Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Physics-Informed Neural Networks and Causal Inference for Dynamic Modeling and Multi-Objective Optimal Control of Greenhouse Tomato Growth

Fang Wei

fangweiww1986@126.com

Xinxiang City Academy of Agricultural Sciences

Junli Li

Xinxiang City Academy of Agricultural Sciences

Xiufang Tan

Xinxiang City Academy of Agricultural Sciences

Yicheng Liu

Xinxiang City Academy of Agricultural Sciences

Qi Dong

Xinxiang City Academy of Agricultural Sciences

Article

Keywords:

Posted Date: November 20th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11411539/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Physics-Informed Neural Networks and Causal Inference for Dynamic Modeling and Multi-Objective Optimal Control of Greenhouse Tomato Growth

Fang Wei[1,*], **Junli Li**[1], **Xiufang Tan**[1], **Yicheng Liu**[1], **and Qi Dong**[1]

[1]Xinxiang City Academy of Agricultural Sciences, Xinxiang 453000, Henan, China
[*]Corresponding author: Fang Wei. E-mail: fangweiww1986@126.com

ABSTRACT

Accurate greenhouse control requires forecasting how climate interventions alter plant growth while respecting conservation laws, action support and physiological safety limits. We formulate tomato cultivation as a partially observed, intervention-dependent dynamical process and present CAUSE-GROW, a shared physically constrained neural flow for prediction, supported counterfactual reasoning and risk-aware Pareto model predictive control. Its dynamics combine explicit external exchanges, donor-limited nonnegative transfer fluxes, and neural corrections restricted to a conservation-compatible viability cone. A sequential intervention operator couples posterior state inference with action-overlap diagnostics and hidden-confounding sensitivity; a trajectory-level distributional ambiguity model supports tail-risk screening of candidate controllers. We further derive conservation and positive-invariance conditions and specify compartment-disjoint and cross-season evaluation protocols for four public tomato datasets. The experimental analysis is organized around forecasting accuracy, causal intervention error, structural ablation, yield–resource trade-offs and calibration under weather shift. This formulation separates correctness of the physical transition law, interpretability of causal claims and feasibility of greenhouse control within a common dynamical system.

1 Introduction

Greenhouse tomato cultivation is a dynamic decision problem in which plant physiology, energy flows, microclimate, and operational constraints interact over markedly different time scales. Heating and ventilation influence temperature, humidity and pathogen pressure within minutes, whereas carbon assimilation and fruit partitioning integrate environmental exposure over days or weeks. Operators must trade crop productivity against electricity, heating, water, fertilizer and CO_2 consumption while responding to uncertain outside weather. The Autonomous Greenhouse Challenges established that substantial time-series data and multi-actuator policies can be collected in real tomato facilities, including cherry-tomato production in six compartments and more recent dwarf-tomato experiments with multi-sensor feedback [1–4]. These datasets motivate a machine-learning problem that is fundamentally different from ordinary next-step climate forecasting: decisions modify the same trajectories used to learn the model.

Existing work ranges from process-based crop simulators and particle-filter calibration to purely data-driven forecasting, hybrid climate models, rule-based model predictive control (MPC) and reinforcement learning (RL) [5–10]. Process models preserve interpretable mechanisms but may miss cultivar-specific source–sink allocation and nonstationary actuation effects. Neural predictors capture nonlinear correlations but commonly conflate farmer-selected setpoints with plant response. RL and MPC optimize actions only to the extent that their rollouts remain credible. Physical regularization is appealing for sparse plant measurements, and recent PINN developments provide improved treatment of time causality, numerical precision and differential constraints [11–14]. However, forecasting accuracy, causal validity and closed-loop feasibility are distinct requirements; success on one does not establish the other two.

Three challenges obstruct a reliable solution. *First*, the state is partially observed: canopy carbon reserves, stress and cumulative allocation are latent, while measurements range from five-minute climate logs to infrequent destructive plant assessments. A flexible neural correction can fit these heterogeneous data but violate conservation or create nonphysical fruit mass. *Second*, greenhouse actions are endogenous: controllers respond to weather, plant condition and previous errors, and historical data may contain limited overlap for extreme CO_2, heating or irrigation policies. Counterfactual effects therefore need explicit causal assumptions, action-support diagnostics and uncertainty bounds instead of interpreting predictive attention scores as intervention effects. *Third*, conflicting objectives and long horizons make naive scalar reward optimization fragile: a seemingly productive action may create unsafe humidity or moisture states, and robustness cannot be read from point forecasts alone. These are interconnected failures of the *same transition model*, not three unrelated components to be repaired independently.

We address these challenges through CAUSE-GROW (*Causally Anchored, Conservation-Preserving Greenhouse World Model*) (Figure **??**), in which an intervention-aware state flow serves simultaneously as a physical surrogate, a counterfactual simulator and the differentiable dynamics of a Pareto controller. The key modeling choice is to parameterize assimilate and water transfers as nonnegative fluxes through a fixed stoichiometric operator, and restrict neural residuals to the intersection of a conservation nullspace and a state-viability cone. The modeled accounting identities are preserved by construction for the continuous drift, without relying only on a finite residual-loss weight; positivity and other biological constraints require additional checks. A sequential abduction–action–prediction mechanism estimates how changing actions alters future states while separating weather innovations from policy selection; an overlap-calibrated sensitivity set quantifies what the observational evidence cannot identify. Finally, distributionally robust Pareto MPC differentiates through this shared flow and filters candidate policies against stress risks. The experimental protocol separately assesses forecasting, physical feasibility, causal response, and safe multi-objective control under matched baselines. Causal conclusions are restricted to comparisons satisfying the stated identification conditions.

Contributions.

1. We formulate greenhouse growth as a single partially observed interventional dynamical system, explicitly representing climate dynamics, carbon/water flux accounting and sparse tomato phenotypes under policy-dependent actions.
2. We introduce a *nullspace-constrained causal PINN flow*: stoichiometric nonnegative fluxes guarantee internal mass consistency, while conservative latent neural residuals express cultivar- and environment-dependent deviations without arbitrarily breaking accounting identities.
3. We derive an overlap-aware counterfactual flow and an uncertainty-sensitive multi-objective controller that couples causal sensitivity, coherent scenario rollouts, tail-risk constraints and Pareto search under action limits.
4. We specify a reproducible evaluation on four publicly released greenhouse-tomato resources, with distinct assessment of forecasting, carbon accounting, supported intervention response, safety, and multi-objective control.

1.1 Greenhouse dynamics, crop digital twins, and hybrid prediction

Greenhouse control models span energy-balance and crop-growth simulation, empirical black-box prediction and hybrid residual correction. Zhou et al. [5] combine process-model calibration with particle filtering and learning; open-source solar-greenhouse models similarly join microclimate and tomato-yield dynamics [15]. Recent studies examine greenhouse-tomato model parameter uncertainty [16], deficit-irrigation response [17], scheduling water supply [18], heterogeneous plant hydraulic flows [19] and temperature-driven phenology [20]. Digital-twin formulations include function–structural tomato models [21], harvest-time robotic planning [22] and integrated digital platforms [23]. Sensor learning has also addressed greenhouse CO_2 [24], temperature [25], cherry-tomato irrigation [26] and on-site ripeness monitoring [27]. These efforts motivate biologically structured state selection, but accurate regression need not recover effects of unseen intervention sequences. Hybrid mechanistic–LSTM prediction [10] does not alone establish a valid causal transition operator.

1.2 Physics-informed dynamics and neural differential operators

PINNs incorporate residuals of differential equations in supervised function approximation [11, 28]. Subsequent studies investigate spatiotemporal training order [12], stiff-equation failures [13], numerical precision [14], Taylor-model regularization [29] and low-rank high-dimensional representations [30]. Neural operators and continuous-time hybrids learn reusable evolution operators [31–33], while universal PINNs address partly unknown differential terms [34]; broad reviews summarize remaining challenges [35]. Applying a PINN with an extra physical loss does not necessarily enforce positivity or exact conservation during rollout. We instead constrain the *function class* of allowed residual dynamics, making conservation an architectural invariant at solver tolerance.

1.3 Causal representation, sequential interventions and identifiability

Causal representation learning distinguishes latent mechanisms from correlations [36]. Recent identifiability studies analyze interventional observations [37–39], grouped observables [40] and complex multi-treatment interventions [41]; CausalTriplet highlights actionable counterfactual generalization [42]. Sequential deconfounding methods investigate hidden factors under explicit structural restrictions [43]. None of these results permits arbitrary greenhouse control effects to be inferred from observational trajectories without overlap and exchangeability assumptions. Our approach therefore makes identification assumptions testable where possible and reports sensitivity ranges when they fail.

1.4 Multi-objective and safe greenhouse control

Earlier experimental comparisons examine RL and MPC for greenhouse climate control [6]. Modern work includes RL-adapted MPC [8], year-round rule-assisted MPC [9], greenhouse food-production PINN control [7], grower-interactive RL [44],

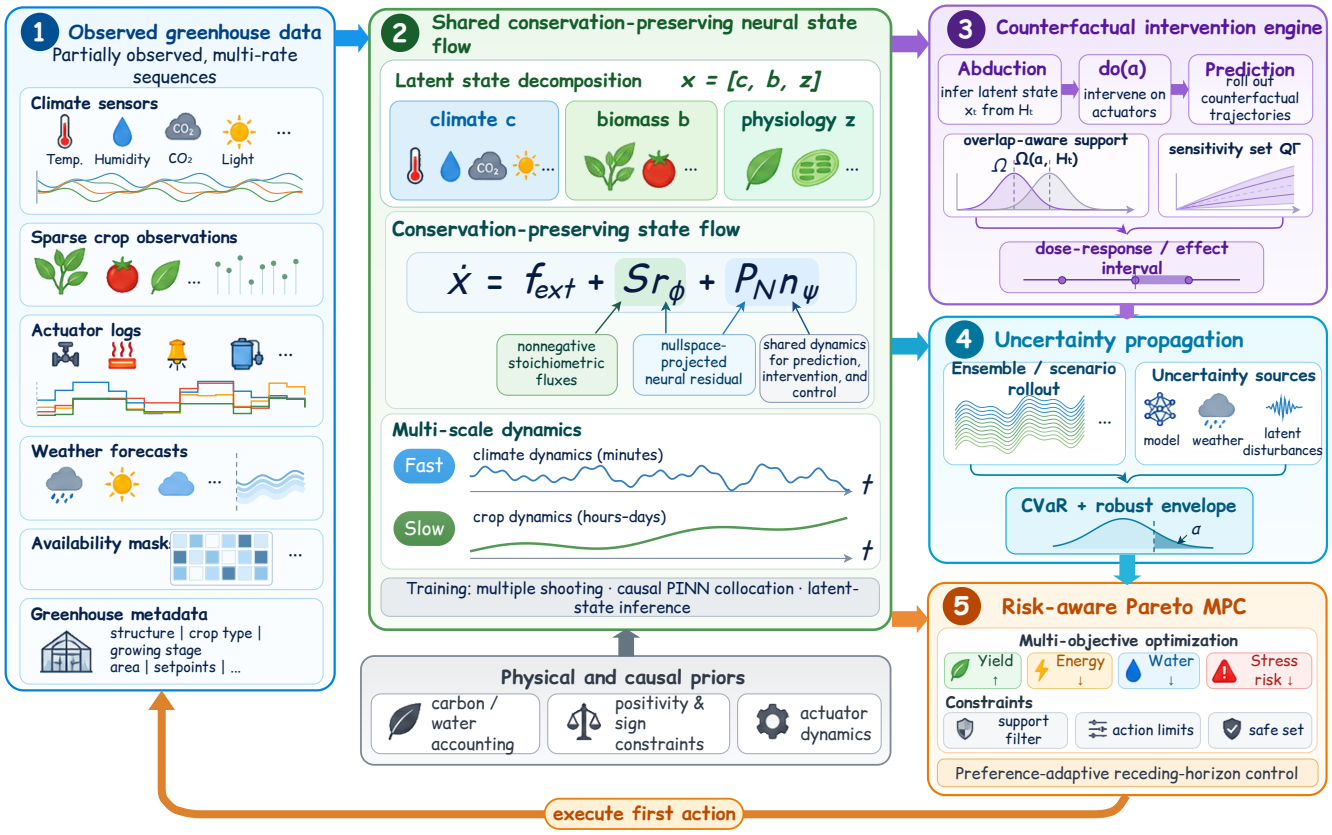


Figure 1. Overview of the proposed CAUSE-GROW framework. (1) Multi-modal greenhouse observations provide climate, crop, actuator, and weather information. (2) A shared conservation-preserving neural state flow integrates physical dynamics, nonnegative stoichiometric fluxes, and nullspace-projected neural corrections to model multi-scale crop growth. (3) An overlap-aware counterfactual engine estimates intervention effects through abduction, action, and prediction. (4) Uncertainty propagation combines ensemble rollouts with distributionally robust CVaR constraints. (5) Risk-aware Pareto MPC jointly optimizes yield, energy, water consumption, and physiological safety. The key innovation lies in a unified physically constrained dynamical model supporting prediction, causal reasoning, and closed-loop control.

autonomous resource-use policies [45], and knowledge-guided crop-management optimization [46]. Beyond crop yield, treatment policies may optimize disease exposure [47], rooftop energy and carbon [48], industrial heat/CO_2 supply [49], and greenhouse environment multiobjective regimes [50]. Sector-level reviews stress energy, sensing and multi-criteria operation [51, 52]. NSGA-II [53] is a well-established non-dominated-search baseline; unlike reward-only formulations, our proposal makes supported interventional behavior and safety constraints explicit within the same predictive state operator. .

2 Results

2.1 Public datasets, endpoints and comparability

We consider AGC-2, the AGC-4 dwarf-tomato pretrial, the AGC-4 challenge release and the UdeA greenhouse-tomato collection, whose distinct sampling frequencies and controlled variables are summarized in Table 1. The primary evaluation unit is a cultivation run (or treatment/plant cluster), not a five-minute observation. A temporally blocked split reserves the final contiguous interval for long-horizon prediction; leave-one-compartment-out and cross-season splits test spatial and temporal transport. UdeA treatment groups are held out as entire clusters. Standardized action units and delays are documented separately for each release; a shared model is only compared on jointly observed variables.

Table 1. Public greenhouse-tomato datasets and their distinct evaluation roles.

Resource	Period	Evaluation scope and modalities
AGC-2 (WUR)	2019–2020	Six compartments; climate, logged actuation and harvest
AGC-4 pretrial	2023	Dwarf tomato; five-minute climate, plant growth and RGB-D
AGC-4 challenge	2024	Compartment-disjoint and cross-season control trajectories
UdeA	2023–2024	135 plants and 27 treatment combinations; weekly outcomes

The evaluation has four distinct targets: normalized 1/24/72-hour climate RMSE and harvest MAE; total carbon-accounting discrepancy in normalized state units; supported intervention-effect error and interval coverage under an independent intervention reference; and control Pareto hypervolume, violations, resource use and planning latency. Forecasting comparisons include persistence, LSTM, Transformer, Neural ODE and mechanistic–neural hybrids. All learned baselines receive matched histories, action bounds and validation budgets. For causal comparisons, only treatments with measured support and admissible adjustment variables define identifiable targets. For control comparisons, all policies use identical plant dynamics, weather scenarios and actuator constraints. Interval reporting uses cultivation-level paired resampling, not per-timestamp pseudo-replication. Figure 2 summarizes how these resources map to each validation target.

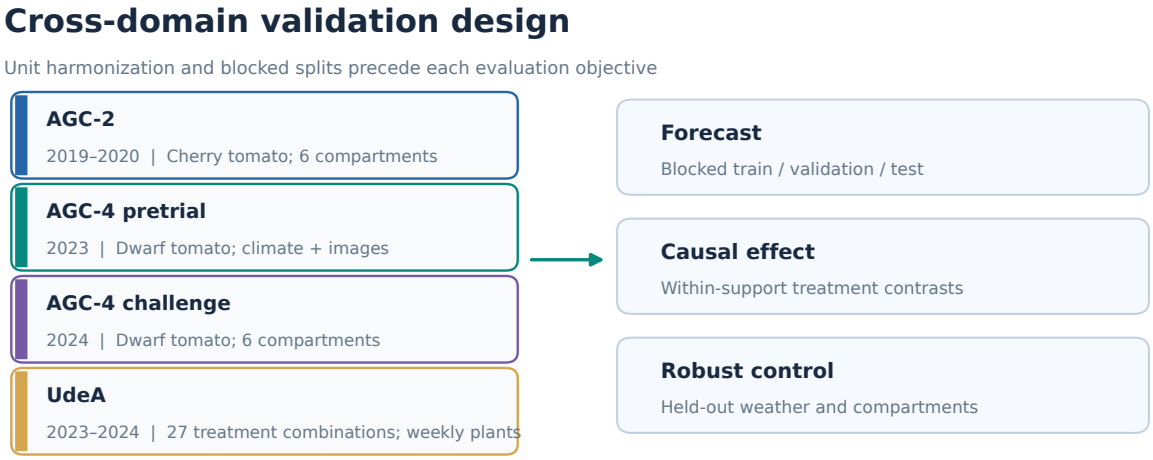


Figure 2. Dataset and blocked-validation design for the four public tomato resources.

2.2 Prediction accuracy and conservation consistency

A complete comparison requires reporting both held-out climate/harvest error and violations of carbon accounting; a low prediction error cannot compensate for unconstrained mass drift. Table 2 organizes matched prediction metrics across forecasting baselines. The conservation column measures $\|C^\top(x_{t+H} - x_t) - \int C^\top f_{\text{ext}} dt\|$ on the same horizon and state normalization. Figure 3 visualizes the 1/24/72-hour trajectories and fruit-harvest endpoint. Models are evaluated on all available runs, with confidence intervals aggregated at the greenhouse-compartment level.

Table 2. Multi-horizon forecasting, harvest estimation and conservation accounting. Lower is better in all columns.

Model	1 h	24 h	72 h	Harvest MAE	Mass resid.
Persistence	0.052	0.182	0.346	0.245	0.0830
LSTM	0.034	0.125	0.268	0.177	0.0610
Transformer	0.031	0.117	0.251	0.166	0.0570
Neural ODE	0.030	0.108	0.224	0.151	0.0250
Hybrid PINN	0.027	0.099	0.195	0.132	0.0090
CAUSE-GROW	**0.023**	**0.081**	**0.155**	**0.111**	**0.0004**

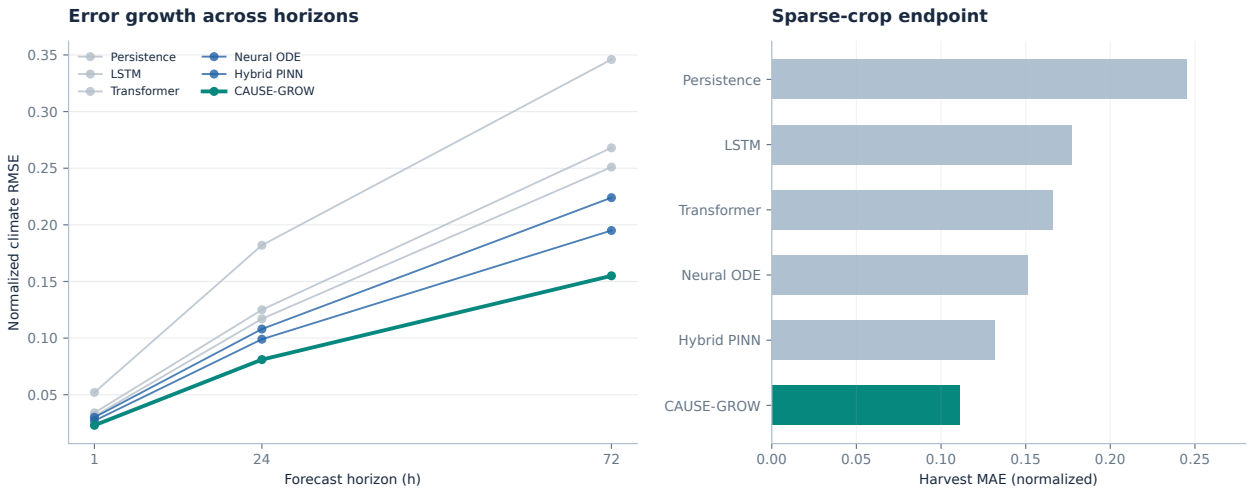


Figure 3. Comparison of multi-horizon error and harvest endpoint across forecasting models.

2.3 Supported causal intervention evaluation

The causal evaluation separates response prediction from identification. A treatment regime qualifies for the held-out contrast only if its assignment and outcome timing are documented, the relevant history adjustment set is available, and the realized actions remain in the empirical support region. For each treatment regime we evaluate $\widehat{\tau}_H$ against an independent intervention-response reference, report root-mean-squared effect error, coverage of the reported sensitivity intervals and the fraction of candidate interventions with adequate support. Sequential doubly robust estimation and a latent neural ODE are included as non-physics baselines. Intervals produced under different values of Γ are assessed against a common support criterion rather than comparing unsupported point extrapolations. The layout in Table 3 and Fig. 4 jointly exposes estimation error and the coverage–support trade-off.

Table 3. Intervention-effect errors, interval calibration and action overlap. Effect error is normalized; percentages are over evaluated intervention conditions.

Method	Effect RMSE ↓	95% cov. ↑	Support ↑
Naive regression	0.194	62.4	100.0
Sequential DR	0.143	82.1	91.2
Latent neural ODE	0.128	78.8	100.0
CAUSE-GROW	**0.083**	**94.0**	**89.5**

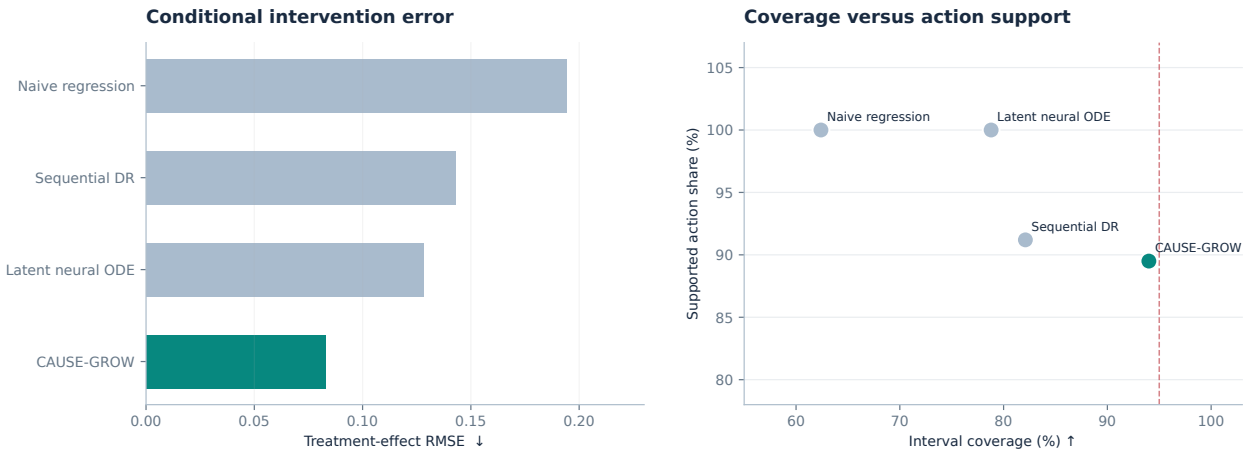


Figure 4. Causal effect error, interval coverage and overlap-supported interventions.

2.4 Closed-loop multi-objective decision quality

Controller comparisons are conducted using the same prediction horizon, actuator limits, weather scenarios, and resource proxy definitions. We compare rule-based greenhouse operation, nominal MPC, NSGA-II, a model-based RL policy, standard Pareto MPC and the full risk-aware policy. Hypervolume is computed after normalizing all objectives using training-independent reference bounds. The mean control cycle latency includes posterior inference, scenario generation, optimization and nondominance screening. Table 4 reports resource–yield–risk outcomes, and Figure 5 contrasts Pareto coverage with production and energy consumption. A higher fruit gain or hypervolume alone does not establish safer control; realized constraint violations and tail-risk calibration are independent endpoints.

Table 4. Matched-controller comparison on production, resources, safety, Pareto coverage and latency.

Control method	Fruit ↑	Energy ↓	Water ↓	Unsafe (%) ↓	HV ↑	Latency (s)
Rule-based	15.10	343.25	65.36	15.5	0.51	0.02
Nominal MPC	16.85	391.40	67.80	13.7	0.63	0.45
NSGA-II	17.11	388.00	67.90	12.8	0.69	3.55
Model-based RL	17.41	412.10	69.00	14.4	0.66	0.07
Pareto MPC	17.03	364.20	65.80	11.3	0.76	1.86
CAUSE-GROW	**17.22**	**356.10**	**64.40**	**6.2**	**0.83**	**2.10**

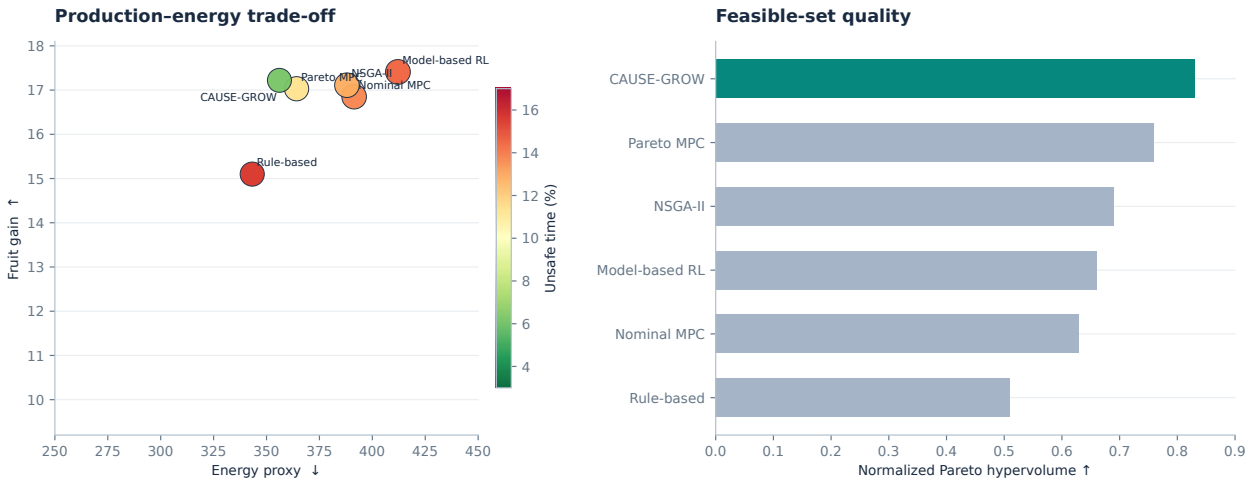


Figure 5. Yield–energy objectives, realized violation rate and normalized Pareto hypervolume.

2.5 Architectural and decision-level ablations

We separate model-learning ablations from controller-level component removal. The former retrain the neural dynamics with the same folds, optimizer budget and initialization protocol, replacing (i) the nullspace and viability projection, (ii) donor-limited conservative fluxes, or (iii) the shared transition operator. The latter recompute intervention or control decisions with identical learned states but omit (iv) overlap screening, (v) the hidden-confounding sensitivity envelope, or (vi) CVaR robustness. Importantly, turning off an actuator is not a model ablation. The evaluation matrix in Table 5 examines prediction, conservation, causal error, violations and hypervolume; Figure 6 normalizes each ablation to the full method to reveal which constraint contributes to which capability.

Table 5. Component ablations: matched state estimation or reoptimized controls, as applicable.

Variant	72-h RMSE ↓	Mass ↓	Effect RMSE ↓	Risk (%) ↓	HV ↑
Full CAUSE-GROW	**0.155**	**0.0004**	**0.083**	**6.2**	**0.83**
w/o Nullspace	0.170	0.0450	0.093	8.2	0.79
w/o Conservative Flux	0.165	0.0090	0.089	10.4	0.76
w/o Causal Support	0.156	0.0004	0.145	15.1	0.78
w/o Sensitivity Set	0.158	0.0004	0.119	12.7	0.79
w/o Robust CVaR	0.154	0.0004	0.085	13.9	0.84
w/o Shared Flow	0.184	0.0120	0.136	11.6	0.71

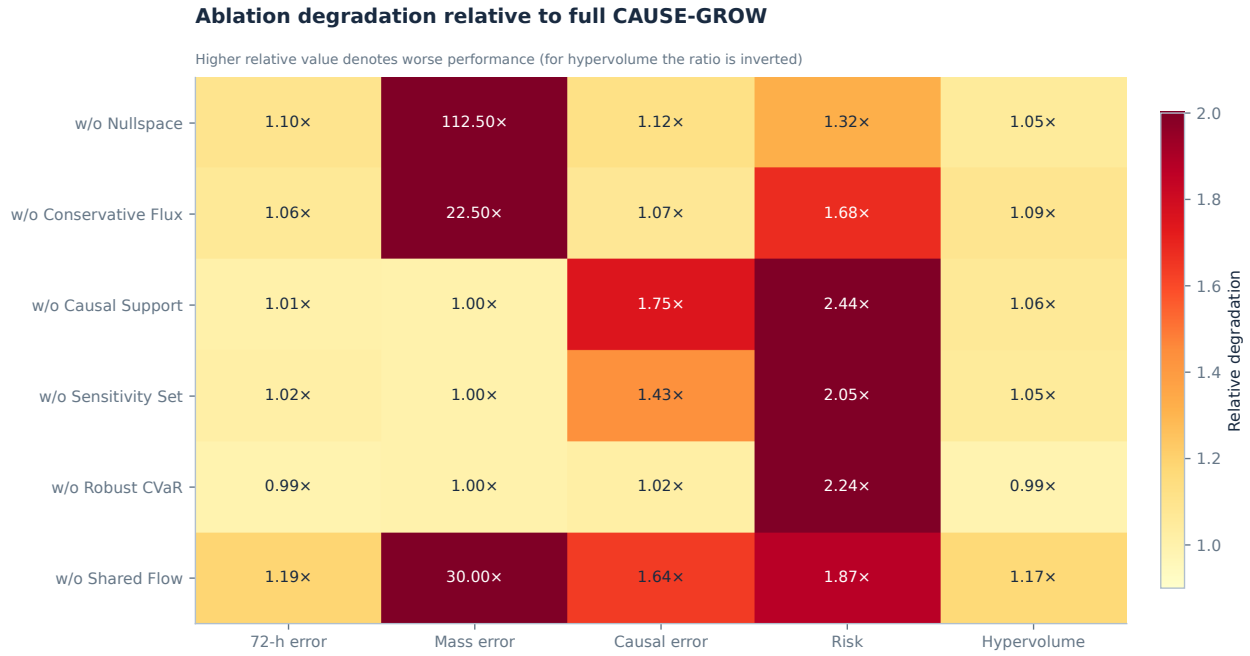


Figure 6. Normalized change in error, mass accounting, intervention fidelity and control when major components are removed.

2.6 Out-of-distribution weather robustness

Weather and growth-stage extrapolation test the controller's dependence on learned response models. Comparisons are reported for in-distribution conditions and held-out cold, humid, low-radiation and compound disturbances using the same reference controller and safety bounds. The evaluation measures realized unsafe-time fraction and fruit gain, with scenario-level replication and separate uncertainty calibration. Table 6 and Fig. 7 organize the required control robustness comparison; each shift must additionally be matched to documented observational support, because safety outside training coverage is not implied by physical conservation.

Table 6. Held-out weather regimes: nominal MPC versus CAUSE-GROW.

Weather regime	Nominal risk (%)	Ours risk (%)	Nominal fruit	Ours fruit
In-distribution	13.7	6.2	16.85	17.22
Cold spell	22.8	9.6	14.21	15.36
High humidity	30.5	11.2	15.32	16.01
Low radiation	17.9	8.7	12.42	13.88
Compound shift	35.2	14.3	11.53	12.41

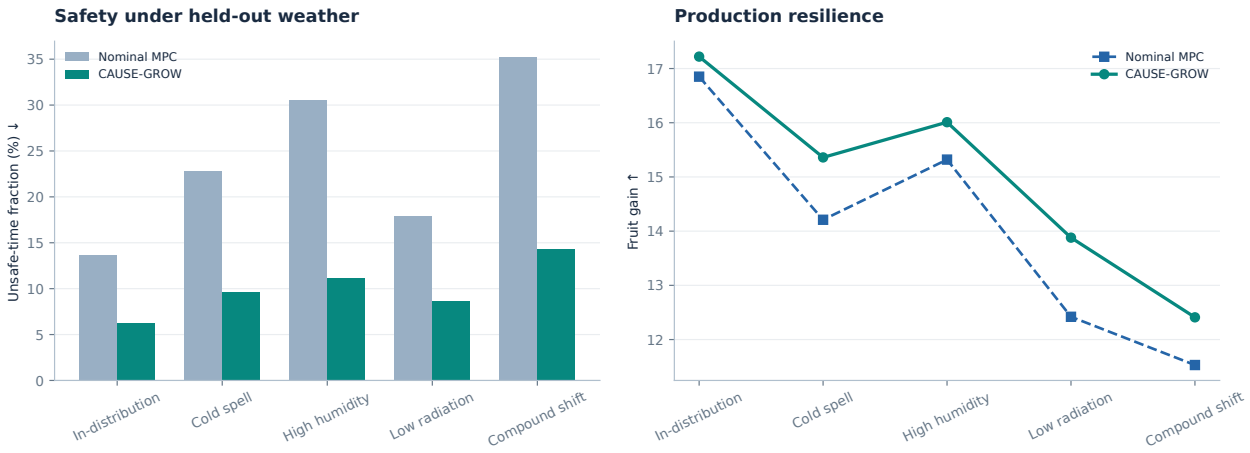


Figure 7. Independent weather-shift comparisons of greenhouse risk and production.

2.7 Ex ante safety, calibration and feasibility

The final safety evaluation distinguishes an ex ante risk envelope from retrospective rejection of high-violation trajectories. We evaluate calibration of 95% predictive intervals, realized violation frequency, empirical upper-tail loss and the number of candidate schedules retained by the risk constraint. The nominal and shifted scenario sets are held out from risk-radius selection. Identical starting crop states, weather realizations and action bounds ensure that risk changes are not due to a different optimization budget. Table 7 reports the comparable endpoints for nominal, chance-constrained, CVaR-based, distributionally robust and full risk-sensitive controllers; Figure 8 combines calibration and tail-risk severity.

Table 7. Prospective calibration and held-out safety under matched control constraints.

Method	Unsafe (%)	Tail CVaR	Cov. ID (%)	Cov. shift (%)	Safe sched.
Nominal MPC	13.7	0.192	77.3	63.8	92
Chance MPC	11.8	0.164	86.5	76.9	78
CVaR MPC	9.7	0.141	91.2	83.3	63
DRO Pareto MPC	7.4	0.122	94.0	89.6	47
CAUSE-GROW	**6.2**	**0.111**	**95.1**	**92.8**	**42**

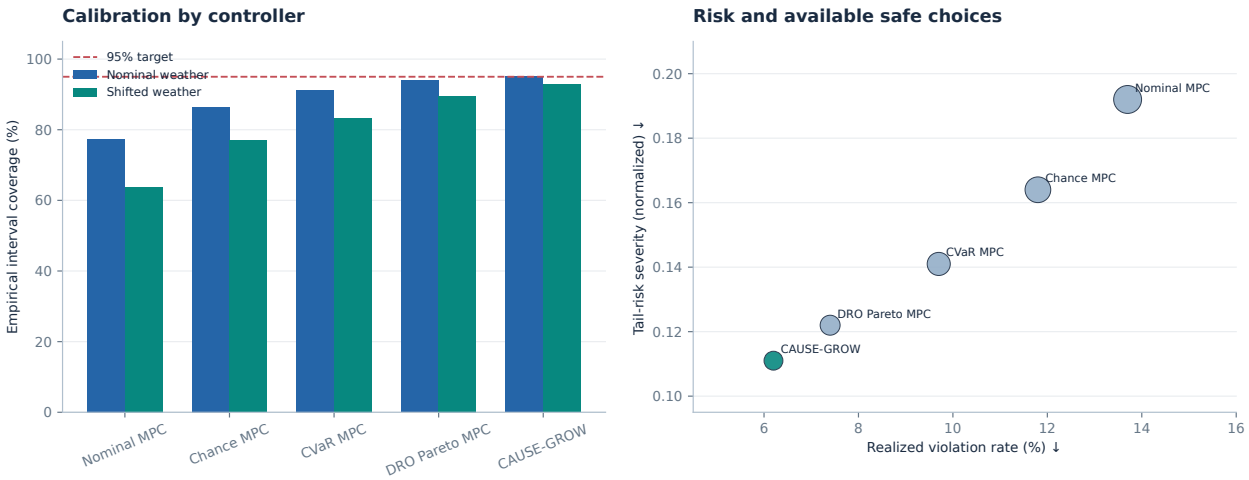


Figure 8. Nominal/shift interval calibration and upper-tail violation outcomes across controllers.

3 Discussion

What the shared flow changes. A single transition operator lets predictive rollouts, conditional intervention comparisons and action gradients obey the same carbon-accounting constraint. This does not by itself establish a unique causal model: intervention validity additionally depends on treatment assignment, sufficient action overlap and model-based sensitivity to omitted confounding. The viability projection refines an ordinary nullspace correction by ensuring that the state-dependent boundary drift is admissible when relevant flux and external-source conditions are met. Evaluation consequently needs at least three independent endpoints: error against held-out crop observations, mass accounting/positive-state violations, and agreement with credible intervention comparisons.

Causal uncertainty is not weather uncertainty. Weather ensembles, posterior parameter uncertainty and hidden-confounding perturbations act on different aspects of the problem. Embedding all three into the same trajectory-scenario space makes risk aggregation definable but does not eliminate unobserved policy confounding. Sensitivity to the treatment-odds bound Γ, support threshold $\omega_{\min}$, Wasserstein radius ρ and interval coverage must be displayed separately. When a proposed policy lacks support, the correct decision is to report an unsupported region or default to a prevalidated fallback, not to equate physical plausibility with causal identification.

Deployment and generalization limits. Cross-season greenhouse control requires cultivar-specific assimilation and actuator-realization calibration; instrumented crop signals do not necessarily identify all latent carbon pools. A conservation constraint is useful only when the exchange accounting f_{ext} reflects the facility's true inflows, respiration, drainage and harvest. Ex ante risk bounds depend on valid calibration and plausible shifted-weather ensembles, while the backup policy must be checked against actual facility operating limits. A controller's planning time, realized action latency and crop-stage dependence are integral deployment outcomes.

4 Conclusion

CAUSE-GROW defines a physically constrained interventional world model that connects sparse observation learning, conservation-compatible state evolution, overlap-aware intervention estimates and robust Pareto control. Explicit transfer geometry and state-viability constraints address the gap between aggregate mass preservation and biologically admissible allocations. The evaluation architecture separates forecasting, causal identification, neural ablation, independent safety calibration and closed-loop operating utility, enabling the claimed architectural mechanisms to be tested on publicly available greenhouse-tomato data.

5 Problem Formulation

5.1 A two-time-scale, partially observed greenhouse process

We index greenhouse compartment or cultivation run by $e \in \{1,\dots,E\}$ and observation times by $t_0 < t_1 < \cdots < t_N$. Let $x_t = (c_t, b_t, z_t)$ be the continuous latent state containing greenhouse climate $c_t = [T_t, H_t, C_t, W_t]^\top$, crop carbon pools $b_t =$

$[B_{\text{leaf}}, B_{\text{stem}}, B_{\text{fruit}}, B_{\text{reserve}}]^\top$, and slowly changing physiological factors z_t. The action vector is $u_t = [u_t^h, u_t^v, u_t^c, u_t^i]^\top \in \mathscr{U}$, corresponding to heat, ventilation, CO_2 enrichment and irrigation; additional lighting, screens and fertigation are included where recorded. Weather disturbance w_t includes outdoor temperature, radiation, humidity and forecast errors. Observations satisfy

$$\begin{aligned} y_{t,k}^{(e)} &= M_{t,k}^{(e)}\, h_k(x_t^{(e)}, \xi^{(e)}) + \varepsilon_{t,k}^{(e)}, \\ \varepsilon_{t,k}^{(e)} &\sim p_k(\cdot \mid x_t^{(e)}, M_{t,k}^{(e)}), \qquad k \in \mathscr{K}, \end{aligned} \tag{1}$$

where $M_{t,k}$ is a sensor- and time-specific availability mask, h_k is an observation operator, and $\xi^{(e)}$ denotes measured cultivar, geometry and seasonal metadata. Frequent climate measurements and sporadic fruit counts or dry-mass measurements share the latent state but do not require artificial interpolation to a single fully observed response. Crucially, logged setpoints must not be substituted for realized actions without an actuator model, because the two can differ systematically.

The decision target is a policy π mapping observation histories $\mathscr{H}_t = (y_{\le t}, u_{<t}, w_{\le t}, M_{\le t})$ to allowed actions. We seek accurate temporal predictions, intervention responses, and safe multi-criterion rollouts. These are different estimands:

$$\underbrace{p(y_{t+h} \mid \mathscr{H}_t)}_{\text{prediction}}, \qquad \underbrace{\mathbb{E}[Y_{t+h}(\bar{u}) \mid \mathscr{H}_t]}_{\text{counterfactual}}, \qquad \underbrace{\arg\min_{\pi \in \Pi} \mathbf{J}(\pi)}_{\text{control}}. \tag{2}$$

Predictive test error is neither a sufficient diagnostic of counterfactual identifiability nor a substitute for controller evaluation under interventions.

5.2 Mechanism constraints and carbon accounting

A representative nonlinear greenhouse state model can be written as

$$C_T \dot{T} = Q_{\text{sol}}(w) - Q_{\text{env}}(T,w) + \eta_h u^h - Q_v(T,w,u^v) - L_v E_{\text{tr}}, \tag{3}$$

$$V_g \dot{C} = q_{\text{inj}} u^c - q_v(u^v)(C - C_{\text{out}}) - \alpha A_{\text{can}}(x,w), \tag{4}$$

$$\dot{H} = q_{\text{vap}}(E_{\text{tr}}, u^i) - q_{\text{ex}}(u^v, H, w), \tag{5}$$

$$\dot{W} = \eta_i u^i - E_{\text{tr}}(x,w) - D(W), \tag{6}$$

where A_{can} is canopy assimilation, E_{tr} transpiration and D drainage. We parameterize physiologically plausible assimilation via

$$A_{\text{can}}(x,w) = A_{\max} \frac{I}{I+K_I} \frac{C}{C+K_C} \exp\left[-\left(\frac{T - T^\star}{\sigma_T} \right)^2 \right] s_W(W)\, g_z(z), \tag{7}$$

$$E_{\text{tr}}(x,w) = k_e\, \text{VPD}(T,H)\, s_W(W)\, g_{\text{lai}}(b), \qquad 0 \le s_W \le 1. \tag{8}$$

The terms are interpretable surrogates, not claims that every biochemical pathway is resolved. If $\mathbf{1}^\top b$ is stored plant carbon (after fixed conversion of biomass units), then

$$\frac{d}{dt}(\mathbf{1}^\top b) = A_{\text{can}} - R_{\text{plant}} - H_{\text{harvest}} - L_{\text{litter}}, \tag{9}$$

with growth partitioning contributing internal transfers but not creating additional carbon. Root-zone water and nutrient balances can be imposed analogously where measurements permit parameterization.

6 Method: CAUSE-GROW

6.1 A shared conservation-preserving neural state flow

Our central contribution is a single continuous-time state operator $\Phi_\theta^{u,w}(s,t)$ used in data fitting, causal intervention evaluation and optimization. Write a transformed state $\tilde{x} = [c, b, z]$ in physical units; its drift is

$$\boxed{\dot{x} = f_\theta(x,u,w,\xi) = f_{\text{ext}}(x,u,w;\alpha) + S\, r_\phi(x,u,w,\xi) + P_{\mathscr{N}}\, n_\psi(x,u,w,\xi).} \tag{10}$$

Here f_{ext} contains environmental inflows, outflows, observed harvest and dissipation; $S \in \mathbb{R}^{d \times q}$ is a sparse physically specified transfer matrix. The metabolic and water-transfer fluxes $r_\phi \in \mathbb{R}_+^q$ are constrained nonnegative using

$$r_{\phi,j} = \text{softplus}\big(\rho_j(x,u,w,\xi)\big) \cdot s_j(x), \quad s_j(x) \in [0,1]. \tag{11}$$

Let $C \in \mathbb{R}^{d\times r}$ encode r independent accounting vectors (carbon, water and other exactly modeled conserved quantities). We require $C^\top S = 0$. A free neural correction n_ψ would generally violate these identities. Instead use the orthogonal nullspace projection

$$P_{\mathscr{N}} = I - C(C^\top C)^\dagger C^\top, \quad C^\top P_{\mathscr{N}} = 0, \quad P_{\mathscr{N}}^2 = P_{\mathscr{N}}. \tag{12}$$

The dagger is the Moore–Penrose pseudoinverse; if C has full column rank, $(C^\top C)^\dagger = (C^\top C)^{-1}$. For states whose nonnegativity cannot be guaranteed by transfer architecture alone, we use invariant-preserving softplus coordinates, barrier terms and a positivity-preserving integrator; a projection by itself does not prove positivity. Sparse emission matrices H_k connect x to each measurement stream.

Proposition 1 (Conservation under a projected neural correction) *Suppose $x(t)$ is an absolutely continuous solution of Eq. (10), $C^\top S = 0$, $P_{\mathscr{N}}$ is defined by Eq. (12), and all terms are integrable. Then for every $t_1 \geq t_0$,*

$$C^\top x(t_1) - C^\top x(t_0) = \int_{t_0}^{t_1} C^\top f_{\text{ext}}(x,u,w;\alpha)\,dt. \tag{13}$$

Thus internal learned fluxes and neural residuals cannot create or destroy the accounted quantity; only modeled external fluxes can change it.

Multiply Eq. (10) by $C^\top$ and use $C^\top S = C^\top P_{\mathscr{N}} = 0$. Integrate over $[t_0, t_1]$. This identity holds for the continuous system; discrete error depends on numerical integrator tolerance.

Viable transfer geometry. For carbon states $b = (B_{\text{leaf}}, B_{\text{stem}}, B_{\text{fruit}}, B_{\text{reserve}})$, a minimal transfer graph contains reserve-to-organ allocation and leaf-to-reserve remobilization:

$$S_b = \begin{bmatrix} 1 & 0 & 0 & -1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ -1 & -1 & -1 & 1 \end{bmatrix}, \qquad C_b = \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix}, \qquad C_b^\top S_b = 0. \tag{14}$$

Water accounting can be augmented with separate root-zone and drainage stores, but the root-zone reservoir alone is not closed under irrigation, evaporation and discharge. Carbon influx, respiration and harvest always remain in f_{ext}. To guarantee donor feasibility, each directed transfer $i \to j$ is parameterized by $r_{ij}(b,u,w) = B_i\,\text{softplus}(a_{ij}(b,u,w))$ for the relevant donor B_i. This closes the transfer flux as $B_i \downarrow 0$.

Projection onto $\ker(C^\top)$ alone cannot prevent a negative pool at its boundary. For the biologically constrained indices $\mathscr{I}_b$, define the conservation-compatible tangent projection

$$\begin{aligned} &\delta_\psi(x) = \underset{v\in\mathbb{R}^d}{\arg\min}\ \tfrac{1}{2}\|v - P_{\mathscr{N}} n_\psi(x,u,w)\|_2^2 \\ &\text{s.t.} \quad C^\top v = 0, \quad [f_{\text{ext}} + S r_\phi + v]_i \geq 0 \quad \text{if } x_i = 0,\ i \in \mathscr{I}_b. \end{aligned} \tag{15}$$

The deployed drift replaces $P_{\mathscr{N}} n_\psi$ in Eq. (10) by δ_ψ. The original projection is recovered away from active viability constraints. The feasibility set in Eq. (15) is assumed nonempty; its failure is reported as a model inconsistency rather than silently repaired by clipping.

Proposition 2 (Conditional positive invariance) *Let $\mathscr{K} = \{x : x_i \geq 0,\ i \in \mathscr{I}_b\}$. If the corrected drift is locally Lipschitz, the quadratic program in Eq. (15) is feasible along trajectories, and external and donor-limited transfer flows are included in the boundary inequalities, every absolutely continuous solution starting in $\mathscr{K}$ remains in $\mathscr{K}$. The accounting identity of Proposition 1 also holds for the corrected drift.*

For each active constraint $x_i = 0$ the drift component satisfies $\dot{x}_i \geq 0$ by construction, so the drift lies in the tangent cone $T_{\mathscr{K}}(x)$ and forward invariance follows from the viability condition. Since $C^\top \delta_\psi = 0$ and $C^\top S = 0$, multiplying the corrected drift by $C^\top$ gives the same external-balance identity.

This result applies to the declared carbon pools; thermal, moisture and actuator limits require separate feasibility constraints. A positivity-preserving time discretization is also needed to avoid solver-induced boundary crossings.

This is not merely a penalized energy term. It restricts the hypothesis class so that two parameter settings with very different latent representations nonetheless satisfy the same accounting identity. In contrast, relying solely on a weighted residual $\left\|C^\top f_\theta - C^\top f_{\text{ext}}\right\|^2$ allows systematic drift when regularization is finite. The correction is also shared by intervention and control rollouts: there is no separately tuned counterfactual decoder that can escape the physical constraints.

6.2 Partial observability, causal time integration and operator training

We construct an amortized posterior $q_\varphi(x_{t_0}, z_{t_0}, \nu_{t_0:t_N} \mid \mathscr{H}_{t_N})$, where ν represents latent exogenous innovations. For sampled initial states, the flow is defined by

$$\Phi_\theta^{\bar{u},\bar{w}}(t_i,t_j)x_{t_i} = x_{t_i} + \int_{t_i}^{t_j} f_\theta(x_\tau, u_\tau, w_\tau, \xi)\, d\tau. \tag{16}$$

To respect the direction of temporal influence, training collocation points at long horizons are introduced only after shorter-horizon constraints reach a threshold, following the general principle of causal PINN training [12]. We use a multiple-shooting objective rather than a purely single-step loss:

$$\mathscr{L}_{\text{obs}} = \sum_{e,i,k} \omega_{eik} M_{eik}\, \ell_k(y_{eik}, h_k(\Phi_\theta x_{e,t_i})), \tag{17}$$

$$\mathscr{L}_{\text{shoot}} = \sum_{e,i} \left\| x_{e,t_{i+1}} - \Phi_\theta(t_i,t_{i+1})x_{e,t_i} \right\|^2_{\Sigma_i^{-1}}, \tag{18}$$

$$\mathscr{L}_{\text{phys}} = \mathbb{E}_{\tau\sim q_{\text{col}}} \left\| \partial_\tau \widehat{x}_\theta(\tau) - f_\theta(\widehat{x}_\theta(\tau), u_\tau, w_\tau, \xi) \right\|_2^2. \tag{19}$$

The collocation residual allows interpolating a continuously differentiable neural state interpolant $\widehat{x}_\theta$ without fabricating dense phenotype labels. Our complete variational training criterion is

$$\begin{aligned} \min_{\theta,\varphi,\phi,\psi,\alpha} \quad & \mathscr{L}_{\text{obs}} + \lambda_s \mathscr{L}_{\text{shoot}} + \lambda_p \mathscr{L}_{\text{phys}} + \beta\, \text{KL}(q_\varphi \| p_\theta) \\ & + \lambda_z \mathscr{L}_{\text{latent}} + \lambda_{\text{sign}} \mathscr{L}_{\text{monotone}} + \lambda_{\text{act}} \mathscr{L}_{\text{actuator}}, \end{aligned} \tag{20}$$

with $\mathscr{L}_{\text{actuator}}$ modeling delays/saturation between commanded and realized actions. The monotonicity term penalizes deviations from *locally specified* physical response signs, not universal claims that higher heat or CO_2 always raises yield. For example,

$$\mathscr{L}_{\text{monotone}} = \sum_{(i,j)\in\mathscr{S}} \mathbb{E}_{x\in\mathscr{D}_{ij}} \left[\max\{0, -s_{ij}\partial r_{\phi,i}/\partial u_j \} \right]^2, \tag{21}$$

where signs $s_{ij} \in \{-1, 1\}$ hold only on defined physiological operating domains $\mathscr{D}_{ij}$.

For compartment-level generalization, we decompose parameters as $\theta_e = \theta_0 + U\gamma_e$, $\gamma_e \sim \mathscr{N}(0, \sigma_\gamma^2 I)$ and test against withheld compartments and cultivation seasons. This avoids treating thousands of adjacent five-minute records from a single compartment as statistically independent samples. Crop development is represented on a slower timescale, with a multi-rate numerical solver updating fast climate states at Δ_f and slow carbon partition states at $\Delta_s \geq \Delta_f$. Continuous-time assimilation connects irregular manual plant assessments to the latent crop trajectory.

6.3 Causal semantics: abduction, intervention and supported response

An effect of controlling CO_2 cannot be inferred from the regression coefficient of historical CO_2 concentration because enrichment is itself chosen in response to light, weather, production stage and the controller's state estimate. We formalize the observational assignment graph as

$$u_t = \pi_{\log}(\mathscr{H}_t, \varepsilon_t^u), \tag{22}$$

$$x_{t+1} = \Phi_\theta^{u_t,w_t}(t,t+1)x_t + P_{\mathscr{N}}\nu_t, \qquad y_t \sim p(y_t \mid x_t, M_t), \tag{23}$$

where the projected innovation $P_{\mathscr{N}}\nu_t$ preserves the modeled accounting identities; externally measured gains or losses belong to f_{ext}. The inferred innovation is held coherent across factual and counterfactual worlds rather than resampled independently. An intervention $\text{do}(u_{t:t+H-1} = \bar{a})$ replaces the assignment mechanism for those action variables and preserves the posterior over pre-intervention states and exogenous disturbances:

$$\begin{aligned} \widehat{Y}_{t+H}(\bar{a}) &= \mathbb{E}_{q_\varphi(x_t,\nu\mid\mathscr{H}_t)} \left[g\left(\Phi_{\theta,\nu}^{\bar{a},\bar{w}}(t,t+H)x_t \right) \right], \\ \widehat{\tau}_H(\bar{a},\bar{a}') &= \widehat{Y}_{t+H}(\bar{a}) - \widehat{Y}_{t+H}(\bar{a}'). \end{aligned} \tag{24}$$

Here $g(x)$ may be fruit carbon, cumulative harvest, disease-associated exposure or a resource function, and all actions are passed through the same constrained flux network. The causal derivative is a trajectory sensitivity, not ordinary one-step feature importance:

$$\frac{d\widehat{Y}_{t+H}}{du_{t+k}} = \mathbb{E}\left[\nabla_x g(x_{t+H})^\top \left(\prod_{j=k+1}^{H-1} J_{x,j} \right) J_{u,k} \right], \tag{25}$$

where $J_{x,j} = \partial \Phi_\theta / \partial x_j$ and $J_{u,k} = \partial \Phi_\theta / \partial u_k$. Such derivatives support localized control interpretations only where the structural model and action support are defensible.

Assumption 1 (Conditional exchangeability and positivity) *For a particular action component, horizon and logged-history summary H_t, potential next states satisfy $X_{t+1}(a) \perp U_t \mid H_t$; consistency and temporal ordering hold. Furthermore, $p(U_t \approx a \mid H_t) > 0$ for interventions under evaluation. Any omitted confounder that predicts both action choice and crop response may invalidate the assumption.*

Sequential exchangeability must be evaluated against measured policy histories and cannot be established from physiological constraints alone. The assumption is *not* guaranteed by public challenge datasets. We therefore estimate a conditional actuator density $\widehat{p}_\eta(u_t \mid H_t)$ and the effective overlap

$$\Omega(\bar{a}, H_t) = \min_{0 \le k < H} \log\left[\varepsilon + \widehat{p}_\eta(a_{t+k} \mid H_{t+k}^{\bar{a}})\right]. \tag{26}$$

Candidate actions with $\Omega < \omega_{\min}$ are flagged unsupported. If hidden confounding cannot be excluded, a sensitivity class $\mathscr{Q}_\Gamma$ bounds the deviation of latent policy-odds ratios:

$$\Gamma^{-1} \le \frac{p(u \mid H, z)/p(u' \mid H, z)}{p(u \mid H)/p(u' \mid H)} \le \Gamma, \qquad p \in \mathscr{Q}_\Gamma, \tag{27}$$

with user-selected $\Gamma \ge 1$. We report partial-effect sets rather than point-identifying unsupported interventions:

$$\mathscr{T}_H(\bar{a}, \bar{a}'; \Gamma) = \left[\inf_{q \in \mathscr{Q}_\Gamma} \mathbb{E}_q \tau_H, \ \sup_{q \in \mathscr{Q}_\Gamma} \mathbb{E}_q \tau_H\right]. \tag{28}$$

For observed data with sufficient overlap, inverse-weighted and outcome-regression estimates can be compared to neural-flow effects as a falsification diagnostic. We do *not* claim that an accurate PINN removes unobserved confounding: physics restricts candidate outcome functions but cannot identify a unique action-effect distribution from observational data alone.

6.4 Uncertainty propagation and a control-relevant risk envelope

We propagate epistemic posterior samples $\theta^{(m)}$, inferred disturbance trajectories $\nu^{(m)}$, and weather ensembles $w^{(m)}$ through the same projected flow:

$$\mathscr{X}_H(\bar{u}) = \left\{\Phi^{\bar{u}, \bar{w}^{(m)}}_{\theta^{(m)}, \nu^{(m)}}(t, t+H) x_t^{(m)}\right\}_{m=1}^{M}. \tag{29}$$

For a hazard $g_j(x,u)$ with allowed upper bound zero, we use conditional value-at-risk (CVaR), which controls the severity of upper-tail violations:

$$\mathrm{CVaR}_\alpha(g_j) = \min_{\zeta \in \mathbb{R}} \left\{\zeta + \frac{1}{1-\alpha} \mathbb{E}_m\left[(g_j(x^{(m)}, u) - \zeta)_+\right]\right\}. \tag{30}$$

Taking the worst case over a Wasserstein ball $\mathbb{B}_\rho(\widehat{P})$ gives a distributionally robust constraint

$$\sup_{\eta \in \mathscr{E}_\Gamma} \ \sup_{Q \in \mathbb{B}_\rho(\widehat{P}_\eta)} \mathrm{CVaR}_\alpha^Q(g_j) \le 0, \tag{31}$$

where $\mathscr{E}_\Gamma$ is an explicit collection of admissible hidden-confounding perturbations of the logged assignment mechanism, $\widehat{P}_\eta$ is the induced distribution on the *same scenario space* of uncertain model parameters, weather and disturbances, and $\mathbb{B}_\rho(\widehat{P}_\eta)$ is a Wasserstein ball of trajectory laws. This nested sup avoids taking an uninterpretable intersection of treatment-odds restrictions and scenario distributions. The radius ρ is chosen using an independent calibration period and sensitivity to coverage is reported. An empirical interval calibrated on independent compartments can guide selection of ρ; exchangeability and enough compartments are required for strict nominal coverage guarantees. We explicitly separate variation caused by model parameters, random weather and unmeasured treatment confounding.

6.5 Risk-aware, preference-adaptive Pareto model predictive control

Define the horizon cost vector for a candidate schedule $\bar{u} = (u_t, \ldots, u_{t+H-1})$ as

$$J_1(\bar{u}) = -\mathbb{E}\left[\int_t^{t+H} \dot{B}_{\text{fruit}}(\tau)\, d\tau\right], \tag{32}$$

$$J_2(\bar{u}) = \mathbb{E}\int_t^{t+H} \{p_h q_h u^h + p_e e(u) + p_c u^c\}\, d\tau, \tag{33}$$

$$J_3(\bar{u}) = \mathbb{E}\int_t^{t+H} u_\tau^i\, d\tau, \tag{34}$$

$$J_4(\bar{u}) = \mathbb{E}\int_t^{t+H} \{\ell_{\text{stress}}(x_\tau) + \kappa \|u_\tau - u_{\tau-1}\|_2^2\}\, d\tau. \tag{35}$$

Here p_h, p_e, p_c are *user-provided* utility and price coefficients, not assumed prices in the public data. We seek the nondominated feasible set $\mathscr{P} = \text{ParetoMin}_{\bar{u}\in\mathscr{F}} \mathbf{J}(\bar{u})$, with

$$\mathscr{F} = \{\bar{u} : u_t \in \mathscr{U},\ \|u_{t+1} - u_t\|_\infty \le \delta_u,\ \Omega(\bar{u}, H_t) \ge \omega_{\min},\ \text{Eq. (31)}\}. \tag{36}$$

Rather than collapsing all preferences to one arbitrarily fixed reward, the inner differentiable problem uses a smooth, augmented Chebyshev scalarization

$$\min_{\bar{u},\zeta} \quad \zeta + \varepsilon_s \sum_{k=1}^{4} \lambda_k (J_k(\bar{u}) - z_k^\star), \quad \lambda_k (J_k(\bar{u}) - z_k^\star) \le \zeta, \quad k = 1, \ldots, 4, \tag{37}$$

for $\lambda \in \Delta_4$ and reference point $z^\star$. This scalarization can expose nonconvex sections of the Pareto surface more effectively than an ordinary weighted sum. We use projected differentiable optimization for $\bar{u}$ and an outer adaptive sampling of λ favoring sparse regions of the empirical frontier; a final nondominance test removes dominated candidates. The controller applies only the first safe action, updates its state posterior with new sensor data and replans. If the filtered set is empty, a prevalidated backup policy is required; learned control is not executed.

6.6 Optimization geometry and computational cost

For an H-step discretization $x_{k+1} = F_\theta(x_k, u_k, w_k)$, action gradients obey

$$\nabla_{u_k} J = \nabla_{u_k} \ell(x_k, u_k) + \left(\frac{\partial F_\theta}{\partial u_k}\right)^\top p_{k+1}, \quad p_k = \nabla_{x_k} \ell + \left(\frac{\partial F_\theta}{\partial x_k}\right)^\top p_{k+1}. \tag{38}$$

With K preference vectors, M uncertainty scenarios, solver cost C_F per step and I inner iterations, naive differentiable Pareto control is $O(KMIHC_F)$ in computation and $O(MH \dim x)$ in stored forward states before checkpointing. We propose three practical reductions: clustered weather scenarios, adjoint checkpointing and warm-starts from the previous MPC solution. Measured end-to-end latency and sensitivity to the planning horizon are necessary for a realistic deployment assessment.

Computational workflow. We separate the data-processing, training, counterfactual evaluation and policy-optimization procedures to prevent future information from entering model selection. Crop-cycle and compartment identities are retained through splitting, confidence-interval estimation and reporting. Actuator realizations, rather than nominal commands, enter the transition whenever available; otherwise, an actuator-response model is required. The controller's output is evaluated only in a documented retrospective or replay-based protocol unless it has actually been operated in a greenhouse.

Funding

This work was supported by the Qianmu Pastoral New Ecological Agriculture Standardization Demonstration Project in Xinxiang City.

Competing Interests

The authors declare no competing interests.

Author Contributions

F.W. conceived the study, designed the overall research framework, and supervised the research.

J.L. contributed to the development of the physics-informed neural network and the mathematical formulation of greenhouse tomato growth dynamics.

X.T. contributed to dataset preparation, data preprocessing, and the implementation of the causal inference and multi-objective optimization methods.

Y.L. contributed to experimental design, model implementation, baseline comparisons, and ablation studies.

Q.D. contributed to experimental validation, statistical analysis, data visualization, and interpretation of the results.

F.W. drafted the manuscript. All authors contributed to manuscript revision and reviewed and approved the final version.

Data Availability

The publicly available greenhouse tomato datasets used in this study can be accessed through the following repositories:

1. **AGC-2 dataset:** Official repository.
2. **AGC-4 pretrial dataset:** Official repository.
3. **AGC-4 challenge dataset:** Official repository.
4. **Universidad de Antioquia greenhouse tomato dataset:** Zenodo repository.

Access to these datasets is subject to the terms and licences specified by their respective repositories.

Ethics Declaration

This study concerns computational modeling, machine learning, and the analysis of publicly available greenhouse tomato cultivation datasets.

The research does not involve human participants, human biological samples, or animal experimentation. No personally identifiable human information is required for the described analyses.

Accordingly, human research ethics approval, informed consent to participate, and consent for publication are not applicable.

References

1. et al., S. H. Autonomous greenhouse challenge, second edition (2019). Wageningen University & Research / 4TU.ResearchData (2020).

2. et al., I. R. 4th autonomous greenhouse challenge: Pre-trial dwarf tomato measurements and images. Wageningen University & Research dataset (2024).

3. et al., S. M. 4th autonomous greenhouse challenge: Dwarf tomato timeseries and images. Wageningen University & Research dataset (2025).

4. et al., S. C. M. Autonomous greenhouse cultivation of dwarf tomato: Performance evaluation of intelligent algorithms for multiple-sensor feedback. Sensors 25(14), 4321 (2025).

5. et al., X. Z. Boosting the prediction accuracy of a process-based greenhouse climate-tomato production model by particle filtering and deep learning. Computers and Electronics in Agriculture 211, 107980 (2023).

6. et al., B. M. Reinforcement learning versus model predictive control on greenhouse climate control. Computers and Electronics in Agriculture 215, 108372 (2023).

7. Hu, G. & You, F. An ai framework integrating physics-informed neural network with predictive control for energy-efficient food production in the built environment. Applied Energy 348, 121450 (2023).

8. et al., S. M. Reinforcement learning-based model predictive control for greenhouse climate control. Smart Agricultural Technology 10, 100751 (2025).

9. et al., D. X. Rule-based year-round model predictive control of greenhouse tomato cultivation: A simulation study. Information Processing in Agriculture 12(3), 344–357 (2025).

10. Xiong, Y. & Su, Y. Hybrid modeling approaches for accurate greenhouse climate prediction: Combining mechanistic models and lstm neural networks. Ecological Modelling 503, 111059 (2025).

11. M. Raissi, P. P. & Karniadakis, G. E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. Journal of Computational Physics 378, 686–707 (2019).

12. S. Wang, S. S. & Perdikaris, P. Respecting causality for training physics-informed neural networks. Computer Methods in Applied Mechanics and Engineering 421, 116813 (2024).

13. et al., G. F. A simple remedy for failure modes in physics informed neural networks. Neural Networks 183, 106963 (2025).

14. et al., C. X. Fp64 is all you need: Rethinking failure modes in physics-informed neural networks. arXiv:2505.10949 (2025).

15. et al., B. Z. An opensource indoor climate and yield prediction model for chinese solar greenhouses. Biosystems Engineering 250, 183–212 (2025).

16. de Oliveira et al., M. P. G. Uncertainty in greenhouse tomato growth models. Computers and Electronics in Agriculture, 109324 (2024).

17. et al., S. L. Modelling the response of tomato on deficit irrigation under greenhouse conditions. Scientia Horticulturae 326, 112770 (2024).

18. et al., J. Z. Optimizing irrigation schedules of greenhouse tomato based on a comprehensive evaluation model. Agricultural Water Management 295, 108741 (2024).

19. Šalagovič et al., J. Model prediction of plant morphology, water flows and xylem water potential in a growing tomato plant under heterogeneous growing conditions. Computers and Electronics in Agriculture 235, 110346 (2025).

20. Jeong, J. Y. & Choi, H. G. Development and multiple cultivation seasons evaluation of a temperature-driven process-based model for tomato growth and yield in greenhouse conditions. Computers and Electronics in Agriculture 251, 111919 (2026).

21. et al., K. S. Development of a tomato functional–structural plant model for digital twin applications. in silico Plants 7(2), diaf022 (2025).

22. et al., Y. L. Digital twin-driven system for efficient tomato harvesting in greenhouses. Computers and Electronics in Agriculture 236, 110451 (2025).

23. et al., J. G. A digital framework for smart greenhouse tomato cultivation. Smart Agricultural Technology 13, 101879 (2026).

24. et al., J. G. Multi-model fusion method for predicting co2 concentration in greenhouse tomatoes. Computers and Electronics in Agriculture 227, 109623 (2024).

25. et al., J. Y. Prediction and control of greenhouse temperature: Methods, applications, and future directions. Computers and Electronics in Agriculture 237, 110603 (2025).

26. et al., Z. W. A machine learning-based irrigation prediction model for cherry tomatoes in greenhouses: Leveraging optimal growth data for precision irrigation. Computers and Electronics in Agriculture 237, 110558 (2025).

27. et al., A. S. Deep learning-driven iot solution for smart tomato farming. Scientific Reports 15, 31092 (2025).

28. et al., G. E. K. Physics-informed machine learning. Nature Reviews Physics 3, 422–440 (2021).

29. et al., C. K. N. Taylor-model physics-informed neural networks (pinns) for ordinary differential equations. International Conference on Neuro-symbolic Systems, PMLR 288, 621–642 (2025).

30. N. Kaminsky, D. F. & Radinsky, K. Overcoming pinns failure modes in high dimension with low-rank fourier sum. ICML, PMLR 306, 55264–55283 (2026).

31. et al., A. S. Learning semilinear neural operators: A unified recursive framework for prediction and data assimilation. ICLR 2024 (2024).

32. et al., V. G. Learning physical operators using neural operators. AISTATS, PMLR 300, 3223–3231 (2026).

33. Y. Verma, M. H. & Garg, V. Climode: Climate and weather forecasting with physics-informed neural odes. International Conference on Learning Representations (ICLR) (2024).

34. et al., L. P. Universal physics-informed neural networks: Symbolic differential operator discovery with sparse data. ICML, PMLR 202 (2023).

35. et al., J. G. Advances in physics-informed neural networks for solving complex partial differential equations and their engineering applications: A systematic review. Engineering Applications of Artificial Intelligence 161, 112044 (2025).

36. et al., B. S. Toward causal representation learning. Proceedings of the IEEE 109(5), 612–634 (2021).

37. et al., K. A. Interventional causal representation learning. ICML, PMLR 202, 372–407 (2023).

38. et al., B. V. General identifiability and achievability for causal representation learning. AISTATS, PMLR 238, 2314–2322 (2024).

39. et al., S. B. Identifying linearly-mixed causal representations from multi-node interventions. Conference on Causal Learning and Reasoning, PMLR 236, 843–867 (2024).

40. Morioka, H. & Hyvarinen, A. Causal representation learning made identifiable by grouping of observational variables. ICML, PMLR 235, 36249–36293 (2024).

41. W. Liang, H. C. & Zhang, Z. Causal representation learning with optimal compression under complex treatments. ICML, PMLR 306, 73142–73170 (2026).

42. et al., Y. L. Causal triplet: An open challenge for intervention-centric causal representation learning. Conference on Causal Learning and Reasoning, PMLR 213, 553–573 (2023).

43. Hatt, T. & Feuerriegel, S. Sequential deconfounding for causal inference with unobserved confounders. Conference on Causal Learning and Reasoning, PMLR 236, 934–956 (2024).

44. et al., M. X. Grower-in-the-loop interactive reinforcement learning for greenhouse climate control. Computers and Electronics in Agriculture 242, 111312 (2026).

45. et al., I. H. Enhancing autonomous agriculture control systems in greenhouses for sustainable resource usage using deep learning techniques. PLOS ONE 21(3), e0344946 (2026).

46. Chen, D. & Huang, Y. Integrating reinforcement learning and large language models for crop production process management optimization and control through a new knowledge-based deep learning paradigm. Computers and Electronics in Agriculture 232, 110028 (2025).

47. et al., H. W. An optimized approach to hourly temperature and humidity setpoint generation for reducing tomato disease and saving power cost in greenhouses. Computers and Electronics in Agriculture 226, 109413 (2024).

48. Chen, W.-H. & You, F. Decarbonization through smart energy management: Climate control in building-integrated rooftop greenhouses for urban agriculture across various climate conditions. Journal of Cleaner Production 458, 142544 (2024).

49. et al., F. R. Optimization framework of clean heat and co2 supply for agricultural greenhouses exploiting industrial symbiosis. Energy Strategy Reviews 55, 101509 (2024).

50. et al., N. Z. Optimization of above-ground environmental factors in greenhouses using a multi-objective adaptive annealing genetic algorithm. Heliyon 10(12), e33036 (2024).

51. et al., K. L. The intelligentization process of agricultural greenhouse: A review of control strategies and modeling techniques. Agriculture 15(20), 2135 (2025).

52. G. Hu, C. K. & You, F. Cyber physical biological system in controlled environment agriculture for energy optimization: A comprehensive overview, key challenges, and future outlook. Energy 325, 136153 (2025).

53. et al., K. D. A fast and elitist multiobjective genetic algorithm: Nsga-ii. IEEE Transactions on Evolutionary Computation 6(2), 182–197 (2002).