains come from irrigation schedules that follow crop demand more closely rather than applying uniform doses.

Panel B compares energy efficiency distributions, expressed as crop yield per unit of energy input (t/ha^{-1} per GJ^{-1}), across climate scenarios and management approaches. The AI-based strategy achieved consistently higher energy efficiency under all climate conditions: under normal scenarios, AI-based management produced 1.766 ± 0.186 t/ha^{-1} per GJ^{-1} compared with 1.391 ± 0.141 t/ha^{-1} per GJ^{-1} for conventional practices, representing a 27.0% improvement. Under drought conditions, energy efficiency declined for both approaches but maintained AI superiority at 1.276 ± 0.145 t/ha^{-1} per GJ^{-1} versus 0.772 ± 0.082 t/ha^{-1} per GJ^{-1} conventionally, corresponding to a 65.3% advantage. Extreme heat scenarios showed similar patterns with AI-based efficiency of 1.364 ± 0.128 t/ha^{-1} per GJ^{-1} compared with 0.845 ± 0.078 t/ha^{-1} per GJ^{-1} for conventional management, yielding a 61.4% improvement. The box plot distributions reveal narrower interquartile ranges for AI-based systems under normal conditions, indicating more consistent energy performance, whereas both approaches show greater variability under climate stress. These energy efficiency gains arise primarily from optimized irrigation pump scheduling during off-peak electricity pricing periods, reduced fertilizer and pesticide application through precision nutrient management guided by the expert system and minimized field operation redundancy coordinated by the multi-agent framework, all achieved without compromising crop yield.

Panel C illustrates carbon efficiency variability, expressed as crop yield per unit of carbon dioxide emissions (t/ha^{-1} per ton CO_2), using violin plots that reveal full distributional characteristics across climate scenarios. Under normal conditions, AI-based management achieved 4.939 ± 0.534 t/ha^{-1} per ton CO_2 compared with 3.726 ± 0.383 t/ha^{-1} per ton CO_2 for conventional practices, representing a 32.6% improvement in carbon efficiency. The performance gap widened substantially under climate stress: drought scenarios showed AI-based carbon efficiency of 3.896 ± 0.412 t/ha^{-1} per ton CO_2 versus 2.330 ± 0.287 t/ha^{-1} per ton CO_2 conventionally (67.2% advantage), while extreme heat conditions yielded 4.112 ± 0.368 t/ha^{-1} per ton CO_2 for AI management compared with 2.601 ± 0.315 t/ha^{-1} per ton CO_2 for conventional approaches (58.1% improvement).

Panel D presents resource input-yield relationships through a scatter plot mapping crop yield against normalized combined water and energy inputs, with data points differentiated by climate scenario (triangles for AI-based, circles for conventional) and color-coded by stress condition (purple for normal, orange for drought, pink for extreme heat). Under drought and extreme heat scenarios, conventional management typically requires 85 to 105 normalized resource units (combining water mm and energy MJ ha^{-1} scaled to comparable dimensionless indices) to achieve yields ranging from 2.5 to 4.0 t/ha^{-1}, forming a dispersed cloud in the lower-right quadrant of the input-yield space. In contrast, the AI-based framework achieves yields of 4.5–6.0 t/ha^{-1} with only 60 to 75 normalized resource units, occupying the upper-left region that represents high productivity with low resource intensity. Normal climate conditions for both management approaches show intermediate clustering between these extremes, though AI-based systems maintain separation toward higher yield-to-input ratios.

3.4. Regional performance variability and spatial adaptation

Regional performance variations in Fig. 5 reflect the distributed agent coordination mechanisms (Eqs. (9)–(11)) and localized learning parameters specified in Table 3.

The hybrid AI framework demonstrated distinct performance patterns across three European agricultural regions, with differential climate adaptation capabilities shaped by baseline productivity levels and regional climate vulnerabilities.

Panel A presents regional yield trajectories under AI-based management across climate scenarios. Central Europe achieved the highest productivity at 5.90 ± 0.15 t/ha^{-1} under normal conditions, followed by Nordic regions at 5.25 ± 0.15 t/ha^{-1} and Mediterranean zones at 4.50 ± 0.10 t/ha^{-1}. All regions experienced yield declines under climate stress, with drought reducing Central Europe to 4.89 ± 0.14 t/ha^{-1} (17.1% decline), Nordic to 4.15 ± 0.23 t/ha^{-1} (20.9% decline), and Mediterranean to 3.64 ± 0.13 t/ha^{-1} (19.1% decline). Under extreme heat, yields partially recovered: Central Europe reached 5.06 ± 0.18 t/ha^{-1}, Nordic 4.48 ± 0.15 t/ha^{-1}, and Mediterranean 3.79 ± 0.17 t/ha^{-1}. The larger variability in Nordic regions under drought (±0.23 t ha^{-1}) likely reflects less frequent exposure to severe water stress and correspondingly lower adaptation capacity.

Panel B displays water efficiency gains as a heatmap, showing percentage reductions in water use relative to conventional management. Under normal conditions, all regions achieved similar savings: Mediterranean 18.1%, Central Europe 17.9%, and Nordic 17.5%. Climate stress amplified these gains substantially. During drought, Mediterranean zones reached 29.9% savings, Central Europe 29.3%, and Nordic 28.1%. Extreme heat produced comparable improvements: Nordic achieved the highest savings at 30.0%, followed by Central Europe at 29.3% and Mediterranean at 29.0%. The consistent intensification of water savings under stress conditions demonstrates that the fuzzy logic irrigation controller and multi-agent coordination provide greater value when water becomes limiting.

Panel C illustrates carbon footprint reduction variability through a scatter plot with error bars. Under normal conditions, reductions were similar across regions: Mediterranean 18.8 ± 1.0%, Central Europe 18.8 ± 0.9%, and Nordic 18.5 ± 1.2%. Drought conditions slightly increased carbon savings, with Central Europe reaching 19.7 ± 0.9%, Nordic 19.2 ± 1.2%, and Mediterranean 19.1 ± 1.0%. Extreme heat yielded the largest reductions: Nordic achieved 20.6 ± 1.2%, Mediterranean 20.4 ± 1.0%, and Central Europe 19.9 ± 0.9%. The upward trend under climate stress reflects the AI framework's ability to maintain productivity through efficient resource use, thereby distributing emissions across higher outputs.

Panel D presents energy efficiency gains as horizontal bars grouped by region and climate scenario. Under normal conditions, savings ranged from 14.5 ± 2.5% (Nordic) to 15.3 ± 5.2% (Central Europe) and 15.2 ± 2.9% (Mediterranean). Drought increased energy savings to 19.5 ± 3.2% in Central Europe, 17.9 ± 3.8% in Nordic, and 17.1 ± 3.8% in Mediterranean zones. Extreme heat produced the highest energy efficiency gains: Nordic reached 23.1 ± 3.6%, Mediterranean 22.6 ± 2.8%, and Central Europe 20.4 ± 1.6%. The pronounced advantage under extreme heat reflects optimized pump scheduling and reduced cooling requirements through precision irrigation timing.

3.5. Temporal dynamics of crop development and resource management

Temporal growth dynamics in Fig. 6 are governed by reinforcement learning reward functions (Eq. (6)) optimized through Actor-Critic

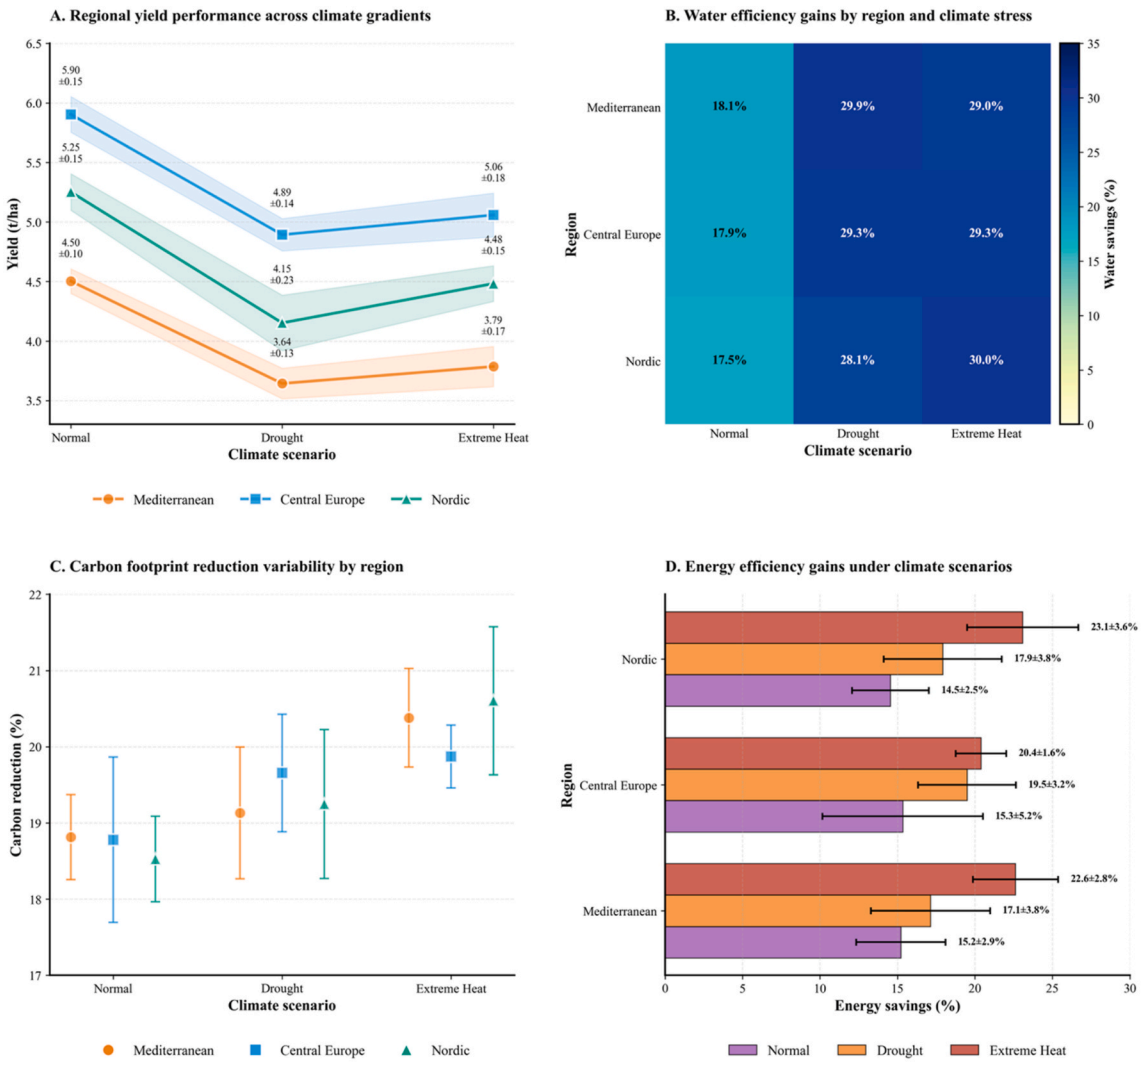

Fig. 5. Regional performance variability across European agricultural zones under climate scenarios.

policy updates (Eq. (7)) integrated with fuzzy logic control outputs (Eqs. (2)–(4)).

The hybrid AI framework produced distinct temporal patterns in crop development and resource use across the 120-day growing season, demonstrating how dynamic optimization influences plant growth and input timing decisions.

Panel A presents aboveground biomass accumulation with phenological stages annotated. Both management approaches followed sigmoidal growth curves typical of winter wheat: gradual increase during emergence (days 0–20), acceleration through tillering and stem elongation (days 20–70), and plateau during flowering and grain filling (days 70–120). The AI-based system reached final biomass of approximately 14,000 kg ha^{-1} compared with 11,700 kg ha^{-1} for conventional management, representing a 19.7% improvement. The divergence became visible around day 60 during stem elongation, when optimized nitrogen timing and irrigation scheduling provided maximum benefit for vegetative growth. Shaded confidence bands show consistent variance throughout the season for both approaches.

Panel B illustrates leaf area index (LAI) dynamics and canopy development. The AI-based system achieved a peak LAI of 7.5 m^2/m^{-2} compared with 6.4 m^2/m^{-2} for conventional management, representing a 17.2% advantage in light interception capacity. Both systems showed similar early development with LAI below 3.0 during the first 50 days as seedlings established. The critical LAI threshold of 3.0 (illustrated by the dashed horizontal line), representing 95% light interception efficiency, was surpassed approximately 10 days earlier by the AI-based system, thereby extending the period of maximum photosynthetic activity. After day 100, both systems showed parallel LAI decline due to natural leaf senescence during grain filling, though the AI-based approach maintained higher values throughout.

Panel C displays water stress index fluctuations alongside irrigation management strategies. The conventional approach exhibited regular oscillations with peaks reaching 0.4–0.5, reflecting fixed-schedule irrigation every 12–15 days regardless of actual crop demand. Reference thresholds for high stress (0.6) and moderate stress (0.3) are indicated by dashed lines. The AI-based system-maintained water stress predominantly below 0.3 through variable irrigation timing guided by fuzzy logic processing of soil moisture, weather forecasts, and crop phenology. The presence of blue triangular markers is indicative of AI irrigation events, characterized by increased frequency but reduced magnitude of applications. These events were observed to prevent the accumulation of stress while concomitantly reducing total seasonal water usage. The reduced stress amplitude under AI management (peaks rarely exceeding 0.3 vs. conventional peaks near 0.5) sustained photosynthetic activity and minimized yield losses during critical reproductive stages.

Panel D presents daily nitrogen uptake rates with fertilizer application timing overlaid. Both systems showed similar uptake patterns until day 35, after which the AI-based approach demonstrated higher daily uptake rates. During the period of maximum nitrogen demand (days 50–80), the AI-based system reached peak uptake rates near 2.6 kg ha^{-1}

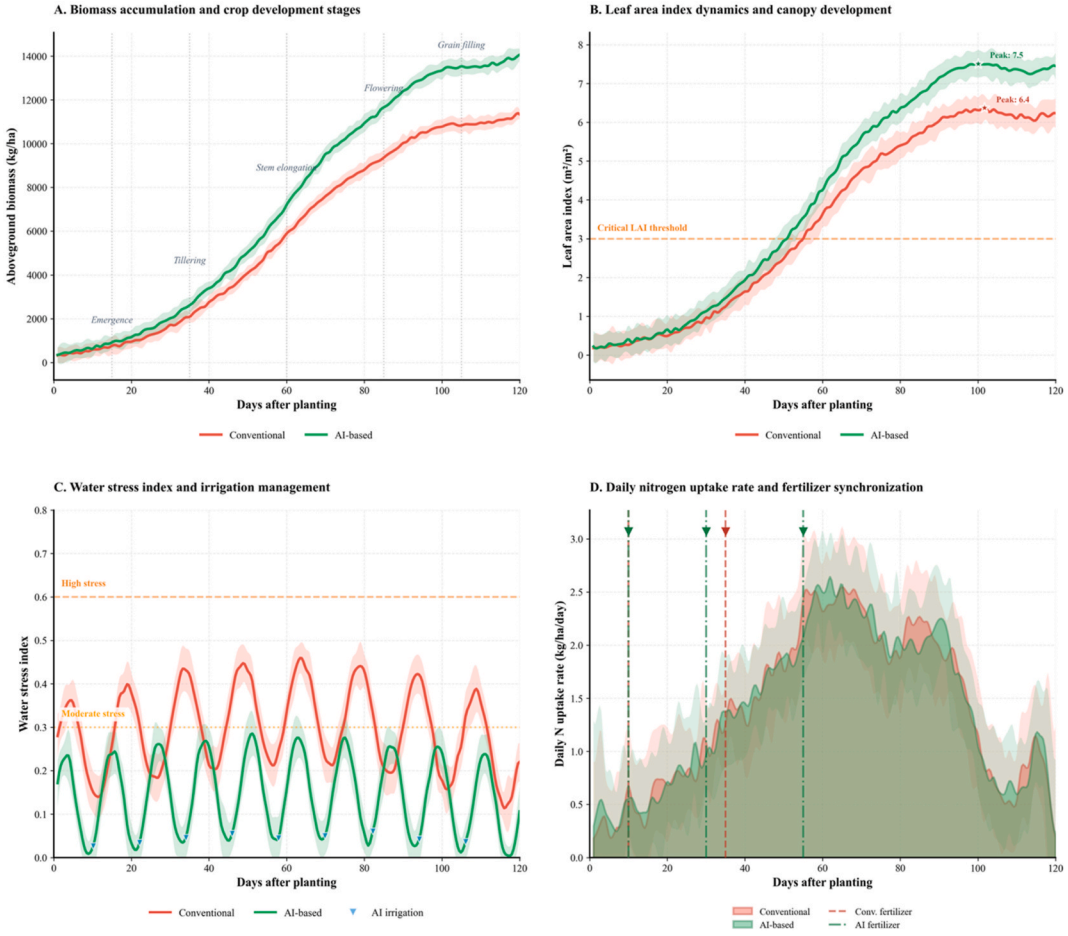

Fig. 6. Temporal dynamics of crop development and resource management.

day^{-1} compared with approximately 2.0 kg ha^{-1} day^{-1} for conventional management. Vertical dashed lines with arrows mark fertilizer applications: conventional applications (red) occurred at fixed intervals around days 10 and 37, following predetermined calendar-based schedules, while AI-based applications (green) were strategically timed at days 10, 35, and 57 based on the expert system's real-time assessment of crop nitrogen status and growth stage requirements. The AI system's fertilizer timing aligned more closely with peak uptake demand periods, particularly evident in the third application at day 57 that precedes the sustained high uptake phase during grain filling. This improved synchronization enhanced nitrogen use efficiency by minimizing losses to leaching and volatilization that occur when fertilizer is applied during periods of low crop demand, while ensuring adequate nitrogen availability during critical growth stages when uptake capacity is highest. The shaded confidence bands reveal consistent variance for both approaches, though the AI-based system maintains more stable elevated uptake rates throughout the mid-to-late season.

3.6. Parameter sensitivity analysis and model robustness

Parameter sensitivity analysis in Fig. 7 examines variations in learning rates and discount factors (Eq. (5) and (8)) across reinforcement learning agents and fuzzy logic membership function parameters (Eq. (2)).

The hybrid AI framework exhibited varying degrees of sensitivity to key operational parameters, revealing which components require precise calibration and which demonstrate robust performance across parameter ranges.

Panel A presents yield sensitivity to parameter variations spanning −30% to +30% from baseline values across five critical parameters: irrigation threshold, fertilizer rate, learning rate, discount factor, and climate variability. The shaded regions represent 95% confidence intervals, with green shading indicating positive yield impact zones and pink shading denoting negative impact regions. Irrigation threshold demonstrated the highest sensitivity, exhibiting an asymmetric response whereby −30% variation (earlier irrigation triggering) produced approximately +15% yield gains, whereas +30% variation (delayed irrigation) caused roughly −10% yield losses. This pattern indicates that lowering the soil moisture threshold enhances productivity by reducing water stress during critical growth periods, though at the trade-off of increased seasonal water consumption. Fertilizer rate showed moderate sensitivity with approximately +8% yield gains under +30% increases and −5% yield reductions under −30% decreases, reflecting the crop's capacity to utilize additional nitrogen inputs up to physiological saturation limits. Learning rate exhibited asymmetric behavior, with −30% variations reducing yields by approximately −9% due to insufficient policy exploration during reinforcement learning training, while +30% variations caused smaller losses around −8% due to training instability and suboptimal convergence. Discount factor displayed the most modest sensitivity, showing approximately +11% yield improvement at −30% variation and −7% yield reduction at +30% variation, suggesting that temporal credit assignment remains robust across a wide parameter range and that the framework tolerates variation in future reward weighting. Climate variability demonstrated −30% variations improving yields by roughly +6.5% through reduced environmental uncertainty, while +30% increases caused −8% yield reductions, reflecting the framework's adaptive capacity to buffer moderate climate fluctuations through intelligent decision-making while remaining

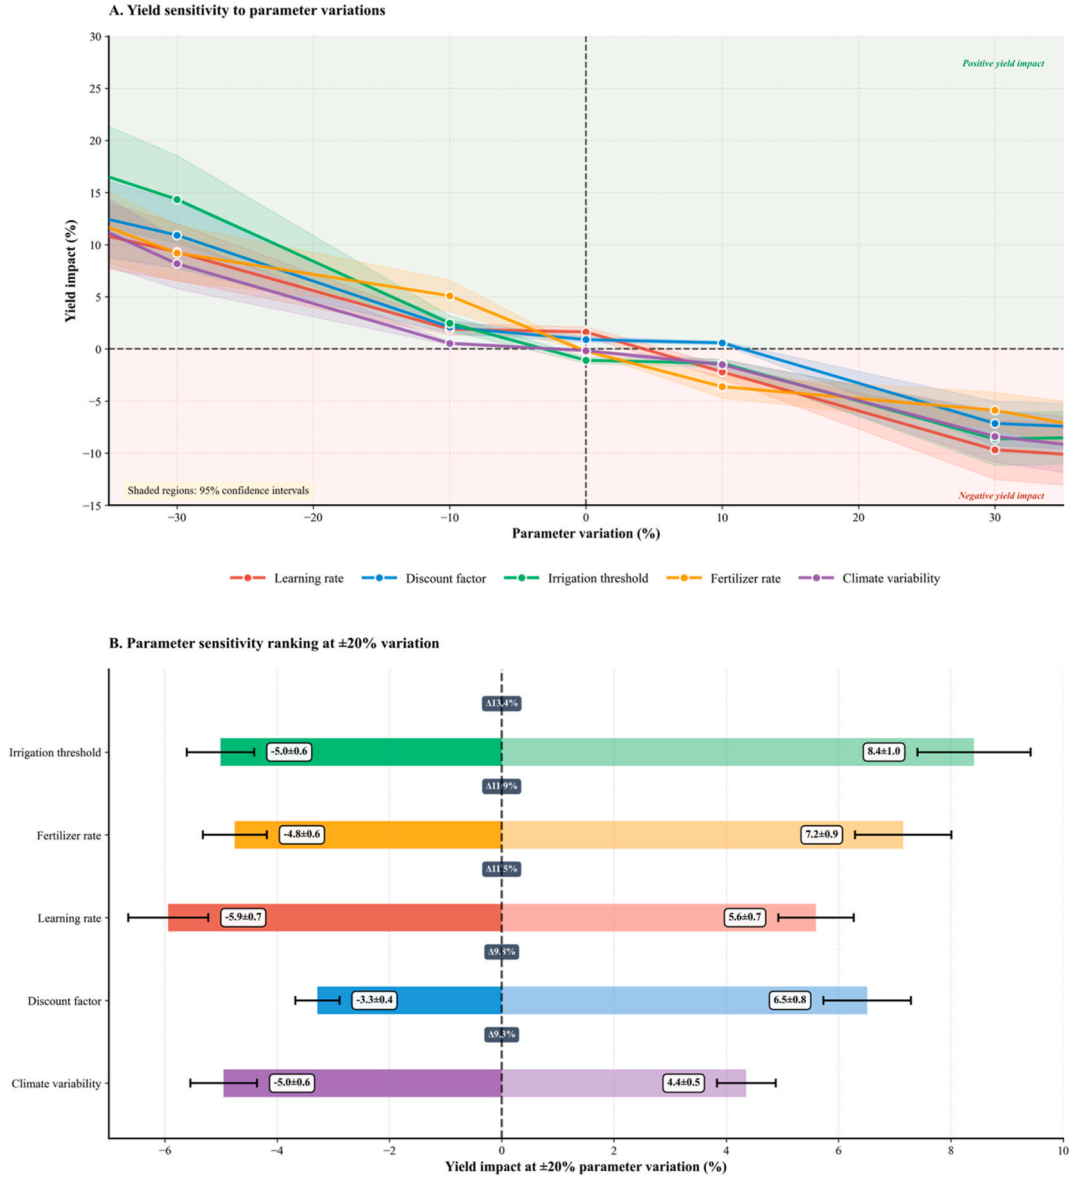

Fig. 7. Parameter sensitivity analysis and model robustness assessment.

vulnerable to extreme stochastic conditions that exceed the system's optimization capabilities.

Panel B presents a tornado diagram ranking parameter sensitivity at ±20% variation, enabling direct comparison of relative parameter importance. Irrigation threshold exhibited the largest total sensitivity range of 13.4 percentage points, with +20% parameter increase producing yield improvements of 8.4 ± 1.0% and −20% decrease causing yield reductions of −5.0 ± 0.6%. Fertilizer rate ranked second with a total range of 11.9 percentage points (+20%: 7.2 ± 0.9%; −20%: −4.8 ± 0.6%). Learning rate showed a sensitivity range of 11.5 percentage points with nearly symmetric impacts (+20%: 5.6 ± 0.7%; −20%: −5.9 ± 0.7%), indicating that both over- and under-parameterization impair performance similarly. Discount factor demonstrated a 9.8 percentage point range (+20%: 6.5 ± 0.8%; −20%: −3.3 ± 0.4%), while climate variability showed the lowest sensitivity at 9.3 percentage points (+20%: 4.4 ± 0.5%; −20%: −5.0 ± 0.6%). The error bars confirm statistical significance for all estimates, with irrigation threshold and fertilizer rate exhibiting the narrowest confidence bands, reflecting consistent sensitivity responses across simulations. These results indicate that irrigation threshold and fertilizer rate require careful calibration, whereas discount factor and climate variability parameters demonstrate acceptable robustness across typical operational ranges.

3.7. Model validation against empirical agricultural data

Model validation in Fig. 8 compares framework predictions against independent empirical observations from European winter wheat production systems over 30 growing seasons (2020–2024), with prediction accuracy quantified using the validation score formulation in Eq. (20). Residual normality is assessed with the Shapiro–Wilk test, where p values greater than 0.05 indicate no statistically significant deviation from a normal distribution, supporting the use of Gaussian error assumptions for uncertainty analysis.

The hybrid AI framework shows adequate predictive skill when evaluated against aggregated empirical data, providing simulation-level validation evidence that establishes a basis for prospective field-level evaluation across contrasting European environments.

Panel A reports crop yield validation as a scatter plot of predicted versus observed yields, with a regression line and 95% confidence band superimposed on the 1:1 reference line. The framework attains

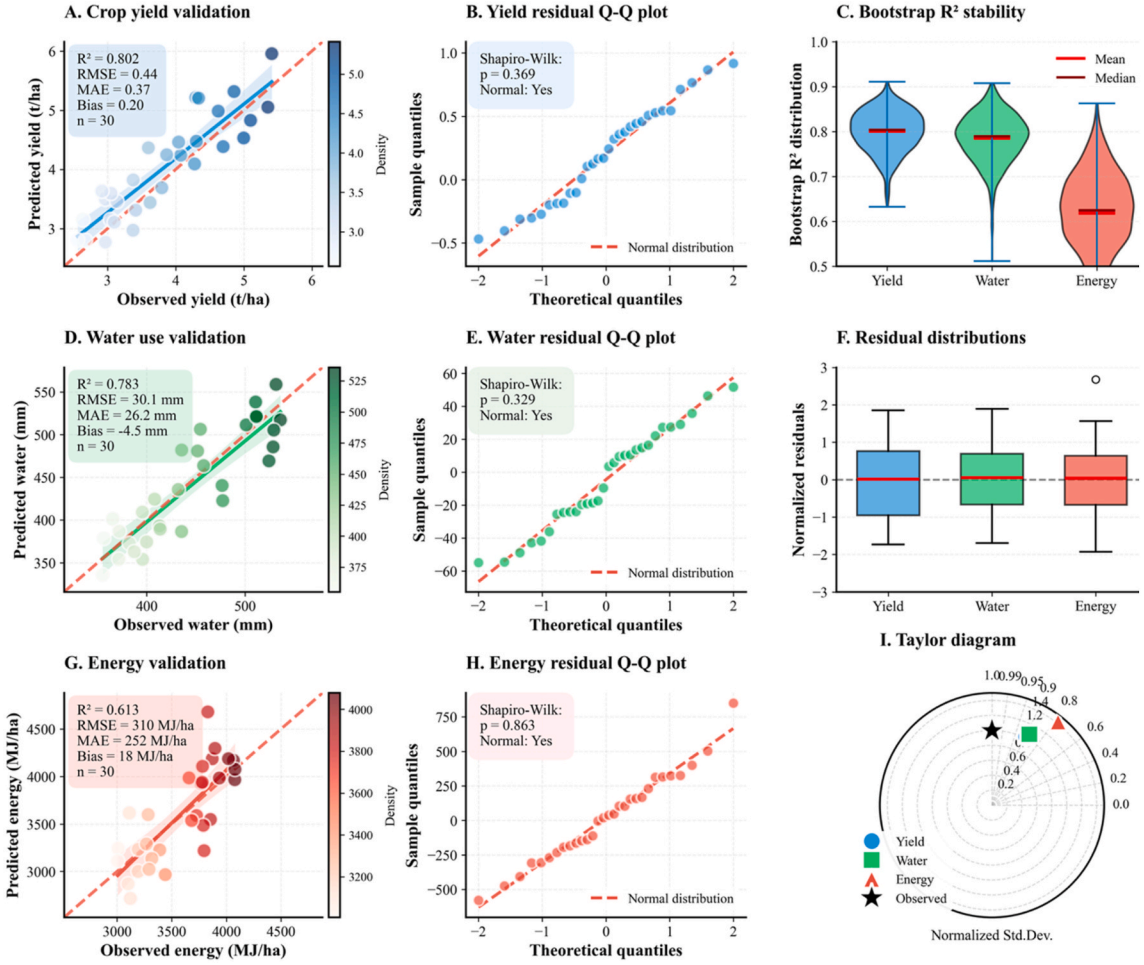

Fig. 8. Empirical validation of the hybrid AI framework against independent observations.

R^2 = 0.802 (95% CI: 0.70–0.89), RMSE = 0.44 t/ha^{-1}, MAE = 0.37 t/ha^{-1}, and a small positive bias of 0.20 t/ha^{-1}, with points tightly clustered along the 1:1 line over the 2.5–6.0 t/ha^{-1} range, indicating robust performance across low-to high-yielding conditions. Panel B shows the corresponding yield residual Q–Q plot; residuals align closely with the theoretical normal quantiles and the Shapiro–Wilk p = 0.369 supports approximate normality, indicating that yield errors are well-behaved and suitable for parametric uncertainty propagation.

Panel C summarizes bootstrap R^2 stability for 1000 resamples, with violin plots for yield, water, and energy. Yield exhibits the highest and narrowest R^2 distribution (median ≈ 0.80), followed by water (≈0.78) and energy (≈0.62), showing that yield predictions are both accurate and stable across resampled subsets, whereas energy predictions are more variable, consistent with their more complex drivers.

Panel D evaluates water use, achieving R^2= 0.783 (95% CI: 0.66–0.88), RMSE = 30.1 mm, MAE = 26.2 mm, and a modest negative bias of −4.5 mm over the 350–550 mm range. The tight clustering around the 1:1 line across both water-limited and non-limiting regimes indicates that the framework reliably captures seasonal water consumption patterns. Panel E presents the water residual Q–Q plot; residuals track the theoretical diagonal with Shapiro–Wilk p = 0.329, again indicating no significant departure from normality over the full residual range and supporting unbiased representation of hydrological variability. Panel F presents normalized residual distributions as boxplots, providing a complementary view of prediction error patterns across the three validated metrics. Yield residuals (blue) are centered near zero with a median close to 0.0 and symmetric interquartile range spanning approximately −0.8 to +0.8, indicating unbiased predictions with balanced overprediction and underprediction across the validation dataset. Water residuals (green) display similar characteristics with a median near zero and interquartile range from approximately −0.7 to +0.7, confirming consistent model performance for irrigation predictions. Energy residuals (red) show a median near zero but exhibit a slightly wider distribution, with interquartile range spanning approximately −0.6 to +0.6 and one notable outlier exceeding +2.5 standard deviations.

Panel G summarizes energy use validation, with R^2 = 0.613 (95% CI: 0.43–0.78), RMSE = 310 MJ ha^{-1}, MAE = 282 MJ/ha^{-1}, and a positive bias of 118 MJ/ha^{-1} across the 3000–4500 MJ/ha^{-1} range. Although less accurate than yield and water, the model captures broad gradients in energy consumption, while the consistent overestimation likely reflects regional differences in energy pricing, equipment efficiency, and management intensity not fully resolved by the input data. Panel H shows that energy residuals also approximate normality (Shapiro–Wilk p = 0.863), but with somewhat heavier tails than yield and water, indicating occasional larger deviations that align with the observed outlier structure.

As illustrated in Panel I, a Taylor diagram has been employed to simultaneously visualize three validation statistics. These are the correlation coefficient, the normalized standard deviation, and the centered root mean square error. The purpose of presenting these is to predict the yield, water use and energy consumption. The black star at the reference point (1.0, 1.0) represents perfect agreement between observed and predicted values. The radial distance from the origin indicates normalized standard deviation of predictions relative to observations, the angular position denotes correlation coefficient (with the arc labeled

from 0.0 to 1.0), and the distance from the reference point represents normalized centered RMSE. Yield predictions (blue circle) achieved the strongest performance with a correlation coefficient of approximately 0.80, normalized standard deviation near 0.90, and minimal distance from the reference point, indicating close alignment between predicted and observed variability patterns. Water use predictions (green square) showed correlation of approximately 0.78 and normalized standard deviation near 1.0, positioning slightly farther from the reference point but maintaining acceptable prediction accuracy. Energy consumption predictions (red triangle) exhibited lower correlation of approximately 0.61 and normalized standard deviation near 1.2, with greater distance from the reference point reflecting the systematic bias and outlier effects observed in Panel G. The clustering of all three metrics within the upper-right quadrant of the diagram confirms that the hybrid artificial intelligence framework maintains positive correlation and reasonable variance agreement across all validated outputs, with yield demonstrating superior predictive skill followed by water use and energy consumption. This integrated visualization corroborates the individual validation statistics reported in Panels A, D, and G, demonstrating that the framework achieves acceptable multi-metric simulation-level validation performance, which constitutes a necessary precondition for agricultural decision support applications; whether these outcomes generalize to field operational contexts requires independent empirical evaluation under actual growing conditions across representative European wheat production environments.

Fig. 8 demonstrates the visual representation of validation performance across the integrated framework. To provide detailed quantitative assessment of individual component contributions, Table 5 presents validation metrics and performance improvement statistics for each AI module (expert system, fuzzy logic controller, reinforcement learning, and multi-agent coordination), enabling systematic evaluation of component-level effectiveness within the hybrid architecture.

Table 5 presents component-level validation metrics assessed using composite F1-scores aggregating accuracy, precision, and recall across multiple decision categories with 70-30 training-test splits and 5-fold cross-validation. Each AI component underwent independent validation against baseline approaches to quantify individual contributions before integration within the hybrid architecture; for the expert system, component performance reflects rule activation through the priority-weighted confidence scoring of Eq. (1), which differentiates decision quality across the 15 production rules by combining confidence level, priority weight, and condition match degree.

A sequential ablation analysis was conducted to quantify the marginal contribution of each AI component to integrated hybrid performance. In each ablation trial, one component was withheld and the remaining three operated under otherwise identical conditions; the resulting hybrid F1-scores are reported alongside the standalone values in the enhanced Table 5. Withholding the reinforcement learning component produced the largest single-component decline, reducing the hybrid F1-score from 0.893 to 0.831, corresponding to a 6.2-percentage-point reduction that reflects the module's specialized capacity for temporal optimization across sequential growing-season decisions. Removing the fuzzy logic controller reduced the score to 0.841 (5.2-point decline), consistent with the controller's primary role in managing measurement uncertainty during continuous irrigation scheduling. Withholding the expert system reduced the score to 0.852 (4.1 points), while removing the multi-agent coordination component produced the smallest single-component decline to 0.864 (2.9 points), reflecting its contribution to distributed spatial coordination rather than individual decision accuracy. The integrated system score of 0.893 exceeded all four single-omission configurations, confirming that hybrid performance arises from complementary contributions across components rather than from dominance of any single module.

Table 5
Component-level validation metrics and performance improvement assessment.

AI component	Hybrid AI (F1)	Baseline (F1)	Improvement (%)	Hybrid without component (F1)
Expert System	0.867	0.617	40.7	0.852
Fuzzy Controller	0.903	0.673	34.2	0.841
Multi-Agent System	0.870	0.630	38.2	0.864
RL Agent	0.867	0.613	41.5	0.831
Overall Hybrid AI	0.893	0.663	34.7	–

Notes:
- Validation metrics calculated using composite F1-Score formulation (Eq. (20)) that aggregates accuracy, precision, and recall across irrigation scheduling, nutrient management, pest control, and harvest timing decisions, computed with 70-30 training-test split and 5-fold cross-validation.
- Baseline comparisons represent conventional practices: binary threshold-based control without fuzzy membership functions (fuzzy logic baseline), fixed calendar-based scheduling without adaptive learning (RL baseline), informal farmer decision-making without formalized rule encoding (expert system baseline), centralized optimization without distributed agent coordination (multi-agent baseline), and integrated conventional management without AI components (overall hybrid baseline).
- Component outputs weighted by decision context following the integration layer formulation (Eq. (12) and (13)): fuzzy logic 0.40 (continuous control), expert system 0.30 (categorical decisions), RL 0.20 (temporal sequences), multi-agent 0.10 (spatial coordination).
- Statistical significance ($p < 0.001$) confirmed via paired Wilcoxon signed-rank tests across 500+ decision instances per category, with effect sizes exceeding Cohen's $d = 0.8$ for all components.

4. Discussion

The hybrid artificial intelligence framework integrating expert systems, fuzzy logic controllers, reinforcement learning, and multi-agent coordination demonstrated substantial performance advantages across multiple agricultural sustainability metrics for winter wheat production systems. The integrated system achieved yield improvements ranging from 7.8% under normal conditions to 34.8% under drought stress and 26.1% under extreme heat, with water savings of 17.8–29.4% and energy reductions of 15.1–22.1% across climate scenarios relative to conventional management practices, exceeding single-technique approaches that reported 8.7–18.5% improvements in recent agricultural engineering literature (Hasan et al., 2023; Waghela et al., 2024). This performance differential arises from complementary integration of four computational intelligence paradigms within a unified control architecture. Expert systems provide interpretable rule-based reasoning essential for farmer trust and regulatory compliance, while fuzzy logic controllers manage measurement uncertainty inherent in agricultural sensor data through linguistic variable processing. Reinforcement learning optimizes temporal decision sequences by adapting policies through iterative environmental interaction, whereas multi-agent coordination enables distributed responsiveness to localized conditions while maintaining system-wide resource allocation efficiency. This integrated approach aligns with farming system design principles that emphasize holistic optimization across production, environmental, and socioeconomic dimensions in cereal production systems, particularly winter wheat cultivation under European pedoclimatic conditions. The simulation-based assessment framework enables comprehensive evaluation of hybrid AI performance across diverse climate scenarios and regional contexts that would be impractical through field experimentation alone, providing systematic evidence for climate-smart wheat management under conditions spanning normal, drought, and extreme heat scenarios.

A notable finding emerges regarding benefit amplification under climate stress conditions for wheat production systems. Drought scenarios revealed 34.8% yield improvements compared to only 7.8% under normal environmental conditions, with water use efficiency gains reaching 29.1% during drought versus 17.8% under normal conditions. Extreme heat scenarios demonstrated 26.1% yield gains with 29.4% water savings and 22.1% energy efficiency improvements. This pattern

suggests that the value proposition of intelligent agricultural control systems increases proportionally with environmental uncertainty and resource scarcity, providing a critical insight for technology deployment strategies in climate-vulnerable wheat production regions (Wang et al., 2024). The fuzzy logic irrigation controller and reinforcement learning temporal optimizer provide greatest value precisely when conventional fixed-schedule approaches break down under non-standard environmental conditions characteristic of drought and heat stress periods. Energy consumption reductions reached 18.3% during drought and 22.1% under extreme heat compared to 15.1% normally, reflecting the multi-agent coordination system's capacity to optimize pump scheduling when electricity demand and irrigation requirements diverge from typical patterns. From a farming systems perspective, these results demonstrate that adaptive control architectures maintain or even enhance effectiveness during environmental perturbations, whereas systems designed for nominal operating conditions exhibit degraded performance under climate extremes that increasingly characterize European wheat production environments.

The reinforcement learning algorithm comparison revealed that Actor-Critic methods achieved superior convergence performance for wheat management optimization, reaching cumulative rewards of 519.7 compared to 409.6 for Q-learning. Crucially, Actor-Critic required more training episodes to reach convergence (episode 441) compared to Q-Learning (episode 371), though it achieved higher final cumulative rewards representing better long-term wheat production outcomes. Actor-Critic exhibited moderate final loss value of 0.284, which carries important implications for agricultural engineering applications. This finding indicates that loss minimization alone constitutes an insufficient predictor of practical decision-making quality because agricultural decisions in wheat production systems balance multiple objectives spanning yield, resource efficiency, economic returns, and risk management across uncertain timeframes. Prediction accuracy metrics correlate imperfectly with field-level agronomic outcomes due to this complex, multi-objective nature. Consequently, the selection of appropriate reinforcement learning algorithms for agricultural cyber-physical systems prioritizes cumulative reward metrics over training loss values when evaluating suitability for integration into climate-smart wheat production decision frameworks; this criterion retains methodological validity across both simulation-based and prospective field-level implementation contexts. Empirical evaluation of algorithmic performance within these frameworks necessitates a focus on long-term agricultural outcomes rather than the convergence rate of the loss function.

Regional performance variations across Mediterranean, Central European, and Nordic wheat production zones demonstrated both framework transferability and the importance of environmental context in determining absolute outcomes. Central Europe achieved the highest wheat yields at 5.90 t/ha under normal conditions, while Mediterranean regions exhibited proportionally larger relative improvements of 23.4% compared to 21.9% in Nordic zones, reflecting baseline differences in winter wheat productivity across contrasting European climates and the framework's capacity to address region-specific resource constraints. This pattern suggests the integrated framework provides greatest value in wheat production environments characterized by greater baseline resource constraints, an important consideration for technology deployment prioritization (Cao et al., 2023). Nevertheless, water efficiency gains remained relatively consistent across regions at 28–30% under drought conditions, indicating that the fuzzy logic controller's capacity to process intermediate soil moisture states and prevent redundant irrigation applications, coordinated through the distributed constraint optimization of Eq. (11) across water resource agents sharing district-level allocations, operates effectively across diverse pedoclimatic contexts relevant to European wheat cultivation (Zhou et al., 2025). Carbon footprint reductions ranged from 18.5% to 20.6% across regions, demonstrating that environmental sustainability benefits extend beyond direct resource consumption to encompass broader climate mitigation objectives relevant to agricultural engineering applications in wheat production systems.

Temporal analysis throughout the 120-day wheat growing season revealed mechanistic insights into how the integrated system achieves improved outcomes; the continuous irrigation output values generated by centroid defuzzification (Eq. (4)) enabled gradual quantity adjustments in response to evolving soil moisture trajectories, in contrast to the step-changes imposed by threshold-based conventional scheduling that produced oscillating water stress patterns across the growing season. Final aboveground biomass accumulation reached 14,000 kg ha^{-1} for AI-based management compared to 11,700 kg ha^{-1} conventionally, representing a 19.7% enhancement in wheat biomass production. Peak leaf area index reached 7.5 for the hybrid system versus 6.4 for conventional approaches, indicating that optimized nitrogen application timing guided by expert system recommendations enables enhanced canopy development during critical vegetative growth phases characteristic of winter wheat phenology (Gao et al., 2021). Water stress management patterns showed that AI-based irrigation-maintained stress indices predominantly below 0.30 throughout the wheat growing season through variable timing and quantities, whereas conventional fixed-schedule irrigation exhibited oscillating stress patterns with peaks reaching 0.45 to 0.50. This demonstrates the practical advantage of dynamic optimization over predetermined management calendars, reflecting a core control engineering principle applied to wheat production systems (Xiong et al., 2020). Daily nitrogen uptake patterns revealed that strategic fertilizer timing at days 10, 35, and 57 for AI-based management aligned more closely with peak wheat nitrogen demand periods compared to fixed conventional applications, improving nitrogen use efficiency and reducing environmental losses.

Sensitivity analysis revealed that irrigation threshold parameters exerted the strongest influence on wheat system performance, with yield impacts ranging from −5.0% to 8.4% at ±20% parameter variations, followed by fertilizer rate sensitivity at −4.8%–7.2%. This parameter importance ranking provides guidance for calibration priorities within the simulation framework and informs parameter selection strategies for future field trials in wheat production systems, indicating that precise calibration of fuzzy logic membership functions for soil moisture thresholds should receive primary attention in any prospective empirical extension of the framework (Mateus Freitas Silveira et al., 2024). In contrast, learning rates and discount factors for reinforcement learning agents demonstrated greater robustness across parameter ranges, with yield impacts remaining within approximately 6% for both positive and negative variations. Climate variability demonstrated asymmetric sensitivity, with reduced variability improving yields by 6.5% while increased uncertainty caused 8% reductions, confirming that the framework provides adaptive capacity to buffer moderate climate fluctuations while remaining vulnerable to extreme stochastic conditions beyond optimization thresholds.

Model validation against independent empirical observations from European wheat production systems spanning 2015–2024 yielded prediction accuracies of $R^2 = 0.802$ for yield, 0.783 for water use, and 0.613 for energy consumption. The lower predictive accuracy for energy reflects greater complexity of factors influencing actual field-level energy use in wheat production, including equipment maintenance status, operator decision-making, field topography variations, and temporal electricity pricing dynamics not fully captured by simulation parameters. The root mean square errors of 0.44 t/ha for yield and 30.1 mm for water use fall within acceptable ranges for agricultural decision support systems deployed in operational contexts (Kuradusenge et al., 2023). Component-level validation demonstrated that the overall hybrid framework improvement of 34.7% derives from complementary strengths rather than simple additive effects; the ensemble synthesis of Eq. (12) dynamically reallocates weights toward the component with the highest contextual confidence score (Eqs. (14) and (15)), so that no single module consistently dominates across all decision categories, confirming the engineering rationale for multi-paradigm integration within unified control architectures for climate-smart agricultural

applications (Ashfaq et al., 2025). From a farming systems perspective, the hybrid framework addresses the fundamental challenge of multi-objective optimization under uncertainty that characterizes climate-smart wheat production systems. The demonstrated simulation-level capacity to simultaneously improve productivity, resource efficiency, and environmental outcomes across diverse European agroecological zones provides a computationally grounded evidence base for wheat producers, agricultural advisors, and policy-makers evaluating the potential of hybrid AI architectures within broader climate adaptation strategies for winter wheat production systems; these findings most appropriately inform the design of field validation trials rather than constitute a direct prescription for immediate operational integration.

5. Conclusion

This research developed and validated a unified computational framework that synergistically integrates expert systems, fuzzy logic controllers, reinforcement learning, and multi-agent coordination for resource-efficient, climate-smart wheat production across contrasting European regions. The novel contribution extends beyond performance metrics to methodological innovation in agricultural engineering. The study provides explicit mathematical formulations for component integration, conflict resolution mechanisms that enable distributed agents to coordinate resource allocation, and context-dependent weighting strategies that generate synergistic performance exceeding individual technique capabilities. Actor-Critic reinforcement learning achieved the highest convergence performance, with cumulative rewards of 519.7 compared to 409.6 for Q-learning; the advantage function $\widehat{A}(s,a,t)$ in the Actor-Critic policy gradient (Eq. (7)) provides a lower-variance credit assignment signal than plain Q-learning by explicitly quantifying each action's improvement over the baseline state value, which explains the algorithm's superior multi-objective performance in agricultural decision contexts where competing yield, water, and carbon objectives must be optimized simultaneously. Regional validation from Mediterranean to Nordic European wheat production zones achieved prediction accuracies of $R = 0.802$ for yield and $R = 0.783$ for water use; these results support simulation-level spatial transferability across diverse climate scenarios while sensitivity analysis identifies irrigation threshold parameters as priority targets for local pedoclimatic calibration. The demonstrated simulation performance characteristics align with the structural scalability inherent in the framework design spanning individual farm to landscape-level contexts; this spatial scope remains consistent with farming-systems research regarding crop modeling and decision support. Translating such structural compatibility into operational planning requires further validation through independent field experiments conducted at the respective spatial scales.

The demonstrated performance characteristics support scalable deployment from individual farm operations to landscape-level climate adaptation planning, aligning with farming-systems research that addresses crop modeling and decision support from farm to regional scales. At the field scale, hourly agent coordination optimized irrigation micro scheduling and real-time stress responses for winter wheat; this optimization prevented resource waste through precise temporal allocation, embodying cleaner production principles of doing more with less while minimizing environmental externalities. At the farm scale, daily decision cycles integrated resource allocation across multiple fields through distributed agent communication protocols. Regional validation across three European agricultural zones spanning 2015–2024 confirmed that the simulation-based framework remains robust under diverse climate regimes and management contexts. The modular Python architecture is designed to accommodate incremental adoption pathways, wherein a standalone fuzzy logic irrigation controller represents the most computationally accessible entry point and could serve as an initial configuration for prospective field evaluation. Subsequent extension to reinforcement learning optimization and multi-agent coordination would be warranted as field validation evidence accumulates and on-farm data infrastructure matures. This structural modularity preserves upward compatibility toward regional advisory configurations spanning multiple farm operations; field-level empirical testing at each scale of integration would nevertheless be necessary before such configurations could be responsibly implemented in operational contexts.

Several limitations constrain the interpretation and generalizability of these findings. The framework was validated exclusively for winter wheat production systems; crop-specific physiological parameters, phenology, and management calendars would require substantial recalibration before the architecture could be transferred to other commodities or mixed cropping configurations. The reward function employs a weighted linear scalarization of five sustainability objectives (Eq. (6)), with weight vectors calibrated via analytic hierarchy process and held-out validation refinement under three climate scenarios; this formulation does not enumerate the complete Pareto frontier across feasible weight combinations, so trade-off surfaces between competing objectives such as yield maximization and carbon minimization are not fully characterized for all plausible operational priorities. The local state transition independence assumed in Eq. (9) provides computational tractability for the 118-agent system but does not capture lateral soil water dynamics between adjacent field units or shared groundwater depletion; inter-agent message passing (Eq. (10)) partially compensates for this spatial coupling, yet fields sharing subsurface irrigation infrastructure may retain residual cross-agent dependencies unresolved by the current architecture. The conflict resolution mechanism (Eq. (16)) prioritizes crop survival under high-uncertainty conditions; under prolonged multi-year drought sequences where the system repeatedly enters conflict-resolution mode, this conservative protocol may systematically suppress efficiency optimization across consecutive growing seasons before the reinforcement learning policy has accumulated sufficient episodes to adapt to the new distributional regime. Validation relied on 30 growing seasons derived from three geographically separated regions, a structure that supported detection of the observed large effect sizes (Cohen's $d > 0.8$; $\alpha = 0.001$) but limits characterization of distributional tails and multi-year compound climate extremes; the ten-season within-region temporal series additionally constrains the effective degrees of freedom available for inference. More broadly, this study constitutes a simulation-based assessment in which field-level variability is represented through synthetically generated scenarios anchored in published agronomic data rather than through independent experimental measurements; consequently, the performance metrics and comparative advantages reported here characterize the behaviour of the computational framework under simulated conditions and should not be extrapolated to actual field outcomes without prospective empirical validation across representative European wheat production environments. The ablation analysis quantified single-component omissions across four trials; systematic evaluation of all $2^4 = 16$ possible subset configurations was not conducted. Uncertainty propagates from three principal sources, namely parameter uncertainty in calibrated regional means, structural uncertainty in model specification choices including the MDP independence assumption and the scalarized reward formulation, and stochastic variability in RL policy convergence; of these, only the first two are reflected in the reported 95% confidence intervals for yield (± 0.44 t ha^{-1}) and water use (± 30.1 mm), whereas structural model uncertainty remains unquantified. The economic assessment focused on water and energy cost savings and carbon footprint reduction without comprehensively incorporating technology adoption costs, computational infrastructure requirements, or farmer training investments that determine life-cycle feasibility for small- and medium-scale operations.

Future research should extend validation to diverse cropping systems including vegetables, perennial crops, and integrated crop-livestock configurations where temporal dynamics and resource flows exhibit

greater complexity than winter wheat monocultures. The modular architecture accommodates livestock agents representing grazing management, feed optimization, and manure nutrient cycling alongside existing crop agents; expert system rules could incorporate crop residue allocation decisions, fuzzy logic controllers could process livestock density indicators jointly with soil moisture and crop phenology data, and the reinforcement learning component could optimize decisions across multi-season rotations that account explicitly for long-term soil fertility and carbon dynamics. Expanding the validation dataset through integration of long-term crop monitoring networks and retrospective yield archives across additional European countries would increase statistical power for tail-event characterization and support more robust assessment under multi-year drought sequences and compound climate extremes. Multi-objective reinforcement learning methods that generate Pareto-optimal policy sets, such as envelope Q-learning, offer a direct extension to the scalarized reward formulation of Eq. (6) and would enable systematic characterization of trade-off surfaces among yield, water use, energy, and carbon objectives across the full feasible weight space (Roijers et al., 2013). Factored Markov decision process representations (Guestrin et al., 2003) or mean-field approximations provide tractable formalisms for relaxing the local state transition independence of Eq. (9), enabling explicit modelling of spatial dependencies arising from shared irrigation infrastructure without incurring the exponential cost of full joint state enumeration. Real-time sensor integration via on-farm Internet-of-Things infrastructure (Li et al., 2018) would replace the simulation-based field-level variability used here with observed in-situ data streams, enabling continuous online policy updating and fully empirical prospective validation under operational deployment conditions.

Informed consent statement

Not applicable. This study does not involve human subjects.

Institutional review board statement

Not applicable. This study is based on mathematical simulations and computational modeling and does not involve human subjects or animal experimentation.

Simulation output files

Regional crop performance data for Mediterranean, Central Europe, and Nordic zones spanning 30 independent growing seasons (2015–2024) across normal, drought, and extreme heat climate scenarios (3600 daily observations total).

Validation and analysis data

Sensitivity analysis results for parameter variations, empirical validation dataset comparing predicted versus observed values for 30 growing seasons, temporal dynamics measurements throughout 120-day growing seasons, and aggregated performance summaries supporting all manuscript tables and figures.

All analysis scripts are implemented in Python 3.14.0 and provided with the dataset.

Hybrid AI system data

Reinforcement learning training performance across 1000 episodes for four algorithms (Q-learning, Deep Q-Network, Policy Gradient, Actor-Critic), daily decision logs from all four AI components (expert system, fuzzy logic, reinforcement learning, multi-agent), and multi-agent coordination records across 118 agents.

Declaration of generative AI and AI-assisted technologies

During preparation of this manuscript, DeepL Write and ScienceDirect AI were utilized exclusively for language refinement, grammar correction, academic tone enhancement, and manuscript formatting. DeepL Write provided sentence-level writing improvements and clarity enhancements. ScienceDirect AI assisted with literature search refinement and related terminology accuracy. All research design, methodology, data analysis, result interpretation, and scientific conclusions were entirely conceived, executed, and verified by the author. Following application of these tools, the manuscript was thoroughly reviewed and edited to ensure all content reflects the author's original work and scientific intent. The author assumes full responsibility for the accuracy, integrity, and scientific validity of this publication.

Funding

This research received no external funding.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Mohammad Fazle Rabbi reports administrative support and article publishing charges were provided by University of Debrecen. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgment

This research was supported by the "University of Debrecen Program for Scientific Publication."

Appendix A. Nomenclature and Abbreviations

Table A1
Complete nomenclature and abbreviations used in this study

Abbreviation/Symbol	Definition
AI	Artificial Intelligence
AHP	Analytic Hierarchy Process
AQUASTAT	FAO Global Information System on Water and Agriculture
CAPRI	Common Agricultural Policy Regional Impact model
CSA	Climate-Smart Agriculture
DCOP	Distributed Constraint Optimization Problem

(*continued on next page*)

Table A1 (*continued*)

Abbreviation/Symbol	Definition
DQN	Deep Q-Network
DSS	Decision Support System
E-OBS	European Observational gridded dataset
ERA5-Land	ECMWF Reanalysis v5 Land surface data
ES	Expert System
FAOSTAT	Food and Agriculture Organization Statistics Database
FLC	Fuzzy Logic Controller
GHG	Greenhouse Gas
GS	Zadoks Growth Stage scale (e.g. GS25, GS31)
IoT	Internet of Things
MAE	Mean Absolute Error
MAS	Multi-Agent System
MDP	Markov Decision Process
MODIS	Moderate Resolution Imaging Spectroradiometer
RMSE	Root Mean Square Error
RL	Reinforcement Learning
RETE	Rete pattern-matching algorithm (Forgy, 1982)
TSK	Takagi-Sugeno-Kang fuzzy inference
α	Learning rate; controls policy update step size
γ	Discount factor; weights future rewards in Eqs. (5)–(9)
ε	Exploration rate in ε-greedy strategy (Eq. (8))
τ	Conflict resolution threshold (Eq. (16)); set to 0.3
$\mu_A(x)$	Membership function for linguistic term A (Eq. (2))
$\hat{A}(s,a,t)$	Advantage function in Actor-Critic policy gradient (Eq. (7))
r_t	Reward signal at time t (Eq. (6))
$Q(s,a)$	Action-value function (Eq. (5))
conf_k	Component-specific confidence score for module k (Eqs. (14) and (15))

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.cesys.2026.100446.

Data availability

The complete dataset supporting the findings of this study is provided as supplementary material submitted with this manuscript

References

Aderele, M.O., Srivastava, A.K., Butterbach-Bahl, K., Rahimi, J., 2025. Integrating machine learning with agroecosystem modelling: current state and future challenges. Eur. J. Agron. 168, 127610. https://doi.org/10.1016/j.eja.2025.127610.

Ajith, S., Vijayakumar, S., Elakkiya, N., 2025. Yield prediction, pest and disease diagnosis, soil fertility mapping, precision irrigation scheduling, and food quality assessment using machine learning and deep learning algorithms. Discov. Food 5, 67. https://doi.org/10.1007/s44187-025-00338-1.

Ashfaq, M., Khan, I., Afzal, R.F., Shah, D., Ali, S., Tahir, M., 2025. Enhanced wheat yield prediction through integrated climate and satellite data using advanced AI techniques. Sci. Rep. 15, 18093. https://doi.org/10.1038/s41598-025-02700-w.

Becker-Reshef, I., Vermote, E., Lindeman, M., Justice, C., 2010. A generalized regression-based model for forecasting winter wheat yields in Kansas and Ukraine using MODIS data. Remote Sens. Environ. 114, 1312–1323. https://doi.org/10.1016/j.rse.2010.01.010.

Blanco, M., Witzke, P., Pérez Domínguez, I., Salputra, G., Martínez, P., 2015. Extension of the CAPRI model with an irrigation sub-module. Bonn. https://doi.org/10.2791/319578.

Botega, L.F. de C., da Silva, J.C., 2020. An artificial intelligence approach to support knowledge management on the selection of creativity and innovation techniques. J. Knowl. Manag. 24, 1107–1130. https://doi.org/10.1108/JKM-10-2019-0559.

Caicedo, A.M.D., Mejia, É.F., Luna, E.G., 2024. Revolutionizing protection dynamics in microgrids: local validation environment and a novel global management control through multi-agent systems. Comput. Electr. Eng. 120, 109748. https://doi.org/10.1016/j.compeleceng.2024.109748.

Cao, M., Yang, R., Choi, C.Y., Rong, L., Zhang, G., Wang, K., Wang, X., 2023. Effects of discharge angle of jet from a slot orifice on cooling performance for a perforated air ducting system in dairy cattle barn. Comput. Electron. Agric. 210, 107890. https://doi.org/10.1016/j.compag.2023.107890.

Cornes, R., van der Schrier, G., van den Besselaar, E.J.M., Jones, P., 2018. E-OBS daily gridded meteorological data for Europe from 1950 to present derived from in-situ observations. https://doi.org/10.24381/cds.151d3ec6.

Durbin, J., Watson, G.S., 1950. Testing for serial correlation in least squares regression. Biometrika 37, 409–428. https://doi.org/10.1093/biomet/37.3-4.409.

EC, 2025. model CAPRI - Common Agricultural Policy Regional Impact Analysis | Modelling Inventory and Knowledge Management System of the European Commission (MIDAS) [WWW Document]. https://web.jrc.ec.europa.eu/policy-model-inventory/explore/models/model-capri/, 11.11.25.

FAO, 2009. How to Feed the World in 2050.

FAO, 2020. Water Scarcity | Land & Water [WWW Document]. https://www.fao.org/land-water/water/water-scarcity/en/, 11.11.25.

FAO, 2024. AQUASTAT - FAO's Global Information System on Water and Agriculture [WWW Document]. https://www.fao.org/aquastat/en/databases/, 11.11.25.

FAO, 2025a. FAO study reveals alarming agricultural land degradation in the Arab region [WWW Document]. https://www.fao.org/newsroom/detail/fao-study-reveals-alarming-agricultural-land-degradation-in-the-arab-region/en, 11.11.25.

FAO, 2025b. FAOSTAT [WWW Document]. https://www.fao.org/faostat/en/#data/QCL, 11.11.25.

Fereres, E., Soriano, M.A., 2007. Deficit irrigation for reducing agricultural water use. J. Exp. Bot. 58, 147–159. https://doi.org/10.1093/jxb/erl165.

Forgy, C.L., 1982. Rete: a fast algorithm for the many pattern/many object pattern match problem. Artif. Intell. 19, 17–37. https://doi.org/10.1016/0004-3702(82)90020-0.

Gao, Y., Shao, G., Yang, Z., Zhang, K., Lu, J., Wang, Z., Wu, S., Xu, D., 2021. Influences of soil and biochar properties and amount of biochar and fertilizer on the performance of biochar in improving plant photosynthetic rate: a meta-analysis. Eur. J. Agron. 130, 126345. https://doi.org/10.1016/J.EJA.2021.126345.

Glorot, X., Bengio, Y., 2010. Understanding the difficulty of training deep feedforward neural networks. In: Artificial Intelligence and Statistics (AISTATS). JMLR Workshop and Conference Proceedings, Quebec, pp. 249–256.

Guestrin, C., Koller, D., Parr, R., Venkataraman, S., 2003. Efficient solution algorithms for factored MDPs. J. Artif. Intell. Res. 19, 399–468. https://doi.org/10.1613/jair.1000.

Hasan, M., Marjan, M.A., Uddin, M.P., Afjal, M.I., Kardy, S., Ma, S., Nam, Y., 2023. Ensemble machine learning-based recommendation system for effective prediction of suitable agricultural crop cultivation. Front. Plant Sci. 14. https://doi.org/10.3389/fpls.2023.1234555.

Hristov, J., Toreti, A., Pérez Domínguez, I., Dentener, F., Fellmann, T., Elleby, C., Ceglar, A., Fumagalli, D., Niemeyer, S., Cerrani, I., Panarello, L., Bratu, M., 2020. Analysis of Climate Change Impacts on EU Agriculture by 2050 JRC PESETA IV project-task 3. https://doi.org/10.2760/121115.

Ikram, A., Ikram, S., El-kenawy, E.-S.M., Hussain, A., Alharbi, A.H., Eid, M.M., 2025. A fuzzy-optimized hybrid ensemble model for yield prediction in maize-soybean intercropping system. Front. Plant Sci. 16. https://doi.org/10.3389/fpls.2025.1567679.

Jang, J.S.R., Sun, C.T., Mizutani, E., 1997. Neuro-Fuzzy and Soft Computing-A Computational Approach to Learning and Machine Intelligence [Book Review]. IEEE

Trans. Automat. Control 42, 1482–1484. https://doi.org/10.1109/TAC.1997.633847.
Kamilaris, A., Prenafeta-Boldú, F.X., 2018. Deep learning in agriculture: a survey. Comput. Electron. Agric. 147, 70–90. https://doi.org/10.1016/j.compag.2018.02.016.
Kuradusenge, M., Hitimana, E., Hanyurwimfura, D., Rukundo, P., Mtonga, K., Mukasine, A., Uwitonze, C., Ngabonziza, J., Uwamahoro, A., 2023. Crop yield prediction using machine learning models: case of Irish potato and maize. Agriculture 13, 225. https://doi.org/10.3390/agriculture13010225.
Li, D., Wang, C., Jiang, H., Peng, Z., Yang, J., Su, Y., Song, J., Chen, S., 2018. Monitoring litchi canopy foliar phosphorus content using hyperspectral data. Comput. Electron. Agric. 154, 176–186. https://doi.org/10.1016/j.compag.2018.09.007.
Li, L., Wang, B., Feng, P., Jägermeyr, J., Asseng, S., Müller, C., Macadam, I., Liu, D.L., Waters, C., Zhang, Y., He, Q., Shi, Y., Chen, S., Guo, X., Li, Y., He, J., Feng, H., Yang, G., Tian, H., Yu, Q., 2023. The optimization of model ensemble composition and size can enhance the robustness of crop yield projections. Commun. Earth Environ. 4, 362. https://doi.org/10.1038/s43247-023-01016-9.
Mamdani, E.H., Assilian, S., 1975. An experiment in linguistic synthesis with a fuzzy logic controller. Int. J. Man Mach. Stud. 7, 1–13. https://doi.org/10.1016/S0020-7373(75)80002-2.
Mateus Freitas Silveira, R., Andréa Evangelista Façanha, D., McManus, C., Bermejo Asensio, L.A., José Oliveira da Silva, I., 2024. Intelligent methodologies: an integrated multi-modeling approach to predict adaptive mechanisms in farm animals. Comput. Electron. Agric. 216, 108502. https://doi.org/10.1016/j.compag.2023.108502.
Melese, M., Getachew, B., Woldemeskel, E., Gunnabo, A.H., 2025. Legume integration in smallholder farming systems for food security and resilience to climate change. PLoS One 20, e0327727.
Mmbando, G.S., 2025. Harnessing artificial intelligence and remote sensing in climate-smart agriculture: the current strategies needed for enhancing global food security. Cogent Food Agric. 11, 2454354. https://doi.org/10.1080/23311932.2025.2454354.
Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A.A., Veness, J., Bellemare, M.G., Graves, A., Riedmiller, M., Fidjeland, A.K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., Hassabis, D., 2015. Human-level control through deep reinforcement learning. Nature 518, 529–533. https://doi.org/10.1038/nature14236.
Moayedi, H., Ahmadi Dehrashid, A., Nguyen Le, B., 2024. A novel problem-solving method by multi-computational optimisation of artificial neural network for modelling and prediction of the flow erosion processes. Eng. Appl. Comput. Fluid Mech. 18, 2300456. https://doi.org/10.1080/19942060.2023.2300456.
Muñoz Sabater, J., 2019. ERA5-Land monthly averaged data from 1950 to present [WWW Document]. https://doi.org/10.24381/cds.68d2bb30.
Musa, S.F.P.D., Ariff Lim, S., 2025. Revitalising agriculture through climate change mitigation: a systematic literature review on smart technologies and sustainable practices. Int. J. Clim. Change Strateg. Manag. 17, 483–501. https://doi.org/10.1108/IJCCSM-05-2024-0071.
Nevo, A., Oad, R., Podmore, T.H., 1994. An integrated expert system for optimal crop planning. Agric. Syst. 45, 73–92. https://doi.org/10.1016/S0308-521X(94)90281-X.
Paudel, D., de Wit, A., Boogaard, H., Marcos, D., Osinga, S., Athanasiadis, I.N., 2023. Interpretability of deep learning models for crop yield forecasting. Comput. Electron. Agric. 206, 107663. https://doi.org/10.1016/J.COMPAG.2023.107663.
Qi, J., Lin, J., 2025. Optimizing the strategy of new media content distribution using fuzzy logic. Discov. Artif. Intell. 5, 130. https://doi.org/10.1007/s44163-025-00385-8.
Rabbi, M.F., 2025a. Unlocking sustainability in the EU food system: a regional analysis of sectoral carbon emission drivers and SDG-12 performance. Sustain. Horiz. 15, 100144. https://doi.org/10.1016/j.horiz.2025.100144.
Rabbi, M.F., 2025b. A machine learning framework for forecasting multidimensional sustainability and informing integrated policy thresholds in the EU. Environ. Dev. Sustain. https://doi.org/10.1007/s10668-025-06762-8.
Rabbi, M.F., 2025c. Cross-framework hybrid artificial intelligence for high-penetration renewable energy integration: Multi-regional forecasting and adaptive control. Appl. Energy 401, 126834. https://doi.org/10.1016/j.apenergy.2025.126834.
Rabbi, M.F., 2025d. Optimizing carbon emissions and SDG-12 performance in the EU food system. Carbon Res. 4. https://doi.org/10.1007/s44246-025-00220-w.
Radwan, M., Alhussan, A.A., Ibrahim, A., Tawfeek, S.M., 2025. Potato leaf disease classification using optimized machine learning models and feature selection techniques. Potato Res. 68, 897–921. https://doi.org/10.1007/s11540-024-09763-8.
Roijers, D.M., Vamplew, P., Whiteson, S., Dazeley, R., 2013. A survey of multi-objective sequential decision-making. J. Artif. Intell. Res. 48, 67–113. https://doi.org/10.1613/jair.3987.
Saaty, T.L., 1980. The Analytic Hierarchy Process. McGraw-Hill International Book Co., New York.
Saccone, D., Vallino, E., 2025. Global food security in a turbulent world: reviewing the impacts of the pandemic, the war and climate change. Agricult. Food Econ. 13, 47. https://doi.org/10.1186/s40100-025-00388-0.
Santos, J.S., Zahn, J.O., Silvestre, E.A., Silva, V.T., Vasconcelos, W.W., 2017. Detection and resolution of normative conflicts in multi-agent systems: a literature survey. Auton. Agent. Multi. Agent. Off. Syst. 31, 1236–1282. https://doi.org/10.1007/s10458-017-9362-z.
Singh, A.R., Sujatha, M.S., Kadu, A.D., Bajaj, M., Addis, H.K., Sarada, K., 2025. A deep learning and IoT-driven framework for real-time adaptive resource allocation and grid optimization in smart energy systems. Sci. Rep. 15, 19309. https://doi.org/10.1038/s41598-025-02649-w.
Sylvester-Bradley, R., Kindred, D.R., 2009. Analysing nitrogen responses of cereals to prioritize routes to the improvement of nitrogen use efficiency. J. Exp. Bot. 60, 1939–1951. https://doi.org/10.1093/jxb/erp116.
UN, 2022. World Population Prospects 2022.
Waghela, A., Makadia, D., Sheth, P., Mangla, M., Sharma, D., 2024. Crop yield prediction in India: a comparative analysis of ensemble techniques. In: Goar, V., Sharma, A., Shin, J., Mridha, M.F. (Eds.), Deep Learning and Visual Artificial Intelligence. Springer Nature Singapore, Singapore, pp. 367–380. https://doi.org/10.1007/978-981-97-4533-3_28.
Wang, L., Cong, J., Ren, N., Ying, J., Wang, X., Liao, Y., Liao, Q., 2024. Influence of surface slope on the seeding performance of air-assisted centralized metering device for rapeseed based on numerical simulation. Comput. Electron. Agric. 218, 108734. https://doi.org/10.1016/j.compag.2024.108734.
Watkins, C.J.C.H., Dayan, P., 1992. Q-learning. Mach. Learn. 8, 279–292. https://doi.org/10.1007/BF00992698.
WEF, 2021. Artificial Intelligence for Agriculture Innovation.
Wing, I.S., De Cian, E., Mistry, M.N., 2021. Global vulnerability of crop yields to climate change. J. Environ. Econ. Manag. 109, 102462. https://doi.org/10.1016/j.jeem.2021.102462.
Xiong, Y., Ge, Y., From, P.J., 2020. An obstacle separation method for robotic picking of fruits in clusters. Comput. Electron. Agric. 175, 105397. https://doi.org/10.1016/j.compag.2020.105397.
Ye, Z., Yin, S., Cao, Y., Wang, Y., 2024. AI-driven optimization of agricultural water management for enhanced sustainability. Sci. Rep. 14, 25721. https://doi.org/10.1038/s41598-024-76915-8.
Zadeh, L.A., 1965. Fuzzy sets. Inf. Control 8, 338–353. https://doi.org/10.1016/S0019-9958(65)90241-X.
Zadeh, L.A., 1973. Outline of a new approach to the analysis of complex systems and decision processes. IEEE Trans. Syst. Man Cybern. SMC-3, 28–44. https://doi.org/10.1109/TSMC.1973.5408575.
Zeng, Y., Hussein, Z.A., Chyad, M.H., farhadi, A., Yu, J., Rahbarimagham, H., 2025. Integrating type-2 fuzzy logic controllers with digital twin and neural networks for advanced hydropower system management. Sci. Rep. 15, 5140. https://doi.org/10.1038/s41598-025-89866-5.
Zhao, J., Fan, S., Zhang, B., Wang, A., Zhang, L., Zhu, Q., 2025. Research status and development trends of deep reinforcement learning in the intelligent transformation of agricultural machinery. Agriculture 15, 1223. https://doi.org/10.3390/agriculture15111223.
Zhou, Y., Han, Q., Sarabi, S., de Vries, B., 2025. Multi-agent systems in climate-resilient land-use planning: a review. Int. J. Digit. Earth 18, 2487051. https://doi.org/10.1080/17538947.2025.2487051.