Archives of Computational Methods in Engineering
https://doi.org/10.1007/s11831-026-10640-5

RESEARCH

# Optimizing Irrigation and Fertilization Strategies for Crop Growth: A Comparative Study of Genetic Algorithm and Model Predictive Control Under Weather Uncertainty

**Carla J. Becker**[1] · **Tarek I. Zohdi**[1]

Received: 19 March 2026 / Revised: 2 May 2026 / Accepted: 5 May 2026


**Abstract**

Rising production costs and increasing weather variability pose significant challenges to agricultural profitability, motivating computational approaches to optimize resource allocation. This paper presents a coupled ordinary differential equation (ODE) crop model that captures two key physiological phenomena often neglected in optimization studies: delayed nutrient absorption and cumulative stress tracking, which we address via finite impulse response (FIR) convolution and exponential moving average (EMA) filtering. We compare two optimization approaches for irrigation and fertilizer scheduling: genetic algorithm (GA) optimization for fixed seasonal strategies and model predictive control (MPC) for adaptive daily decision-making. Applied to corn production in Iowa across 21 stochastic weather scenarios spanning normal conditions to extreme drought and heat stress, our results yield a surprising conclusion. GA optimization achieves 35% higher mean revenue than farmer best practices across all 21 scenarios ($842 vs. $626 per acre), and contrary to our initial hypothesis, the fixed GA strategies outperform adaptive MPC in all 21 scenarios. The GA trained under normal weather conditions achieves the highest mean revenue with the lowest risk (coefficient of variation 14.7%), while MPC, despite its daily adaptation, cannot match the GA's full-season optimization horizon. However, MPC offers a sustainability advantage: it achieves competitive returns while using dramatically less water and fertilizer than all other strategies. The practical conclusion is that for seasonal agricultural planning, pre-optimization for expected conditions outperforms adaptive control, and the choice of which GA strategy to deploy depends simply on whether extreme weather is anticipated.

## 1 Introduction

The agriculture sector in the United States faces significant challenges as the number of farms declines and the cost of farming continues to rise [1], . Rising production expenses for equipment, seeds, and labor, coupled with elevated interest rates and declining commodity prices, have made farming increasingly expensive. To navigate this challenging landscape, farmers are employing strategies such as cost management and operational optimization. One promising approach is to use modeling and simulation to optimize farm operations without substantial capital investment. Recent advances in computational methods have enabled sophisticated digital-twin frameworks for precision agriculture [2–5], and machine learning techniques have been applied to optimize sensor placement and resource delivery in agricultural systems [6–8].

Mathematical modeling of crop growth has a rich history, with models ranging from simple empirical relationships to complex mechanistic simulations. Logistic growth models, first proposed by Verhulst in 1838, remain widely used due to their interpretability and ability to capture resource-limited growth dynamics [9]. More sophisticated crop models such as DSSAT [10], APSIM [11], and WOFOST [12] simulate detailed physiological processes but require extensive parameterization and may be computationally expensive for optimization applications. In contrast, reduced-order models that capture essential dynamics while remaining tractable for optimization have gained attention in precision agriculture [13].

✉ Carla J. Becker
carlabecker@berkeley.edu

Tarek I. Zohdi
zohdi@berkeley.edu

[1] Department of Mechanical Engineering, University of California, 6141 Etcheverry Hall, Berkeley 94720, CA, USA

Published online: 20 May 2026

Optimization of irrigation and fertilizer application has been studied using various approaches, including linear and nonlinear programming [14], dynamic programming [15], and metaheuristic algorithms [16]. Genetic algorithms (GAs) are particularly well-suited for this domain because they can handle nonlinear, non-convex objective functions and do not require gradient information [17]. Previous work has applied GAs to irrigation scheduling [18] and fertilizer optimization [19], but these studies often use simplified plant response models that do not capture the delayed, cumulative effects of resource application.

Climate change is increasing the frequency and severity of extreme weather events, posing significant challenges to agricultural production [20]. Traditional approaches to irrigation and fertilizer scheduling rely on fixed strategies—either following agronomic best practices or strategies optimized for expected conditions—that cannot adapt when actual weather deviates from assumptions. As drought, heat waves, and other extremes become more common, there is growing interest in adaptive resource management strategies that can respond to observed conditions in real time.

Model Predictive Control (MPC) offers an alternative approach: a closed-loop control strategy that repeatedly solves an optimization problem over a finite horizon, applies only the first control action, then re-optimizes based on updated state and disturbance information [21]. This receding-horizon structure enables the controller to adapt to changing conditions while maintaining optimality over the planning horizon. The central question motivating this work is whether MPC's adaptive capabilities can outperform fixed strategies that are pre-optimized for specific conditions.

This paper makes the following contributions:

(1) A coupled ODE crop growth model that captures delayed nutrient absorption via FIR convolution and cumulative stress tracking via EMA filtering, providing a physiologically-motivated yet computationally tractable framework for optimization.
(2) A genetic algorithm approach that optimizes fixed irrigation and fertilizer strategies for specific weather conditions, achieving 35% revenue improvement over farmer best practices, on average.
(3) An MPC framework with Bayesian-optimized parameters that adapts resource allocation daily based on observed plant state and weather forecasts.
(4) A stochastic weather scenario generator that creates realistic test conditions ranging from normal variability to extreme drought and heat stress.
(5) A comprehensive comparison of four strategies—farmer best practices, GA optimized for normal conditions, GA optimized for drought, and adaptive MPC—across 21 weather scenarios, revealing that pre-optimization outperforms adaptive control for seasonal agricultural planning.

The remainder of this paper is organized as follows. Section 2 presents the coupled ODE crop model with FIR convolution and EMA filtering. Section 3 describes the stochastic weather scenario generation process. Sections 4 and 5 detail the genetic algorithm and model predictive control approaches, respectively. Section 6 presents results from the four-way strategy comparison across 21 weather scenarios. Section 7 interprets the findings and discusses limitations. Section 8 concludes with practical recommendations.

## 2 Generalized, Coupled-ODE Crop Model

This paper adopts a control systems perspective on crop growth, treating the plant as a dynamical system with internal states $\mathbf{x} = [x_h, x_A, x_N, x_c, x_P]^\top$, comprising plant height (m), leaf area ($m^2$), number of leaves (count), flower spikelet count, and fruit biomass (kg). Control inputs $\mathbf{u} = [u_w, u_f]^\top$, irrigation (in) and fertilizer (lbs), are decision variables chosen by the farmer. Disturbance inputs $\mathbf{d} = [d_S, d_T, d_R]^\top$, precipitation (in), temperature (°C), and radiation ($W/m^2$), affect the system but cannot be directly controlled. The internal dynamics follow logistic growth with time-varying growth rates and carrying capacities that depend on environmental conditions and resource availability. The output is a function of the state of the crop at the end of the season. This framing, visualized in Figure 1, naturally motivates an optimization question: how can we choose an irrigation and fertilizer strategy that optimizes crop revenue at the end of the growing season? We consider two optimal control paradigms: open-loop control via genetic algorithm, which pre-computes a fixed input trajectory, and closed-loop model predictive conrol, which adapts inputs based on state feedback.

### 2.1 Effects of Inputs on Plant Growth

Different inputs affect different aspects of plant growth, and such relationships are crop dependent. In this study, we use corn to demonstrate the framework. For corn, Tables 1 and 2 summarize these relationships based on agronomic literature [27–29]. We note that fertilizer-assisted growth encourages a taller stalk and larger leaf area via both increased carrying capacity and higher growth rate. Irrigation-assisted growth encourages larger leaf area, number of tassels (flower size), and ear size (fruit biomass) via increased carrying capacity.

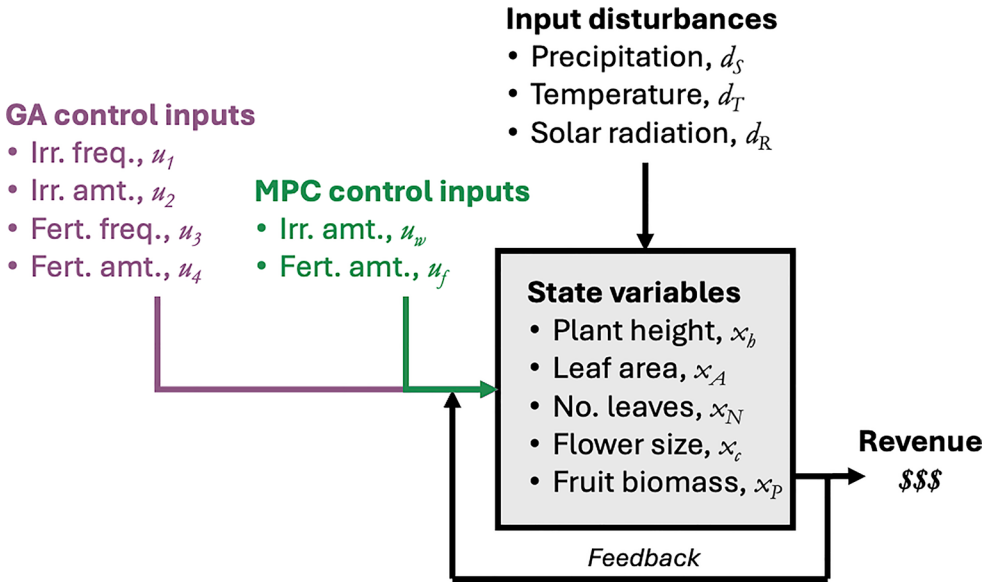


**Fig. 1** This paper treats a crop as a control system with its own internal state variables and dynamics, affected by input controls, input disturbances, and feedback. Output is a function of the state variables

**Table 1** Effects of irrigation and fertilizer on growth dynamics. "+" indicates positive effect, "∼" indicates a negligible effect. [22–25]

| State variable | Irrigation on growth rate | Fertilizer on growth rate | Irrigation on capacity | Fertilizer on capacity |
|---|---|---|---|---|
| Plant height $h$ | ∼ | + | ∼ | + |
| Leaf area $A$ | ∼ | + | + | + |
| Number of leaves $N$ | ∼ | ∼ | ∼ | ∼ |
| Flower size $c$ | ∼ | ∼ | + | ∼ |
| Fruit biomass $P$ | ∼ | ∼ | + | + |

**Table 2** Effects of temperature and solar radiation on growth dynamics. "+" indicates positive effect, "∼" indicates a negligible effect, "–" indicates a negative effect. For flower size, excess heat and radiation reduce flower development, hence negative effects. [26]

| State variable | Temp. on growth rate | Temp. on capacity | Radiation on growth rate | Radiation on capacity |
|---|---|---|---|---|
| Plant height $h$ | + | + | + | + |
| Leaf area $A$ | + | + | + | + |
| Number of leaves $N$ | ∼ | + | ∼ | + |
| Flower size $c$ | – | – | – | – |
| Fruit biomass $P$ | + | + | + | + |

## 2.2 Growth Dynamics

In our approach, each state variable follows logistic growth with time-varying growth rates and carrying capacities. These parameters are modulated by nutrient factors—dimensionless quantities in $[0, 1]$ that translate cumulative input history into growth effects, as detailed in Section 2.5. The logistic growth takes the general form:

$$\frac{dx}{dt} = \hat{a}_x(t) \cdot x(t) \left(1 - \frac{x(t)}{\hat{k}_x(t)}\right) \tag{1}$$

where $\hat{a}_x(t)$ is the effective growth rate and $\hat{k}_x(t)$ is the effective carrying capacity, both functions of the nutrient factors.

The effective parameters are proportional to their typical values by the geometric mean of the relevant nutrient factors, reflecting multiplicative rather than additive effects. This choice is motivated by the observation that growth rates compound over time, making geometric averaging appropriate [30].

**Plant height** growth rate and carrying capacity are modulated by fertilizer, temperature, and radiation through their respective nutrient factors as follows:

$$\hat{a}_h(t) = a_h(\nu_f \nu_T \nu_R)^{1/3}, \quad \hat{k}_h(t) = k_h(\nu_f \nu_T \nu_R)^{1/3} \tag{2}$$

where $a_h$ is the typical plant height growth rate and $k_h$ is the typical plant height carrying capacity, both chosen to reflect what is observed in corn development literature.

**Leaf area** responds to fertilizer, temperature, radiation, and, additionally, to water. Leaf area is also coupled to plant height, representing vegetative growth. We capture those relationships with

$$\hat{a}_A(t) = a_A(\nu_f \nu_T \nu_R)^{1/3}, \quad \hat{k}_A(t) = k_A \left(\nu_w \nu_f \nu_T \nu_R \frac{\hat{k}_h}{k_h}\right)^{1/5} \tag{3}$$

where $a_A$ and $k_A$, respectively, are the growth rate and carrying capacity both chosen to reflect typical literature values.

**Number of leaves** depends only on temperature and radiation through the carrying capacity:

$$\hat{a}_N(t) = a_N, \quad \hat{k}_N(t) = k_N(\nu_T \nu_R)^{1/2} \tag{4}$$

**Flower size** (spikelet count) exhibits inverse dependence on temperature and radiation, as excess heat and light reduce flowering:

$$\hat{a}_c(t) = a_c \left(\frac{1}{\nu_T}\frac{1}{\nu_R}\right)^{1/2}, \quad \hat{k}_c(t) = k_c \left(\nu_w \frac{1}{\nu_T}\frac{1}{\nu_R}\right)^{1/3} \tag{5}$$

**Fruit biomass** depends on all inputs and is coupled to both vegetative growth and flowering:

$$\hat{a}_P(t) = a_P \left(\frac{1}{\nu_T}\frac{1}{\nu_R}\right)^{1/2}, \quad \hat{k}_P(t) = k_P \left(\nu_w \nu_f \nu_T \nu_R \frac{\hat{k}_h}{k_h}\frac{\hat{k}_A}{k_A}\frac{\hat{k}_c}{k_c}\right)^{1/7} \tag{6}$$

The coupling terms $\hat{k}_h/k_h$, $\hat{k}_A/k_A$, and $\hat{k}_c/k_c$ encode physiological dependencies: taller plants with more leaf area can support larger fruit, while larger tassels (more spikelets) may compete with ear development. [31–34]. Table 3 in Section 2.7 details the actual values used for typical (unmodulated) growth rates and carrying capacities used in this study.

The inverse terms $1/\nu_T$ and $1/\nu_R$ appearing in Equations 5 and 6 merit clarification. The nutrient factors are defined as exponentials of cumulative divergence (Section 2.5), so they are strictly positive but can become arbitrarily small under prolonged temperature or radiation stress. As $\nu_T \to 0$ or $\nu_R \to 0$, the inverse terms grow without bound, which would produce unrealistically large growth rates and carrying capacities for flower size and fruit biomass. To keep the model in a physiologically reasonable range, the implementation clamps all modulated growth rates to $[0,\ 2a_x]$ and carrying capacities to $[x(t),\ 2k_x]$, where $a_x$ and $k_x$ are the typical values. The upper bound of twice the typical value reflects the fact that even when conditions favor accelerated flowering or fruit development, the plant's own biological machinery limits how fast it can grow. The lower bound on carrying capacities is set to the current state $x(t)$ to prevent the carrying capacity from falling below what the plant has already achieved.

For crops exhibiting distinct vegetative and reproductive phases, double logistic models can capture two successive S-curves—rapid vegetative expansion followed by a second sigmoid for fruit or grain filling. However, since our optimization objective is end-of-season yield (and hence revenue), the precise shape of intermediate growth dynamics has limited impact on the optimal input strategy. We therefore retain the single logistic formulation as a reduced-order model that balances fidelity with computational tractability.

## 2.3 Delayed Absorption via FIR Convolution

Plants do not immediately process applied nutrients; instead, there is a physiologically-mediated delay between application and utilization. We model this delayed absorption using finite impulse response (FIR) convolution with Gaussian kernels.

If cumulative nutrient uptake follows a sigmoid trajectory—with slow initial uptake due to transport lag, rapid increase once metabolic pathways activate, and eventual saturation—then instantaneous nutrient absorption follows a bell-shaped curve. A Gaussian kernel is the least assumptive choice for a bell curve, requiring only two parameters: the temporal spread $\sigma$ (characterizing absorption duration) and the peak delay $\mu$.

Given only the temporal spread $\sigma$ for each nutrient type, we determine $\mu$ (hours after application where peak absorption occurs) such that 95% of the kernel mass lies within $[0, 2\mu]$. This requires solving

$$\operatorname{erf}\left(\frac{\mu}{\sigma\sqrt{2}}\right) = 0.95, \tag{7}$$

which yields $\mu \approx 1.96\sigma$. The Gaussian FIR kernel is then

$$G[k] = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left\{-\frac{1}{2}\frac{(k-\mu)^2}{\sigma^2}\right\}. \tag{8}$$

for $k \geq 0$ hours.

The FIR horizon $L^*$ is chosen as the minimum length capturing 95% of the kernel mass:

**Table 3** Model parameters derived from values in table . Initial conditions are set to$k_x/K$where$K \approx 2900$is the season length in hours (about 121 days)

| State | Growth rate | Carrying capacity | Initial condition | Units |
|---|---|---|---|---|
| Height$h$ | $a_h = 0.010\text{hr}^{-1}$ | $k_h = 3.0$ | $h_0 = 1.0 \times 10^{-3}$ | m |
| Leaf area$A$ | $a_A = 0.0105\text{hr}^{-1}$ | $k_A = 0.65$ | $A_0 = 2.2 \times 10^{-4}$ | $\text{m}^2$ |
| Leaves$N$ | $a_N = 0.011\text{hr}^{-1}$ | $k_N = 20$ | $N_0 = 6.9 \times 10^{-3}$ | count |
| Spikelets$c$ | $a_c = 0.010\text{hr}^{-1}$ | $k_c = 1000$ | $c_0 = 3.5 \times 10^{-1}$ | count |
| Fruit$P$ | $a_P = 0.005\text{hr}^{-1}$ | $k_P = 0.25$ | $P_0 = 8.6 \times 10^{-5}$ | kg |

$$L^* = \min_L \left\{ L : \frac{\sum_{k=0}^{L-1} G[k]}{\sum_{k=0}^{K-1} G[k]} \geq 0.95 \right\} \tag{9}$$

where $K$ is the simulation length in hours. Figure 2 panel 1 visualizes this transformation.

Absorbed nutrient signals are obtained from applied nutrients by convolving the applied signal with the corresponding Gaussian kernel. Different nutrients have different metabolic timescales, so we use $\sigma_w = 30$ hours for water (rapid uptake), $\sigma_f = 300$ hours for fertilizer (slow uptake reflecting root absorption dynamics), and $\sigma_T = \sigma_R = 30$ hours for temperature and radiation (immediate physiological effects with short memory). In future additions to the model, rate of water uptake from soil could be incorporated here.

## 2.4 Cumulative Stress Tracking via EMA Filtering

While FIR convolution captures delayed absorption, plants also accumulate stress from sustained deviations from optimal conditions. Intuitively, a plant "expects" to receive nutrients at a steady rate throughout the season; prolonged under- or over-supply relative to this expectation causes stress. We formalize this by comparing cumulative absorbed nutrients to a linear reference trajectory based on seasonal requirements from agronomic literature. To capture the delayed effects from the stress, we then additionally apply an exponential moving average (EMA) filter, which is equivalent to a first-order infinite impulse response (IIR) filter.

The EMA filter with memory parameter $\beta \in [0, 1)$ has the recursive form:

$$y[k] = (1 - \beta)x[k] + \beta y[k - 1] \tag{10}$$

where larger $\beta$ values correspond to longer memory (slower response to changes). This formulation preserves constant signals ($x[k] = c$ implies $y[k] \to c$) while smoothing transient fluctuations.

## 2.5 Nutrient Factor Calculation

We now describe the complete transformation pipeline that converts raw input signals into nutrient factors $\nu \in [0, 1]$ that modulate plant growth. For a nutrient signal $f$ (could represent water, fertilizer, temperature, or solar radiation):

**Step 1: Delayed absorption.** Convolve the raw signal $f[k]$ with the Gaussian FIR kernel:

$$\bar{f}[k] = \sum_{n=0}^{L-1} G[n] f[k - n] \tag{11}$$

**Step 2: Cumulative absorption.** Compute the running sum of absorbed nutrient:

$$F[k] = \sum_{n=0}^{k} \bar{f}[n] \tag{12}$$

**Step 3: Instantaneous anomaly.** Compare actual cumulative absorption to expected levels:

$$\delta[k] = \left| \frac{k \cdot f_{\text{typ}} - F[k]}{k \cdot f_{\text{typ}} + \epsilon} \right| \tag{13}$$

where $f_{\text{typ}}$ is the typical hourly nutrient level the plant "expects" and $\epsilon$ is a small constant preventing division by zero.

**Step 4: Cumulative divergence.** Apply EMA smoothing to track sustained anomalies:

$$\Delta[k] = \beta_\Delta \Delta[k - 1] + (1 - \beta_\Delta)\delta[k] \tag{14}$$

where $\beta_\Delta = 0.95$ provides long memory.

**Step 5: Nutrient factor.** Convert divergence to a nutrient factor via exponential decay with additional EMA smoothing:

$$\nu[k] = \beta_\nu \nu_f[k - 1] + (1 - \beta_\nu) \exp\{-\alpha \Delta[k]\} \tag{15}$$

where $\alpha = 3$ ensures $\nu \approx 0.05$ when $\Delta = 1$ (complete divergence from expected levels), and $\beta_\nu = 0.05$.

The nutrient factor $\nu[k]$ equals 1 when nutrient application perfectly matches expected levels and decays toward 0 under sustained over- or under-application. This captures the intuition that plants are resilient to brief deviations but suffer cumulative damage from prolonged stress. Figure 2 illustrates the FIR convolution and EMA smoothing operations that constitute the metabolic transformation pipeline.

## 2.6 Simulation

The logistic ODEs (1) admit a closed-form solution, enabling exact time-stepping without numerical integration error. Given state $x(t)$ at time $t$, the state at $t + \Delta t$ is:

$$x(t + \Delta t) = \frac{\hat{k}_x(t)}{1 + \left( \frac{\hat{k}_x(t)}{x(t)} - 1 \right) \exp(-\hat{a}_x(t)\Delta t)} \tag{16}$$

where $\hat{a}_x(t)$ and $\hat{k}_x(t)$ are treated as constant over the time step. This closed-form approach is more accurate than forward Euler integration and avoids instability issues that can arise with explicit methods at larger time steps.

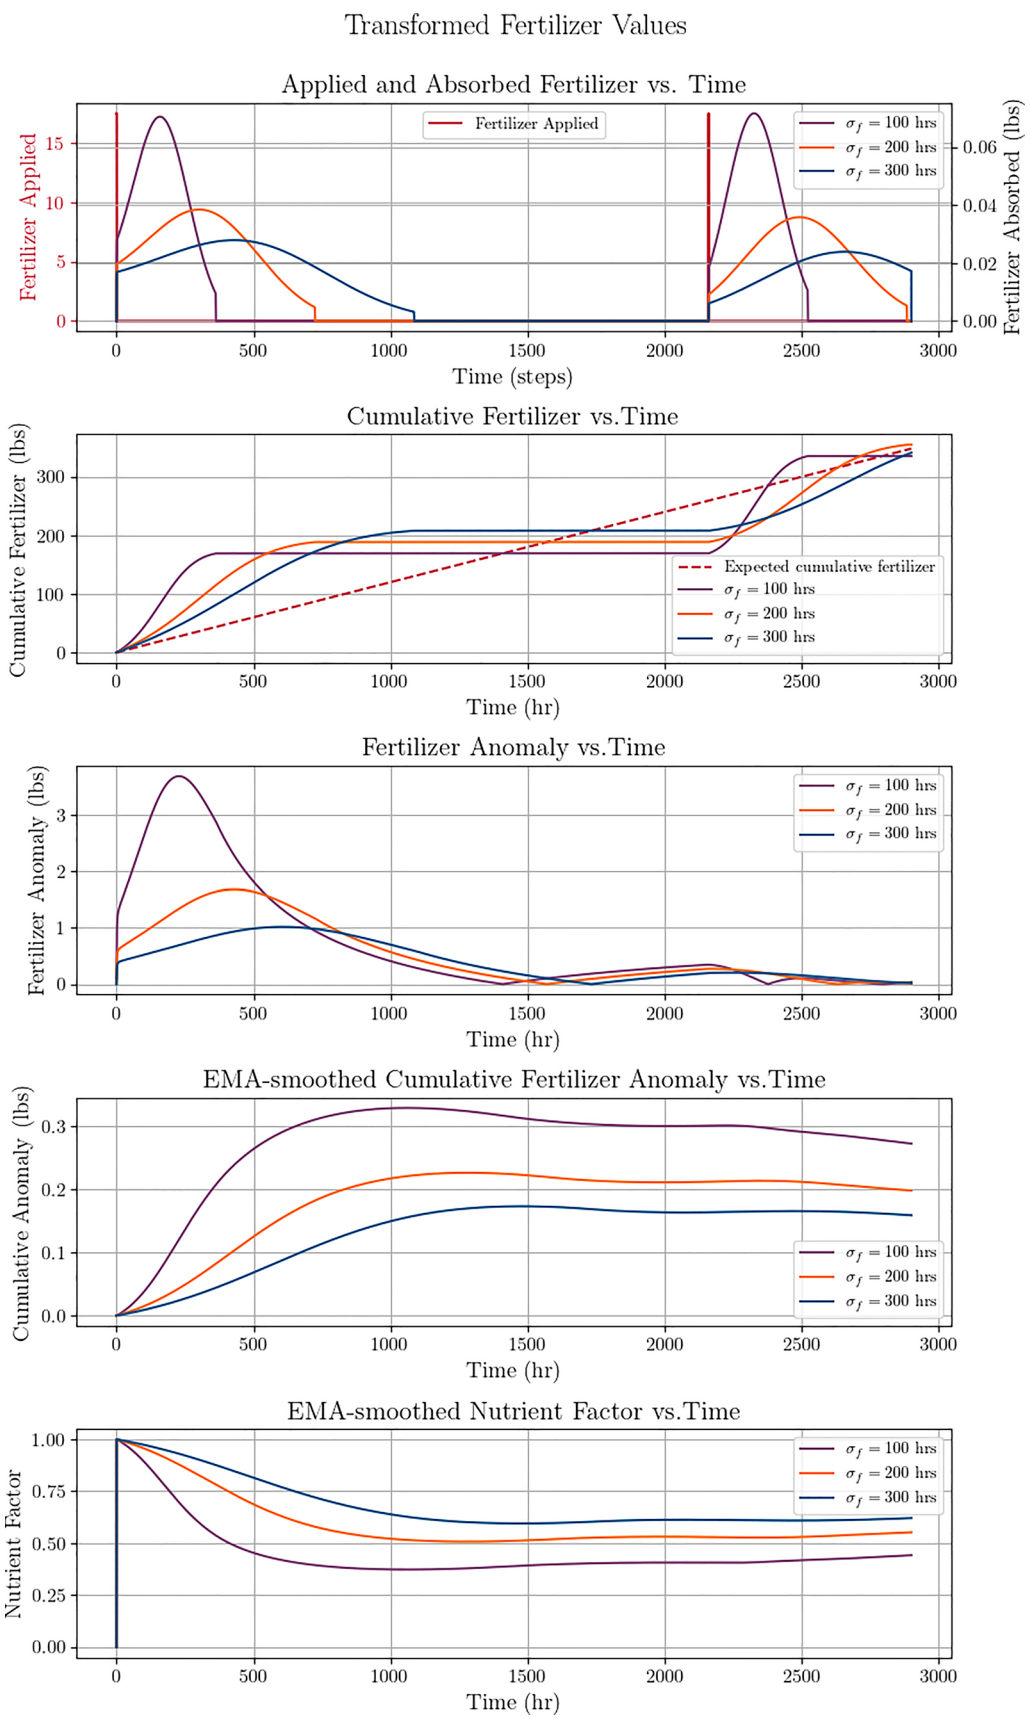
Transformed Fertilizer Values
Applied and Absorbed Fertilizer vs. Time
Fertilizer Applied
Fertilizer Absorbed (lbs)
Time (steps)
Cumulative Fertilizer vs.Time
Cumulative Fertilizer (lbs)
Expected cumulative fertilizer
Time (hr)
Fertilizer Anomaly vs.Time
Fertilizer Anomaly (lbs)
EMA-smoothed Cumulative Fertilizer Anomaly vs.Time
Cumulative Anomaly (lbs)
EMA-smoothed Nutrient Factor vs.Time
Nutrient Factor
$\sigma_f = 100$ hrs
$\sigma_f = 200$ hrs
$\sigma_f = 300$ hrs

**Fig. 2** Illustration of the metabolic transformation pipeline using fertilizer as an example. Applied fertilizer is treated as a finite impulse, while absorbed fertilizer is modeled by convolving the applied signal with a Gaussian FIR kernel. Panel 1 shows how larger temporal spreads, which define the kernel, result in slower absorption. Panel 2 shows the cumulative fertilizer absorbed over time. Panel 3 shows the fertilizer anomaly, which compares the actual cumulative fertilizer absorbed to the expected level. Panel 4 shows how EMA filtering smooths the effects of cumulative anomalies. Panel 5 shows the nutrient factors that result from the preceding transformations

We simulate the growing season at hourly resolution ($\Delta t = 1$ hour), yielding approximately 2900 time steps for a typical corn season (late April to early October). At each step, we: (1) update the nutrient factors based on cumulative inputs and divergences, (2) compute effective growth rates and carrying capacities, and (3) advance each state variable using Equation 16.

### 2.7 Case Study: Corn in Iowa

We demonstrate the framework using corn, the most widely planted crop in the United States with over 90 million acres harvested annually [35]. The case study uses historical weather data from Fairfax, Iowa (41.76°N, 91.87°W), a representative location in the Corn Belt (USDA climate zones 4b–5b).

A single simulation covers a typical growing season from late April to early October (approximately 2900 hours). Environmental inputs are:

- **Temperature and radiation:** Hourly data from NSRDB for Fairfax, IA. Mean temperature is 22.8°C; mean solar radiation is 580 W/m$^2$ [36].
- **Precipitation:** Daily data from NOAA, interpolated to hourly resolution [37].
- **Typical nutrient expectations:** Based on agronomic recommendations [38], the model expects 28 inches of water and 355 lbs of NPK fertilizer over the season ($u_{w,\text{typ}} \approx 0.01$ in/hr, $u_{f,\text{typ}} \approx 0.12$ lb/hr).

Table 4 summarizes expected corn development timelines used to calibrate model parameters. The baseline growth rates and carrying capacities are crop-specific parameters that can be estimated from field data or literature values. For corn, we use the values in Table 3, calibrated to match typical development timelines where the corn reaches full vegetative size around 65–70 days after sowing and grain fill completes around 125 days [39]. For the GA cost function weights to be specific to corn, we use the values in Table 5, which were also taken from the literature.

## 3 Stochastic Weather Scenario Generation

To evaluate optimizer robustness, we generate a suite of stochastic weather scenarios from historical data. Each scenario applies perturbations representing different climate conditions, from normal variability to extreme events. A single scenario can comprise multiple weather events. In this study, the possible types of weather events are droughts/downpours $\mathcal{D}$ and heat waves/cold snaps $\mathcal{H}$. The process of adding perturbations to the historical baseline data in order to obtain synthetic, stochastic, and at times extreme, weather is as follows: 1) scale the precipitation and radiation data by fixed scaling factors, offset the temperature data by a fixed value, 2) add white noise to each time step, independent from other time steps, 3) inject all events of type $\mathcal{D}$, 4) inject all events of type $\mathcal{H}$, and 5) check that synthetic data falls within reasonable physical bounds. Details of those transformations follow below.

### 3.1 Defining a Stochastic Weather Scenario

Let the historical hourly time series for the input disturbances be

$$d_S^{\text{hist}}[k] \geq 0 \tag{16}$$

$$-20 \leq d_T^{\text{hist}}[k] \leq 50 \tag{17}$$

$$d_R^{\text{hist}}[k] \geq 0 \tag{19}$$

for $k = 0, \ldots, K-1$ where $K$ is the length of the growing season in hours, respectively representing precipitation (inches), temperature (°C), and solar radiation (W/m$^2$). Then, let a stochastic weather scenario be defined by the parameters

$$\mathcal{W} = (s_S, s_T, s_R, \eta, D, H) \tag{20}$$

where

- $s_S \in [0, 2]$ is the precipitation adjustment (scaling factor)
- $s_T \in \mathbb{R}$ is the temperature adjustment (offset)
- $s_R \in [0, 2]$ is the radiation adjustment (scaling factor)
- $\eta \in [0, 1]$ is the relative noise level (a fraction of the historical standard deviation)
- $D$ is a set of drought or downpour events, $\mathcal{D}$
- $H$ is a set of heatwave or cold snap events, $\mathcal{H}$

Each event type ($\mathcal{D}$ or $\mathcal{H}$) is defined by three parameters:

$$[k_0, \kappa, \iota] \tag{21}$$

**Table 4** Corn development timeline and typical final values from agronomic literature. [39–50]

| State variable | Days to maturity | Hours to maturity | Typical final value |
|---|---|---|---|
| Plant height $h$ | 65–70 | 1560–1680 | 2.7–3.7 m |
| Leaf area $A$ | 55–65 | 1320–1560 | 0.6–0.7 $m^2$ |
| Number of leaves $N$ | 65 | 1560 | 18–20 |
| Spikelets $c$ | 65–70 | 1560–1680 | ∼1000 |
| Fruit biomass $P$ | 125 | 3000 | 0.15–0.36 kg |

**Table 5** Economic weights for the GA objective function. The fruit biomass weight accounts for approximately 28,350 plants per acre at \$0.157/kg. [51–65]

| Parameter | Value | Derivation |
|---|---|---|
| $w_w$ | \$2.00/inch | Typical irrigation cost |
| $w_f$ | \$0.61/lb | Weighted NPK cost |
| $w_h$ | \$35/m | Silage value proxy |
| $w_A$ | \$215/$m^2$ | Silage value proxy |
| $w_P$ | \$4,450/kg | \$4/bushel × plant density |

where

- $k_0$ is the hour when the event begins
- $\kappa$ is the duration of the event in number of hours
- $\iota \in [0, 1]$ is the intensity, with 1 being the most intense

### 3.1.1 Global Scaling and Offset

We initialize the synthetic environmental disturbance time series by applying the global adjustments $(s_S, s_T, s_R)$ for $k = 0, \ldots, K-1$ as below

$$\begin{aligned} d_S^{(1)}[k] &= s_S d_S^{\text{hist}}[k] \quad \text{(scaling)} \\ d_T^{(1)}[k] &= d_T^{\text{hist}}[k] + s_T \quad \text{(offset)} \\ d_R^{(1)}[k] &= s_R d_R^{\text{hist}}[k] \quad \text{(scaling)} \end{aligned} \tag{22}$$

### 3.1.2 Add White Noise

If $\eta > 0$, we draw independent white noise sequences

$$\varepsilon_T[k] \sim \mathcal{N}(0, (\eta\sigma_T)^2) \tag{23}$$

$$\varepsilon_R[k] \sim \mathcal{N}(0, (\eta\sigma_R)^2) \tag{24}$$

where $\sigma_T$ and $\sigma_R$ are the standard deviations of the historical temperature and radiation time series, respectively.

We then smooth the white noise with a moving average over $m_T = 24$ hours for temperature and $m_R = 12$ hours for radiation, as we will only apply the noise to the radiation during the day time (when it is not as close to zero).

$$\bar{\varepsilon}_T[k] = \frac{1}{m_T} \sum_{n=0}^{m_T-1} \varepsilon_T[k-n] \tag{25}$$

$$\bar{\varepsilon}_R[k] = \frac{1}{m_R} \sum_{n=0}^{m_R-1} \varepsilon_R[k-n] \tag{26}$$

The smoothed noise is then added to the signals from section 3.1.1.

$$d_T^{(2)}[k] = d_T^{(1)}[k] + \bar{\varepsilon}_T[k]$$

$$d_R^{(2)}[k] = d_R^{(1)}[k] + \mathcal{M}_{\text{day}} \bar{\varepsilon}_R[k] \tag{28}$$

where

$$\mathcal{M}_{\text{day}} = \{k | d_R^{\text{hist}}[k] > d_R^{\text{day}}\} \tag{29}$$

and we choose the threshold $d_R^{\text{day}} = 10$ W/$m^2$. As our historical data for precipitation is daily, not hourly, precipitation remains unchanged in this step, so

$$d_S^{(2)}[k] = d_S^{(1)}[k] \tag{30}$$

If $\eta = 0$, then we let $(d_S^{(2)}, d_T^{(2)}, d_R^{(2)}) = (d_S^{(1)}, d_T^{(1)}, d_R^{(1)})$.

### 3.1.3 Drought/Downpour Injection

A drought or downpour event is

$$\mathcal{D} = [k_0, \kappa, \iota] \tag{31}$$

and drought events only affect synthetic precipitation. Each event takes place for a limited amount of time, so only some time steps are affected. For each event, we define the affected hourly index set as

$$\mathcal{I} = \{k_0, k_0 + 1, \ldots, \min(k_0 + \kappa - 1, K - 1)\}, \tag{32}$$

so either $\kappa$ steps or the end of the season, whichever comes first, and then apply the intensity scaling to those indices with

$$d_S^{(3)}[k] \leftarrow (1 - \iota) d_S^{(2)}[k] \quad \text{for all } k \in \mathcal{I} \tag{33}$$

Temperature and radiation are unchanged in this step, so

$$d_T^{(3)}[k] = d_T^{(2)}[k] \quad \text{and} \quad d_R^{(3)}[k] = d_R^{(2)}[k] \tag{34}$$

#### 3.1.4 Heat Wave/Cold Snap Injection

A heat wave or cold snap event affects only synthetic temperature and is defined by

$$\mathcal{H} = [k_0, \kappa, \iota] \tag{35}$$

where $k_0$ and $\kappa$ have the same meanings they did for the drought event, but now $\iota$ is the peak/valley temperature difference (an offset rather than a scaling factor). We then construct a ramp-up, hold, and ramp-down for the event. Let

$$\kappa_{\text{ramp}} = \max(1, 0.1\kappa) \tag{36}$$

i.e. either one tenth of the heat wave duration or at least one hour. Then, let

$$\kappa_{\text{hold}} = \kappa - 2\kappa_{\text{ramp}} \tag{37}$$

to account for the ramp-up and ramp-down.

We can then define

- ramp-up: $w_{\text{up}}[k] = \frac{k}{\kappa_{\text{ramp}}-1}$ for $k = 0, \ldots, \kappa_{\text{ramp}} - 1$
- ramp-hold: $w_{\text{hold}}[k] = 1$
- ramp-down: $w_{\text{down}}[k] = \frac{k}{\kappa_{\text{ramp}}-1}$ for $k = 0, \ldots, \kappa_{\text{ramp}} - 1$

and concatenate to form the heat wave window

$$w_{\text{heat}} = [w_{\text{up}}, w_{\text{hold}}, w_{\text{down}}] \tag{38}$$

Applying the event to the temperature time series, we obtain

$$d_T^{(4)}[k] \leftarrow d_T^{(3)}[k] \pm \iota w_{\text{heat}}[k - k_0] \tag{39}$$

for $k = k_0, k_0 + 1, \ldots, \min(k_0 + \kappa - 1, K - 1)$, where $\iota$ is added in the case of a heat wave and subtracted in the case of a cold snap. Precipitation and radiation are unchanged in this step, so

$$d_S^{(4)}[k] = d_S^{(3)}[k] \quad \text{and} \quad d_R^{(4)}[k] = d_R^{(3)}[k] \tag{40}$$

#### 3.1.5 Clipping to Physical Bounds

Finally, we check to ensure that the transformations above have not violated the bounds on the input disturbances specified in equations 19. If they have, we clip the values as below

$$d_S^{\text{syn}}[k] = \max(0, d_S^{(4)}[k]) \tag{41}$$

$$d_T^{\text{syn}}[k] = \min(50, \max(-20, d_T^{(4)}[k])) \tag{42}$$

$$d_R^{\text{syn}}[k] = \max(0, d_R^{(4)}[k]) \tag{43}$$

for all $k \in \{0, \ldots, K - 1\}$.

### 3.2 Extremity Index

To systematically compare strategy performance across diverse conditions, we define an extremity index $\mathcal{E}$ that quantifies how far a scenario deviates from baseline conditions. This index serves two purposes: (1) analyzing whether different optimization methods perform better at different extremity levels, and (2) ensuring that MPC parameter tuning covers a representative range of conditions. The index aggregates contributions from precipitation anomalies, temperature anomalies, and discrete weather events, with each component scaled so that a value of 1 represents a meaningful agronomic threshold—i.e., a deviation large enough to require management intervention or cause measurable yield impact.

**Precipitation extremity.** The precipitation scaling factor $s_S$ determines seasonal water availability, with $s_S = 1$ representing baseline conditions. Deviations in either direction—drought ($s_S < 1$) or downpour ($s_S > 1$)—contribute to extremity:

$$\mathcal{E}_{\text{precip}} = 2|1 - s_S| \tag{44}$$

The factor of 2 is chosen so that a 50% reduction in precipitation ($s_S = 0.5$) yields $\mathcal{E}_{\text{precip}} = 1$. This threshold reflects agronomic practice: a 50% precipitation deficit during the growing season typically triggers supplemental irrigation and is considered a moderate drought requiring active management [66].

**Temperature extremity.** Seasonal temperature offsets contribute to extremity based on their magnitude:

$$\mathcal{E}_{\text{temp}} = \frac{|s_T|}{5} \tag{45}$$

The $\pm 5°\text{C}$ threshold reflects corn physiology: sustained temperatures above 35°C or below 10°C significantly reduce photosynthesis and pollen viability [67]. Since mean baseline temperature is approximately 23°C, a 5°C offset moves conditions into ranges where yield impacts become measurable.

**Drought event extremity.** Discrete drought events contribute based on their duration and intensity:

$$\mathcal{E}_{\text{drought}} = \sum_{(k_0, \kappa, \iota) \in \mathcal{D}} \frac{\kappa}{500} \iota \tag{46}$$

**Table 6** Revenue ($/acre) for all 21 weather scenarios sorted by extremity index. Cell shading indicates revenue level (darker green = higher revenue). GA (normal) dominates in normal/moderate conditions; GA (drought) excels in drought scenarios; all strategies struggle in extreme combined scenarios

| Scenario | Extremity | Farmer | GA (Drought) | GA (Normal) | MPC |
|---|---|---|---|---|---|
| normal 1 | 0.00 | 587 | 784 | 914 | 850 |
| moderate variable | 0.00 | 586 | 784 | 913 | 847 |
| baseline | 0.00 | 587 | 784 | 915 | 848 |
| normal 2 | 0.20 | 622 | 819 | 945 | 877 |
| normal 3 | 0.20 | 550 | 739 | 883 | 815 |
| normal 4 | 0.40 | 652 | 842 | 973 | 886 |
| normal 5 | 0.40 | 515 | 696 | 845 | 774 |
| moderate cool | 0.60 | 485 | 652 | 804 | 736 |
| moderate wet | 0.70 | 453 | 629 | 762 | 716 |
| moderate warm | 0.70 | 614 | 791 | 925 | 846 |
| moderate dry | 0.80 | 721 | 913 | 945 | 796 |
| wet year | 1.10 | 385 | 551 | 669 | 633 |
| cool wet | 1.10 | 393 | 550 | 680 | 641 |
| mild drought | 1.30 | 853 | 997 | 843 | 717 |
| late drought | 1.52 | 833 | 991 | 869 | 700 |
| early drought | 1.64 | 817 | 993 | 855 | 727 |
| heat stress | 1.70 | 669 | 855 | 948 | 795 |
| summer drought | 1.70 | 752 | 944 | 888 | 746 |
| multiple heatwaves | 2.26 | 651 | 837 | 959 | 829 |
| extreme drought heat | 4.97 | 768 | 836 | 609 | 518 |
| worst case | 8.38 | 659 | 734 | 527 | 445 |

The 500-hour (approximately 21-day) normalization reflects the timescale over which soil moisture depletion becomes critical. Corn can typically tolerate 7–10 days without rainfall before showing stress symptoms; by three weeks, yield losses become significant without irrigation intervention [66]. A full-intensity drought lasting 21 days thus represents the threshold for serious crop impact.

**Heat wave and cold snap extremity.** Temperature events contribute based on both duration and intensity:

$$\mathcal{E}_{\text{heat}} = \sum_{(k_0,\kappa,\iota)\in\mathcal{H}} \frac{\kappa}{200}\frac{\iota}{5} \tag{47}$$

The 200-hour (approximately 8-day) duration threshold reflects the timescale over which sustained heat stress affects corn reproductive development. During tasseling and silking, even brief heat waves can reduce kernel set, but cumulative effects become severe after about one week

**Table 7** Weather scenario categories with extremity index ranges

| Category | Count | Extremity Range |
|---|---|---|
| Normal | 5 | 0.03–0.18 |
| Moderate | 5 | 0.35–0.90 |
| Drought | 4 | 1.20–2.45 |
| Wet/Cool | 2 | 0.95–1.10 |
| Heat Stress | 2 | 2.10–3.50 |
| Extreme | 2 | 5.60–8.38 |

[67]. The 5°C intensity threshold again corresponds to the deviation that moves temperatures into physiologically stressful ranges.

**Aggregate extremity.** The total extremity index is the sum of all components:

$$\mathcal{E} = \mathcal{E}_{\text{precip}} + \mathcal{E}_{\text{temp}} + \mathcal{E}_{\text{drought}} + \mathcal{E}_{\text{heat}} \tag{48}$$

The additive form assumes that stressors contribute independently to overall extremity. In practice, compound events—such as simultaneous drought and heat stress—can interact synergistically, producing impacts greater than the sum of their individual effects. The extremity index should therefore be interpreted as a lower bound on effective stress severity for compound-event scenarios (e.g., "extreme drought heat" and "worst case" in Table 6), and revenue declines in these scenarios may be steeper than the index alone would suggest.

### 3.3 Weather Scenario Suite

We evaluate strategies across 21 stochastic weather scenarios generated from the baseline historical data. Table 7 summarizes the scenario categories.

## 4 Genetic Algorithm Optimization

Given the nonlinear, delay-affected dynamics of the crop model, gradient-based optimization is challenging. The delayed effects of inputs create a non-convex landscape with potentially many local optima. We therefore employ a genetic algorithm (GA), a population-based metaheuristic inspired by natural selection, that can effectively explore complex search spaces [17].

### 4.1 Algorithm Description

The GA maintains a population of $M$ candidate solutions and iteratively improves them through selection, crossover, and mutation over $G$ generations. Figure 3 presents the complete procedure.

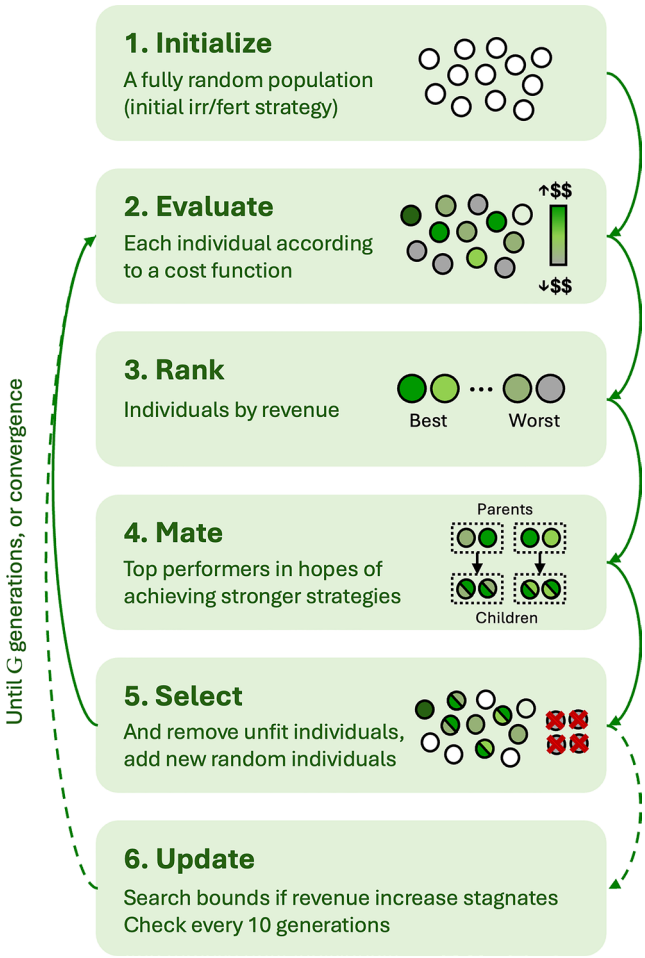


**Fig. 3** Genetic algorithm visualized. Initialization, followed by evaluation, ranking, mating, and selection, repeated for 100 generations with an added check for stagnation every 10 generations

**Selection.** After each generation, population members are ranked by cost. In our application, we use negative revenue as the cost, as this is a minimization problem and we seek to maximize revenue. The top $P$ members survive as "parents" for the next generation. This selection ensures the best solutions are never lost.

**Crossover.** New "children" $C$ are created by blending two parent solutions. For each child, we randomly select two parents, $\mathbf{u}^{(a)}$ and $\mathbf{u}^{(b)}$, and compute a weighted average:

$$\mathbf{u}^{(\text{child})} = \phi \cdot \mathbf{u}^{(a)} + (1-\phi) \cdot \mathbf{u}^{(b)} \tag{49}$$

where $\phi \sim \text{Uniform}(0,1)$ under normal operation. This crossover can produce children anywhere along the line segment connecting the parents, enabling smooth exploration of the search space.

**Mutation and Diversity.** To maintain population diversity and escape local optima, the remaining $M - P - C$ population slots are filled with randomly generated solutions. Additionally, if the best cost stagnates (changes by less than \$0.01) for 10 consecutive generations, we switch to aggressive crossover with $\phi \sim \text{Uniform}(-1,2)$. This allows children to lie outside the convex hull of their parents, promoting exploration of new regions.

**Default Parameters.** We use $M = 128$ members, $P = 16$ parents, $C = 16$ children, and $G = 100$ generations. The large population relative to generations ensures diversity for exploration while preventing premature convergence.

## 4.2 Decision Variables and Search Space: GA

Each candidate solution encodes a complete irrigation and fertilization strategy as a four-dimensional vector:

$$\mathbf{u} = \begin{bmatrix} u_1 \\ u_2 \\ u_3 \\ u_4 \end{bmatrix} = \begin{bmatrix} \text{irrigation frequency (hours)} \\ \text{irrigation amount (inches)} \\ \text{fertilizer frequency (hours)} \\ \text{fertilizer amount (lbs)} \end{bmatrix} \tag{50}$$

The frequencies specify application intervals: $u_1 = 168$ means irrigate every 168 hours (weekly). The amounts specify the quantity applied at each event. This parameterization assumes regular, periodic application—a simplification that captures common agricultural practice while keeping the search space tractable.

The GA searches over the following bounds:

- Irrigation frequency: 100–700 hours (4–29 days between applications)
- Irrigation amount: 0.5–5.0 inches per application
- Fertilizer frequency: 700–2900 hours (29–121 days, i.e., 1–4 applications per season)
- Fertilizer amount: 100–500 lbs per application

## 4.3 Genetic Algorithm Objective Function

The objective is to maximize net revenue, defined as crop value minus input costs.

$$\text{Revenue}(\mathbf{u}) = \text{Crop Value}(\mathbf{x}) - \text{Input Costs}(\mathbf{u}) \tag{51}$$

The crop value depends on final plant state at harvest:

$$\text{Crop Value}(\mathbf{x}) = w_h x_h[K] + w_A x_A[K] + w_P x_P[K] \tag{52}$$

where $K$ is the final time step and $w_h$, $w_A$, $w_P$ are economic weights (dollars per unit) for height, leaf area, and fruit biomass respectively.

The input costs accumulate over the season:

$$\text{Input Costs}(\mathbf{u}) = w_w \sum_{k=0}^{K} u_2[k] + w_f \sum_{k=0}^{K} u_4[k] \tag{53}$$

where $w_w$ and $w_f$ are costs per unit of irrigation and fertilizer.

For corn, the economic weights are derived from market prices and typical yields (Table 5). The fruit biomass weight dominates, reflecting that grain yield is the primary economic output.

## 5 Model Predictive Control

### 5.1 Optimization and Decision Epochs

Model Predictive Control (MPC) frames resource allocation as a sequence of optimization problems. At each *decision epoch*—in our case, the start of each day—the controller observes the current system state and solves an optimization problem to determine the best actions over a finite planning horizon. The solution specifies how much irrigation and fertilizer to apply, balancing immediate costs against long-term crop development.

This approach differs fundamentally from the genetic algorithm method, which optimizes a single fixed strategy over the entire growing season. MPC instead solves a smaller optimization problem at each decision epoch, using updated state and forecast information to adapt to changing conditions.

### 5.2 Constrained Finite-Time Optimal Control

At each decision epoch, model predictive control solves a constrained finite-time optimal control (CFTOC) problem over a planning horizon of $K_d$ days. Discrete-time optimal control is concerned with choosing an optimal input sequence over the horizon $K_d$

$$\mathcal{U}_{0\to K_d} = \{\mathbf{u}[k]\} \quad \text{for} \quad k = 0, \ldots, K_d - 1 \tag{54}$$

with respect to some objective function over a finite or infinite time horizon in order to apply it to a system with a given initial state $\mathbf{x}[0]$. The objective function is often defined as a sum of stage costs $q(\mathbf{x}[k], \mathbf{u}[k])$ and when the horizon has finite length, a terminal cost $p(\mathbf{x}[K_d])$. That is

$$J_{0\to K_d}(\mathbf{x}[0], \mathcal{U}_{0\to K_d}) = p(\mathbf{x}[K_d]) + \sum_{n=0}^{K_d-1} q(\mathbf{x}[n], \mathbf{u}[n]) \tag{55}$$

where the states $\mathcal{X}_{0\to K_d} = \{\mathbf{x}[k]\} \quad \text{for} \quad k = 0, \ldots, K_d - 1$ must satisfy the initial condition and system dynamics

$$\begin{cases} \mathbf{x}[0] = \mathbf{x}_0 \\ \mathbf{x}[k+1] = g(\mathbf{x}[k], \mathbf{u}[k]) \quad \text{for } k = 0, \ldots, K_d - 1 \end{cases} \tag{56}$$

and there may be other state or input constraints formulated as inequalities

$$h(\mathbf{x}[k], \mathbf{u}[k]) \leq 0 \quad \text{for} \quad k = 0, \ldots, K_d - 1 \tag{57}$$

In the finite horizon case, there may also be a terminal constraint requiring the final state to lie in some terminal set $\mathbf{x}[K] \in \mathcal{X}_{\text{final}}$.

In our specific case, we construct the problem as a minimization

$$J^*_{0\to K_d}(\mathbf{x}[0]) = \min_{\mathcal{U}_{0\to K_d}} J_{0\to K_d}(\mathbf{x}[0], \mathcal{U}_{0\to K_d}) \quad \text{s.t.} \\ \begin{cases} \mathbf{x}[0] = \mathbf{x}_0 \\ \mathbf{x}[k+1] = g(\mathbf{x}[k], \mathbf{u}[k]) \quad \text{for } k = 0, \ldots, K_d - 1 \\ \mathbf{x} \in \mathbb{R}^+ \\ \mathbf{u} \in \mathcal{U} \end{cases} \tag{58}$$

where the stage cost penalizes input usage and nutrient anomalies:

$$q(\mathbf{x}[k], \mathbf{u}[k]) = \underbrace{\omega_w \left(\frac{u_w[k]}{u_{w,\text{typ}}}\right)^2 + \omega_f \left(\frac{u_f[k]}{u_{f,\text{typ}}}\right)^2}_{\text{input usage}} + \underbrace{\omega_{\Delta w}(\Delta_w[k])^2 + \omega_{\Delta f}(\Delta_f[k])^2}_{\text{nutrient anomalies}} \tag{59}$$

Here $\omega_w$ and $\omega_f$ are weights on normalized irrigation and fertilizer inputs, $u_{w,\text{typ}}$ and $u_{f,\text{typ}}$ are typical application amounts per time step, and $\omega_{\Delta w}$ and $\omega_{\Delta f}$ penalize the cumulative nutrient anomalies $\Delta_w$ and $\Delta_f$ defined in (14). The quadratic form is not merely a convenience: an MPC formulation using a linear revenue objective mirroring the GA's cost function (Eqs. 51–53) produced degenerate solutions that front-loaded all irrigation and fertilizer into the earliest time steps, applying nothing for the remainder of the season. This "bang-bang" behavior is a well-known consequence of linear objectives in finite-horizon optimal control [68]. The squared terms resolve this by penalizing large instantaneous applications, naturally spreading resource use over time and producing schedules more similar to what a farmer might actually implement.

The terminal cost rewards crop development:

$$p(\mathbf{x}[K_d]) = -\omega_h \frac{x_h[K_d]}{k_h} - \omega_A \frac{x_A[K_d]}{k_A} - \omega_P \frac{x_P[K_d]}{k_P} \tag{60}$$

where $\omega_h$, $\omega_A$, and $\omega_P$ weight the normalized final height, leaf area, and fruit biomass, respectively, and we treat a reward as a negative cost. There is no terminal set constraint because we want the plant to grow as much as possible.

**Table 8** Bayesian optimization search space for MPC parameters

| Parameter | Lower | Upper | Scale | Optimal |
|---|---|---|---|---|
| $\omega_w$(irrigation cost) | 0.001 | 10.0 | log | 1.43 |
| $\omega_f$(fertilizer cost) | 0.0001 | 1.0 | log | $4.81 \times 10^{-3}$ |
| $\omega_{\Delta w}$(water anomaly) | 0.001 | 10.0 | log | 1.88 |
| $\omega_{\Delta f}$(fertilizer anomaly) | 0.001 | 10.0 | log | 0.450 |
| $\omega_h$(height value) | 10.0 | 1000.0 | linear | 323 |
| $\omega_A$(leaf area value) | 10.0 | 1000.0 | linear | 629 |
| $\omega_P$(fruit biomass value) | 100.0 | 10000.0 | linear | $1.95 \times 10^3$ |
| $K_d$(horizon, days) | 3 | 14 | integer | 9 |

### 5.3 Decision Variables and Search Space: MPC

The MPC controller uses a 9-day planning horizon ($K_d = 9$ ) with daily re-optimization. Control bounds are set to $u_w \in [0, 0.7]$ inches/hour and $u_f \in [0, 12]$ lbs/hour, reflecting practical application rate limits.

The CFTOC problem is formulated in Pyomo [69] and solved using the IPOPT interior-point nonlinear optimizer [70] with the MUMPS linear solver. The nonlinearity arises from the logistic dynamics, FIR convolution, EMA filtering, and exponential nutrient factor computation.

### 5.4 MPC Parameter Tuning via Bayesian Optimization

The MPC cost function contains seven tunable weights: $\omega_w, \omega_f, \omega_{\Delta w}, \omega_{\Delta f}, \omega_h, \omega_A, \omega_P$, plus the horizon length $K_d$ . These parameters significantly affect controller behavior, and their optimal values depend on the weather scenario ensemble. We use Bayesian Optimization (BO) to tune these parameters for robust performance.

Bayesian Optimization (BO) is well-suited to hyperparameter tuning problems where function evaluations are expensive and gradients are unavailable [71]. The approach maintains a probabilistic model of the objective function and uses this model to decide where to sample next, balancing exploration of uncertain regions against exploitation of promising regions.

Tree-structured Parzen Estimation (TPE) [72] models the posterior (probability of particular MPC parameter set, given a particular revenue), rather than the likelihood (revenue, given a parameter set), using two density functions $\ell(x)$ and $g(x)$. Observations are partitioned by a quantile threshold $\gamma$ (typically 0.1–0.25) into "good" and "bad" sets, with corresponding densities $\ell(x) = p(x|y \leq y^-)$ and $g(x) = p(x|y > y^-)$, where $y^-$ is the $\gamma$-quantile of observed values. Starting from expected improvement and applying Bayes' rule, one can show that maximizing expected improvement is equivalent to maximizing $\ell(x)/g(x)$, which serves as the TPE acquisition function.

TPE approximates the densities $\ell(x)$ and $g(x)$ using kernel density estimation (KDE), placing Gaussian kernels around observed data points. For multivariate parameters, TPE assumes independence across dimensions, and the "tree-structured" aspect accommodates conditional parameter dependencies. TPE offers computational advantages over Gaussian-process-based methods and handles mixed parameter spaces (continuous, discrete, categorical) making it well-suited for MPC parameter tuning where the objective landscape may contain discontinuities from solver failures and the search space includes both continuous weights and integer horizon length.

Table 8 defines the search space for MPC parameter tuning. Stage cost weights use log-scale sampling to span multiple orders of magnitude, while terminal cost weights use linear sampling (as ask prices for crops do not fluctuate over orders of magnitude). Integer sampling is used for the horizon.

For robust parameter tuning, we optimize average performance across multiple weather scenarios:

$$\boldsymbol{\theta}^* = \arg\max_{\boldsymbol{\theta}} \frac{1}{|\mathcal{S}|} \sum_{\mathcal{S}_i \in \mathcal{S}} \text{Revenue}(\text{MPC}(\boldsymbol{\theta}, \mathcal{S}_i)) \tag{61}$$

where $\mathcal{S}$ is a representative subset of 5 weather scenarios spanning the extremity range. We use $N = 100$ trials with $N_0 = 20$ random startup trials before engaging the TPE sampler. The optimal MPC parameter values resulting from the Bayesian optimization procedure are also included in Table 8.

### 5.5 Receding-Horizon Algorithm

MPC implements the CFTOC in a receding-horizon fashion, re-optimizing daily based on updated state and weather observations. At each decision day $k$, it

(1) Observes current plant state $\mathbf{x}[k]$ and obtain $K_d$-day weather forecast.
(2) Solves the CFTOC problem (58) to obtain optimal control sequence.

**Algorithm 1** Model Predictive Control for Irrigation and Fertilization

**Input:** Initial state $x[0]$, horizon $K_d$, season length $K$ days, weather data
**Output:** Control history $\{u[k]\}_{k=0}^{K-1}$, final state $x[K]$

Initialize FIR kernel buffers, EMA state variables
**for** $k = 0$ to $K - 1$ **do**
  Aggregate hourly weather forecast into daily averages for days $k, \ldots, k+K_d-1$
  Solve CFTOC$(x[k], \{(\hat{d_S}, \hat{d_T}, \hat{d_R})_{k+i}\}_{i=0}^{K_d-1}) \rightarrow \{u^*[k+i]\}_{i=0}^{K_d-1}$
  Extract first control: $(u_w[k], u_f[k]) \leftarrow u^*[k]$
  **for** $t = 0$ to 23 **do** ▷ Hourly simulation within day $k$
    Apply $(u_w[k]/24, u_f[k]/24)$ with actual weather $(d_{S,24k+t}, d_{T,24k+t}, d_{R,24k+t})$
    Update FIR buffers, EMA states, nutrient factors
    Advance plant state using closed-form logistic solution (16)
  **end for**
  $x_{k+1} \leftarrow$ current plant state
**end for**
**return** $\{u[k]\}_{k=0}^{K-1}$, $x[K_d]$

(3) Applies only the first control $\mathbf{u}^*[k] = (u_w[k], u_f[k])$ uniformly over day $k$.
(4) Simulates hourly plant dynamics over the day using actual (not forecast) weather.
(5) Advances to day $k + 1$ and repeat.

The key advantage of MPC over open-loop optimization (such as GA) is its closed-loop nature. By re-solving the optimization problem each day with updated state and weather information, MPC can:

- **Correct for forecast errors:** If yesterday's weather differed from the forecast, today's optimization accounts for the actual plant state.
- **Adapt to changing conditions:** If a heat wave arrives unexpectedly, MPC can adjust irrigation to maintain nutrient factors.
- **Exploit updated forecasts:** Longer-range forecasts become more accurate as the event approaches.

# 6 Results

## 6.1 Simulation Configuration

All simulations use a 121-day growing season representative of corn production in Iowa, with hourly resolution for plant dynamics and daily resolution for closed-loop control decisions. We generate 21 weather scenarios spanning normal conditions to extreme drought and heat stress using the stochastic weather generation procedure described in Section 3.

The crop model parameters are calibrated to literature values with carrying capacities of 3.0 m for plant height, 0.65 $m^2$ for leaf area, 20 leaves, 1000 flower spikelets, and 0.25 kg for fruit (ear) biomass. Absorption kernels use $\sigma_w = 30$ hr for water and $\sigma_f = 300$ hr for fertilizer, reflecting faster water uptake relative to slower fertilizer absorption. Economic parameters follow current USDA estimates: corn price of \$4/bushel, yield conversion of 200 bushels/acre at maximum biomass, water cost of \$2/acre-inch, and NPK fertilizer cost of \$0.61/lb. Net revenue is computed as gross revenue minus input costs.

## 6.2 Baseline Performance: Farmer Best Practices

To establish a performance baseline, we simulate crop growth using conventional farmer practices: weekly irrigation at 1 inch per application [66] and monthly fertilization at 90 lbs per application [73]. These values reflect standard agronomic recommendations for corn in the Corn Belt region [74–76].

Figure 4 shows the environmental disturbances and control inputs over the growing season under historical ("normal") conditions, as well as the periodic irrigation and fertilizer applications from farmer best practices.

Figure 5 shows the resulting plant state trajectories. Under normal conditions with conventional management, the plant reaches suboptimal final values: height of 2.6 m (vs. 3.0 m capacity), leaf area of 0.45 $m^2$ (vs. 0.65 $m^2$), and fruit biomass of 0.16 kg (vs. 0.25 kg). The net revenue under this scenario is \$587/acre. The baseline scenario demonstrates how the model captures cumulative stress effects: despite regular irrigation, the mismatch between applied water and the plant's metabolic expectations under drought conditions leads to sustained nutrient factor depression and reduced growth potential. Supplementary Figures 12 – 15 provide detailed visualizations of the nutrient absorption dynamics and stress factor evolution.

For reference, current Iowa corn yields average approximately 200 bushels per acre [77] at prices near $4 per bushel [78]. Variable costs for irrigation and fertilizer are approximately $220 per acre [79], yielding estimated net revenue of roughly $580/acre under favorable conditions. Our simulated farmer best practices achieves $587/acre, suggesting reasonable model calibration.

### 6.3 Genetic Algorithm Optimization

We optimize irrigation and fertilization strategies using a genetic algorithm with 100 generations and population size of 128. Each individual encodes four decision variables: irrigation frequency (hours), irrigation amount (inches), fertilizer frequency (hours), and fertilizer amount (lbs). Fitness is evaluated as net revenue over the 121-day growing season.

To assess robustness, we executed 10 independent GA runs with different random seeds under normal weather conditions. Figure 6 shows the convergence of all 10 runs, with each curve representing the best revenue achieved at each generation. All runs converge to similar final values (within 3% of each other), demonstrating that the GA reliably finds near-optimal solutions despite the stochastic nature of the search.

The best GA strategy under normal conditions achieves $915/acre net revenue on the baseline weather scenario, compared to $587/acre for farmer best practices—a 56% improvement on that scenario. Figure 7 shows the plant state trajectories for the best member from each of the 10 independent GA runs. All optimized strategies achieve similar, near-optimal growth trajectories, confirming that different GA runs converge to similar optimal strategies.

To verify that the model appropriately discriminates between effective and ineffective strategies, we examined the plant state trajectories for the *worst* member in each GA run's final population (Supplementary Figure 16). These strategies—the least fit survivors after 100 generations of evolution—achieve almost no growth and underperform the farmer best practices. This confirms that: (1) the GA successfully eliminates poor strategies through evolutionary pressure, and (2) arbitrary irrigation and fertilization schedules are indeed worse than conventional farmer practices, confirming the model is appropriately sensitive to irrigation and fertilizer input.

### 6.4 Four-Way Strategy Comparison

We compare four irrigation and fertilization strategies across all 21 weather scenarios:

**Farmer Best Practices.** Standard agronomic recommendations: weekly irrigation at 1 inch per application and monthly fertilization at 90 lbs per application.

**GA (Drought).** A strategy optimized via genetic algorithm for drought conditions (50% of historical precipitation). The optimal strategy applies irrigation every 52 days at 5.0 inches per application, and fertilizer every 33 days at 77 lbs per application.

**GA (Normal).** A strategy optimized via genetic algorithm for normal weather conditions. The optimal strategy applies irrigation every 81 days at 2.1 inches per application, and fertilizer every 9 days at 30 lbs per application.

**Adaptive MPC.** Model Predictive Control that re-optimizes daily based on observed weather and current plant state, using a 9-day planning horizon as determined by Bayesian optimization.

Table 9 presents summary statistics for all four strategies. Figure 8 visualizes the statistics from Table 9. Choosing a strategy with less variance in revenue across weather scenarios could be considered less risk averse, in which case GA (Normal) would be the ideal choice.

All three optimization-based methods outperform farmer best practices, with mean advantages of $170/acre for GA (Drought), $215/acre for GA (Normal), and $123/acre for MPC. Figure 9 visualizes revenue across all 21 scenarios sorted by increasing extremity index.

Figure 10 shows revenue comparison grouped by weather category. GA (Normal) dominates under normal and moderate conditions, while GA (Drought) provides better performance under drought and extreme scenarios. MPC provides consistent but not optimal performance across all categories.

Table 6 presents detailed results for all 21 scenarios. GA (Drought) wins all 21 scenarios against farmer best practices, GA (Normal) wins 18 of 21, and MPC wins 15 of 21. When comparing optimization methods directly, GA (Normal) outperforms MPC in all 21 scenarios, GA (Drought) outperforms MPC in 9 of 21 scenarios, and GA (Normal) outperforms GA (Drought) in 15 of 21 scenarios.

Figure 11 compares mean resource usage across strategies. While the GA strategies achieve higher revenue, MPC does so with substantially less water and fertilizer. MPC uses an average of 1.6 inches of irrigation and 240 lbs of fertilizer across all scenarios—roughly one-tenth the water of farmer best practices (18 inches) and one-tenth of GA (Drought) (15 inches), and less than half the fertilizer of any other strategy. Despite this dramatically lower resource footprint, MPC still outperforms farmer best practices in 15 of 21 scenarios, demonstrating that adaptive optimization can deliver meaningful economic gains alongside substantial sustainability benefits through reduced water and fertilizer use.

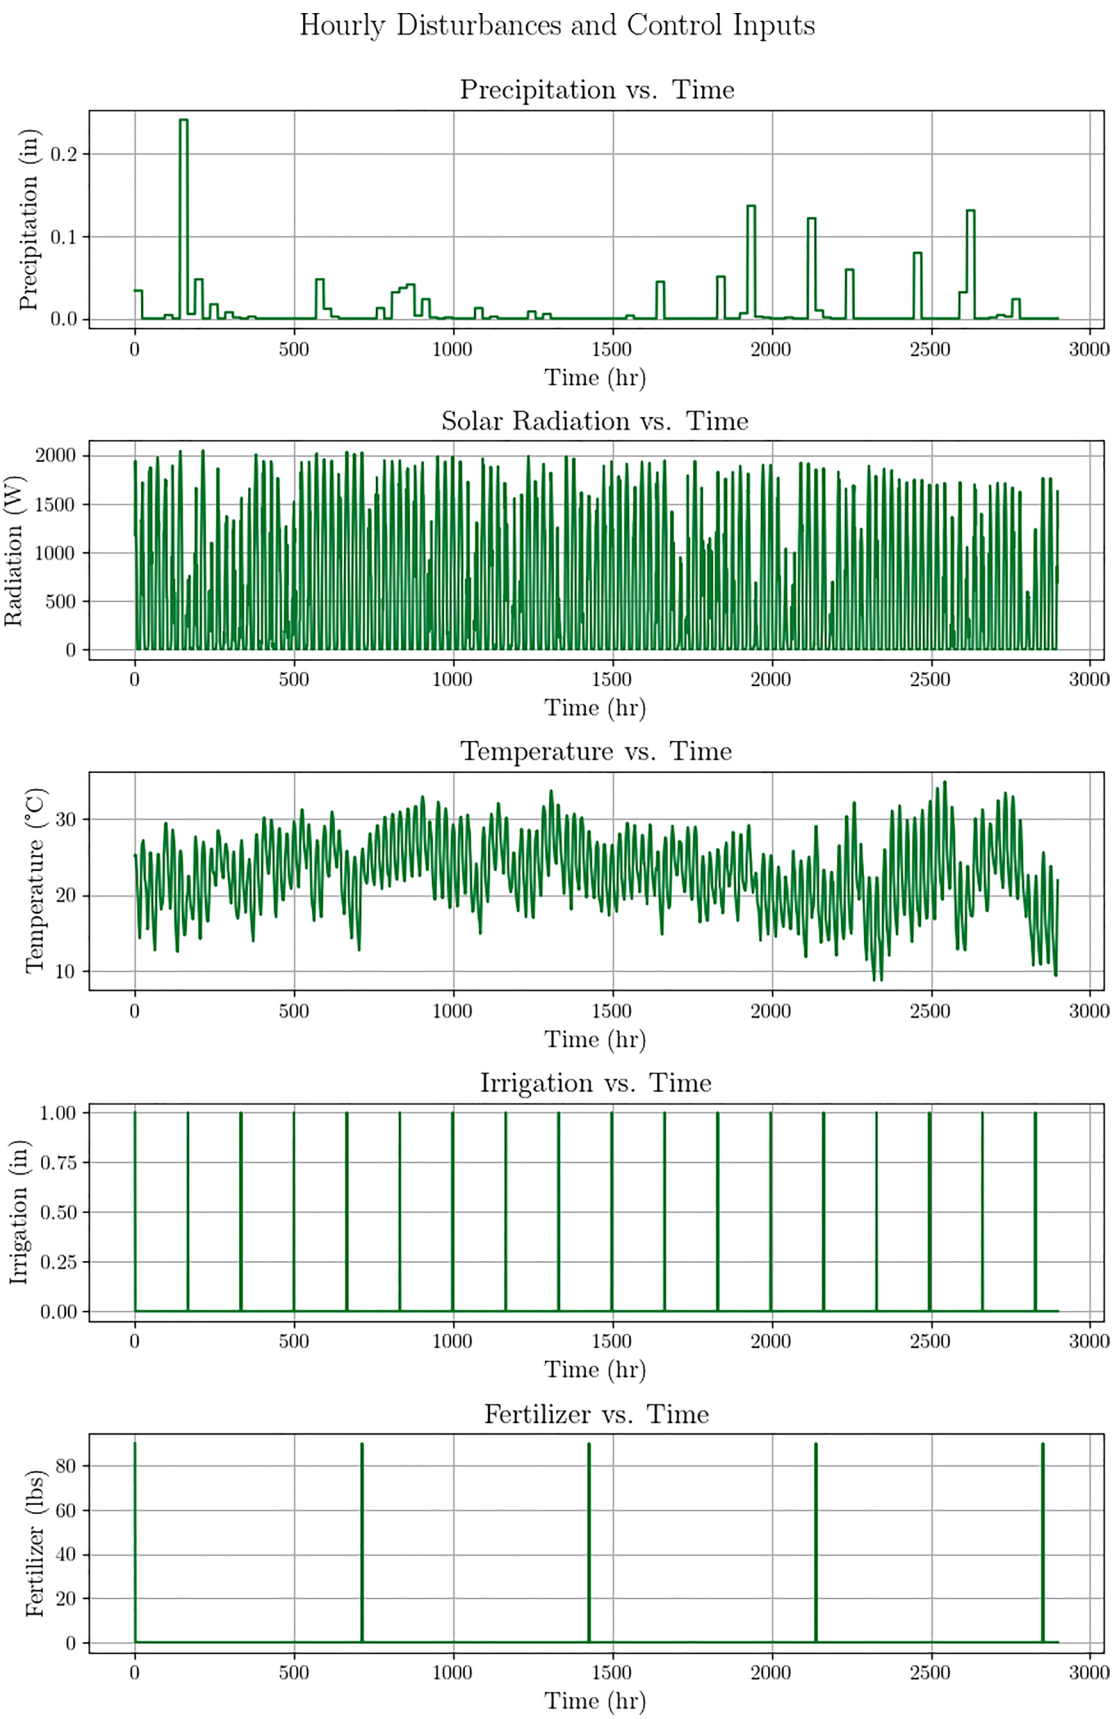
Hourly Disturbances and Control Inputs
Precipitation vs. Time
Precipitation (in)
0.0
0.1
0.2
0
500
1000
1500
2000
2500
3000
Time (hr)
Solar Radiation vs. Time
Radiation (W)
0
500
1000
1500
2000
Time (hr)
Temperature vs. Time
Temperature (°C)
10
20
30
Time (hr)
Irrigation vs. Time
Irrigation (in)
0.00
0.25
0.50
0.75
1.00
Time (hr)
Fertilizer vs. Time
Fertilizer (lbs)
0
20
40
60
80
Time (hr)

**Fig. 4** Environmental disturbances and control inputs for the baseline ("normal") scenario. Top three panels show hourly precipitation, solar radiation, and temperature from historical Iowa data. Bottom two panels show the farmer's irrigation (weekly, 1 inch) and fertilizer (monthly, 90 lbs) application strategy

## 7 Discussion

### 7.1 Interpretation of Results

The GA-optimized strategies differ from conventional wisdom in several notable ways. The algorithm discovers that less frequent but larger irrigation events, combined with reduced total fertilizer input, can outperform conventional uniform application schedules. This counterintuitive result emerges from the model's delayed absorption dynamics: the plant's metabolic expectations are calibrated to typical nutrient availability over time, and the GA discovers that strategically-timed resource inputs better maintain alignment with these metabolic expectations than aggressive compensation through frequent, uniform applications.

Under delayed absorption dynamics, timing matters more than total quantity. The FIR convolution introduces a temporal spread of 30 hours for water and 300 hours for fertilizer, meaning that today's application affects plant growth over the coming days to weeks. The GA can exploit this by front-loading or back-loading resources to align cumulative absorption with plant expectations throughout the season. Farmer best practices, with their rigid weekly/monthly schedules, cannot achieve this alignment.

MPC does not outperform the GA strategies, contrary to our initial hypothesis. The GA strategies, despite being "fixed," are optimized for their target conditions over the entire 121-day growing season. MPC re-optimizes daily but only sees 9 days ahead—insufficient foresight given the slow physiological dynamics. With fertilizer absorption spreading over 300 hours (12+ days), MPC's horizon barely captures a single absorption cycle. This limited foresight prevents MPC from making the long-term resource allocation decisions that the GA can pre-compute. In agricultural systems with slow dynamics—where nutrient absorption takes days to weeks—the GA's full-season optimization horizon appears more valuable than MPC's adaptive but myopic decision-making. We note, however, that this horizon-based explanation is not the only factor; structural and formulation differences between the two methods also contribute to the performance gap, as we discuss in Section 7.2.

Each GA strategy performs best when conditions match its optimization assumptions. GA (Normal) outperforms GA (Drought) in 12 of 21 scenarios, while GA (Drought) wins in the 9 scenarios with more extreme conditions. This is unsurprising: each strategy was optimized for its target conditions. The practical implication is that farmers should select the appropriate GA strategy based on seasonal forecasts—if drought is anticipated, use GA (Drought); otherwise, use GA (Normal).

Despite lower peak revenue, MPC offers a sustainability advantage worth noting. MPC achieves competitive returns (outperforming farmer best practices in 15 of 21 scenarios) while using dramatically less water (1.3–2.0 inches vs. 15–18 inches for other strategies) and fertilizer (216–307 lbs vs. 450 lbs for farmer best practices). For operations prioritizing environmental impact or input cost reduction, MPC may be the preferred choice despite lower revenue.

Our MPC formulation assumes perfect weather forecasts over the planning horizon, which represents MPC's best-case operating condition. Because MPC is the only method that consumes forecast information at runtime—the GA strategies are fixed and do not adapt to observed or predicted weather—degrading forecast quality can only reduce MPC's performance while leaving the GA unaffected. The GA's robustness to weather uncertainty is already tested by a different mechanism: each GA strategy is optimized for one weather condition (e.g., baseline) but evaluated across all 21 scenarios, including conditions it was not tuned for. Introducing forecast noise into MPC would therefore widen the performance gap in favor of the GA, reinforcing rather than undermining the conclusions drawn above. Stochastic or robust MPC formulations that explicitly account for forecast uncertainty could partially mitigate this degradation, but such extensions would also increase the computational cost of each daily solve and are left to future work.

### 7.2 Asymmetries in the Optimization Strategies

The GA-versus-MPC comparison involves asymmetries beyond the horizon mismatch discussed above: the decision space structure, the objective formulation, and the computational effort. Disentangling these is important for understanding whether the results reflect a fundamental property of open-loop versus closed-loop control in slow-dynamics systems, or are partly a consequence of the problem design.

**Decision space asymmetry.** The GA searches a 4-dimensional space (Eq. 50), where each candidate encodes a periodic application schedule for the full 121-day season. MPC, by contrast, solves a sequence of 18-dimensional nonlinear programs (2 control inputs $\times$ 9-day horizon) at each of 121 daily decision epochs. The GA's low-dimensional parameterization acts as an implicit regularization: by restricting the policy space to periodic schedules, it excludes erratic strategies and biases the search toward smooth, implementable policies. MPC's richer decision space gives it the *potential* to outperform any fixed periodic strategy, since it can express arbitrary time-varying input trajectories. That it does not realize this potential is attributable to the horizon

limitation: a full-horizon MPC formulation (if computationally feasible) would subsume the GA's search space and could only do at least as well.

**Objective function asymmetry.** The GA directly maximizes end-of-season revenue (Eqs. 51–53), a linear function of final crop state minus cumulative input costs. MPC minimizes a sum of quadratic stage costs (Eq. 59) plus a linear terminal reward (Eq. 60). These are not equivalent objectives. The quadratic input penalties in MPC's stage cost penalize large applications disproportionately—doubling an input quadruples its cost contribution—which actively discourages aggressive resource use. This largely explains MPC's dramatically lower water consumption (1.3–2.0 vs. 15–18 inches): the quadratic formulation imposes an implicit conservation preference absent from the GA's linear cost structure. Importantly, this formulation difference is not an arbitrary design choice: as noted in Section 5.2, an MPC formulation using the GA's linear revenue objective produced degenerate "bang-bang" solutions that front-loaded all resources. The quadratic formulation is *necessary* for MPC to produce physically sensible schedules in this setting. The GA, by contrast, can use a simple linear revenue objective because its periodic parameterization inherently regularizes the solution—the 4-parameter structure prevents bang-bang behavior by construction. This asymmetry is therefore structural: MPC's flexible decision space requires regularization through the objective function, and that regularization comes at the cost of revenue.

To further investigate these effects, we ran a resource-constrained GA with total seasonal inputs capped at MPC's maximum usage levels (2.0 inches water, 310 lbs fertilizer). Under these constraints, the GA found a strategy using 0.76 inches of water and 283 lbs of fertilizer—comparable to MPC's resource usage—and achieved a mean revenue of \$759/acre across all 21 scenarios, compared with MPC's \$750/acre. The constrained GA wins 11 of 21 scenarios, with MPC winning in drought conditions and the constrained GA winning in normal and wet conditions. This near-parity could be interpreted as evidence that the two methods are equivalent at equal resource levels but it could also suggest that the revenue cost of MPC's *required* quadratic regularization approximately offsets the benefit of its daily adaptive capability: MPC gains from adapting to observed conditions but pays a comparable price for the objective formulation it needs to produce sensible schedules. The unconstrained GA's \$863/acre mean reflects its structural advantage—the periodic parameterization provides free regularization, allowing the GA to pursue revenue directly without the front-loading behavior that forces MPC into a more conservative formulation.

**Computational effort.** The methods also differ in computational cost and the design choices driven by tractability. The GA evaluates approximately 12,800 full-season simulations (128 population $\times$ 100 generations) to optimize 4 parameters. While conceptually simple, a pure Python implementation required hours per run; only by migrating the inner loop to C++ did we reduce GA optimization time to approximately one minute. MPC solves 121 daily nonlinear problems via the IPOPT (Interior Point OPTimizer) library, each embedding the full crop model as equality constraints, and completes a full-season simulation in comparable time. Several MPC-specific design choices were driven by tractability. We use a two-tier temporal scheme: plant dynamics are simulated at hourly resolution (matching physiological timescales), but decisions are made daily—solving 2,900 hourly optimizations would be prohibitive. Similarly, the planning horizon is constrained to 3–14 days, as longer horizons would increase the CFTOC problem size and solver time. Despite these differences, both methods were given sufficient computational resources to converge: the GA's consistency across 10 independent runs (Section 6.3) and MPC's Bayesian-optimized hyperparameters (Section 5.4) both indicate convergence.

### 7.3 Limitations and Future Extensions

Several model limitations suggest directions for future work:

**Growth model.** The logistic equation (1) assumes symmetric growth around the inflection point. Richards growth generalizes this with a shape parameter $v$:

$$\frac{dx}{dt} = a_x x \left[1 - \left(\frac{x}{k_x}\right)^v\right] \tag{62}$$

where $v > 1$ produces steeper early growth and $n < 1$ produces steeper late growth.

**Absorption kernels.** Gaussian kernels are symmetric, but physiological absorption often exhibits fast activation followed by slow decay. Log-normal or Gamma kernels could better capture this asymmetry.

**Saturation.** The current model does not explicitly limit nutrient uptake. Modeling nutrient saturation would provide more realistic response to over-application.

**Per-event operational costs.** The objectives penalize total resource quantity but not the number of application events. In practice, each irrigation or fertilization event incurs operational costs (labor, equipment, fuel) regardless of the amount applied. Including per-event costs would differentially affect strategies with varying application frequencies.

**Single-point model.** The crop model represents a single plant without spatial heterogeneity. Field-scale

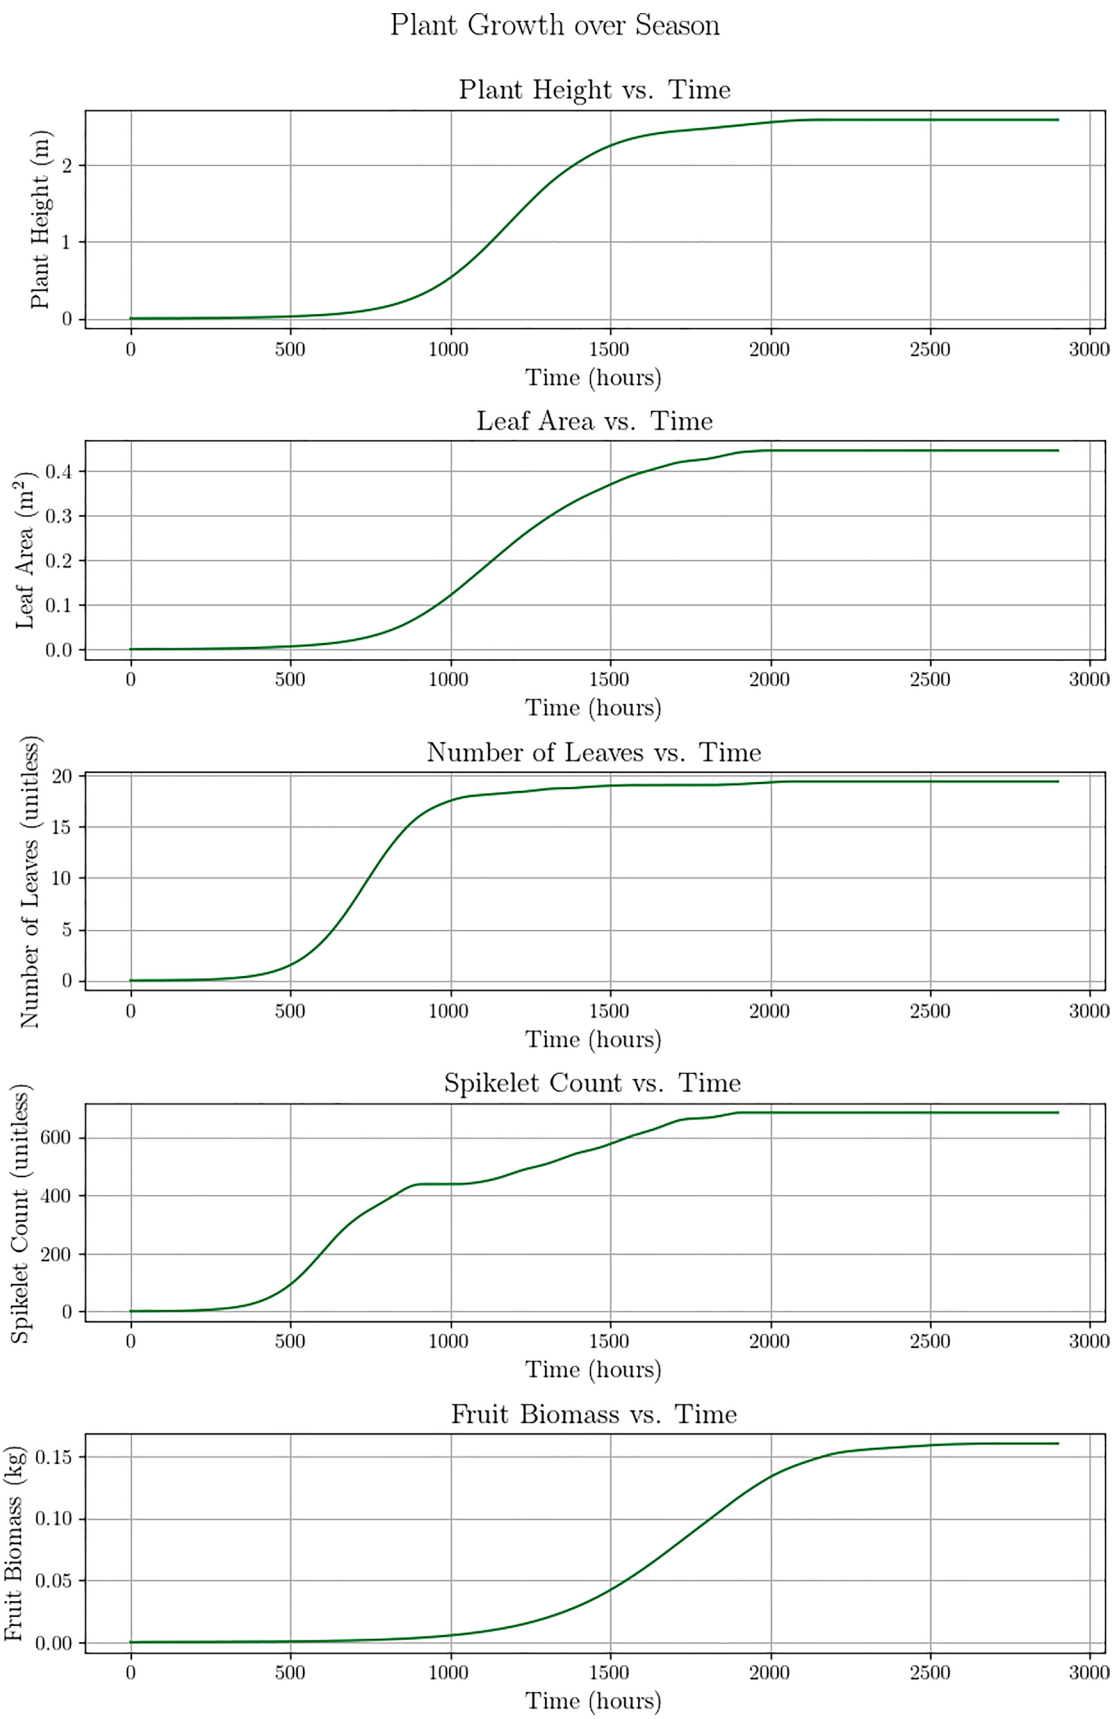


**Fig. 5** Plant state variable trajectories under the baseline scenario (farmer best practices during mild drought). All state variables reach suboptimal final values due to cumulative water stress. This strategy yields $587/acre in revenue

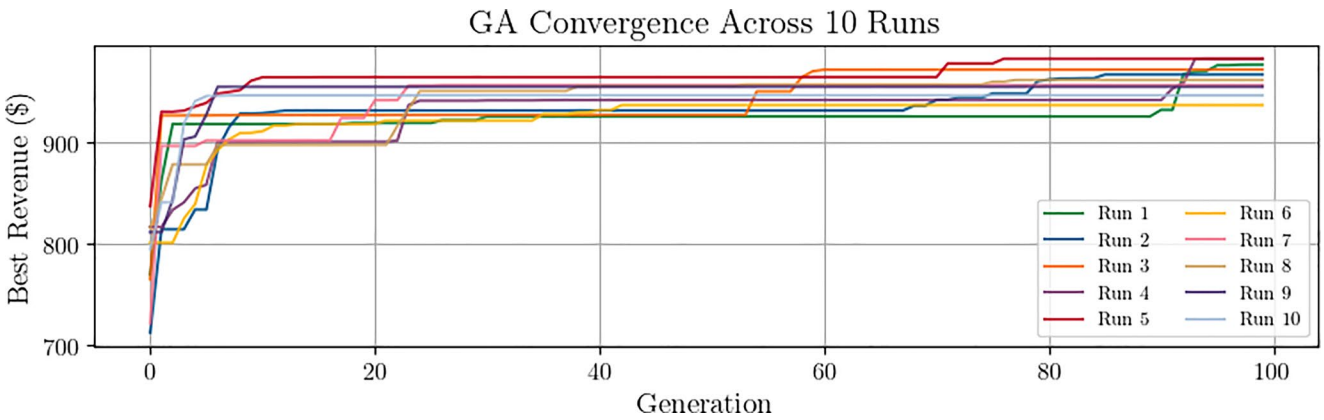


**Fig. 6** GA convergence across 10 independent runs. Each curve shows the best revenue at each generation. All runs exhibit rapid improvement in early generations followed by convergence to near-optimal solutions. The consistency across runs demonstrates algorithm robustness

implementation would require accounting for spatial variation in soil properties, drainage, and microclimate.

**Soil moisture dynamics.** The current model treats irrigation as directly available water without explicitly modeling soil moisture content, drainage, or field capacity. This omission likely favors the GA's infrequent, large-volume irrigation strategy, but, in practice, applying several inches in a single event would exceed the soil's water-holding capacity, and the excess would be lost to deep drainage or runoff. MPC's frequent, small applications would suffer less from such losses. Incorporating a soil water balance model could therefore narrow the GA–MPC performance gap and is a natural extension of this work.

**GA–MPC comparison isolation.** Although we partially disentangle the effects of resource conservation and horizon myopia through a resource-constrained GA ablation (Section 7.2), further ablation studies—such as increasing the GA's temporal resolution beyond periodic schedules—would more fully isolate the effect of the control paradigm (open-loop versus closed-loop) from the differing objective formulations and decision space structures.

## 8 Conclusion

This paper presented a coupled ODE model for crop growth that captures two physiological phenomena often neglected in optimization studies: delayed nutrient absorption and cumulative stress tracking. The model tracks five state variables (plant height, leaf area, number of leaves, flower size, and fruit biomass) with time-varying growth rates modulated by nutrient factors derived from a series of transformations which capture the physiology of the stress and delayed absorption.

We compared two optimization approaches for irrigation and fertilizer scheduling: genetic algorithm (GA) optimization for fixed seasonal strategies and model predictive control (MPC) for adaptive daily decision-making. Applied to corn production in Iowa across 21 stochastic weather scenarios, both optimization-based methods outperform farmer best practices. GA optimization achieves 35% higher mean revenue than conventional practices ($842 vs. $626 per acre), demonstrating the value of computational optimization in agriculture.

Contrary to our initial hypothesis that adaptive control would outperform fixed strategies under weather uncertainty, the results show that MPC does not outperform the GA strategies. GA (Normal) outperforms MPC in all 21 scenarios, while GA (Drought) outperforms MPC in 9 of 21 scenarios. The explanation lies primarily in the mismatch between MPC's planning horizon (9 days) and the crop model's physiological timescales (fertilizer absorption spreads over 300+ hours). Although the 9 day horizon was found by Bayesian optimization to be optimal, MPC's daily adaptation cannot substitute for the GA's full-season optimization horizon when system dynamics (especially fertilizer absorption dynamics) are slow. The comparison also involves decision space and objective function asymmetries. The GA optimizes four scalar parameters under a linear revenue objective while MPC optimizes daily with quadratic penalties on irrigation and fertilizer usage. A resource-constrained ablation study, in which the GA's total seasonal inputs were capped at MPC's usage levels, shows that the constrained GA achieves near-parity with MPC ($759 vs. $750/acre mean, winning 11 of 21 scenarios). However, this near-parity does not necessarily imply equivalence: MPC's quadratic formulation is not an optional design choice but a necessary regularization to prevent solutions which front load nutrient application and fail long term. The GA's periodic parameterization, which was born out of computational tractability, actually provides this regularization structurally, allowing it to optimize revenue directly.

The practical conclusion is straightforward: for seasonal agricultural planning, pre-optimization for expected

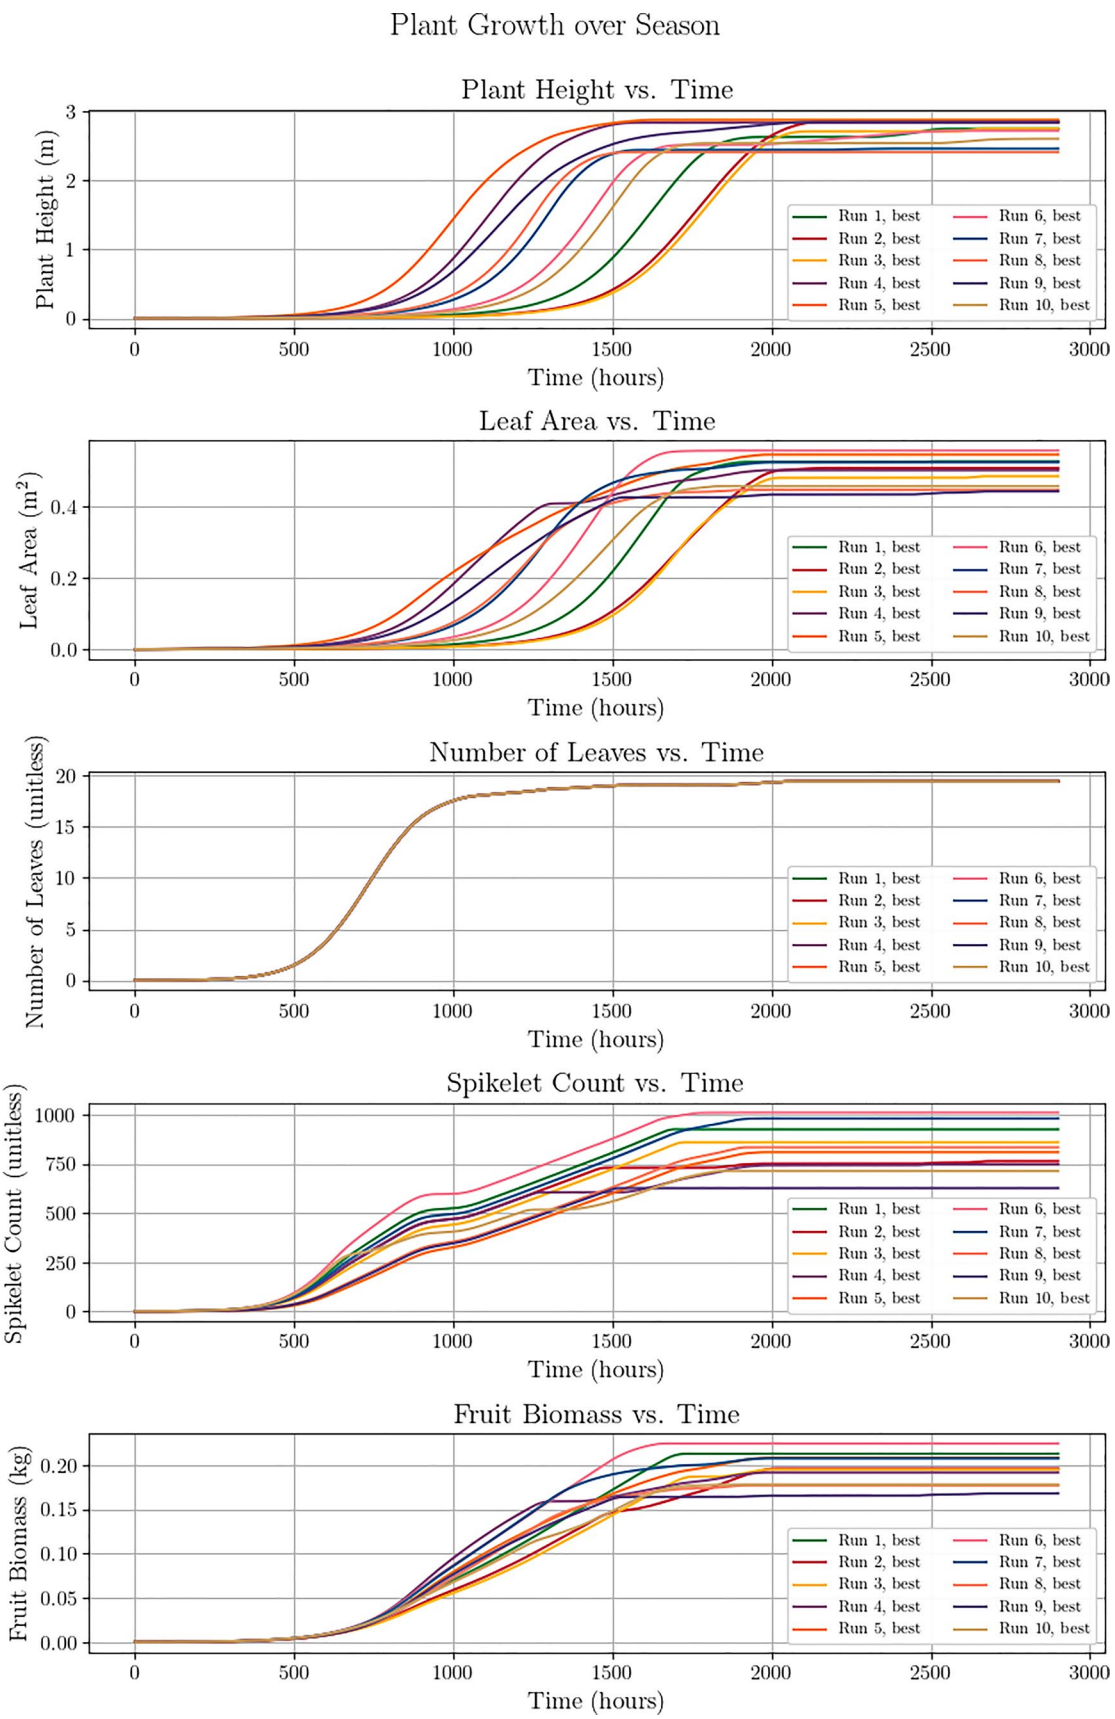


**Fig. 7** Plant state variable trajectories for the best member from each of the 10 independent GA runs. The tight clustering of trajectories demonstrates that different GA runs converge to similar optimal strategies

**Table 9** Summary statistics for the four strategies across 21 weather scenarios. CV = coefficient of variation (standard deviation / mean)

| Strategy | Mean ($/acre) | Std Dev | CV (%) | Min | Max |
|---|---|---|---|---|---|
| Farmer Best Practices | 626 | 135 | 21.6 | 385 | 853 |
| GA (Drought) | 796 | 132 | 16.6 | 550 | 997 |
| GA (Normal) | 842 | 124 | 14.7 | 527 | 973 |
| Adaptive MPC | 750 | 115 | 15.4 | 445 | 886 |

conditions outperforms adaptive control. If normal weather is expected, use GA (Normal); if drought is anticipated, use GA (Drought); if a different weather condition is expected, a GA can be tuned to that expectation. However, MPC offers a sustainability advantage worth considering: it achieves competitive returns while using dramatically less water and fertilizer than all other strategies, which may be valuable for

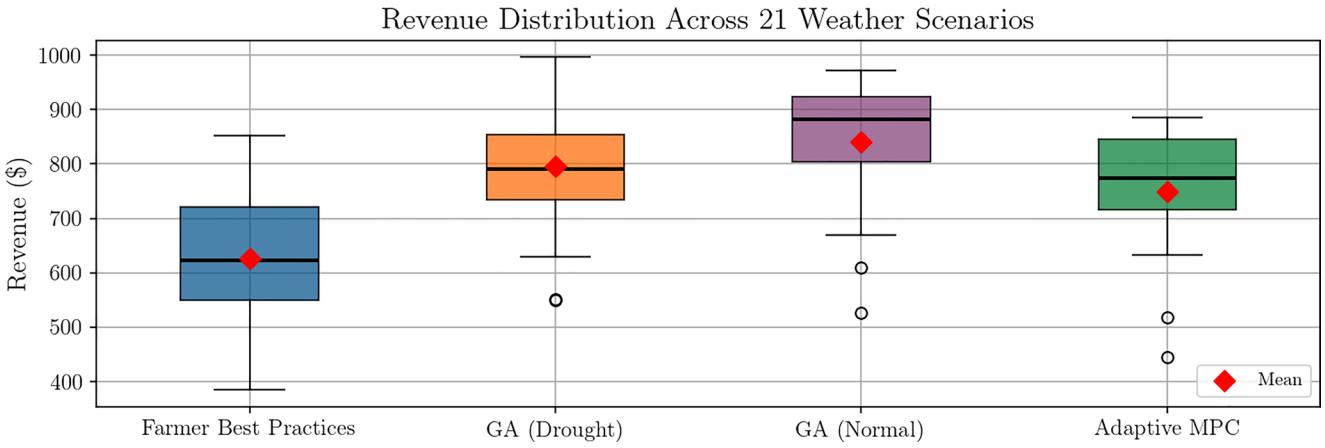


**Fig. 8** Revenue distribution across all 21 weather scenarios. Box plots show median, interquartile range, and outliers for each strategy. GA (normal) achieves the highest median revenue but also exhibits the widest range, while MPC shows the lowest variability

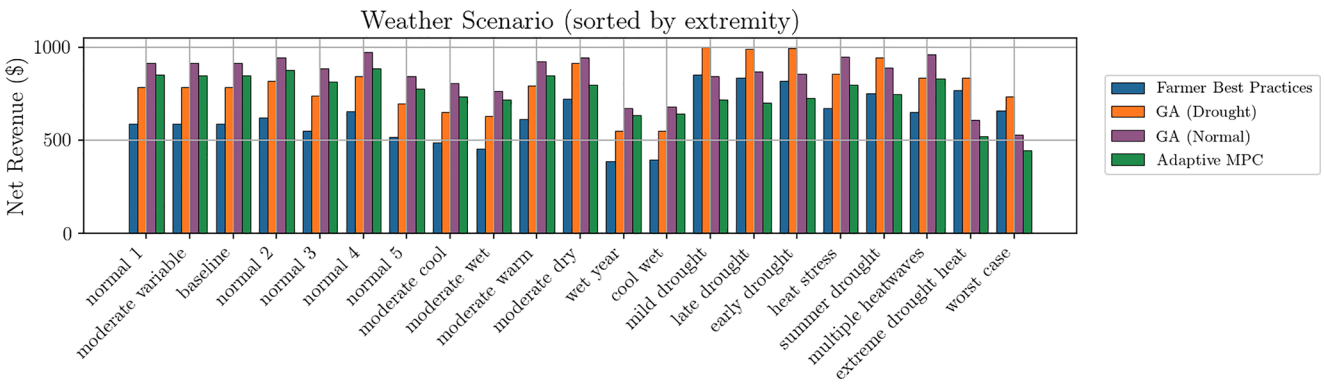


**Fig. 9** Revenue comparison across all 21 weather scenarios, sorted by extremity index. All optimization-based strategies outperform farmer best practices in most scenarios, with GA (normal) achieving the highest revenues under normal conditions and GA (drought) excelling under extreme conditions

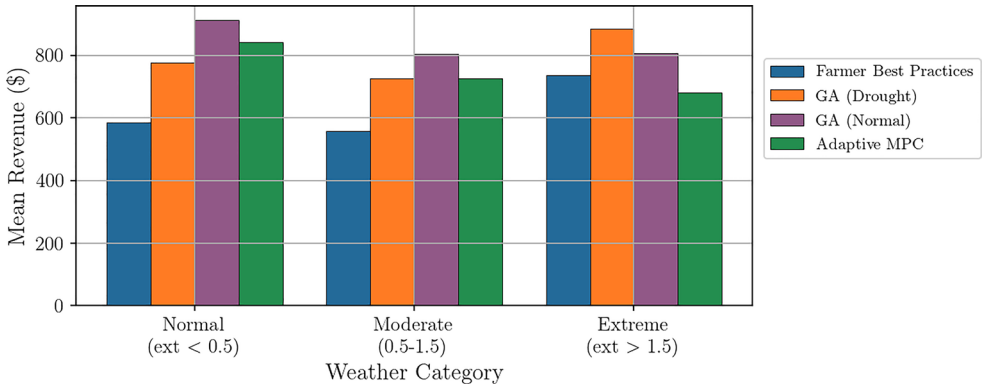


**Fig. 10** Revenue comparison by weather category. GA (normal) achieves the highest revenues under normal conditions, while GA (drought) is better suited for extreme weather scenarios

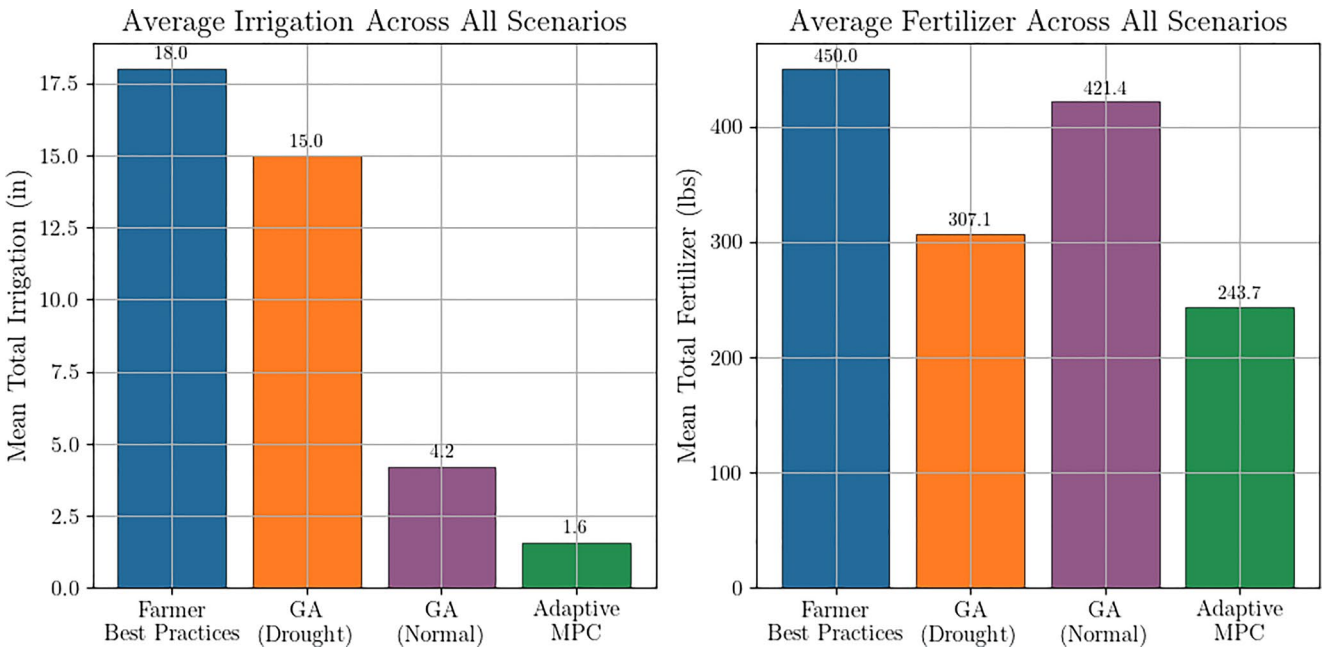


**Fig. 11** Mean total irrigation (left) and fertilizer (right) across all 21 weather scenarios for each strategy. MPC achieves competitive revenue while using dramatically less water and fertilizer than all other approaches

operations prioritizing environmental impact or input cost reduction.

Future work should explore inclusion of soil nutrient effects, crop interaction with a field-scale model, and field validation of the model. Though this study calibrated model parameters to corn, the framework generalizes readily to other crops by adjusting the growth parameters and input-effect relationships.

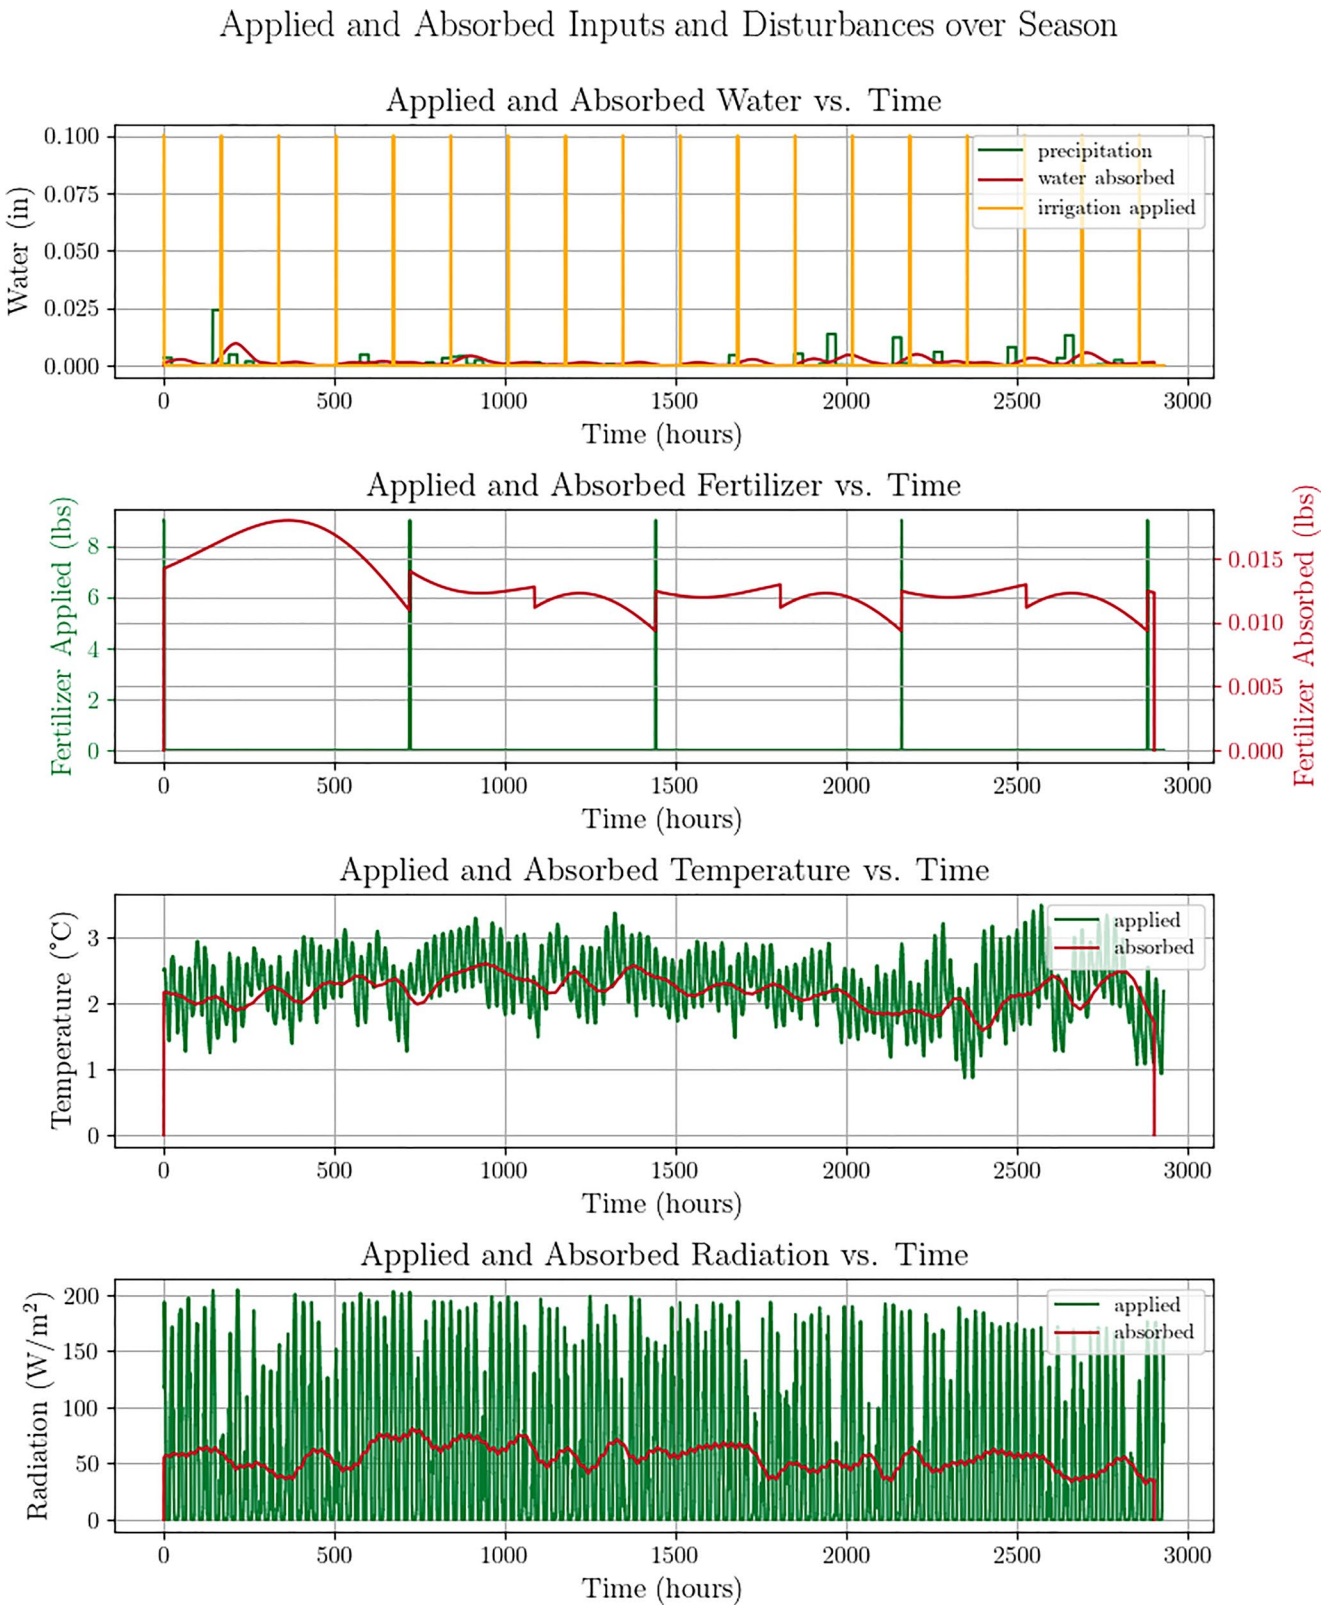


**Fig. 12** Applied versus absorbed nutrients under the baseline farmer strategy. The delayed absorption dynamics are evident in the lag between applied inputs and the smoothed absorbed signals. Water absorption ($\sigma_w = 30$ hr) responds more quickly than fertilizer absorption ($\sigma_f = 300$hr)

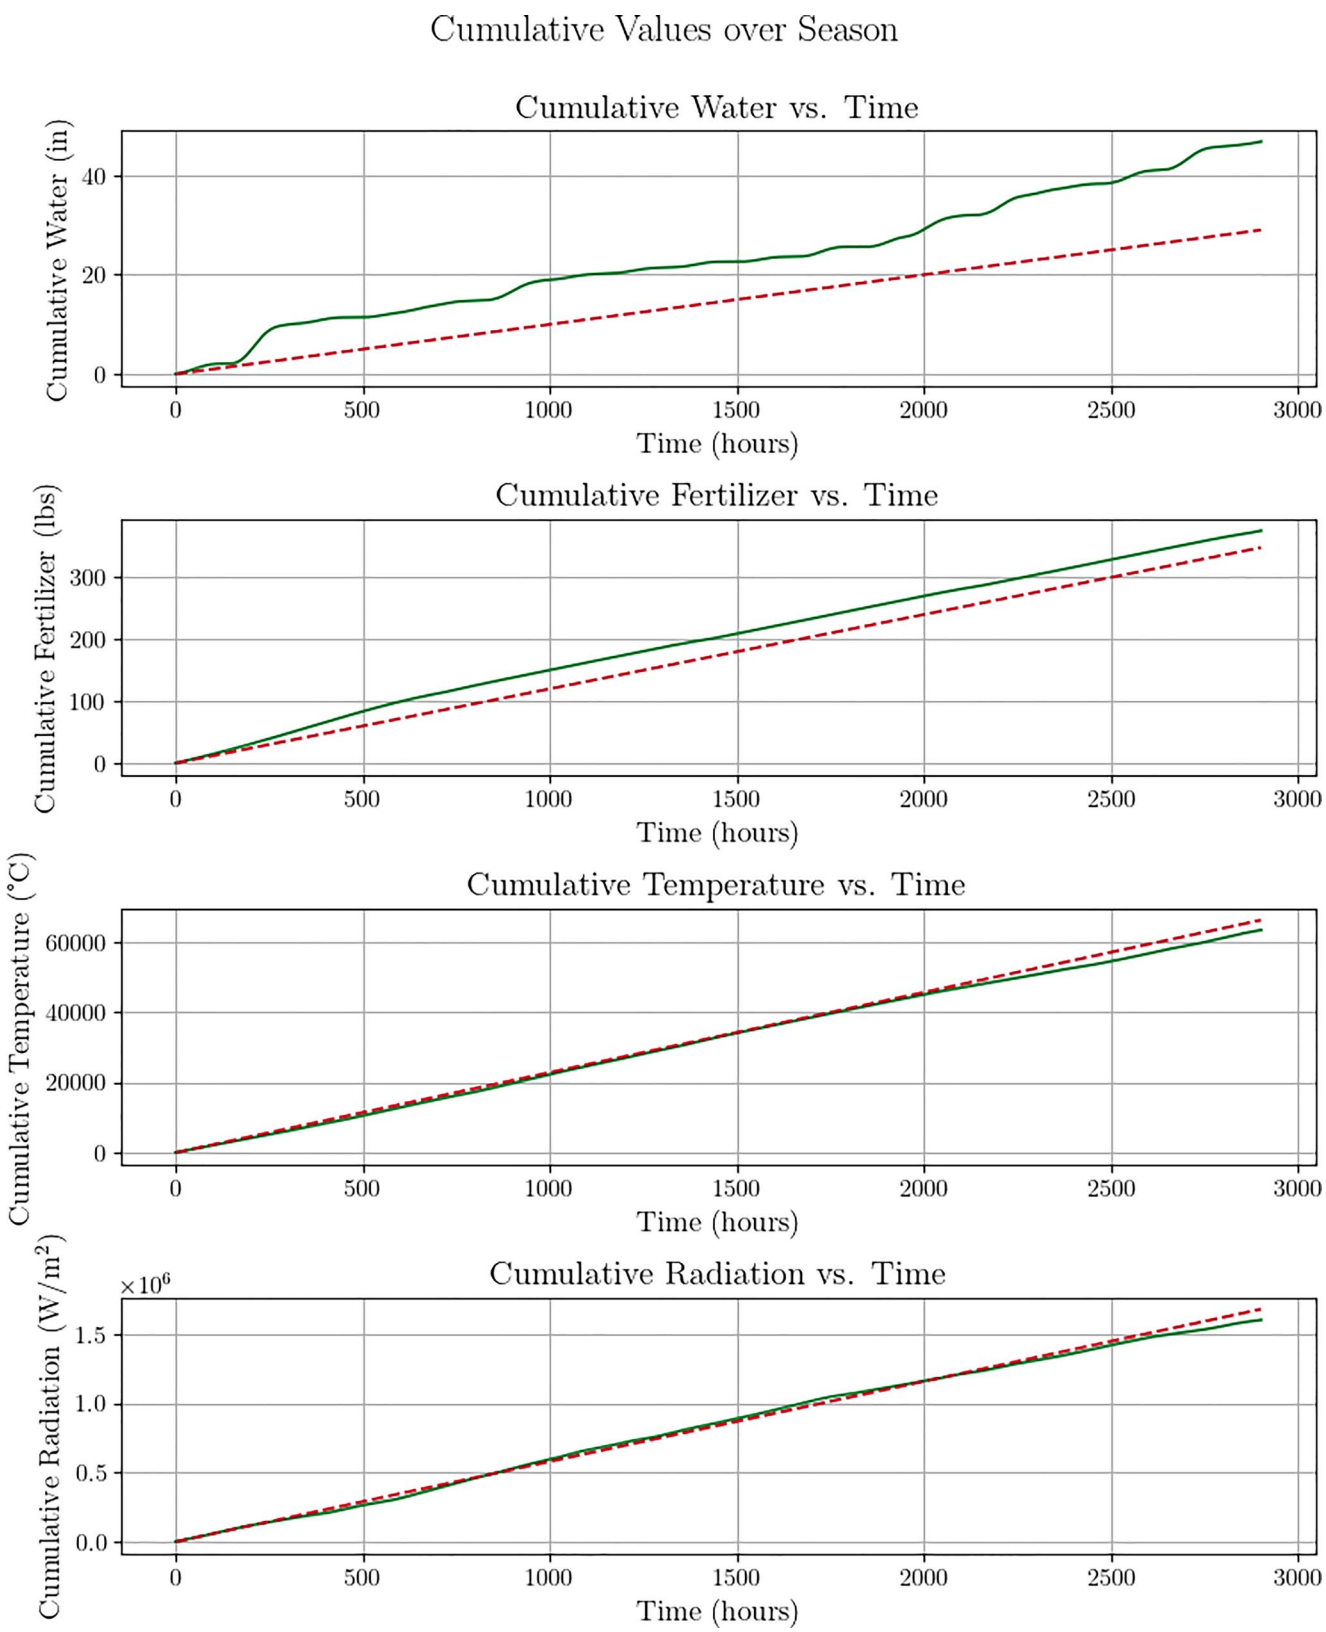


**Fig. 13** Cumulative absorbed nutrients (solid) versus expected levels (dashed red). The growing gap between actual and expected water absorption reflects the drought stress accumulating over the season

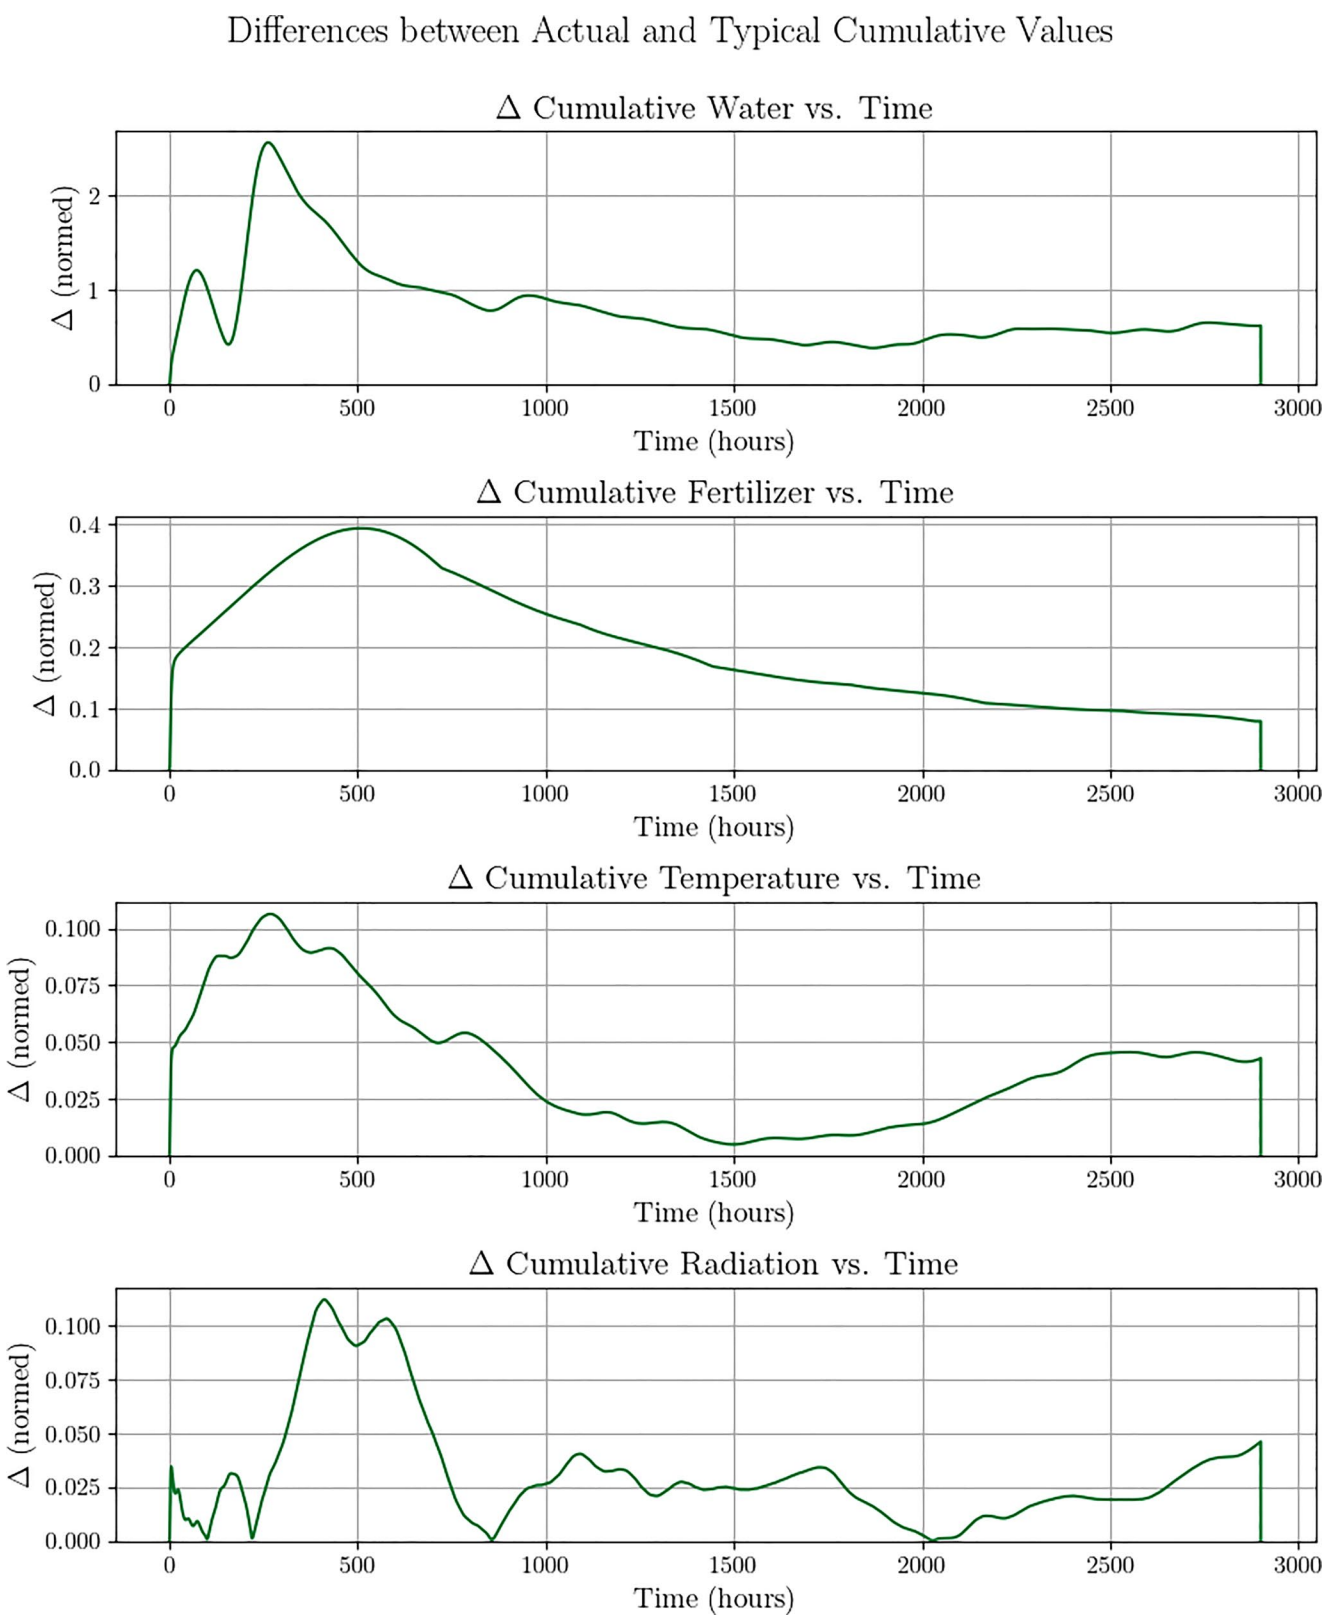


**Fig. 14** Instantaneous divergence from expected cumulative nutrient levels. Higher divergence indicates greater stress

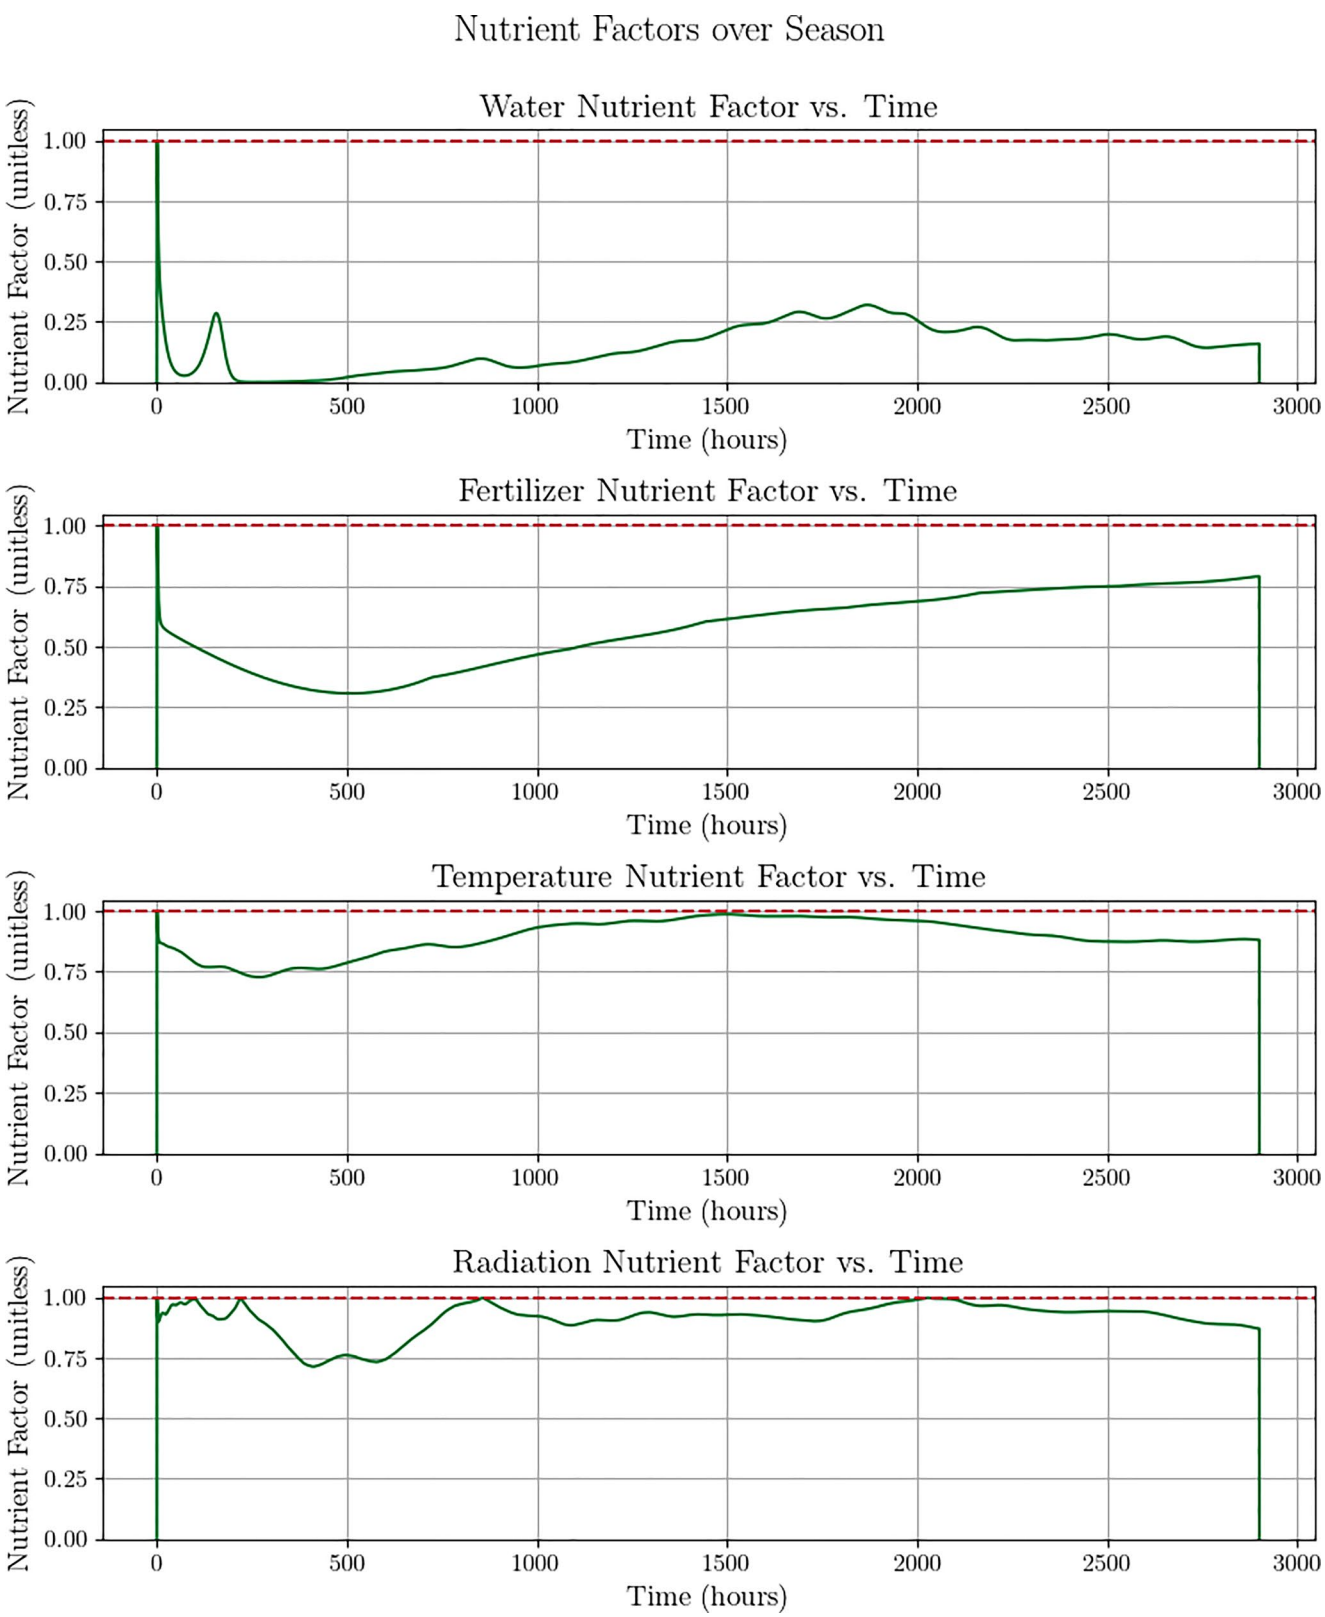


**Fig. 15** Nutrient factor evolution under the baseline scenario. The water nutrient factor $\nu_w$ declines due to cumulative drought stress, while fertilizer, temperature, and radiation factors remain closer to 1.0 (no stress). The declining $\nu_w$ (which later recovers) reduces effective growth rates and carrying capacities throughout the season

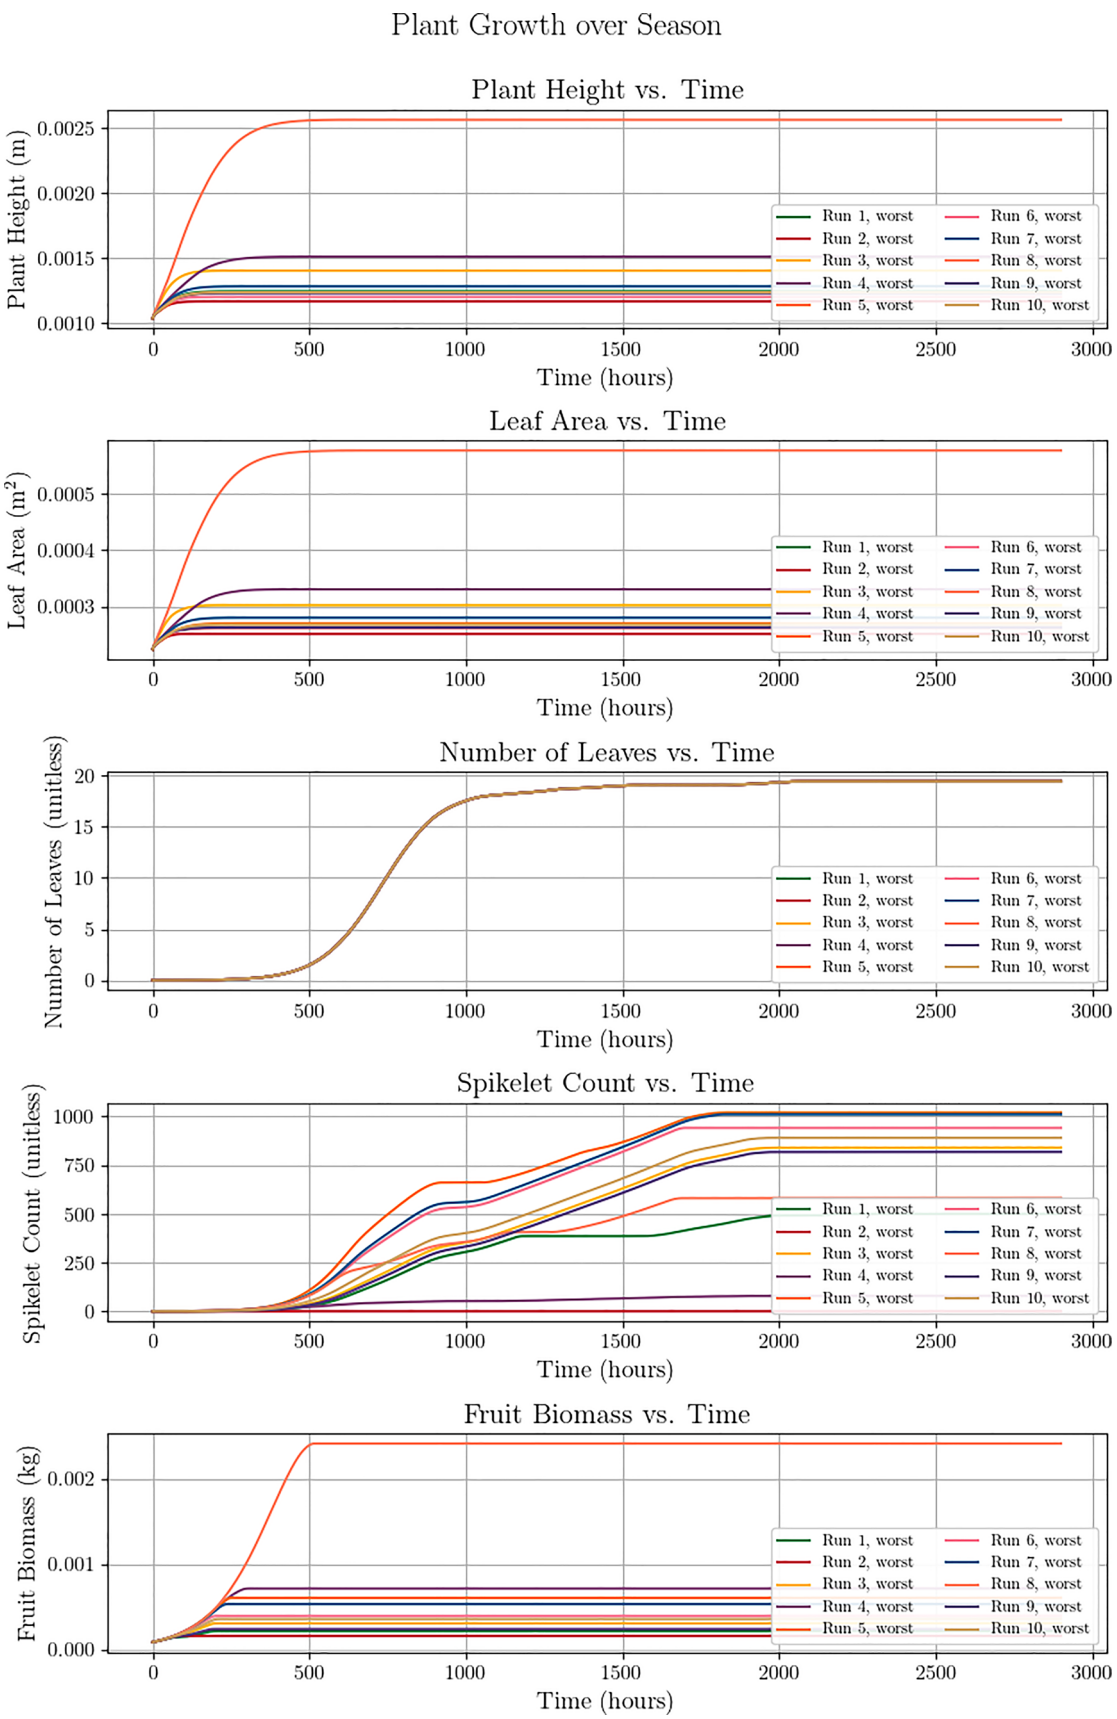


**Fig. 16** Plant state variable trajectories for the worst member from each GA run's final population. These suboptimal strategies achieve almost no growth, confirming that the model does not guarantee excellent growth and appropriately penalizes poor resource allocation strategies

**Author contributions** Carla J. Becker: Conceptualization, Methodology, Software, Formal Analysis, Investigation, Visualization, Writing – Original Draft. Tarek I. Zohdi: Funding Acquisition, Supervision, Writing – Review & Editing.

**Funding** This work has been partially supported by the UC Berkeley College of Engineering and the USDA AI Institute for Next Generation Food Systems (AIFS), USDA award number 2020–67021-32855.

**Code Availability** The source code used for this study is archived on GitHub at https://github.com/carlacupcake/smartfarm.

## Declarations

**Competing Interests** The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.



## References

1. Economic Research Service (2024) Farm income and wealth statistics. US Department of Agriculture.
2. Zohdi TI (2024) A machine-learning enabled digital-twin framework for next generation precision agriculture and forestry. Comput Methods Appl Mech Eng 431:117250. https://doi.org/10.1016/j.cma.2024.117250.
3. Mengi E, Becker CJ, Sedky M, Yu S, Zohdi TI (2023) A digital-twin and rapid optimization framework for optical design of indoor farming systems. Comput Mech 72:953–970. https://doi.org/10.1007/s00466-023-02421-9.
4. Isied R, Mengi E, Zohdi TI (2022) A digital-twin framework for genomic-based optimization of an agrophotovoltaic greenhouse system. Proc R Soc A 478(2267). https://doi.org/10.1098/rspa.2022.0414.
5. Mengi E, Samara OA, Zohdi TI (2023) Crop-driven optimization of agrivoltaics using a digital-replica framework. Smart Agric Technol 4:100150. https://doi.org/10.1016/j.atech.2022.100168.
6. Goodrich P, Betancourt O, Arias A, Zohdi TI (2023) Placement and drone flight path mapping of agricultural soil sensors using machine learning. Comput Electron Agriculture 205:107591. https://doi.org/10.1016/j.compag.2022.107591.
7. Tagkopoulos I, Brown SF, Liu X, Zhao Q, Zohdi TI, Earles JM, Nitin N, Runcie DE, Lemay DG, Smith AD, Ronald PC, Feng H, Youtsey GD (2022) Special report: AI institute for next generation food systems (AIFS). Comput Electron Agriculture 196:106819. https://doi.org/10.1016/j.compag.2022.106819.
8. Betancourt JO, Li I, Mengi E, Corrales L, Zohdi TI (2024) A computational framework for precise aerial agricultural spray delivery processes. Arch Computat Methods Eng. https://doi.org/10.1007/s11831-024-10106-6.
9. Verhulst P-F (1838) Notice sur la loi que la population suit dans son accroissement. Correspondance Mathématique et physique 10:113–126. https://doi.org/10.1007/BF02309004.
10. Jones JW, Hoogenboom G, Porter CH, Boote KJ, Batchelor WD, Hunt L, Wilkens PW, Singh U, Gijsman AJ, Ritchie JT (2003) The dssat cropping system model. Eur J Agron 18(3–4):235–265. https://doi.org/10.1016/S1161-0301(02)00107-7.
11. Holzworth DP, Huth NI, DeVoil PG, Zurcher EJ, Herrmann NI, McLean G, Chenu K, Oosterom EJ, Snow V, Murphy C, Moore AD, Brown H, Whish JPM, Verrall S, Fainges J, Bell LW, Peake AS, Poulton PL, Hochman Z, Thorburn PJ, Gaydon DS, Dalgliesh NP, Rodriguez D, Cox H, Chapman S, Doherty A, Teixeira E, Sharp J, Cichota R, Vogeler I, Li FY, Wang E, Hammer GL, Robertson MJ, Dimes JP, Whitbread AM, Hunt J, van Rees H, McClelland T, Carberry PS et al. (2014) Apsim–evolution towards a new generation of agricultural systems simulation. Environ Modell Softw 62:327–350. https://doi.org/10.1016/j.envsoft.2014.07.009.
12. Van Diepen C, Wolf J, Van Keulen H, Rappoldt C (1989) Wofost: a simulation model of crop production. Soil Use Manag 5(1):16–24. https://doi.org/10.1111/j.1475-2743.1989.tb00755.x.
13. Gebbers R, Adamchuk VI (2010) Precision agriculture and food security. Science 327(5967):828–831. https://doi.org/10.1126/science.1183899.
14. Singh A (2012) An overview of the optimization modelling applications. J Hydrol 466-467:167–182. https://doi.org/10.1016/j.jhydrol.2012.08.004.
15. Epperson JE, Hook JE, Mustafa YR (1993) Dynamic programming for improving irrigation scheduling strategies of maize. Agric Syst 42(1–2):85–101. https://doi.org/10.1016/0308-521X(93)90070-I. Applications of Dynamic Optimization Techniques to Agricultural Problems.
16. Wang W, Xu Y, Xu J (2020) Realization of wireless charging in intelligent greenhouse with orthogonal coil system uniform magnetic field. Comput Electron Agriculture 175:105524. https://doi.org/10.1016/j.compag.2020.105524.
17. Goldberg DE (1989) Genetic algorithms in search, optimization, and machine learning. Addison-Wesley, Reading, MA.
18. Wardlaw R, Bhaktikul K (2004) Application of genetic algorithms for irrigation water scheduling. Irrig Drain 53(4):397–414. https://doi.org/10.1002/ird.121.
19. Yu Z, Fu W (2025) Optimization of nitrogen fertilization strategies for drip irrigation of cotton in large fields by dssat combined with a genetic algorithm. Appl Sci 15(7):3580. https://doi.org/10.3390/app15073580.
20. IPCC (2021) Climate change 2021: the physical science basis. contribution of working group i to the sixth assessment report of the intergovernmental panel on climate change. Technical report, Cambridge University Press (https://doi.org/10.1017/9781009157896.
21. Rawlings JB, Mayne DQ, Diehl M (2017) Model predictive control: theory, computation, and design, 2nd edn. Nob Hill Publishing, Madison, WI.
22. Azi A et al. (2025) Growth and yield of sweet corn (zea mays var. rugosa) in response to bio-n and different levels of inorganic fertilizers. Asian J Agric Hortic Res). Accessed 2026-03-08.
23. Lum MS, Aldam NFB, Chin CFS (2024). Accessed 2026-03-08 The effect of azolla (azolla pinnata) extract on the growth and yield of sweet corn under reduced nitrogen fertilization practice. SJA 40(Issue s1). https://doi.org/10.17582/journal.sja/2024/40/s1.101.109.
24. Liu Y, Wang N, Jiang C, Wang Y (2024) Effects of irrigation type and fertilizer application rate on growth, yield, and water and fertilizer use efficiency of silage corn in the north China plain.

PeerJ 12:18315. https://doi.org/10.7717/peerj.18315. Accessed 2026-03-08.
25. Khan MA, Marwat KB, Gul B, Wahid F, Khan H, Hashim S (2014) Pistia stratiotes l. (araceae): phytochemistry, use in medicines, phytoremediation, biogas and management options. Pak J Botany 46(3):851–860). Accessed 2026-03-08.
26. Dodd J, Sweby PK, Mayes S, Murchie EH, Karunaratne AS, Massawe F, Tindall MJ (2023) A multiscale mathematical model describing the growth and development of bambara groundnut. J. Theor. Biol 560:111373. https://doi.org/10.1016/j.jtbi.2022.111373.
27. Sawyer JE, Nafziger ED, Randall GW, Bundy LG, Rehm GW, Joern BC (2006) Nitrogen management for corn production. Iowa State Univ Ext Publ PM 1714.
28. Payero JO, Tarkalson DD, Irmak S, Davison DR, Petersen JL (2008) Effect of irrigation amounts applied with subsurface drip irrigation on corn evapotranspiration, yield, water use efficiency, and dry matter production in a semiarid climate. Agric Water Manag 95(8):895–908. https://doi.org/10.1016/j.agwat.2008.02.015.
29. Sánchez B, Rasmussen A, Porter JR (2014) Temperatures and the growth and development of maize and rice: a review. Glob Change Biol 20(2):408–417. https://doi.org/10.1111/gcb.12389.
30. Lewontin RC, Cohen D (1969) On population growth in a randomly varying environment. Proc Natl Acad Sci USA 62(4):1056–1060. https://doi.org/10.1073/pnas.62.4.1056.
31. Ren H, Liu M, Zhang J, Li P, Liu C (2022) Effects of agronomic traits and climatic factors on yield and yield stability of summer maize (zea mays L) in the huang-huai-hai plain in China. Front. Plant Sci 13:1042107. https://doi.org/10.3389/fpls.2022.1050064. Accessed 2026-03-08.
32. Stubbs CJ, Kunduru B, Bokros N, Verges V, Porter J, Cook DD, DeBolt S, McMahan C, Sekhon RS, Robertson DJ (2023) Moving toward short stature maize: the effect of plant height on maize stalk lodging resistance. Field Crops Res 300:109008. https://doi.org/10.1016/j.fcr.2023.109008. Accessed 2026-03-08.
33. Wang Y, Bao J, Wei X, Wu S, Fang C, Li Z, Qi Y, Gao Y, Dong Z, Wan X (2022) Genetic structure and molecular mechanisms underlying the formation of tassel, anther, and pollen in the male inflorescence of maize (zea mays l.). Cells 11(11):1753. https://doi.org/10.3390/cells11111753. Accessed 2026-03-08.
34. Silveira DL, Cargnelutti Filho A, Souza M, Trivisiol VS, Somavilla FM (2024) Adaptability and stability of grain yield and maize tassel traits. Cienc. Rural 54(9):20230319. 10.1590/0103-8478cr20230395. Accessed 2026-03-08.
35. US Department of Agriculture. (2023). 2023 acreage data as of august 9, 2023. Fsa crop acreage data reported to fsa, US Department of Agriculture Farm Service Agency.
36. Sengupta M, Xie Y, Lopez A, Habte A, Maclaurin G, Shelby J (2018) The national solar radiation data base (nsrdb). Renewable Sustain Energy Rev 89:51–60. https://doi.org/10.1016/j.rser.2018.03.003.
37. National Oceanic and Atmospheric Administration. (2024). Climate data online. https://www.ncdc.noaa.gov/cdo-web/.
38. Sawyer J, Nafziger E, Randall G, Bundy L, Rehm G, Joern B (2006) Concepts and rational for regional nitrogen rate guidelines for corn. University Extension Iowa State University Ames.
39. Abendroth LJ, Elmore RW, Boyer MJ, Marlay SK (2011) In corn growth and development. Iowa State Univ Ext Publ PMR 1009.
40. Schierbaum F (2004) Book reviews: corn: chemistry and technology, 2nd edition. By pamela j. white and lawrence a. johnson (editors). Starch Stärke 56(6):263–264. https://doi.org/10.1002/star.200490027.
41. Bayer Crop Science (2018) Height difference between traditional corn products and short stature corn hybrids. https://www.cropscience.bayer.us/articles/traits/height-difference-between-traditional-corn-products-and-short-stature-corn-hybrids. Accessed 2026-03-08.
42. Keller's Farmstand (2018) Sweet corn vs. Field corn height. https://kellersfarmstand.com/sweet-corn-vs-field-corn-height/. Accessed 2026-03-08.
43. Nanda D (2017) Number of corn leaves provides insight into maturity. https://www.farmprogress.com/corn/number-of-corn-leaves-provides-insight-into-maturity. Accessed 2026-03-08.
44. Nielsen RL (2025) Tassel emergence and pollen shed. https://www.agry.purdue.edu/ext/corn/news/timeless/tassels.html. Accessed 2026-03-08.
45. Urban Farmer Seeds (UF Seeds) (2024) Pencil cob corn seed. https://www.ufseeds.com/product/pencil-cob-corn-seed/COPCB.html. Accessed 2026-03-08.
46. Bayer Crop Science (2020) Corn growth stages and GDU requirements. https://www.cropscience.bayer.us/articles/bayer/corn-growth-stages-and-gdu-requirements. Accessed 2026-03-08.
47. University of California, Davis: Crop Development (2024) https://corn.ucdavis.edu/crop-development. Accessed 2026-03-08.
48. Sisson AJ, Mueller D, Conley SP, Gerber CK (2020) Corn growth stages. https://cropprotectionnetwork.org/web-books/crop-scouting-basics-for-corn-and-soybean?section=21-corn-growth-stages. Accessed 2026-03-08.
49. Western Cooperative Company (2024) Corn growth. https://www.westco.coop/pages/custom.php?id=33409. Accessed 2026-03-08.
50. Corn States (2018) Corn growth stages and GDU requirements. https://www.corn-states.com/app/uploads/2018/07/corn-growth-stages-and-gdu-requirements.pdf. Accessed 2026-03-08.
51. Eisenthal J (2025) Fertilizer prices put economic squeeze on farmers. https://www.farms.com/news/fertilizer-prices-put-economic-squeeze-on-farmers-235057.aspx. Accessed 2026-03-08.
52. National Corn Growers Association (2025) What is driving high production costs in corn? https://www.morningagclips.com/what-is-driving-high-production-costs-in-corn/. Accessed 2026-03-08.
53. Swanson K (2025) High production cost series Part 3. https://ncga.com/stay-informed/media/the-corn-economy/article/2025/08/high-production-cost-series-part-3. Accessed 2026-03-08.
54. Farmonaut (2025) Average cost to clear land, plant corn, farm insurance. https://farmonaut.com/usa/average-cost-to-clear-land-plant-corn-farm-insurance. Accessed 2026-03-08.
55. Newman M (2023) Irrigating more U.S. Crops by mid-century will Be worth the investment. https://natsci.msu.edu/news/2023-08-irrigating-more-us-crops-by-mid-century-will-be-worth-the-investment.aspx. Accessed 2026-03-08.
56. Rathore LS, Kumar M, McNider RT, Magliocca N, Ellenburg W (2024) Contrasting corn acreage trends in the midwest and southeast: the role of yield, climate, economics, and irrigation. J Agric Food Sys Community Dev Food Res 18:101373. https://doi.org/10.1016/j.jafr.2024.101373.
57. Berger A (2020) Estimating bushels of corn on the ground by counting ears prior to grazing with cattle. https://beef.unl.edu/estimating-bushels-corn-ground-counting-ears-prior-grazing-cattle/. Accessed 2026-03-08.
58. Groenke B (2018) Estimating corn yield potential. https://www.canr.msu.edu/news/estimating-corn-yield-potential. Accessed 2026-03-08.
59. Bayer Crop Science (2023) Estimating corn yield. https://www.cropscience.bayer.us/articles/channel/estimating-corn-yield. Accessed 2026-03-08.
60. Licht M, Clemens Z (2020) Corn yield estimates. https://crops.extension.iastate.edu/encyclopedia/corn-yield-estimates. Accessed 2026-03-08.
61. USDA Agricultural Marketing Service (2026) Iowa daily Cash grain bids. https://www.ams.usda.gov/mnreports/ams_2850.pdf. Accessed 2026-03-08.

62. Sucu E (2020) Optimizing plant density improves corn silage quality. https://dellait.com/optimizing-plant-density-improves-corn-silage-quality/. Accessed 2026-03-08.
63. Dhaliwal DS, Ainsworth EA, Williams MM (2021) Historical trends in sweet corn plant density tolerance using era hybrids (1930–2010s). Front. Plant Sci 12:685066. https://doi.org/10.3389/fpls.2021.707852. Accessed 2026-03-08.
64. Hanna Instruments (2024) Weight of corn. https://www.hanna-one.com/weight-corn.html. Accessed 2026-03-08.
65. Hommood MI, Isaak M (2020) Physical and Mechanical properties of sweet corn plant. Agric Eng Int: CIGR J 22(2):138–147). Accessed 2026-03-08.
66. Kranz WL, Irmak S, Donk SJ, Yonts CD, Martin DL (2008) Irrigation management for corn. Technical Report G1850, Univ Neb–Linc Ext. https://extensionpubs.unl.edu/publication/g1850/.
67. Hatfield JL, Prueger JH (2015) Temperature extremes: effect on plant growth and development. Weather Clim Extremes 10:4–10. https://doi.org/10.1016/j.wace.2015.08.001.
68. Liberzon D (2012) Calculus of variations and optimal control theory: a concise introduction. Princeton University Press, Princeton, NJ.
69. Bynum ML, Hackebeil GA, Hart WE, Laird CD, Nicholson BL, Siirola JD, Watson J-P, Woodruff DL (2021) Pyomo—optimization modeling in python, 3rd edn. Springer, ???. https://doi.org/10.1007/978-3-030-68928-5.
70. Wächter A, Biegler LT (2006) On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. Math Program 106(1):25–57. https://doi.org/10.1007/s10107-004-0559-y.
71. Shahriari B, Swersky K, Wang Z, Adams RP, Freitas N (2016) Taking the human out of the loop: a review of Bayesian optimization. Proc. IEEE 104(1):148–175. https://doi.org/10.1109/JPROC.2015.2494218.
72. Bergstra J, Bardenet R, Bengio Y, Kégl B. In: Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K.Q. (eds.) Advances in Neural Information Processing Systems, vol. 2011) Algorithms for hyper-parameter optimization, vol 24. Curran Associates, Inc..
73. Davies B, Coulter JA, Pagliari PH (2020) Timing and rate of nitrogen fertilization influence maize yield and nitrogen use efficiency. PLoS One 15(5):0233674. https://doi.org/10.1371/journal.pone.0233674.
74. Gouthami B, Ashok kumar B (2022) Water management in maize. Just Agriculture 2(7). Article ID: 059, e-ISSN: 2582-8223.
75. Roy RN, Finck A, Blair GJ, Tandon HLS (2006). Chap. 8. FAO document a0443e. Nutrient management guidelines for some major field crops. In: Plant Nutrition for Food Security: A Guide for Integrated Nutrient Management. FAO Fertilizer and Plant Nutrition Bulletin, pp. Food and Agriculture Organization of the United Nations, Rome, pp 235–274. https://www.fao.org/4/a0443e/a0443e04.pdf.
76. Mallarino A, Barbieri P, Pecinovsky K (2015). Winter 2015. Long-term phosphorus and potassium fertilization effects on yields of corn and soybean grown in rotation. Technical Report ISRF14-13, RFR-A14104, Iowa State Univ, Northeast Res Demonstration Farm, Nashua, IA. https://www.agronext.iastate.edu/soilfertility/info/NE%20farm%20report_PK%20Winter%202015-Long%20term%20P-K%20effects.pdf.
77. Robinson A (2025) USDA crop Report Continues large 2025 crop production estimates. https://www.iowafarmbureau.com/Article/USDA-Crop-Report-Continues-Large-2025-Crop-Production-Estimates. Iowa Farm Bureau. Accessed: 2026-03-08.
78. USDA Agricultural Marketing Service (2026) Iowa daily Cash grain bids. https://www.ams.usda.gov/mnreports/ams_2850.pdf. Livestock, Poultry and Grain Market News Division. Accessed: 2026-03-08.
79. Swanson K (2025) High production cost series, part 3. https://ncga.com/stay-informed/media/the-corn-economy/article/2025/08/high-production-cost-series-part-3. National Corn Growers Association. Accessed: 2026-03-08.