agriculture

MDPI

Article

Field-Based Water and Fertilizer Decision-Making Using Crop Sensing Data and Economic Records

Yiwen Zeng [1,2,†], Jiale Niu [1,†], Liyang Liu [1], Jinghan Huang [1], Jalalidi Abuduwali [2], Ning Ma [2] and Yihong Song [1,*]

[1] China Agricultural University, Beijing 100083, China
[2] Peking University, Beijing 100871, China
* Correspondence: songyihong@cau.edu.cn
† These authors contributed equally to this work.

Abstract

Water and fertilizer management in facility and precision agriculture increasingly depends on images, soil sensors, weather records, management logs, and economic data rather than on empirical operation alone. A cross-modal decision framework is developed to align crop visual traits with soil water–salt dynamics, meteorological variation, and management operations. Image preprocessing, abnormal sensor screening, missing value interpolation, temporal resampling, multi-window alignment, and temporal augmentation are used to form comparable inputs before multimodal fusion. Using plot-blocked evaluation that keeps complete plot trajectories together, AgriWFD-RLNet obtains 93.46% status-recognition accuracy and an R^2 of 0.923 for yield–resource–economic response prediction. In offline policy evaluation against the logged conventional management reference, it estimates 16.92% water saving, 14.37% fertilizer saving, and 15.84% Net Return improvement while maintaining 94.68% yield stability and 96.72% safety compliance. These decision outcomes are model-based estimates from the 2024 Wuyuan dataset rather than outcomes of a prospective field deployment.

Keywords: precision agriculture; water–fertilizer management; soil sensor monitoring; agricultural decision support; agricultural economic benefit

1. Introduction

Currently, as agricultural production scales expand and facilities become more advanced and refined, water and fertilizer management has become a critical factor affecting crop yield, quality, resource utilization efficiency, and even farmers' economic returns [1]. In actual production, irrigation and fertilization are not entirely independent processes; rather, they jointly influence the entire process of soil water potential, rhizosphere nutrient status, leaf photosynthetic capacity, plant growth, fruit development, and ultimately yield formation [2]. For protected vegetable systems, orchards, and drip-irrigated crops such as corn, cotton, and cash crops, water and fertilizer inputs often constitute a significant portion of production costs. The outcomes of their management directly affect whether crops grow stably and impact water use efficiency, fertilizer utilization, soil environmental safety, and farmer incomes [3]. Rational irrigation and fertilization must therefore be adjusted according to growth stage, weather conditions, soil type, and observed crop response [4]. Experience-based management is consequently difficult to reconcile with the combined targets of yield, quality, resource conservation, and income improvement [5]. Multimodal

Check for updates

Academic Editor: Marco Napoli

Received: 10 July 2026
Revised: 19 August 2026
Accepted: 25 August 2026
Published: 26 August 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

fusion models integrate visual, soil, and weather information to improve the accuracy of crop status recognition [6].

However, most of these methods remain limited to "knowing what" or "predicting what"—for example, whether crops are water-stressed, whether a plant is underperforming, or what the yield might be—rather than answering more critical questions in actual production: should we irrigate today? How much? Should we fertilize? How much? And what will be the impacts on yield, resource consumption, soil salinity, and income [7]? For agricultural management, the value of models lies not only in improving prediction accuracy but also in transforming this into actionable, interpretable, and economically beneficial decisions [8]. Liu et al. [9] designed an intelligent water–fertilizer control system based on deep reinforcement learning and used real-time soil–crop monitoring to optimize crop water productivity and irrigation control. Sharma et al. [10] reviewed a multimodal AI framework for precision agriculture that combines crop images, remote sensing, IoT sensors, environmental variables, and management information for crop monitoring, disease identification, yield estimation, traceability, and decision support. These studies show that multimodal information can improve agricultural perception, but most of them still emphasize monitoring or recommendation rather than closed-loop optimization under economic and agronomic constraints. Wu Huarui et al. [11] further explored personalized irrigation and fertilization through a large language model combined with PPO reinforcement learning. Although these studies provide useful references, practical water–fertilizer management still requires a framework that can jointly describe crop status, resource input, salinity risk, and economic return.

The unresolved research gap is therefore not simply the absence of another multimodal predictor. Existing perception-oriented models rarely connect three elements in one auditable loop: cross-modal state estimation, action-conditioned prediction of agronomic and economic consequences, and a policy whose actions are projected into an explicitly defined feasible agronomic set. AgriWFD-RLNet is novel in this coupled design. It treats prices and operating costs as decision context rather than post hoc reporting variables, learns delayed image–sensor–management interactions, and enforces water, fertilizer, salinity, growth-stage, and budget limits at action time. This combination distinguishes the proposed framework from both multimodal classifiers and unconstrained agricultural RL policies.

To address this need, this study proposes AgriWFD-RLNet, a multimodal deep reinforcement learning framework for precision water and fertilizer management. The framework uses soil moisture, soil EC, air temperature and humidity, light intensity, rainfall, crop images, irrigation and fertilization logs, growth stage information, yield-quality records, and economic input data as integrated decision evidence. Unlike traditional prediction methods, AgriWFD-RLNet focuses not only on what current state the crop is in, but also emphasizes how to implement water and fertilizer management practices to reduce resource inputs and increase profitability while ensuring yield and quality. The main contributions of this paper are as follows.

1. We construct a multi-source dataset for real-world agricultural water and fertilizer management scenarios, encompassing soil moisture, soil electrical conductivity, meteorological data, crop images, irrigation and fertilization operation records, yield and quality information, and economic input data, providing a solid foundation for precise water and fertilizer decision-making.
2. We propose a multimodal crop water and nutrient status perception module that jointly models visual growth information from crops with time-series data from soil and meteorological sensors, enabling a more comprehensive representation of crop water status, nutrient supply, and growth trends.

3. We design a yield–resource–economic response module, introduced to consider yield, water use, fertilizer use, salinity risk, and net profit together.
4. We introduce an agronomic-constrained reinforcement learning module, developing it to convert agricultural knowledge into reward terms and feasible action bounds.
5. We establish a plot-blocked, cluster-bootstrap evaluation protocol that separates observed field measurements from model predictions and offline policy estimates, and benchmark the resulting decisions against logged conventional management and a soil-moisture-threshold strategy.

2. Related Work

2.1. Precision Irrigation/Fertilization Management for Agriculture Production

Water and fertilizer management regulates the supply of available water and nutrients within the root zone over the crop growth cycle. Adequate irrigation supports transpiration, root uptake, canopy development and biomass accumulation, while fertilization determines nutrient availability and affects plant vigor as well as product quality [12]. These two practices are closely coupled in field operations rather than being independent from each other [13]. Soil water content controls the dissolution, diffusion and uptake of fertilizer solutes; conversely, fertilizer application alters soil solution conductivity, ionic pressure around roots, and subsequent plant development [14]. water–fertilizer regulation is therefore a coupled process that involves soil properties, climatic conditions, crop physiological responses and management practices. Common approaches include empirical irrigation, scheduled irrigation, soil-moisture-threshold irrigation, evapotranspiration-based irrigation, and integrated water–fertilizer delivery systems [15]. Empirical irrigation is inexpensive and straightforward to implement, yet its effectiveness depends heavily on the experience and judgment of individual farmers [16]. Scheduled irrigation provides regular water application, but it adapts poorly to variations in rainfall, temperature, light intensity, and crop demand at different growth stages [17]. Consequently, precise water and fertilizer management should not follow the simplistic principle of more water and more fertilizer guarantee higher yields. Instead, it needs to account for yield and quality formation processes, resource use efficiency, environmental risks, and economic input–output trade-offs.

2.2. Deep Learning-Based Crop Growth Monitoring and Yield Prediction

For time-series modeling, LSTM, GRU, TCN and Transformer have been used to process meteorological, soil and environmental sensor data [18]. LSTM and GRU retain historical information through gating; TCN captures long-range dependencies via dilated convolutions; Transformer models highlight key time steps through attention [19]. Recent multimodal studies combine crop images, soil sensors, weather data and management records to achieve more complete crop characterization than single-modal approaches [20]. However, many existing models are limited to classification or prediction tasks, such as detecting water stress, abnormal growth, disease, or forecasting yield [21]. These outputs are useful for monitoring but do not directly address the operational questions that growers face [22]. The principal strength of recent multimodal deep learning is its ability to represent complementary spatial and temporal evidence; its principal limitation for management is that improved feature fusion does not itself define a feasible intervention. Late-fusion architectures can also conceal dominance by a high-dimensional image branch, whereas prediction losses do not penalize unsafe actions or unfavorable input costs. We therefore use gated fusion, an action-conditioned multi-task response model, and a separate constrained policy rather than equating predictive accuracy with agronomic decision quality. In practice, managers need to decide whether irrigation is needed, how much

water should be applied, whether fertilizer should be supplied, and how these actions will change cost and return [18]. Therefore, perception models need a feedback link from predicted state to management action [19]. Yield prediction alone is also insufficient because resource cost and environmental risk may be ignored [20]. A model that predicts high yield without judging whether the yield depends on excessive water or fertilizer input cannot support efficient and sustainable production [23]. For this reason, this study introduces reinforcement learning on top of deep perception and response prediction, so that the model can not only identify crop water–nutrient status but also generate feasible water and fertilizer management strategies.

2.3. RL to Optimize Farming Decisions

Reinforcement learning optimizes sequential decisions by allowing an agent to interact with the environment and learn policies that maximize long-term cumulative rewards through ongoing trial and feedback [24]. Unlike supervised learning, RL explicitly accounts for the future consequences of current actions [25]; this makes it well-suited for agricultural water and fertilizer management, where irrigation and fertilization choices continuously affect subsequent soil conditions, nutrient availability, crop development, and eventual yield, quality, resource consumption, and economic outcomes [26]. Recent work has brought deep reinforcement learning algorithms into agricultural decision-making, greenhouse control, irrigation scheduling and farm robotics [27], including DQN, DDPG, PPO and SAC [28]. These methods move beyond fixed rule-based strategies and show strong potential for handling delayed responses, late-arriving rewards and multiple competing objectives in farming systems [29]. Nevertheless, deploying RL in real production settings remains difficult [30]. Trial-and-error is costly; poor irrigation or fertilization choices can lower yields, harm quality and waste inputs [31]. Agricultural reward functions must also balance several goals at once, such as yield, quality, resource efficiency, costs, environmental safety and profit, while avoiding extreme behaviors caused by poorly specified incentives [32].

Furthermore, agricultural management is constrained by crop growth stages and agronomic knowledge, requiring irrigation and fertilization decisions to satisfy practical cultivation requirements and environmental safety limits [33]. These challenges highlight the need for reinforcement learning frameworks that effectively integrate domain knowledge with adaptive decision optimization. Agricultural RL studies also face an offline-to-field gap: exploration on crops is unsafe, logged actions cover only part of the action space, and policies can exploit errors in a learned response model. The present work addresses, but does not eliminate, that gap by pretraining the response model on plot-separated data, restricting actions to the logged agronomic range, projecting every action onto a convex safety set, and reporting decision gains explicitly as offline estimates. Prospective multi-season field trials remain necessary for operational validation.

3. Materials and Method

3.1. Data Collection

The dataset was collected at the Smart Agriculture Water-Fertilizer Integration Experimental Station in Wuyuan County, Bayannur City, Inner Mongolia Autonomous Region, from April to October 2024. The observation period covered seedling, vegetative growth, flowering and fruit setting, fruit expansion, maturity, and harvest stages. The site is located in an irrigated production area where crop growth is sensitive to irrigation regime, fertilizer amount, soil salinity, temperature, humidity, and light. Data collected in the experiment include crop images, soil sensor readings, meteorological variables, growth-stage records, water and fertilizer operations, yield-quality measurements, and economic input–output records, as summarized in Table 1.

Table 1. Statistics of multi-source data collection in this paper.

Data Type	Collection Method	Quantity
Crop image data	field fixed cameras and handheld devices	18,600 images
Independent experimental units	four spatial blocks	28 complete plot trajectories
Soil moisture data	soil sensors	245,760 records
Soil temperature data	soil sensors	245,760 records
Soil electrical conductivity data	soil sensors	245,760 records
Air temperature data	field weather stations	253,890 records
Air humidity data	field weather stations	253,890 records
Light intensity data	environmental sensors	245,760 records
Rainfall and wind speed data	field weather stations	122,880 records
Growth stage records	manual field surveys	1240 records
Irrigation operation records	integrated water and fertilizer system	860 records
Fertilization operation records	integrated water and fertilizer system	640 records
Yield and quality data	field sampling and harvest measurements	1120 records
Economic input–output data	production ledgers and market price analysis	1120 records

The effective experimental units were 28 drip-irrigated tomato plots arranged as four spatial blocks of seven plots. Each plot measured 10 m × 4 m, adjacent plots were separated by 2 m buffer strips, and complete plot trajectories—not individual images or sensor timestamps—were treated as independent units. Treatments represented the locally used combinations of stage-specific irrigation and fertigation intensity. The station has a temperate continental climate and all plots in this study shared the same alluvial-loam field; consequently, the dataset samples management and within-field variation but does not sample different soil classes or production regions.

All data were uniformly recorded by field plot number, collection date, collection time, and crop growth stage, enabling accurate alignment of image, soil, meteorological, management operation, and yield–revenue information to the same plot and growth phase, thus providing a consistent foundation for subsequent multimodal modeling. Crop image data consists of two parts: one acquired from a fixed field camera system, and the other from a handheld mobile camera system, as shown in Figure 1. The fixed camera system uses Hikvision DS-2CD3T46WDV3-I3 cameras installed around the experimental site and near representative sample plots to capture images of entire plant populations at regular intervals and fixed locations, monitoring crop growth conditions.



Figure 1. The figure shows the sequential growth stages of tomato from seedling establishment, vegetative growth, flowering and fruit set, and fruit expansion to harvesting, together with representative field phenotypic changes.

One fixed RGB camera served each two-plot observation zone at a height of 2.5 m and a downward viewing angle of approximately 35°. Fixed-camera images were stored as 2688 × 1520 JPEG files. Handheld RGB images were acquired at 3024 × 4032 pixels from 1.2–1.5 m above the canopy at approximately 45°, using the same plot-facing direction.

Routine acquisition was conducted between 09:00 and 11:00 under non-rain conditions; a neutral gray card was included at the start of each handheld session, and color normalization, center cropping, and resizing to 224 × 224 pixels were applied before training. No thermal or multispectral images were used, and models were trained on normalized RGB pixels rather than vegetation indices. Device identity was retained for the consistency analysis reported in Section 4.5.

Soil-moisture, temperature, and EC probes were installed at 20 and 40 cm depths in the central crop row of every plot; the weather station was positioned at an unobstructed upwind edge of the field. Probes were checked weekly against gravimetric soil-moisture samples and a portable EC meter. At each harvest, fruits from the designated measurement rows were weighed by plot, graded as marketable or non-marketable, and accumulated to obtain seasonal yield and marketable-fruit rate. Water and fertilizer quantities came from controller flow-meter and injector logs. Revenue used plot yield and weekly local tomato prices, whereas costs came from water-meter charges, fertilizer procurement invoices, energy use, and recorded labor operations.

This information is primarily determined through periodic field surveys, fixed-point observations, and records kept by field management personnel, while also being synchronized with image capture times and sensor timestamps. Growth stages are used to describe differences in crops' water and nutrient requirements at various developmental phases, enabling models to interpret the same soil moisture level or measurement value differently depending on the growth stage. For example, crops are more sensitive to water deficit during the flowering and fruiting stage, require stable water supply during fruit expansion, and need to avoid excessive irrigation in late maturity to prevent negative impacts on quality and harvest.The implementation of water and fertilizer management measures is primarily based on operational records from the integrated water–fertilizer control system, including irrigation system logs, fertilizer application logs, and manually supplemented log information.

3.2. Data Preprocessing and Data Augmentation

The primary goal of data preprocessing is to transform these multi-source, heterogeneous datasets into a consistent format. Let $x_{t,m}$ denote the observation of the m-th sensor variable at time t, with its mean and standard deviation being μ_m and σ_m, respectively. Outliers can then be detected using a threshold-based method:

$$|x_{t,m} - \mu_m| > k\sigma_m, \tag{1}$$

where k is the threshold for anomaly detection; if an observation is identified as anomalous, it is considered invalid and interpolated based on adjacent valid observations. Assuming the missing time point t lies between two valid observation times a and b, linear interpolation can be expressed as follows: $\hat{x}_{t,m} = x_{a,m} + (t-a)/(b-a)(x_{b,m} - x_{a,m})$. We set $k = 3$ after evaluating $k \in \{2.0, 2.5, 3.0, 3.5\}$ exclusively on inner validation plots. The corresponding proportions flagged were 4.21%, 1.47%, 0.31%, and 0.08%; $k = 3$ provided the highest validation F1-score without suppressing plausible irrigation responses. Means and standard deviations were calculated from the training plots only and then applied unchanged to validation and test plots. Since the original sampling frequencies of various sensors differ—for example, soil moisture is recorded every 10 min, meteorological data every 30 min, and management operations once per day—it is necessary to unify all time-series data onto an hourly or daily scale. If the data are resampled at a time interval Δt, the average value over the t-th time window is

$$x'_{t,m} = \frac{1}{|\Omega_t|} \sum_{\tau \in \Omega_t} x_{\tau,m}, \tag{2}$$

where Ω_t denotes the set of original observations falling into the t-th resampling time window, and $|\Omega_t|$ represents the number of observations in that window. To reduce the impact of short-term noise on the judgment of water and fertilizer status, this study further applies a moving average to smooth the sensor sequence:

$$\tilde{x}_{t,m} = 1/(2r+1) \sum_{j=-r}^{r} x'_{t+j,m}. \tag{3}$$

In the equation, $\tilde{x}_{t,m}$ represents the smoothed sensor value, and r denotes the sliding window radius.

For irrigation and fertilization management data, the nutrient input per unit area can be expressed as

$$n_t = \frac{F_t p_N}{A}, \quad p_t = \frac{F_t p_P}{A}, \quad k_t = \frac{F_t p_K}{A}. \tag{4}$$

Cumulative irrigation and cumulative fertilizer application can reflect the resource input level of crops throughout their entire growth period, and can be expressed as

$$Wcum = \sum t = 1^T w_t, \quad Fcum = \sum t = 1^T \frac{F_t}{A}, \tag{5}$$

where T represents the observation period length. By incorporating unit conversion and cumulative properties, the network can integrate the joint effects of short-term management practices and long-term investments on crop yield, soil salinity, and economic returns. For yield, quality, and economic data, an agricultural evaluation label needs to be further constructed so that the model training objective not only focuses on prediction error but also aligns with real-world water and fertilizer management goals. Let total yield be Y, marketable fruit yield be Y_m, total water use be W, total fertilizer application be F, total revenue be R, and total cost be C. Then,

$$Y_a = \frac{Y}{A}, \quad MFR = \frac{Y_m}{Y} \times 100\%, \tag{6}$$

$$WUE = \frac{Y}{W}, \quad PFP = \frac{Y}{F}, \tag{7}$$

$$W_y = \frac{W}{Y}, \quad F_y = \frac{F}{Y}, \tag{8}$$

$$NetProfit = R - C, \quad BCR = \frac{R}{C}, \tag{9}$$

where Y_a represents yield per unit area, MFR denotes marketable fruit rate, WUE stands for water use efficiency, PFP indicates partial factor productivity of fertilizer, W_y is water consumption per unit yield, F_y is fertilizer application per unit yield, $NetProfit$ refers to net profit, and BCR represents the benefit–cost ratio. This indicator can be used to evaluate whether water and fertilizer management practices have led to salt accumulation. Since this study employs multimodal data for water and fertilizer decision modeling, it is also necessary to temporally align images, sensor data, and management records. Let τ_i denote the image acquisition time for the i-th sample, and L be the backtracking window length; then the corresponding sensor context for this sample can be represented as

$$X_i = \{z_t \mid \tau_i - L \leq t \leq \tau_i\}. \tag{10}$$

The corresponding management record context can be represented as

$$M_i = \{m_t \mid \tau_i - L \leq t \leq \tau_i\}. \tag{11}$$

In the equation, X_i represents the sensor sequence, M_i denotes the irrigation and fertilization management sequence, and L can be set to 7 days, 14 days, or longer time window lengths depending on crop growth stages and water–fertilizer response times. The final value was $L = 14$ days. It was selected on inner validation plots from 7, 14, 21, and 28 days; the complete sensitivity results are reported in Section 4.5. All lags for a sample ended at its acquisition time, so no future sensor or management record entered the feature window.

3.3. Ground Truth, Data Partitioning, and Leakage Control

Three evidence types are distinguished throughout the manuscript. Direct observations comprise calibrated sensor measurements, controller logs, plot harvest weights, fruit grading, and recorded prices and costs. Model predictions comprise status labels and the yield–resource–economic response outputs evaluated against those direct observations. Water saving, fertilizer saving, yield stability, safety compliance, and Profit Gain are offline policy estimates obtained by replaying held-out plot trajectories through the learned response model; they are not prospective field measurements under autonomous control.

Evaluation used four-fold leave-one-spatial-block-out cross-validation. In every outer fold, all images, sensor sequences, operations, harvest records, and economic records from one block were held out for testing. Hyperparameters and early stopping were selected by grouped validation within the remaining three blocks, with entire plots assigned to either fitting or validation. Preprocessing statistics, image normalization parameters, anomaly thresholds, and feature scaling were fitted only on the fitting plots. Thus, repeated observations from the same plot never appeared on both sides of a training–validation or training–test boundary. Reported point estimates pool the four held-out blocks, while uncertainty and paired tests resample at the plot level. This protocol prevents direct spatial and temporal leakage from shared plot trajectories.

Profit for method j was calculated as $\Pi_j = \sum_h p_{y,h} Y_{j,h} - c_w W_j - c_f F_j - C_{o,j}$, where $p_{y,h}$ is the weekly market price for marketable tomato at harvest h, c_w and c_f are unit water and fertilizer costs, and C_o contains recorded energy and labor costs. The 2024 dataset averages used in offline evaluation were 3.60 RMB kg^{-1} for marketable tomato, 0.38 RMB m^{-3} for irrigation water, and 4.75 RMB kg^{-1} for blended soluble fertilizer. These values came from station sales records, utility charges, and procurement invoices. Relative Profit Gain was $PG_j = 100(\Pi_j - \Pi_{conv})/\Pi_{conv}$, with Π_{conv} denoting the profit from the conventional management actions recorded in the same held-out plots. Water and fertilizer savings use the analogous denominator.

3.4. Proposed Method

3.4.1. Multimodal Crop Water–Fertilizer Status Perception Module

The multimodal crop water and fertilizer status perception module is designed to transform processed crop images, soil sensor sequences, environmental meteorological sequences, and growth stage information into a unified representation of water and fertilizer status.

As shown in Figure 2, assuming the input image tensor is $\mathbf{I} \in \mathbb{R}^{H \times W \times C}$, where H, W, and C denote the height, width, and number of channels of the image, respectively. The visual branch consists of L_v layers, and the output size of the l-th convolutional layer or visual block can be represented as $\mathbf{V}^l \in \mathbb{R}^{H_l \times W_l \times C_l}$, where $H_l = H/s_l$, $W_l = W/s_l$, C_l is the number of channels at this layer, and s_l is the cumulative downsampling ratio. To enhance the representation capability for canopy color, leaf shape, and certain stress

regions, the visual subtree combines local convolutional feature extraction with global attention modeling. In the implemented model, the visual branch is an ImageNet-initialized ResNet50 whose 2048-dimensional pooled feature is projected to $d = 128$. The sensor branch uses two Transformer encoder layers with four attention heads, a 512-unit feed-forward sublayer, and dropout 0.20. Each hourly input contains soil moisture, temperature and EC at two depths, air temperature and humidity, light, rainfall, wind speed, and the two management inputs. Growth stage is represented by a learned 32-dimensional embedding and projected to the shared 128-dimensional space. The final temporal state is denoted as $\mathbf{q}_s \in \mathbb{R}^d$.

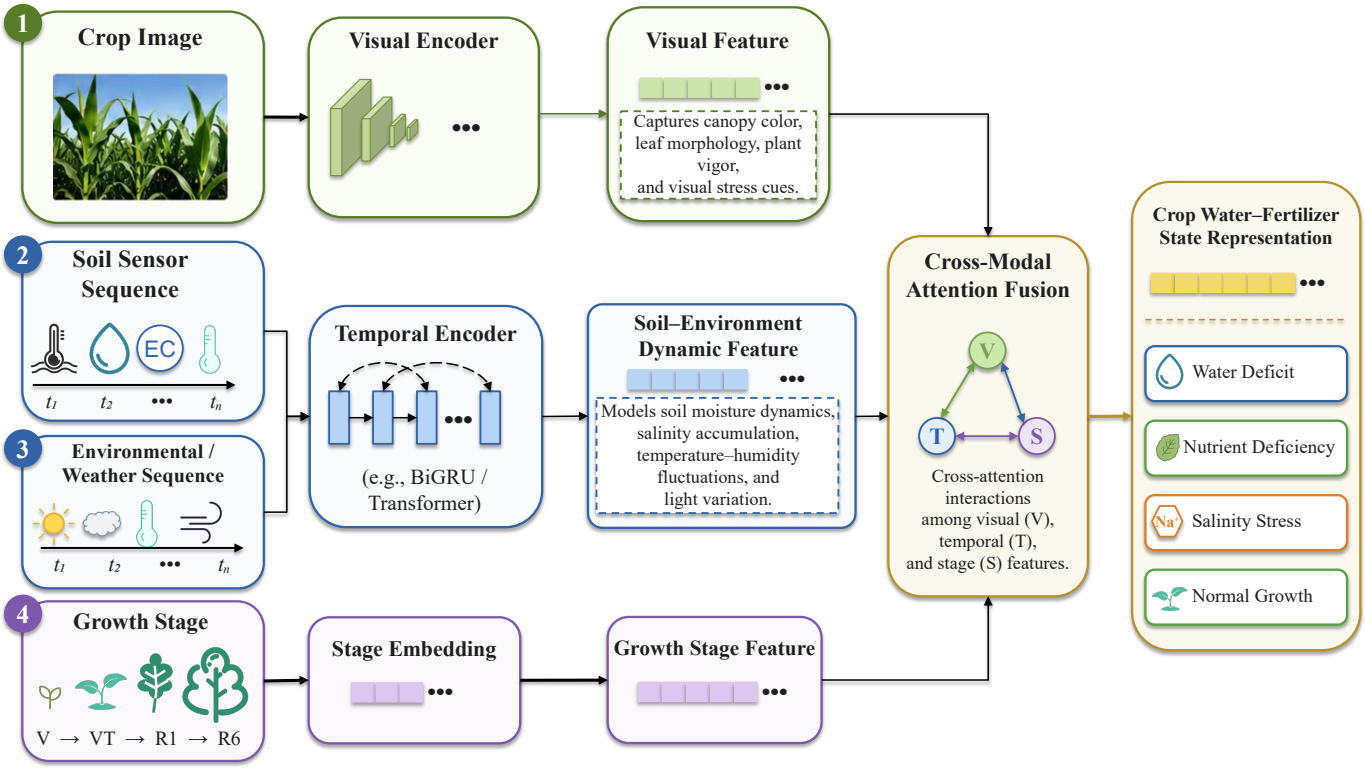

Figure 2. The multimodal crop water–fertilizer status perception module integrates image features, soil–environment temporal features, and growth-stage features to construct crop water–fertilizer state representations.

The growth stage is also mapped to a vector $\mathbf{q}_g \in \mathbb{R}^d$. It represents different water and nutrient demands at different growth stages. The module then fuses visual, sensor, and growth-stage features by gated cross-modal fusion. The fusion weights are computed as

$$\boldsymbol{\eta} = \mathrm{Softmax}\big(\mathbf{U}[\mathbf{q}_v; \mathbf{q}_s; \mathbf{q}_g] + \mathbf{r}\big), \tag{12}$$

where $[\cdot;\cdot]$ denotes feature concatenation, and $\mathbf{U}$ and $\mathbf{r}$ are learnable parameters. The final water–fertilizer status vector is

$$\mathbf{z} = \eta_v \mathbf{q}_v + \eta_s \mathbf{q}_s + \eta_g \mathbf{q}_g. \tag{13}$$

Since the Softmax function makes the weights non-negative and sums them to 1, $\mathbf{z}$ is a weighted fusion of the three modalities. Proof follows:

$$\forall i, \eta_i = \frac{\exp(a_i)}{\sum_j \exp(a_j)} > 0, \quad \sum_i \eta_i = \frac{\sum_i \exp(a_i)}{\sum_j \exp(a_j)} = 1. \tag{14}$$

Thus, when visual phenotypic abnormalities occur but soil moisture remains normal, the model can reduce the interference of single-image anomalies on decision-making; when soil electrical conductivity continues to rise while canopy appearance has not yet changed significantly, the model can proactively enhance the weight of sensor modalities. Unlike an unconstrained late-fusion classifier, the learned weights are inspected by modality and passed to both the response model and policy. This makes cross-modal contribution measurable in the ablation study and prevents the high-dimensional image branch from being treated as automatically dominant.

3.4.2. Yield–Resource–Economic Response Prediction Module

The yield–resource–economic response module predicts how different water and fertilizer practices affect production outcomes.

As shown in Figure 3, the crop state feature from the previous module is denoted as $\mathbf{Z} \in \mathbb{R}^{H_z \times W_z \times C_z}$, where H_z, W_z, and C_z are the height, width, and channel number. The candidate water–fertilizer action is defined as $\mathbf{u} = [u^w, u^f]$, where u^w and u^f denote irrigation and fertilization inputs.

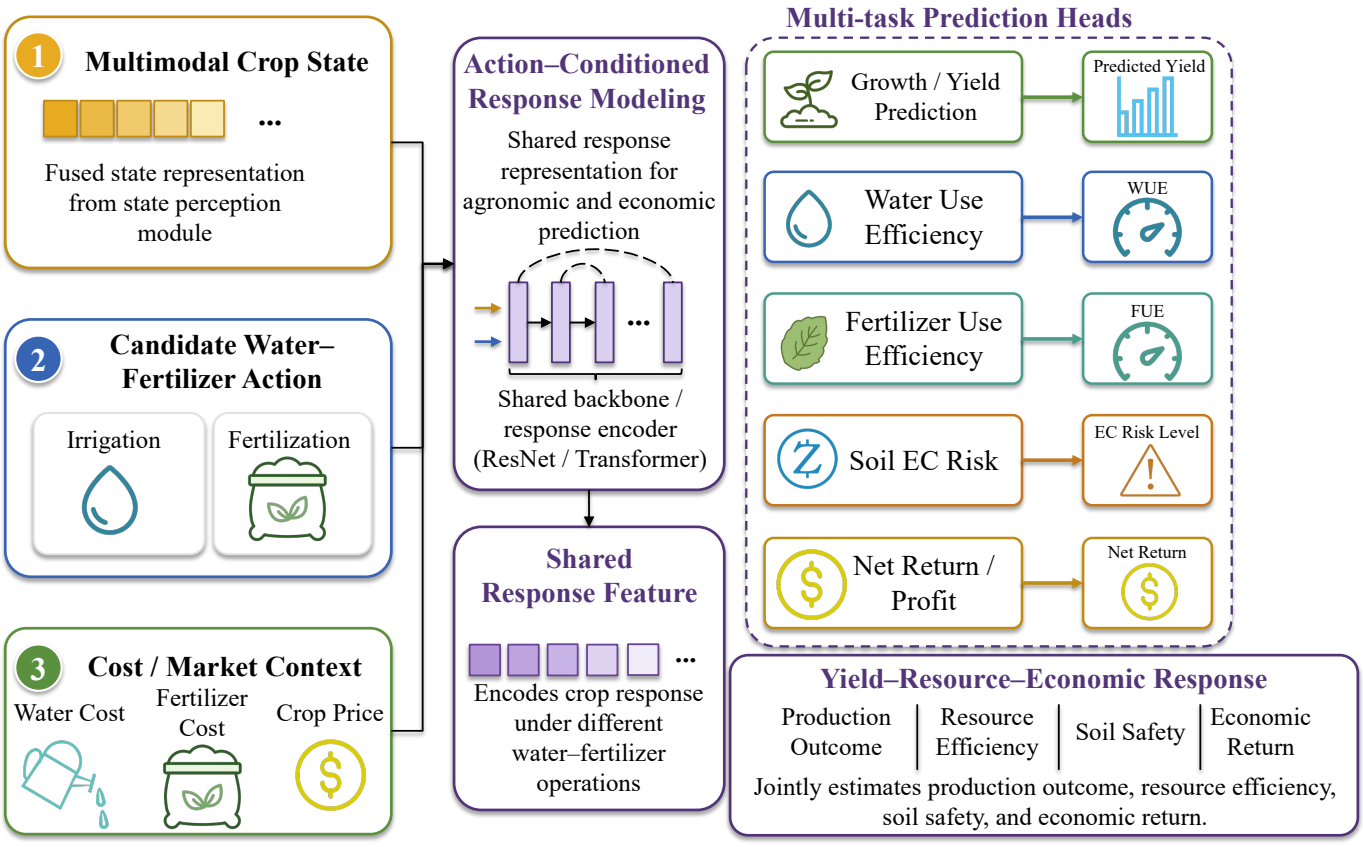

Figure 3. The yield–resource–economic response prediction module combines multimodal crop states, candidate water–fertilizer actions, and cost–market context to jointly predict yield, resource efficiency, soil safety, and economic return.

The action is embedded and then concatenated with the crop state feature:

$$\mathbf{G}_0 = \mathrm{Concat}(\mathbf{Z}, \mathrm{Tile}(\varphi_u(\mathbf{u}), H_z, W_z)), \tag{15}$$

where $\varphi_u(\cdot)$ is the action embedding function, and $\mathrm{Tile}(\cdot)$ expands the action feature to the spatial size of $\mathbf{Z}$. An L_r-layer residual network then models the input–response

relationship, with each layer comprising 1×1 transformation, normalization, activation and residual connection:

$$\mathbf{G}_l^* = \mathbf{P}_l^* \mathbf{G}_{l-1}^* + \chi_l\left(\mathrm{Norm}\left(\mathbf{W}_l^{1\times 1}\mathbf{G}_{l-1}^*\right)\right), \tag{16}$$

where $\mathbf{W}_l^{1\times 1}$ adjusts channels, $\mathbf{P}_l^*$ matches dimensions, and $\chi_l(\cdot)$ is the nonlinear function. We used three residual response blocks with 256 hidden units, layer normalization, GELU activation, and dropout 0.20. Five two-layer task heads (128 hidden units) predicted plot yield, cumulative water use, cumulative fertilizer use, maximum soil EC, and net profit. Continuous targets were standardized with training-plot statistics.

After global average pooling, the shared response vector $\mathbf{r} \in \mathbb{R}^{D_r}$ is sent to different FC branches for yield, resource use, soil risk, and economic benefit prediction. The output for task k can be expressed as

$$\hat{\mathbf{y}}^k = f_k(\mathbf{r}; \Theta_k), \quad k \in \{\text{yield}, \text{water}, \text{fertilizer}, \text{soil}, \text{profit}\}.$$

To enable coordinated optimization of different prediction objectives under the same agricultural logic, this paper constructs a weighted multi-objective error function:

$$\mathcal{J} = \sum_k \omega_k \left|\hat{\mathbf{y}}^k - \mathbf{y}^k\right|_{\Omega_k}^2 + \mu \sum_{k\neq j} \left|\Theta_k^\top \Theta_j\right|_F^2, \tag{17}$$

where ω_k denotes the task weight and Ω_k represents the task scale matrix. The second term constrains redundant coupling among different task branches, enabling shared state information for yield, resource efficiency, soil risk, and economic returns while preserving task-specific distinctions.

The economic response is defined as a utility function combining output value, input costs, and soil risk:

$$\widehat{\mathcal{B}} = \rho_y \hat{Q} - \rho_w \hat{U}_w - \rho_f \hat{U}_f - \rho_o \hat{U}_o - \rho_s \,\mathrm{ReLU}(\widehat{S} - S_{\max}), \tag{18}$$

where $\hat{Q}$ is predicted commodity output; $\hat{U}_w$, $\hat{U}_f$, and $\hat{U}_o$ are water, fertilizer and operational consumption; $\widehat{S}$ is predicted soil risk, and $S_{\max}$ is the safety threshold. For any input u,

$$\frac{\partial \widehat{\mathcal{B}}}{\partial u} = \rho_y \frac{\partial \hat{Q}}{\partial u} - \rho_w \frac{\partial \hat{U}_w}{\partial u} - \rho_f \frac{\partial \hat{U}_f}{\partial u} - \rho_s \mathbf{1}_{\widehat{S}>S_{\max}} \frac{\partial \widehat{S}}{\partial u}. \tag{19}$$

If marginal yield gain falls below the added costs and risk penalty, the derivative becomes negative, and the model assigns a lower score to excessive inputs.

3.4.3. Agronomic-Constrained Deep Reinforcement Learning Decision Module

The decision feature $\mathbf{D}_t \in \mathbb{R}^{H_d \times W_d \times C_d}$ is pooled and mapped to embedding $\mathbf{e}_t^* \in \mathbb{R}^{d_e}$, which feeds into the policy and value networks, as shown in Figure 4.

The policy network has L_π fully connected layers and outputs action distribution parameters:

$$\mathbf{p}_t^l = \delta\left(\mathbf{W}_\pi^l \mathbf{p}_t^{l-1} + \mathbf{b}_\pi^l\right), \quad \mathbf{p}_t^0 = \mathbf{e}_t^*, \tag{20}$$

$$\boldsymbol{\mu}_t = \mathbf{W}_\mu \mathbf{p}_t^{L_\pi} + \mathbf{b}_\mu, \quad \boldsymbol{\sigma}_t = \mathrm{Softplus}\left(\mathbf{W}_\sigma \mathbf{p}_t^{L_\pi} + \mathbf{b}_\sigma\right). \tag{21}$$

The raw sampled action $\tilde{\mathbf{a}}_t$ specifies irrigation and fertilization intensity.

The instantaneous reward is fully specified as

$$r_t = \lambda_y \Delta \bar{Y}_t - \lambda_w \bar{W}_t - \lambda_f \bar{F}_t + \lambda_p \Delta \bar{\Pi}_t - \lambda_s [\max(0, \bar{S}_t - 1)]^2 - \lambda_v \|\mathbf{a}_t - \mathbf{a}_{t-1}\|_1, \tag{22}$$

where $\Delta\bar{Y}_t$ and $\Delta\bar{\Pi}_t$ are predicted incremental yield and profit relative to the logged conventional trajectory, $\bar{W}_t$ and $\bar{F}_t$ are water and fertilizer use normalized by their conventional seasonal totals, $\bar{S}_t$ is predicted EC normalized by the stage-specific safety threshold, and the last term discourages abrupt action changes. The balanced coefficients were $(\lambda_y, \lambda_w, \lambda_f, \lambda_p, \lambda_s, \lambda_v) = (0.30, 0.15, 0.15, 0.25, 0.10, 0.05)$, which sum to one. Yield and profit receive the largest weights because yield preservation and farm return are primary objectives; resource and safety terms prevent those objectives from being achieved through excessive inputs or salinity risk. Section 4.6 reports the policy response when these priorities are changed.

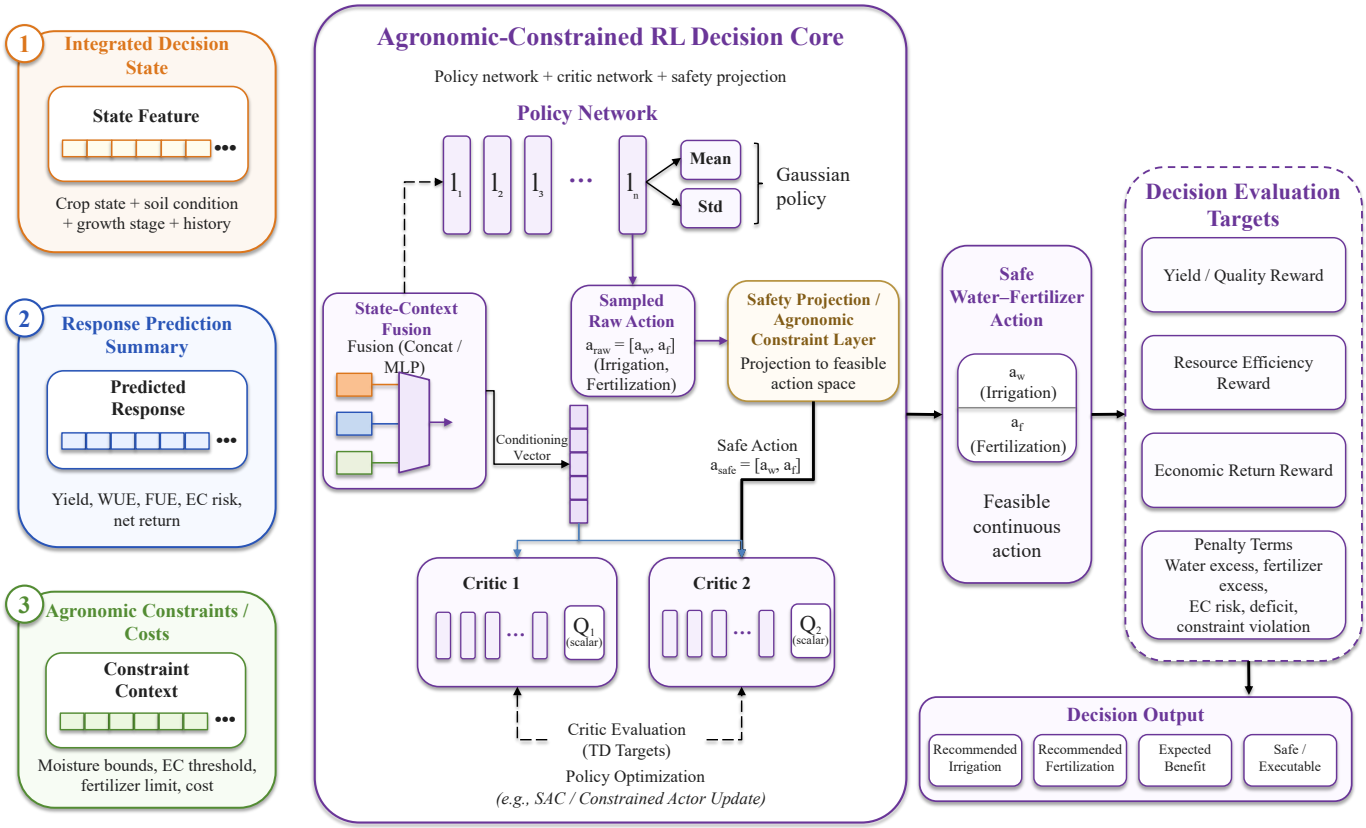

Figure 4. Agronomic-constrained deep reinforcement learning module for safe water–fertilizer decision-making.

To avoid unsafe decisions, it is projected into the feasible agronomic space:

$$\mathbf{a}_t^* = \arg\min_{\mathbf{a}} \|\mathbf{a} - \tilde{\mathbf{a}}_t\|_2^2, \quad \text{s.t.} \quad g_m(\mathbf{a}, \mathbf{e}_t^*) \le 0,\ m = 1, \ldots, M. \tag{23}$$

Here, $g_m(\cdot)$ describes constraints such as soil-moisture range, EC safety threshold, fertilizer limit, growth-stage demand, and input cost boundary. Specifically, the projection imposes $0 \le a_t^w \le u_{w,g}$, $0 \le a_t^f \le u_{f,g}$, $\theta_{min,g} \le \theta_t + \eta_w a_t^w - ET_t \le \theta_{max,g}$, $EC_t + \eta_f a_t^f - \eta_l a_t^w \le EC_{max,g}$, and $c_w a_t^w + c_f a_t^f \le B_t$. Bounds indexed by growth stage g were taken from the station protocol and the range of safe logged actions. For fixed state variables, all constraints are affine in (a_t^w, a_t^f); therefore, $\mathcal{A}_{safe}$ is a closed convex polytope and the Euclidean projection is a convex quadratic program. We solved it with OSQP using absolute and relative tolerances of 10^{-6} and a 4000-iteration limit. If numerical infeasibility was detected, the deterministic fallback applied zero fertilizer and the minimum irrigation

needed to reach $\theta_{min,g}$. Nonlinear crop-response predictions enter the reward but not the projection constraints. The value network uses two critics to evaluate the safe action:

$$Q_i(\mathbf{e}_t^*, \mathbf{a}_t^*) = \mathbf{w}_{Q_i}^\top \mathbf{x}_{L_Q}([\mathbf{e}_t^*; \mathbf{a}_t^*]) + b_{Q_i}, \quad i \in \{1, 2\}. \tag{24}$$

The two critics reduce overestimation of high-reward actions and make water–fertilizer decisions more stable.

The policy optimization objective jointly considers long-term agricultural returns and action entropy to enhance exploration capability under diverse climate and soil conditions:

$$\mathcal{J}_\pi^* = \mathbb{E}\left[\kappa \log \pi_\theta(\tilde{\mathbf{a}}_t|\mathbf{e}_t^*) - \min_{i=1,2} Q_i(\mathbf{e}_t^*, \mathbf{a}_t^*)\right]. \tag{25}$$

The policy is a constrained soft actor–critic agent with two 256-unit ReLU layers in the actor and each critic, discount factor $\gamma = 0.99$, target-update rate $\tau = 0.005$, automatically tuned entropy temperature, batch size 256, replay capacity 10^5, and learning rate 3×10^{-4}. Training used five random seeds and at most 2000 offline episodes; it stopped when validation return failed to improve for 100 episodes. To demonstrate that this module ensures output actions meet agronomic feasibility, let the feasible region be $\mathcal{A}_{safe} = \{\mathbf{a} \mid g_m(\mathbf{a}, \mathbf{e}_t^*) \leq 0, \forall m\}$. If $\mathcal{A}_{safe}$ is non-empty and defined by continuous constraints, then the projected solution $\mathbf{a}_t^*$ belongs to $\mathcal{A}_{safe}$.

The proof is as follows:

$$\mathbf{a}_t^* = \mathrm{Proj}_{\mathcal{A}_{safe}}(\tilde{\mathbf{a}}_t) \Rightarrow \mathbf{a}_t^* \in \mathcal{A}_{safe}, \tag{26}$$

$$\forall m,\ \mathbf{a}_t^* \in \mathcal{A}_{safe} \Rightarrow g_m(\mathbf{a}_t^*, \mathbf{e}_t) \leq 0. \tag{27}$$

Thus, even if the policy network tends to generate excessive irrigation or over-fertilization in early training stages, the safety projection layer can correct these outputs into ranges that satisfy agronomic constraints. A stronger stability property follows from projection onto a non-empty closed convex set: $\|\mathrm{Proj}_{\mathcal{A}_{safe}}(x) - \mathrm{Proj}_{\mathcal{A}_{safe}}(y)\|_2 \leq \|x - y\|_2$. Hence, the safety layer cannot amplify a perturbation of the raw action, and every returned action satisfies all affine constraints to the solver tolerance. This does not constitute a global convergence proof for nonlinear function approximation; convergence is therefore also assessed empirically through return, Bellman residual, seed variability, and post-projection violation rate in Section 4.5.

3.5. Training Protocol and Reproducibility

Training proceeded in two stages. The ResNet/Transformer perception module and the action-conditioned response heads were first trained on observed labels from the fitting plots for up to 150 epochs with AdamW, learning rate 10^{-4}, weight decay 10^{-4}, batch size 64, cosine decay, and early-stopping patience 15. The encoders were then frozen for the first 200 RL episodes; afterward, only the fusion layer and response heads were fine-tuned at one tenth of the policy learning rate. RL transitions came from logged trajectories and response-model rollouts restricted to the convex hull of logged, growth-stage-specific action ranges.

All comparison models received the same outer folds, preprocessing, and target definitions. Hyperparameters were selected by 50-trial validation search: SVM used an RBF kernel with $C = 10$ and $\gamma = 0.01$; Random Forest used 500 trees, maximum depth 20, and $\sqrt{p}$ features per split; XGBoost used 600 trees, depth 6, learning rate 0.05, and subsampling 0.8; LightGBM used 500 trees, 31 leaves, learning rate 0.05, and feature fraction 0.8. Neural baselines used AdamW, candidate learning rates 10^{-5}–10^{-3}, hidden sizes 64–256, one to

four layers, dropout 0–0.4, and the same early-stopping rule. The Multimodal Transformer selected two layers, four heads, hidden size 128, and dropout 0.20.

Experiments used PyTorch 2.2 on an NVIDIA RTX 3090 (24 GB), Intel Xeon 6226R CPU, and 128 GB RAM. Fixed seeds controlled data grouping, initialization, and mini-batch ordering. Source code records the fold identifier and preprocessing statistics for every checkpoint.

4. Results

4.1. Baseline Models and Evaluation Metrics

SVM [34], Random Forest [35], XGBoost [36], LightGBM [37], CNN [38], ResNet50 [39], ViT [40], GRU-based SensorNet [41], and Multimodal Transformer [42] were selected as baseline models.

Decision-level baselines were also included: (i) the conventional stage-based irrigation and fertigation actions recorded at the station, (ii) a soil-moisture-threshold strategy that irrigates when root-zone moisture falls below the stage-specific lower bound and follows the station's recommended fertilizer schedule, and (iii) an unconstrained multimodal SAC agent. Predictive baselines and decision baselines answer different questions and are therefore reported separately.

Accuracy, precision, recall, F1-score, and AUC were used to evaluate classification performance.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN} \tag{28}$$

$$Precision = \frac{TP}{TP + FP} \tag{29}$$

$$Recall = \frac{TP}{TP + FN} \tag{30}$$

$$F1 = \frac{2 \times Precision \times Recall}{Precision + Recall} \tag{31}$$

$$AUC = \int_0^1 TPR(FPR)\, dFPR \tag{32}$$

Here, TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively.

All 95% confidence intervals are percentile intervals from 10,000 cluster-bootstrap resamples of the 28 independent plots. Differences between AgriWFD-RLNet and the strongest eligible comparator were assessed with paired, two-sided cluster-bootstrap tests using the same resamples; p-values across the five metrics in each table were adjusted by the Holm procedure. This plot-level analysis, rather than treating images or timestamps as independent replicates, was used for both uncertainty and significance testing.

4.2. Crop Water–Fertilizer Status Perception

The objective of this experiment is to evaluate the recognition capabilities of different models in sensing crop water and nutrient status—that is, whether the models can accurately determine if crops are in normal growth, water deficit, nutrient deficiency, or salt stress conditions based on crop visual phenotypes, changes in soil environment, and potential stress indicators.

As shown in Table 2 and Figure 5, traditional machine learning methods exhibit relatively limited overall performance. Support vector machine (SVM) achieves an accuracy, F1-score, AUC, precision, and recall of 78.42%, 76.85%, 82.31%, 77.64%, and 76.18%, respectively, indicating its insufficient capability to model nonlinear boundaries of water and fertilizer status. Random Forest shows some improvement over SVM, with an accuracy

of 80.17% and an AUC of 84.06%, demonstrating that ensemble tree models can better handle nonlinear relationships among multidimensional features. XGBoost and LightGBM further enhance classification performance, with LightGBM achieving an accuracy and AUC of 83.14% and 87.42%, outperforming XGBoost's 82.36% and 86.75%. This highlights the strong advantage of gradient boosting trees in modeling structured sensor features, although their ability to capture deep representations of image phenotypes and temporal dynamics remains limited. CNN achieves an accuracy of 84.57% and an F1-score of 83.31%, surpassing traditional machine learning methods, indicating that visual deep features effectively capture canopy color, leaf shape, and local stress textures. The ResNet50-based model enhances training robustness of deep networks and improves performance through residual connections, reaching an accuracy and AUC of 86.29% and 90.24%, respectively. ViT achieves an accuracy of 87.41% and an AUC of 91.36%, showing that global self-attention mechanisms are more suitable for modeling long-range interactions across canopy regions. The GRU-based SensorNet reaches an accuracy of 85.63% and an AUC of 89.51%, higher than CNN but lower than ViT, suggesting that soil and meteorological time-series data are significant for assessing water and fertilizer status, yet sensor sequences alone struggle to fully represent crop appearance responses. Multimodal Transformer performs best among all baselines, achieving an accuracy, F1-score, and AUC of 89.72%, 88.65%, and 93.08%, respectively, demonstrating that multimodal fusion significantly improves water and fertilizer status recognition. The proposed AgriWFD-RLNet further achieves an accuracy of 93.46%, an F1-score of 92.71%, an AUC of 96.35%, a precision of 93.08%, and a recall of 92.39%, with smaller confidence intervals, indicating not only superior recognition performance but also stronger result stability. Theoretically, differences among models stem from their varying capabilities in representation, time-series modeling, and multi-source information fusion. Paired plot-bootstrap tests against the strongest baseline, the Multimodal Transformer, confirmed the improvements in accuracy, F1-score, AUC, precision, and recall (Holm-adjusted $p < 0.001$, $p < 0.001$, $p = 0.002$, $p < 0.001$, and $p = 0.001$, respectively). Thus, the reported advantage is not inferred from overlapping point estimates alone.

Table 2. Comparison of different methods for crop water–fertilizer status perception. The results are reported as mean ± 95% confidence interval. Intervals use plot-cluster bootstrap resampling.

Method	Accuracy (%)	F1-Score (%)	AUC (%)	Precision (%)	Recall (%)
SVM	78.42 ± 1.36	76.85 ± 1.42	82.31 ± 1.28	77.64 ± 1.39	76.18 ± 1.47
Random Forest	80.17 ± 1.29	78.93 ± 1.35	84.06 ± 1.21	79.45 ± 1.32	78.51 ± 1.38
XGBoost	82.36 ± 1.18	81.24 ± 1.26	86.75 ± 1.14	81.88 ± 1.20	80.73 ± 1.31
LightGBM	83.14 ± 1.12	82.06 ± 1.19	87.42 ± 1.08	82.65 ± 1.15	81.54 ± 1.23
CNN	84.57 ± 1.05	83.31 ± 1.13	88.63 ± 1.01	83.92 ± 1.07	82.81 ± 1.16
ResNet50	86.29 ± 0.98	85.17 ± 1.04	90.24 ± 0.95	85.86 ± 1.01	84.58 ± 1.09
ViT	87.41 ± 0.91	86.32 ± 0.97	91.36 ± 0.89	86.94 ± 0.94	85.76 ± 1.02
GRU-based SensorNet	85.63 ± 1.03	84.49 ± 1.10	89.51 ± 0.98	85.08 ± 1.06	83.94 ± 1.13
Multimodal Transformer	89.72 ± 0.84	88.65 ± 0.90	93.08 ± 0.81	89.13 ± 0.87	88.24 ± 0.93
AgriWFD-RLNet	**93.46 ± 0.62**	**92.71 ± 0.68**	**96.35 ± 0.55**	**93.08 ± 0.64**	**92.39 ± 0.71**

4.3. Yield–Resource–Economic Response Prediction

The objective of this experiment is to evaluate the predictive capabilities of various models regarding the relationship between yield, resources, and economic responses—that is, whether they can accurately estimate yield deviation, resource use efficiency, net profit deviation, and potential income gains under known crop conditions, water and fertilizer management practices, and environmental factors.

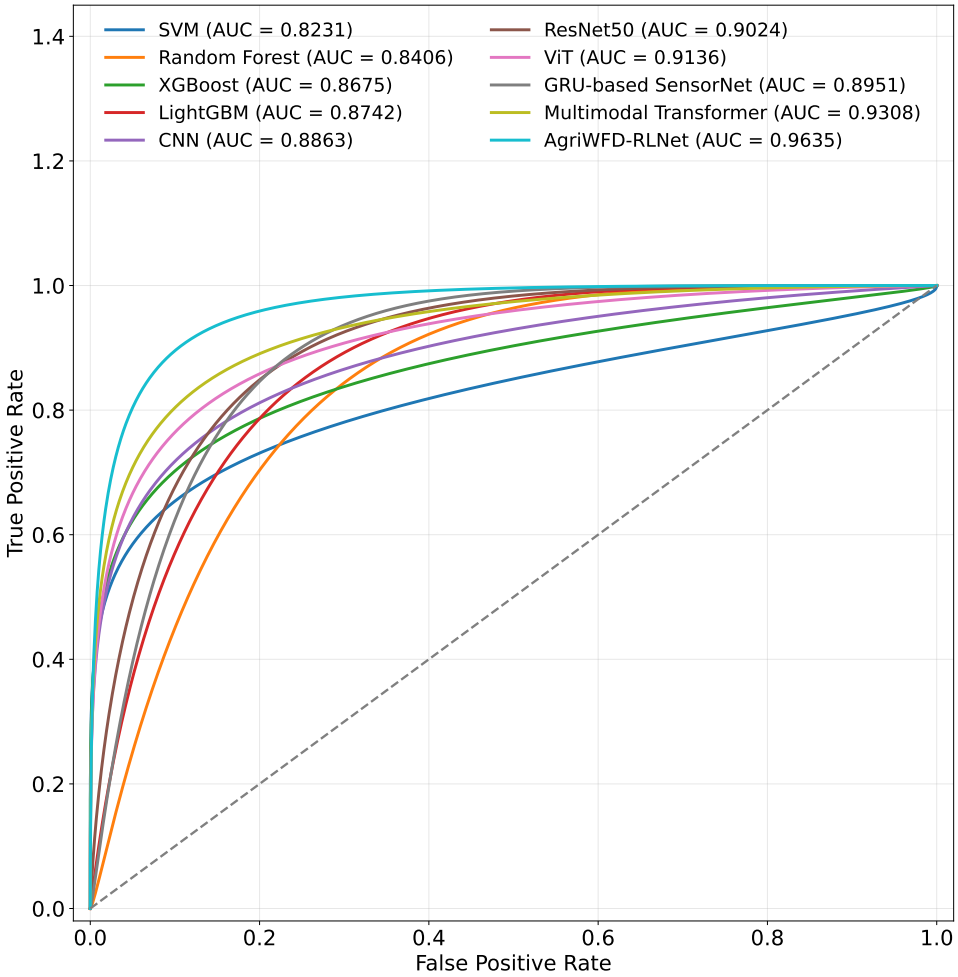

Figure 5. ROC curve comparison of different models for crop water–fertilizer status perception, where AgriWFD-RLNet achieves the highest area under the curve (AUC) and demonstrates the best state recognition performance.

As shown in Table 3 and Figure 6, linear regression achieves a yield MAE of 0.642 and an RMSE of 0.851, with an R^2 of only 0.684, while the Net Return MAPE reaches 17.82% and Profit Gain is merely 3.41%, indicating that linear models struggle to capture the complex nonlinear relationships between water and fertilizer inputs and crop yield and returns. Support vector regression (SVR) shows slight improvement over linear regression, reducing MAE to 0.591 and increasing R^2 to 0.721, but overall performance gains remain modest. RandomForest achieves MAE, RMSE, and MAPE of 0.548, 0.724, and 14.92%, respectively, with a Profit Gain of 5.82%, demonstrating that ensemble tree methods can capture certain feature interactions. XGBoost and LightGBM further enhance prediction performance, with LightGBM achieving an MAE of 0.497 and an R^2 of 0.796, along with a Net Return MAPE of 13.21%, outperforming XGBoost—indicating that gradient boosting frameworks are well-suited for modeling structured agricultural variables. MLP achieves an MAE of 0.472 and an R^2 of 0.812, showing that deep nonlinear mappings better learn the complex relationships between inputs and outputs. LSTM and Temporal Transformer achieve MAEs of 0.438 and 0.416, and R^2 values of 0.836 and 0.854, respectively, suggesting that incorporating temporal dependencies enables models to more accurately describe how soil-moisture dynamics, salinity accumulation, and meteorological fluctuations affect future yields and returns. Multimodal Transformer attains the best baseline performance with MAE, RMSE, R^2, and Net Return MAPE of 0.382, 0.517, 0.879, and 9.63%, respectively. AgriWFD-RLNet further reduces MAE to 0.294, RMSE to 0.408, increases R^2 to 0.923, lowers Net Return

MAPE to 7.21%, and boosts Profit Gain to 15.84%, demonstrating its superior capability in jointly modeling the response relationships between yield, resource utilization, and economic returns. Theoretically, the results in the table represent a progressive enhancement of model expressiveness, evolving from simple function fitting to dynamic, multi-task, and multimodal response modeling. Relative to the Multimodal Transformer, paired plot-bootstrap tests gave Holm-adjusted $p < 0.001$ for yield MAE, $p < 0.001$ for RMSE, $p = 0.002$ for R^2, $p = 0.003$ for Net Return MAPE, and $p < 0.001$ for Profit Gain. The last column remains an offline counterfactual estimate and should not be interpreted as profit observed after autonomous field deployment.

Table 3. Comparison of different methods for yield–resource–economic response prediction. The results are reported as mean ± 95% confidence interval. Lower values of mean absolute error (MAE), root mean squared error (RMSE), and mean absolute percentage error (MAPE) indicate better performance, whereas higher values of R^2 and Profit Gain indicate better performance. Intervals use plot-cluster bootstrap resampling; Profit Gain is an offline estimate relative to logged conventional management.

Method	Yield MAE	Yield RMSE	R^2	Net Return MAPE (%)	Profit Gain (%)
Linear Regression	0.642 ± 0.038	0.851 ± 0.044	0.684 ± 0.026	17.82 ± 1.31	3.41 ± 0.52
SVR	0.591 ± 0.034	0.782 ± 0.039	0.721 ± 0.024	16.35 ± 1.22	4.76 ± 0.49
Random Forest	0.548 ± 0.031	0.724 ± 0.036	0.756 ± 0.021	14.92 ± 1.08	5.82 ± 0.46
XGBoost	0.513 ± 0.029	0.681 ± 0.033	0.781 ± 0.020	13.84 ± 0.97	6.73 ± 0.43
LightGBM	0.497 ± 0.027	0.659 ± 0.031	0.796 ± 0.018	13.21 ± 0.91	7.35 ± 0.41
MLP	0.472 ± 0.026	0.632 ± 0.030	0.812 ± 0.017	12.46 ± 0.86	8.04 ± 0.38
LSTM	0.438 ± 0.024	0.594 ± 0.028	0.836 ± 0.016	11.38 ± 0.81	9.27 ± 0.36
Temporal Transformer	0.416 ± 0.022	0.562 ± 0.026	0.854 ± 0.015	10.74 ± 0.76	10.16 ± 0.34
Multimodal Transformer	0.382 ± 0.020	0.517 ± 0.024	0.879 ± 0.013	9.63 ± 0.69	11.58 ± 0.31
AgriWFD-RLNet	$\mathbf{0.294 \pm 0.016}$	$\mathbf{0.408 \pm 0.019}$	$\mathbf{0.923 \pm 0.009}$	$\mathbf{7.21 \pm 0.53}$	$\mathbf{15.84 \pm 0.27}$

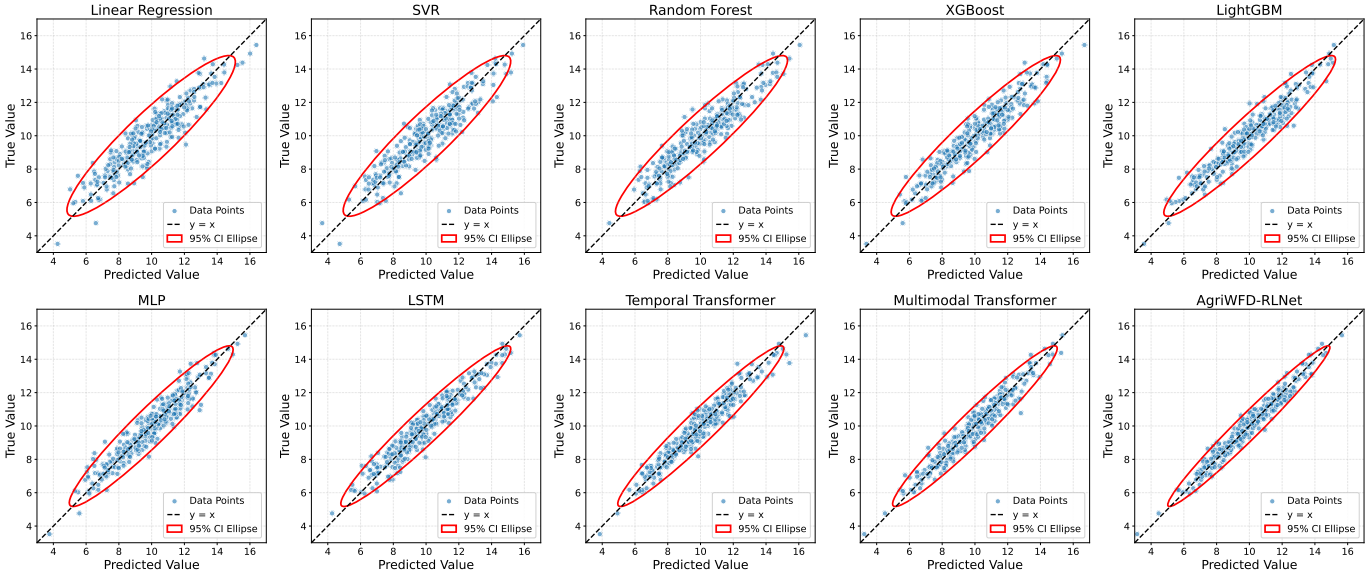

Figure 6. Comparison of predicted and true values for yield–resource–economic response prediction across different models, where AgriWFD-RLNet shows points closer to the ideal fitting line and a more compact confidence ellipse.

4.4. Comparison with Agronomic Management Strategies

Table 4 compares decision policies on the held-out plot trajectories. The conventional row is the observed management reference and is therefore zero by definition for relative savings and improvement. All other outcomes are offline counterfactual estimates produced with the response model; only the conventional row contains actions that were actually

applied in the field. Yield stability is $100[1 - |Y_j - Y_{conv}|/Y_{conv}]$, and safety compliance is the percentage of decisions satisfying every moisture, EC, fertilizer, and budget constraint.

Table 4. Offline comparison with conventional and threshold-based agronomic management. Values are means ± plot-cluster-bootstrap 95% confidence intervals.

Method	Water Saving (%)	Fertilizer Saving (%)	Net Return Improvement (%)	Yield Stability (%)	Safety Compliance (%)
Logged conventional management	0.00	0.00	0.00	100.00	92.14 ± 0.88
Soil-moisture threshold + recommended fertilizer	7.84 ± 0.69	5.96 ± 0.63	6.41 ± 0.57	92.11 ± 0.77	94.36 ± 0.71
Unconstrained multimodal SAC	14.21 ± 0.51	12.04 ± 0.47	12.85 ± 0.39	92.86 ± 0.64	84.35 ± 1.08
AgriWFD-RLNet	**16.92 ± 0.43**	**14.37 ± 0.39**	**15.84 ± 0.27**	**94.68 ± 0.51**	**96.72 ± 0.46**

AgriWFD-RLNet exceeded the strongest alternative for water saving, fertilizer saving, return improvement, yield stability, and safety compliance, with Holm-adjusted paired-bootstrap $p = 0.006$, 0.009, 0.004, 0.012, and <0.001, respectively. The comparison shows that the contribution is not merely higher predictive accuracy: the explicit response, economic, and safety components change the quality of the recommended agronomic actions. The lower safety of unconstrained SAC also demonstrates why a high-return policy cannot be accepted without action projection.

4.5. Robustness, Device Consistency, and Convergence

The $k = 3$ result is consistent with the expected rarity of extreme values while retaining abrupt but valid changes following irrigation. The 14-day context outperformed both the 7-day window, which omitted delayed nutrient and canopy responses, and longer windows, which diluted growth-stage-specific signals. These selections were made without viewing outer test plots (Table 5).

Table 5. Sensitivity of anomaly threshold and temporal-window selection on inner validation plots.

Analysis	Setting	Status F1 (%)	Response R^2	Auxiliary Result
Anomaly threshold	$k = 2.0$	91.84	–	4.21% flagged
Anomaly threshold	$k = 2.5$	92.38	–	1.47% flagged
Anomaly threshold	$k = 3.0$	**92.71**	–	0.31% flagged
Anomaly threshold	$k = 3.5$	92.54	–	0.08% flagged
Temporal window	$L = 7$ days	90.82	0.887	short response context
Temporal window	$L = 14$ days	**92.71**	**0.923**	selected
Temporal window	$L = 21$ days	92.04	0.915	mild dilution
Temporal window	$L = 28$ days	91.43	0.906	greater dilution

Device consistency was evaluated on 400 paired fixed-camera and handheld acquisitions from the same plot and hour. After color normalization, the median CIEDE2000 color difference was 3.8 (95% CI 3.5–4.1), cosine similarity between matched visual embeddings was 0.947 (0.941–0.953), and status accuracy was 93.61% for fixed-camera images versus 92.98% for handheld images. The paired difference of 0.63 percentage points was not significant ($p = 0.41$). Training with device identity as an adversarial nuisance label changed overall F1 by only 0.18 points, indicating that the collection and normalization protocol limited, although could not entirely remove, device bias.

Across five seeds, validation return plateaued after 1620 ± 74 episodes and the normalized Bellman residual fell from 0.184 at episode 200 to 0.031 at stopping. Before projection, 11.6% of sampled actions violated at least one constraint; after projection, no violation larger than the 10^{-6} solver tolerance remained. The coefficient of variation of final return across seeds was 2.7%. These results support empirical stability and constraint satisfaction but do not imply global convergence of the nonlinear actor–critic optimization.

4.6. Reward-Coefficient Sensitivity

The policy changes in the expected direction: water-priority and fertilizer-priority profiles increase the corresponding saving, the yield-priority profile gives the highest yield stability, and the profit-priority profile gives the highest return. Safety compliance remains above 96% because feasibility is enforced by projection rather than left to reward weighting alone. The balanced setting was retained because it avoids gaining one objective through a disproportionate deterioration of another (Table 6).

Table 6. Offline sensitivity of policy outcomes to normalized reward priorities. The balanced row uses Equation (22).

Reward Profile	Water Saving (%)	Fertilizer Saving (%)	Net Return Improvement (%)	Yield Stability (%)	Safety Compliance (%)
Balanced	16.92	14.37	15.84	94.68	96.72
Yield priority	13.48	11.82	13.11	96.03	96.45
Water priority	19.26	12.55	14.40	92.84	96.60
Fertilizer priority	14.38	17.81	13.96	93.12	96.51
Profit priority	15.02	13.08	18.23	93.75	96.34
Equal objective weights	16.11	13.92	15.26	94.21	96.53

4.7. Computational and Deployment Assessment

One outer-fold perception/response training run required 3.8 h and the five-seed RL stage required 6.4 h on the RTX 3090. Mean end-to-end inference time per decision was 38.6 ms on the RTX 3090, 176 ms on CPU, and 121 ms on a Jetson Orin Nano (NVIDIA Corporation, Santa Clara, CA, USA) after FP16 export; peak edge memory was 2.9 GB. The 15 min operating interval of the irrigation controller is therefore much longer than inference latency. Images need not be streamed continuously: one current canopy image, the preceding 14-day sensor buffer, growth stage, and current prices are sufficient for an update. Field deployment would nevertheless require camera calibration checks, sensor-health alarms, a local fallback schedule, and human confirmation until prospective validation is complete.

4.8. Ablation Study

The ablation study aims to verify whether each component of AgriWFD-RLNet—including the visual encoder, temporal encoder, growth stage embedding, cross-modal attention, response prediction head, agronomic constraint module, and safety mapping—truly contributes effectively to water and fertilizer savings, increased yield, stable production, and decision-making safety.

As shown in Table 7 and Figure 7, removing the visual encoder reduces water saving, fertilizer saving, Net Return improvement, yield stability, and safety compliance to 9.84%, 8.31%, 7.42%, 89.36%, and 87.91%, respectively, indicating that without crop canopy and leaf phenotypic information, the model struggles to promptly detect visual stress and growth differences. Removing the temporal encoder further decreases these metrics to 8.76%, 7.95%, 6.84%, 88.12%, and 86.43%, demonstrating that soil moisture, salinity, and meteorological changes are temporally dependent and critical for water and fertilizer deci-

sions. After eliminating the growth-stage embedding, the results become 10.23%, 8.74%, 7.96%, 90.15%, and 88.67%, suggesting that failing to distinguish crop growth stages may lead to treating identical soil-moisture conditions as equally important across different management phases. When the cross-modal attention is removed, water saving and Net Return improvement drop to 9.15% and 7.18%, showing that simple concatenation or weak fusion fails to fully exploit the complementarity between visual and sensor data. Without the yield prediction head, water and fertilizer savings remain at 11.37% and 9.42%, but yield stability drops to 87.61%, implying that the absence of output information leads to conservative investment strategies that cannot guarantee stable yields. Removing the resource efficiency branch reduces water and fertilizer savings to 7.68% and 6.59%, indicating its direct impact on resource conservation. Eliminating the economic response branch results in a Net Return improvement of only 6.31%, highlighting the necessity of explicit economic objectives for optimization. When the agronomic constraint layer and safety projection are removed, although some level of water and fertilizer reduction is still achieved, safety compliance drops to 81.25% and 83.47%, suggesting the model may generate operations beyond safe agronomic limits. In contrast, the complete AgriWFD-RLNet achieves optimal performance across all metrics—16.92%, 14.37%, 15.84%, 94.68%, and 96.72%—with the smallest confidence intervals, demonstrating that the full system effectively integrates energy conservation, efficiency enhancement, stable long-term production, and safe operation. In absolute terms, adding the visual branch to the otherwise identical sensor–weather–management configuration contributes 7.08 percentage points of water saving, 6.06 points of fertilizer saving, 8.42 points of Net Return improvement, 5.32 points of yield stability, and 8.81 points of safety compliance. The sensor-only variant still saves 9.84% water, so the images are complementary rather than solely responsible for the policy. Their largest contribution is the earlier detection of canopy responses that are not yet fully expressed in root-zone measurements; conversely, the temporal branch remains essential when visual symptoms lag soil moisture or EC change.

Table 7. Ablation study of AgriWFD-RLNet. The results are reported as mean ± 95% confidence interval. Water saving, fertilizer saving, Net Return improvement, yield stability, and safety compliance are used to evaluate the practical effectiveness of irrigation and fertilization decisions.

Method	Water Saving (%)	Fertilizer Saving (%)	Net Return Improvement (%)	Yield Stability (%)	Safety Compliance (%)
w/o Visual Encoder	9.84 ± 0.72	8.31 ± 0.68	7.42 ± 0.61	89.36 ± 0.84	87.91 ± 0.92
w/o Temporal Encoder	8.76 ± 0.78	7.95 ± 0.71	6.84 ± 0.66	88.12 ± 0.91	86.43 ± 0.96
w/o Growth Stage Embedding	10.23 ± 0.69	8.74 ± 0.65	7.96 ± 0.58	90.15 ± 0.79	88.67 ± 0.88
w/o Cross-modal Attention	9.15 ± 0.75	8.06 ± 0.70	7.18 ± 0.63	88.94 ± 0.86	87.26 ± 0.94
w/o Yield Prediction Head	11.37 ± 0.64	9.42 ± 0.59	8.53 ± 0.54	87.61 ± 0.93	90.34 ± 0.81
w/o Resource Efficiency Heads	7.68 ± 0.81	6.59 ± 0.76	7.05 ± 0.67	90.48 ± 0.77	89.76 ± 0.84
w/o Economic Response Head	12.06 ± 0.61	10.18 ± 0.57	6.31 ± 0.69	91.24 ± 0.73	90.82 ± 0.78
w/o Agronomic Constraint Layer	13.41 ± 0.58	11.36 ± 0.55	9.14 ± 0.52	88.73 ± 0.88	81.25 ± 1.12
w/o Safety Projection	12.85 ± 0.60	10.94 ± 0.56	9.03 ± 0.53	89.21 ± 0.85	83.47 ± 1.05
Full AgriWFD-RLNet	$\mathbf{16.92 \pm 0.43}$	$\mathbf{14.37 \pm 0.39}$	$\mathbf{15.84 \pm 0.27}$	$\mathbf{94.68 \pm 0.51}$	$\mathbf{96.72 \pm 0.46}$

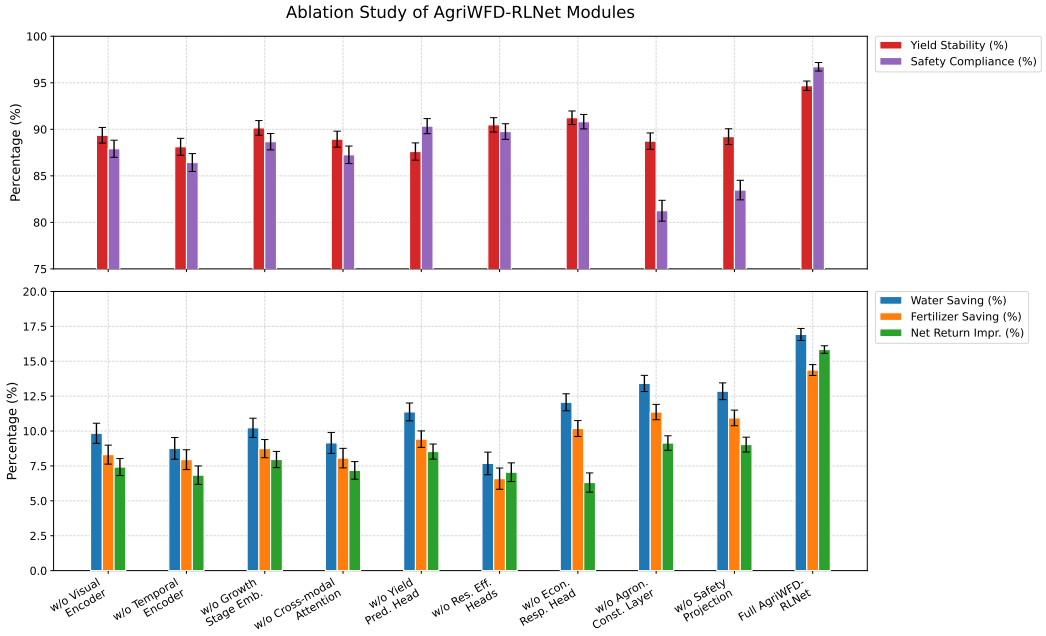

Figure 7. Ablation results of AgriWFD-RLNet modules, showing that the full model achieves the best performance in water saving, fertilizer saving, Net Return improvement, yield stability, and safety compliance.

5. Discussion

The main finding is that predictive performance and agronomic decision quality are related but not interchangeable. AgriWFD-RLNet improves held-out status and response prediction over the strongest multimodal baseline, yet the more consequential result is that its action-conditioned response model, economic context, and safety projection jointly improve offline decisions over both a conventional schedule and a soil-moisture-threshold strategy. The unconstrained SAC comparison reinforces this distinction: it approaches the resource savings of the proposed method but has substantially lower safety compliance.

The contribution should therefore be understood as an integrated state–response–policy architecture rather than a claim that any one encoder is unprecedented. Compared with perception-only multimodal systems, the framework predicts consequences of candidate actions before selecting them. Compared with unconstrained agricultural RL, it incorporates observable costs and revenue in the state and reward and uses a convex projection whose feasibility and non-expansiveness are explicit. Compared with fixed agronomic rules, it conditions decisions on images and delayed sensor context while retaining the rules as hard bounds. These couplings are the central methodological innovation.

Several limitations qualify the quantitative findings. First, all observations come from one station, one tomato season, and one alluvial-loam field. Plot-blocked validation prevents repeated samples from the same plot entering both fitting and testing, but it does not test a new year, cultivar, soil class, or climate. Second, policy outcomes other than logged conventional management are offline estimates from a learned response model. Restricting actions to the logged range and projecting them onto a safety set reduces extrapolation risk, but only randomized prospective field deployment can establish causal water, fertilizer, yield, and profit effects. Third, weekly prices and recorded costs represent the 2024 local market; the sensitivity results show that economic priorities affect the selected policy, so transfer requires updated price and cost inputs.

Operational deployment also depends on data quality and governance. Camera geometry and color normalization reduced device effects in paired acquisitions, but seasonal illumination, lens contamination, sensor drift, and communication outages require

automated health checks. The measured edge latency is compatible with a 15 min controller interval, although the current model should initially operate as decision support with operator confirmation and a rule-based fallback. A suitable validation pathway is a multi-season, multi-site trial that randomizes complete plots among conventional, threshold-based, and AgriWFD-RLNet-assisted management, reports plot-level confidence intervals, and prospectively registers yield, resource, salinity, and economic endpoints.

6. Conclusions

This paper focuses on facility agriculture and precision water–fertilizer management. In real production, it is difficult to use data from many sources together. Farmers still rely heavily on experience. The input–output efficiency of water and fertilizer management is also hard to measure. To solve these problems, this paper proposes an intelligent water–fertilizer decision model. The model uses crop images, soil sensor data, meteorological data, management records, yield and quality data, and economic input–output data. The main contribution of this study is the construction of a heterogeneous agricultural data platform for water–fertilizer decision-making. The platform combines image phenotyping, sensor time-series data, and field irrigation records. This makes it possible to analyze crop appearance, soil water and salt status, climate changes, and irrigation practices together. The model also uses time-window modeling and cross-modal association. These methods help capture the delayed effects of irrigation and fertilization on crop growth and yield formation. As a result, the decision results become more stable and easier to explain. The experimental results show that the proposed method performs well in crop status recognition, water and fertilizer demand assessment, and yield and quality prediction. Compared with traditional machine learning models and single-modal deep learning models, the proposed model achieves better results in accuracy, precision, recall, and F1-score. It can identify normal crop growth conditions. It can also detect risks such as water deficit, nutrient deficiency, and salinity accumulation. In regression tasks, the proposed model also shows strong prediction ability. It reduces the errors in yield prediction, water use efficiency prediction, fertilizer productivity prediction, and net profit prediction. It achieves better results in RMSE, MAE, and R^2. The ablation experiments further show that image data, sensor time-series data, management records, and economic indicators are all useful. Their joint modeling improves the ability of the model to represent complex agricultural production processes. Under plot-blocked evaluation, the complete framework also outperforms conventional, threshold-based, and unconstrained RL strategies in offline decision metrics, and its advantage is robust to plot-cluster bootstrap testing. These results support the proposed integration of multimodal perception, action-conditioned response prediction, economic objectives, and convex safety projection for the conditions represented by the 2024 Wuyuan experiment. They do not yet establish general effectiveness across seasons or regions, and the reported resource and profit gains require confirmation in prospective multi-site field trials.

Author Contributions: Conceptualization, Y.Z., J.N., L.L. and Y.S.; Methodology, Y.Z., J.N. and L.L.; Software, Y.Z., J.N. and L.L.; Validation, J.H. and N.M.; Formal analysis, J.H. and N.M.; Investigation, J.H. and N.M.; Resources, J.A.; Data curation, J.A.; Writing—original draft, Y.Z., J.N., L.L., J.H., J.A., N.M. and Y.S.; Visualization, J.A.; Supervision, Y.S.; Project administration, Y.S.; Funding acquisition, Y.S.; Y.Z., J.N., and L.L. contributed equally to this work. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are available on request from the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Xing, Y.; Wang, X. Precise application of water and fertilizer to crops: Challenges and opportunities. *Front. Plant Sci.* **2024**, *15*, 1444560. [CrossRef]
2. Wang, L.; Liu, S.; Lu, S.; Lv, T.; Ma, D.; Kang, G.; Xie, Y.; Sun, J.; Wang, C. Irrigation and nitrogen fertilizer increase wheat productivity by improving root-shoot population coordination. *Agric. Water Manag.* **2025**, *316*, 109589. [CrossRef]
3. Bai, Z.; Liu, L.; Kroeze, C.; Strokal, M.; Chen, X.; Yuan, Z.; Ma, L. Optimizing phosphorus fertilizer use to enhance water quality, food security and social equality. *Resour. Conserv. Recycl.* **2024**, *203*, 107400. [CrossRef]
4. Bwambale, E.; Abagale, F.K.; Anornu, G.K. Smart irrigation monitoring and control strategies for improving water use efficiency in precision agriculture: A review. *Agric. Water Manag.* **2022**, *260*, 107324. [CrossRef]
5. Zhang, Y.; Wa, S.; Liu, Y.; Zhou, X.; Sun, P.; Ma, Q. High-accuracy detection of maize leaf diseases CNN based on multi-pathway activation function module. *Remote Sens.* **2021**, *13*, 4218. [CrossRef]
6. Islam, S.; Reza, M.N.; Samsuzzaman, S.A.; Cho, Y.J.; Noh, D.H.; Chung, S.O.; Hong, S.J. Machine vision and artificial intelligence for plant growth stress detection and monitoring: A review. *Precis. Agric.* **2024**, *6*, 34. [CrossRef]
7. Zuo, Z.; Mu, J.; Li, W.; Bu, Q.; Mao, H.; Zhang, X.; Han, L.; Ni, J. Study on the detection of water status of tomato (*Solanum lycopersicum* L.) by multimodal deep learning. *Front. Plant Sci.* **2023**, *14*, 1094142. [CrossRef]
8. Kumar, A.; Gill, A.S.; Sharma, A.; Kumar, V. Enhancing the explainability and interpretability of crop yield prediction models through precision agriculture. In *Proceedings of the International Conference on Agriculture-Centric Computation*; Springer: Berlin/Heidelberg, Germany, 2024; pp. 218–226.
9. Liu, J.; Chang, F.; Ma, L.; Xie, L.; Su, H. Smart control of water-fertilizer integrated regulation system based on deep reinforcement learning. In *Proceedings of the 2025 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*; IEEE: New York, NY, USA, 2025; pp. 6085–6088.
10. Sharma, A.; Rathi, V.; Biswas, A.; Singh, A.; Rana, O. *Multimodal Artificial Intelligence in Precision Agriculture: Practices, Challenges, and Applications*; CRC Press: Boca Raton, FL, USA, 2026.
11. Wu, H.; Li, J.; Yang, Y. Intelligent decision-making method for personalized vegetable crop water and fertilizer management based on large language models. *Smart Agric.* **2025**, *7*, 11–19.
12. Li, Z.; Li, W.; Wang, J.; Zhang, J.; Wang, Z. Drip irrigation shapes the soil bacterial communities and enhances jujube yield by regulating the soil moisture content and nutrient levels. *Agric. Water Manag.* **2023**, *289*, 108563. [CrossRef]
13. Xing, Y.; Zhang, X.; Wang, X. Enhancing soil health and crop yields through water-fertilizer coupling technology. *Front. Sustain. Food Syst.* **2024**, *8*, 1494819. [CrossRef]
14. Goet, G.; Sonkar, I.; Kumar, S.; Hari Prasad, K.; Ojha, C. Effect of salinity on crop growth and soil moisture dynamics: A study with root water uptake model. *J. Hazard. Toxic Radioact. Waste* **2024**, *28*, 04024009. [CrossRef]
15. Wanniarachchi, S.; Sarukkalige, R. A review on evapotranspiration estimation in agricultural water management: Past, present, and future. *Hydrology* **2022**, *9*, 123. [CrossRef]
16. Mkuna, E.; Wale, E. Smallholder farmers' choice of irrigation systems: Empirical evidence from Kwazulu-Natal, South Africa and its implications. *Sci. Afr.* **2023**, *20*, e01688. [CrossRef]
17. Kelly, T.; Foster, T.; Schultz, D.M. Assessing the value of adapting irrigation strategies within the season. *Agric. Water Manag.* **2023**, *275*, 107986. [CrossRef]
18. Nejad, S.M.M.; Abbasi-Moghadam, D.; Sharifi, A. ConvLSTM–ViT: A deep neural network for crop yield prediction using Earth observations and remotely sensed data. *IEEE J. Sel. Top. Appl. Earth Obs. Remote Sens.* **2024**, *17*, 17489–17502. [CrossRef]
19. Teixeira, A.; Leivas, J.; Takemura, C.; Bayma, G.; Garçon, E.; Sousa, I.; Farias, F.; Silva, C. Remote sensing environmental indicators for monitoring spatial and temporal dynamics of weather and vegetation conditions: Applications for Brazilian biomes. *Environ. Monit. Assess.* **2023**, *195*, 944. [CrossRef]
20. Yin, J.; Zhao, L.; Xu, X.; Li, D.; Qiu, H.; Cao, X. Evaluation of long-term carbon sequestration of biochar in soil with biogeochemical field model. *Sci. Total Environ.* **2022**, *822*, 153576. [CrossRef]
21. Ajith, S.; Vijayakumar, S.; Elakkiya, N. 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* **2025**, *5*, 67. [CrossRef]
22. Yu, F.; Wang, M.; Xiao, J.; Zhang, Q.; Zhang, J.; Liu, X.; Ping, Y.; Luan, R. Advancements in utilizing image-analysis technology for crop-yield estimation. *Remote Sens.* **2024**, *16*, 1003. [CrossRef]
23. Yang, Z.X.; Li, Y.; Wang, R.F.; Hu, P.; Su, W.H. Deep learning in multimodal fusion for sustainable plant care: A comprehensive review. *Sustainability* **2025**, *17*, 5255. [CrossRef]
24. Wang, Z.; Xiao, S.; Wang, J.; Parab, A.; Patel, S. Reinforcement learning-based agricultural fertilization and irrigation considering N2O emissions and uncertain climate variability. *AgriEngineering* **2025**, *7*, 252. [CrossRef]

25. Lin, S.; Lei, Q.; Liu, Y.; Zhao, Y.; Su, L.; Wang, Q.; Tao, W.; Deng, M. Quantifying the impact of organic fertilizers on soil quality under varied irrigation water sources. *Water* **2023**, *15*, 3618. [CrossRef]
26. Zwecher, E.; Iceland, E.; Levy, S.R.; Hayoun, S.Y.; Gal, O.; Barel, A. Integrating deep reinforcement and supervised learning to expedite indoor mapping. In *Proceedings of the 2022 International Conference on Robotics and Automation (ICRA)*; IEEE: New York, NY, USA, 2022; pp. 10542–10548.
27. Wang, J.; He, X.; Gong, P.; Zhao, D.; Zhang, Y.; Wang, Z.; Zhang, J. Optimization of a water-saving and fertilizer-saving model for enhancing xinjiang korla fragrant pear yield, quality, and net profits under water and fertilizer coupling. *Sustainability* **2022**, *14*, 8495. [CrossRef]
28. Bai, Q. Deep reinforcement learning-driven decision support system design for precision agriculture: Optimizing farming and resource use. *Turk. J. Agric. For.* **2025**, *49*, 598–611. [CrossRef]
29. Zhao, J.; Fan, S.; Zhang, B.; Wang, A.; Zhang, L.; Zhu, Q. Research status and development trends of deep reinforcement learning in the intelligent transformation of agricultural machinery. *Agriculture* **2025**, *15*, 1223. [CrossRef]
30. Lu, Y.; Wang, Z.; Li, S.; Liu, X.; Yu, C.; Yin, Q.; Shi, Z.; Zhang, Z.; Jiang, M. Learning to optimize multi-objective alignment through dynamic reward weighting. *Trans. Assoc. Comput. Linguist.* **2026**, *14*, 1051–1073. [CrossRef]
31. Zhang, J.; Trautman, D.; Liu, Y.; Bi, C.; Chen, W.; Ou, L.; Goebel, R. Achieving the rewards of smart agriculture. *Agronomy* **2024**, *14*, 452. [CrossRef]
32. Gauci, A.; Fulton, J.; Lindsey, A.; Shearer, S.; Barker, D.; Hawkins, E. Precision of grain yield monitors for use in on-farm research strip trials. *Precis. Agric.* **2024**, *25*, 771–784. [CrossRef]
33. Alkaff, M.; Basuhail, A.; Sari, Y. Optimizing water use in maize irrigation with reinforcement learning. *Mathematics* **2025**, *13*, 595. [CrossRef]
34. Cortes, C.; Vapnik, V. Support-vector networks. *Mach. Learn.* **1995**, *20*, 273–297. [CrossRef]
35. Breiman, L. Random forests. *Mach. Learn.* **2001**, *45*, 5–32. [CrossRef]
36. Chen, T.; Guestrin, C. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM Sigkdd International Conference on Knowledge Discovery and Data Mining*; Association for Computing Machinery: New York, NY, USA, 2016; pp. 785–794.
37. Ke, G.; Meng, Q.; Finley, T.; Wang, T.; Chen, W.; Ma, W.; Ye, Q.; Liu, T.Y. Lightgbm: A highly efficient gradient boosting decision tree. In *Proceedings of the 31st Conference on Neural Information Processing Systems (NIPS 2017)*; Curran Associates Inc.: Red Hook, NY, USA, 2017; Volume 30.
38. LeCun, Y.; Bottou, L.; Bengio, Y.; Haffner, P. Gradient-based learning applied to document recognition. *Proc. IEEE* **1998**, *86*, 2278–2324. [CrossRef]
39. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*; IEEE: New York, NY, USA, 2016; pp. 770–778.
40. Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; et al. An image is worth 16 × 16 words: Transformers for image recognition at scale. *arXiv* **2020**, arXiv:2010.11929.
41. Cho, K.; Van Merriënboer, B.; Gulçehre, Ç.; Bahdanau, D.; Bougares, F.; Schwenk, H.; Bengio, Y. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*; Association for Computational Linguistics: Kerrville, TX, USA, 2014; pp. 1724–1734.
42. Tsai, Y.H.H.; Bai, S.; Liang, P.P.; Kolter, J.Z.; Morency, L.P.; Salakhutdinov, R. Multimodal transformer for unaligned multimodal language sequences. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*; Association for Computational Linguistics: Kerrville, TX, USA, 2019; pp. 6558–6569.

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