



DPMF: Dual-Path Multi-timescale Fusion with Phenology-Gated Prediction for Tomato State and Daily Water-Nutrient Flux Estimation

Chenlei Fang^{1,2,3,4,5}, Youlin Cui^{2,3,4,5}, Yisheng Miao^{2,3,4,5,*} and Chunjiang Zhao^{1,2,3,4,5,*}

¹School of Computer Science, Northwestern Polytechnical University, Xi'an 710129, China

²National Engineering Research Center for Information Technology in Agriculture, Beijing 100097, China

³Information Technology Research Center, Beijing Academy of Agriculture and Forestry Sciences, Beijing 100097, China

⁴Key Laboratory of Digital Village Technology, Ministry of Agriculture and Rural Affairs, Beijing 100097, China

⁵Beijing Research Institute of Agricultural Artificial Intelligence and Robotics, Beijing 100097, China

Abstract

Full-cycle management of greenhouse tomato requires continuous information on crop state together with estimates of daily water and nutrient use, yet destructive biomass measurements are sparse, and a single fixed model shared across resources is prone to stage-dependent negative transfer. This study proposes Dual-Path Multi-timescale Fusion (DPMF), a development-aware framework for reconstructing crop state from sparse observations and estimating daily resource fluxes. The state path combines a process-based dry-matter prior with a gated recurrent unit trained on three kinds of supervision: point states, biomass increments between adjacent

samples, and cumulative harvested-fruit dry matter. The resource path uses smooth phenology gating to weight task-specific and shared ExtraTrees predictions across six crop-development stages, so that information is shared between resources only where sharing does not degrade accuracy. On 64 destructive samples from seven light-temperature treatments, DPMF achieved a root mean square error (RMSE) of 132.821 g m^{-2} and a mean absolute error (MAE) of 65.246 g m^{-2} , 6.09% and 5.31% lower in RMSE than interpolation-based and point-only supervision, respectively. On 154 greenhouse test dates, DPMF reduced RMSE relative to a fixed shared model by 4.24%, 5.34%, 2.51% and 0.67% for water, nitrogen (N), phosphorus (P) and potassium (K), respectively. DPMF thus offers a single modelling basis for crop-state reconstruction and coordinated water-nutrient estimation.

Keywords: greenhouse tomato, crop-state reconstruction, multi-timescale supervision, phenology, multi-task learning, precision fertigation, water and nutrient flux.



Academic Editor:

Chunxi Yang

Submitted: 09 July 2026

Revised: 11 October 2026

Accepted: 15 October 2026

Published: 23 October 2026

Vol. 3, No. 4, 2026.

10.62762/TIS.2026.964269

*Corresponding authors:

✉ Yisheng Miao

Nercita_SFX_Y_01@126.com

✉ Chunjiang Zhao

Nercita_SFX_Y_01@126.com

Citation

Fang, C., Cui, Y., Miao, Y., & Zhao, C. (2026). DPMF: Dual-Path Multi-timescale Fusion with Phenology-Gated Prediction for Tomato State and Daily Water-Nutrient Flux Estimation. *ICCK Transactions on Intelligent Systematics*, 3(4), 209–228.



© 2026 by the Authors. Published by Institute of Central Computation and Knowledge. This is an open access article under the CC BY license (<https://creativecommons.org/licenses/by/4.0/>).

1 Introduction

Greenhouse tomato (*Solanum lycopersicum* L.) production increasingly relies on dense climate sensing, automated irrigation and data-driven decision tools. Controlled-environment systems record radiation, temperature, humidity, carbon dioxide, water supply and drainage at daily or sub-daily resolution [1], and machine-learning and deep-learning methods are now widely used in agriculture for tasks ranging from yield prediction [2] to crop classification and general management [3, 4]. Crop variables, however, are not observed at the temporal resolution of the greenhouse climate. Environmental measurements are dense, leaf area and organ measurements are intermittent, destructive biomass observations are sparse, and harvested fruit accumulates over time. A daily crop estimator therefore has to learn from observations with different temporal supports instead of forcing every measurement into an artificial daily point target.

Two modelling issues follow. First, interpolating between destructive samples yields dense labels, but it treats unobserved intermediate states as if they had been measured. Deep time-series models can learn sequential dependence [5], and gated recurrent units (GRUs) offer a compact recurrent representation [6]; the training objective, however, must still respect the temporal support of each crop observation. Second, water, N, P and K share environmental and developmental drivers while differing in physiological and management responses. Multi-task learning can exploit common structure [7–10], and multi-gate mixture-of-experts models can adapt task relations [11]. A single fixed shared model nevertheless imposes the same task coupling throughout the production cycle, although resource relationships change between establishment and harvest.

This study develops and evaluates Dual-Path Multi-timescale Fusion (DPMF), which addresses the two issues with two complementary modelling paths that are trained and evaluated independently. The state path combines point, interval and cumulative supervision to reconstruct total dry matter. The resource path combines task-specific and shared tree ensembles through a gate conditioned on crop phenology. Tree ensembles are stable nonlinear predictors for tabular agricultural data [12, 13], and the phenology gate sets how much independent and shared information contributes at each crop stage. The framework is evaluated on public Autonomous

Greenhouse Challenge datasets [14–17].

The contributions of this study are as follows.

1. A multi-timescale supervision objective that uses destructive point observations, dry-matter increments between adjacent samples and cumulative harvested-fruit dry matter within one recurrent state-reconstruction model.
2. A phenology-conditioned fusion mechanism that combines independent and shared resource predictions across six overlapping crop stages.
3. An accuracy-retention criterion that measures how closely each joint modelling strategy preserves the accuracy of independent single-task models. DPMF reduces the equal-weight mean relative RMSE gap to 1.91%, compared with 2.89% for global fusion and 5.28% for a fixed shared model.

The remainder of this paper is organised as follows. Section 2 reviews process-based and hybrid crop models, multi-task and stage-adaptive learning, and temporal crop-state estimation. Section 3 presents the datasets and the DPMF methodology. Section 4 describes the experimental protocols, reports the results and discusses them. Section 5 concludes the paper.

2 Related Work

2.1 Process-Based Crop Models and Hybrid Methods

Process-based crop models give biologically interpretable descriptions of crop development, water use, carbon assimilation, dry-matter allocation and nutrient uptake. FAO-56 relates reference evapotranspiration to crop coefficients [18], HortSyst provides an analogous discrete-time framework for greenhouse tomato, simulating dry matter production, N uptake, leaf area index and transpiration [19], with uncertainty quantified across the full N, P, K and water variable set [20]. Sequential Bayesian state updates via data assimilation have been used to correct crop-model trajectories: Kalman-family filters have been applied to greenhouse tomato growth models to update dry matter and leaf area estimates from in-season observations [21]. Process-based models have also been adapted to simulate tomato responses under deficit irrigation [22], illustrating how calibrated physiological relationships can represent crop behaviour under management-induced stress. TOMSIM describes tomato dry-matter production and partitioning among organs [23, 24], while the

Wageningen crop models provide broader frameworks for mechanistic crop simulation [25], and GreenLight extends this tradition to integrated greenhouse climate, energy and crop production [26].

For greenhouse water and nutrient management, the VegSyst family is particularly relevant. VegSyst has been validated for crop growth, evapotranspiration and N uptake for greenhouse tomato [27], and has been built into decision-support systems that estimate the water and N requirements of tomato [28, 29]. The broader context of dry matter production and growth modelling for greenhouse tomatoes is reviewed in Higashide [30]. Later versions modelled macronutrient uptake by greenhouse tomato, supported prescriptive-corrective management, and were evaluated in different regions, substrates and soilless systems [31–36]. These studies show the value of crop-development information and process-based demand indicators for greenhouse resource management.

Process-based models nevertheless depend on calibrated parameters, predefined physiological relationships and assumptions about crop response. Realised trajectories vary with cultivar, light and temperature treatment, greenhouse operation, substrate and management history. This matters when destructive biomass observations are sparse, because a departure from the calibrated trajectory can persist between sampling dates. It also matters for joint resource modelling, because the relationships among water, N, P and K demand change as the crop develops.

Hybrid process-data methods combine biological structure with observation-driven correction. Data assimilation [37] and deep learning have been combined to improve the predictions of a greenhouse climate-tomato production model [38], showing that sequential Bayesian state updates can effectively correct process-based trajectories. Existing hybrid approaches mainly correct point predictions or latent states; they give little guidance on learning jointly from destructive point observations, interval biomass changes and cumulative harvested output. Process correction also does not say how information should be shared among resource tasks whose relationships change across crop stages. These two gaps motivate the two paths of DPMF.

2.2 Multi-Task and Stage-Adaptive Learning

Multi-task learning lets related prediction tasks share representations or parameters and can improve data efficiency when the tasks carry complementary information [7, 8]. Adaptive weighting accounts for differences in task uncertainty and optimisation difficulty [9], and surveys show that multi-task performance depends strongly on how shared and task-specific information is organised [10]. In greenhouse resource estimation, water, N, P and K respond to common environmental and developmental drivers, which provides a basis for sharing.

A central difficulty is negative transfer: patterns that help one task can reduce the accuracy of another when the shared representation is too restrictive or task relationships vary across samples. Multi-gate mixture-of-experts architectures address this by learning task-specific combinations of shared experts [11]. Such methods demonstrate the value of selective sharing, but they are formulated as end-to-end neural networks and do not encode crop-development stages. A fixed shared model or a time-invariant fusion coefficient therefore cannot represent resource relationships that change between establishment and harvest.

Several crop applications illustrate selective information sharing. Gong et al. [39] fused a biophysical Tomgro model with a CNN-RNN machine learning model for greenhouse tomato yield prediction, showing that adaptive fusion of complementary models outperformed either component alone. Florakis et al. [40] integrated light interception models with statistical approaches to predict daily water consumption of tomato in a hydroponic greenhouse, illustrating how developmental stage and radiation drive resource demand. Simm et al. [41] proposed a tree-ensemble multi-task learning framework based on Extremely Randomized Trees that jointly trains shared and task-specific components to reduce negative transfer, demonstrating the value of combining shared and independent tree predictors in a multi-task setting. In the crop applications above [39, 40], the targets are seasonal yields, phenology is used to align inputs, and gating combines expert representations. Daily water-nutrient estimation instead needs a task-specific combination of independent and shared predictions whose weights evolve over the crop cycle.

Tree ensembles are a practical choice for agricultural

prediction from tabular measurements. Random forests combine bagging with random feature selection [12], and extremely randomised trees (ExtraTrees) add further randomisation in feature and split selection [13]. For greenhouse tomato, Ge et al. [42] compared eight regression methods on three years of meteorological and evapotranspiration data and found XGBoost to give the most accurate daily evapotranspiration estimates. Extending such prediction from water alone to water, N, P and K adds a further requirement: deciding how task-specific and shared information should be weighted at each developmental stage. DPMF meets this requirement by combining independent and shared ExtraTrees predictions through phenology-conditioned coefficients. Six overlapping stage functions give a smooth representation of development, so the contributions of task-specific and shared predictions can change over the crop cycle.

2.3 Crop Temporal State Estimation

Continuous crop-state estimation supports growth monitoring and resource-demand assessment, but greenhouse observations have heterogeneous temporal supports. Climate and management variables may be recorded daily or more often, leaf-area measurements are intermittent, destructive biomass measurements are sparse, and fruit harvest accumulates over time. A temporal estimator must reconstruct crop development from observations that describe different aspects and time spans of the underlying state.

Deep temporal models learn nonlinear dependence from sequential inputs [5]. GRUs use update and reset gates to retain relevant history while controlling changes in the recurrent state [6], which suits the accumulation of crop responses to environmental sequences. Autonomous greenhouse datasets provide dense sensor records and repeated crop observations for learning such relationships [14–17]. Gong et al. [43] combined temporal convolutional network and recurrent neural network architectures to predict greenhouse tomato yield from historical environmental sequences including temperature, humidity, carbon dioxide and radiation, demonstrating that sequential environmental inputs carry strong predictive information for crop output. Their target was accumulated yield rather than a daily biomass state, which is a different temporal problem from reconstructing dry-matter trajectories between destructive sampling dates. For greenhouse tomato,

Lin et al. [44] evaluated recurrent architectures, including GRU, for temperature forecasting and used the predicted growing degree days in sigmoid models of leaf area index, fresh fruit weight and aboveground dry matter, linking environmental sequences with crop development through thermal time.

Data assimilation techniques, including Kalman-family filters applied to the Reduced-State TOMGRO model, have also been used to update greenhouse tomato state variables—dry matter and leaf area index—from in-season sensor observations, offering an alternative route to continuous state estimation [21].

The remaining difficulty lies in the training objective. Point-only training anchors predictions at destructive sampling dates but weakly constrains the change between them. Linear interpolation supplies dense targets by imposing an assumed trajectory. Cumulative harvested-fruit records carry information on reproductive output, but they describe an accumulated quantity rather than an instantaneous state; treating them as interchangeable point targets discards this distinction.

DPMF organises these signals by temporal support. Point supervision anchors biomass magnitude at destructive observations, interval supervision constrains dry-matter accumulation between adjacent samples, and cumulative supervision represents harvested-fruit production. A fold-calibrated VegSys trajectory supplies a process prior, and the GRU learns treatment-dependent corrections from environmental and developmental sequences. No artificial daily biomass labels are created.

3 Methodology

3.1 Overall DPMF Framework

DPMF contains two complementary modelling paths that are trained and evaluated independently (Figure 1). Calendar and phenology variables give the developmental position of each record. Environmental inputs describe the radiation, temperature and greenhouse conditions that drive growth and transpiration. Crop observations provide biomass, leaf-area and harvested-fruit information, and root-zone and management records describe the history of water and nutrient supply.

The upper path reconstructs crop state with a GRU and multi-timescale supervision, producing a daily trajectory of accumulated tomato dry matter from

sparse observations. The lower path estimates daily net fluxes of water and elemental N, P and K. It combines four task-specific ExtraTrees predictors with one shared multi-output ExtraTrees predictor, and a smooth phenology gate changes their relative contributions over the production cycle.

3.2 Public Datasets and Variable Preparation

3.2.1 WUR Dwarf-Tomato Dataset

The state-reconstruction experiment used the Wageningen University & Research (WUR) dwarf-tomato dataset from the fourth Autonomous Greenhouse Challenge [16, 17]. It contains two cultivars, Orange and Red, grown under seven light-temperature treatments, giving 14 complete cultivar-treatment sequences and 64 destructive sampling records. The treatments span two experimental years, and time is expressed as days after sowing (DAS). Environmental and crop-development variables were aligned to daily sequence positions, whereas destructive biomass measurements kept their original sampling dates.

Total dry matter was defined as standing plant dry matter plus cumulative harvested-fruit dry matter, so that reproductive output removed at harvest remains part of the state. Light-temperature treatments, not individual sampling dates, were the held-out experimental units.

3.2.2 AGC2 Greenhouse Dataset

The daily flux experiment used the six greenhouse compartments of the second Autonomous Greenhouse Challenge (AGC2) [14, 15]. Each compartment contains 166 daily records. After applying the common validity mask defined by the water-flux reference, 154 test dates were retained for all four resource tasks.

Daily resource references were calculated as recorded supply minus drainage. Water was expressed in mm d^{-1} , and N, P and K as elemental $\text{kg ha}^{-1} \text{d}^{-1}$. Where fertiliser records were given in oxide form, elemental quantities were obtained as $m_P = 0.4364 m_{P_2O_5}$ and $m_K = 0.8301 m_{K_2O}$ using standard stoichiometric conversion factors. All methods used the same dates, response definitions, historical inputs and update protocol. Missing predictor values were imputed with medians estimated from the corresponding fitting set.

3.3 Multi-timescale Crop-State Reconstruction

3.3.1 State Definition and Temporal Representation

Tomato total dry matter on day t is

$$B_t = B_t^{\text{st}} + B_t^{\text{hf}}, \quad (1)$$

where B_t^{st} is standing plant dry matter and B_t^{hf} is cumulative harvested-fruit dry matter, both in g m^{-2} . The state path is a process-prior recurrent model. For each cultivar-treatment sequence, 12 daily environmental, temporal and cultivar descriptors were used: mean, minimum and maximum air temperature; mean relative humidity; mean CO_2 concentration; daily photosynthetically active radiation (PAR); DAS; daily and cumulative thermal time; nominal treatment temperature; nominal daily light integral (DLI); and a cultivar indicator. These variables were standardised with the mean and standard deviation of the training fold, and missing values were replaced by the training-fold mean before standardisation.

Crop development in the state path was represented by six normalised Gaussian basis functions $\pi_{k,t}^{\text{st}}$ of $g_t^{\text{st}} = \text{clip}(TT_t^{\text{cum}}/2600, 0, 1)$, where $\text{clip}(x, a, b) = \max(a, \min(b, x))$ denotes clamping to $[a, b]$ and TT_t^{cum} is cumulative thermal time ($^{\circ}\text{C d}$), with centres (0.04, 0.18, 0.36, 0.55, 0.74, 0.93) and a common bandwidth $\sigma^{\text{st}} = 0.13$. A fold-calibrated VegSyst model supplied the process prior p_t , the predicted daily dry-matter increment. VegSyst parameters were calibrated using only the training treatments of each leave-one-treatment-out fold.

The 12 standardised predictors, the six phenology weights and $\log(1 + p_t)$ formed the 19-dimensional input to a single-layer unidirectional GRU with 32 hidden units, whose hidden state on day t is h_t . Six phenology-gated linear experts map h_t to a bounded correction of the process-based daily increment and to a fruit-allocation fraction:

$$d_t = \sum_{k=1}^6 \pi_{k,t}^{\text{st}} (W_k h_t + c_k), \quad r_t = 2.5 \tanh(d_{t,1}), \quad (2)$$

$$\widehat{\Delta B}_t = \text{softplus}(\text{softplus}^{-1}(\max(p_t, 10^{-4})) + r_t), \\ \widehat{B}_t = \sum_{u \leq t} \widehat{\Delta B}_u, \quad (3)$$

where W_k and c_k are the weight matrix and bias of expert k , and d_t is the phenology-weighted expert output with components $d_{t,1}$ and $d_{t,2}$. The inverse softplus was implemented numerically as




Figure 1. DPMF architecture. The upper path reconstructs total dry matter with a GRU trained on point, interval, cumulative-fruit, leaf area index (LAI) and process-residual terms. The lower path combines independent and shared ExtraTrees predictions through a smooth six-stage phenology gate to estimate daily net water and elemental N, P and K fluxes.

$\text{softplus}^{-1}(p_t) = \log\{\max[\exp(\max(p_t, 10^{-4})) - 1, 10^{-6}]\}$ to maintain numerical stability near zero. The second expert output $d_{t,2}$ defines the fruit-allocation fraction through a sigmoid, and LAI is estimated from h_t through a separate softplus linear readout $\hat{L}_t$. Total dry matter is therefore reconstructed by correcting and accumulating the fold-calibrated process prior rather than by an independent biomass output head.

3.3.2 Point, Interval and Cumulative Supervision

Let s index cultivar-treatment sequences and Ω_B be the set of pairs (s, t) with a destructive biomass observation $B_{s,t}$. The point-state loss anchors reconstructed biomass magnitude in log space:

$$\mathcal{L}_{\text{state}} = \frac{1}{|\Omega_B|} \sum_{(s,t) \in \Omega_B} \left[\log(1 + \hat{B}_{s,t}) - \log(1 + B_{s,t}) \right]^2. \quad (4)$$

LAI supervision uses a masked mean-squared error over the set Ω_L of dates with finite LAI observations, $\mathcal{L}_{\text{LAI}} = |\Omega_L|^{-1} \sum_{(s,t) \in \Omega_L} (\hat{L}_{s,t} - L_{s,t})^2$. The process-residual term is $\mathcal{L}_{\text{process}} = |\Omega_D|^{-1} \sum_{(s,t) \in \Omega_D} r_{s,t}^2$, where Ω_D is the set of valid daily positions. This term pulls the learned correction towards zero and keeps the reconstructed daily increment close to the fold-calibrated VegSys prior; it is not a loss that matches biomass or increments to VegSys.

Let $\mathcal{P}_B$ be the set of triples (s, t_-, t_+) in which $t_- < t_+$ are adjacent destructive sampling dates of sequence s . The interval loss constrains dry-matter accumulation over each observed interval:

$$\mathcal{L}_{\text{interval}} = \frac{1}{|\mathcal{P}_B|} \sum_{(s,t_-,t_+) \in \mathcal{P}_B} \left[\log(1 + \hat{B}_{s,t_+} - \hat{B}_{s,t_-}) - \log(1 + B_{s,t_+} - B_{s,t_-}) \right]^2. \quad (5)$$

The cumulative harvested-fruit state is $\hat{\Delta F}_t = \hat{\Delta B}_t \text{sigmoid}(d_{t,2})$ and $\hat{F}_t = \sum_{u \leq t} \hat{\Delta F}_u$. Over the set Ω_F of dates with recorded cumulative harvested-fruit dry matter $F_{s,t}$, the fruit loss is

$$\mathcal{L}_{\text{fruit}} = \frac{1}{|\Omega_F|} \sum_{(s,t) \in \Omega_F} \left[\log(1 + \hat{F}_{s,t}) - \log(1 + F_{s,t}) \right]^2. \quad (6)$$

The common training objective is

$$\mathcal{L}_{\text{common}} = \mathcal{L}_{\text{state}} + 0.4 \mathcal{L}_{\text{LAI}} + 0.01 \mathcal{L}_{\text{process}}, \quad (7)$$

and the DPMF multi-timescale objective is

$$\mathcal{L}_{\text{DPMF}} = \mathcal{L}_{\text{common}} + 0.5 \mathcal{L}_{\text{interval}} + 0.35 \mathcal{L}_{\text{fruit}}. \quad (8)$$

Point, Interpolation and DPMF share the same recurrent architecture, environmental inputs,

fold-calibrated VegSyst prior, LAI objective and process-residual regularisation. Point uses only the observed destructive point states for biomass supervision. Interpolation replaces the sparse point-state target with daily linear pseudo-labels generated only between adjacent observations of a sequence, without extrapolation before the first or after the last observation. DPMF keeps the observed point states and adds interval and cumulative-fruit supervision: point states anchor biomass magnitude, interval increments constrain stage-wise change, and cumulative fruit constrains reproductive output.

3.4 Phenology-Conditioned Daily Water and Nutrient Fusion

3.4.1 Independent and Shared Prediction Paths

The daily resource target is

$$\mathbf{y}_t = [y_{W,t}, y_{N,t}, y_{P,t}, y_{K,t}]. \quad (9)$$

Both prediction paths used the same 31 predictors: days after planting; daily and cumulative growing degree days above 10°C (GDD10); phenology progress; greenhouse mean air temperature, relative humidity, CO₂ concentration, humidity deficit, total PAR and lamp PAR; outdoor temperature, relative humidity, wind speed and global radiation; current root-zone water content and electrical conductivity; current drainage; cumulative yield; process-based LAI, crop dry matter and daily dry-matter increment; process-based water and N, P and K demand indicators; one-day-lagged water, N, P and K fluxes; and one-day-lagged root-zone water content and electrical conductivity. Features were not standardised, and missing values were replaced by medians estimated from the corresponding fitting set.

For both paths, each target was transformed as $z_j = \log[1 + \max(y_j, 0)]$ before fitting, and predictions were back-transformed as $\hat{y}_j = \max[\exp(\hat{z}_j) - 1, 0]$. No further target standardisation was applied.

The Independent path fits one single-output ExtraTrees model per resource and serves as the zero-sharing accuracy reference. Each model used 260 trees, `min_samples_leaf` = 3 and `max_features` = 0.8, with the remaining parameters at their scikit-learn defaults. The Shared path fits one four-output ExtraTrees model on the same 31 predictors to capture response patterns common to water and nutrient fluxes; it used 260 trees, `min_samples_leaf` = 4 and `max_features` = 0.75.

For resource j , the two paths produce $\hat{y}_{j,t}^{\text{ind}}$ and $\hat{y}_{j,t}^{\text{sh}}$. Fixed shared uses only the shared multi-output predictor. Global fusion combines the Independent and Shared predictions with one task-specific, time-invariant pair of non-negative coefficients. DPMF uses the same two candidate predictions but lets their relative contributions vary with phenological development.

3.4.2 Smooth Phenology Gating

Phenological development in the flux path was represented by GDD10. Daily GDD10 was $\delta_t = \max(T_t^{\text{gh}} - 10, 0)$, where T_t^{gh} is the daily mean greenhouse air temperature (°C), and cumulative GDD10 was $C_t = \sum_{u \leq t} \delta_u$. Within each compartment the phenology coordinate was $g_t = \text{clip}(C_t/C_{\text{max}}, 0, 1)$, where C_{max} is the maximum cumulative GDD10 over the complete 166-day covariate record of that compartment.

Six overlapping Gaussian basis functions were centred at $\mu_k \in (0.04, 0.19, 0.36, 0.52, 0.70, 0.91)$, corresponding to establishment, vegetative growth, flowering, fruit set, fruit filling and harvest, with a common bandwidth $\sigma = 0.115$. The responses were normalised to sum to one on every date:

$$\pi_k(g_t) = \frac{\exp[-(g_t - \mu_k)^2/(2\sigma^2)]}{\sum_{l=1}^6 \exp[-(g_t - \mu_l)^2/(2\sigma^2)]}. \quad (10)$$

Global fusion provides the time-invariant comparison. For each task j , non-negative least squares (NNLS) was applied to the Independent and Shared out-of-fold (OOF) predictions over all training OOF rows, and the two coefficients were normalised to sum to one, giving a fixed task-specific mixture over the whole crop cycle.

DPMF instead estimates stage-specific coefficients. For stage k the fitting subset was $I_k = \{t : \pi_k(g_t) > 0.20\}$. Within each subset, NNLS combined the Independent and Shared OOF predictions and the resulting pair was normalised to sum to one. If fewer than eight observations were available for a stage, the global coefficients were used. To limit negative transfer, a stage-specific blend was kept only if its OOF mean squared error was at least 3% lower than that of the Independent predictor on the same subset; otherwise the stage coefficients were set to (1, 0).

The stage-specific coefficient pairs $(\alpha_{k,j}, \beta_{k,j})$ were interpolated with the phenology weights to obtain daily effective weights, and the final DPMF prediction is

$$w_{j,t}^{\text{ind}} = \sum_{k=1}^6 \pi_k(g_t) \alpha_{k,j}, \quad w_{j,t}^{\text{sh}} = \sum_{k=1}^6 \pi_k(g_t) \beta_{k,j}, \quad (11)$$

$$\hat{y}_{j,t} = \max \left\{ 0, w_{j,t}^{\text{ind}} \hat{y}_{j,t}^{\text{ind}} + w_{j,t}^{\text{sh}} \hat{y}_{j,t}^{\text{sh}} \right\}. \quad (12)$$

Because $\alpha_{k,j} + \beta_{k,j} = 1$ for every stage and $\sum_k \pi_k(g_t) = 1$, it follows that $w_{j,t}^{\text{ind}} + w_{j,t}^{\text{sh}} = 1$.

The gate is thus a Gaussian-stage interpolation of stage-specific, simplex-normalised NNLS coefficients, not a separately trained neural gating network. Fusion coefficients were estimated only from leave-one-compartment-out predictions within the four training compartments; the validation compartment was not used to fit either the global or the stage-specific coefficients.

4 Experiments

4.1 Experimental Protocols and Evaluation Criteria

4.1.1 State-Reconstruction Protocol

Seven leave-one-treatment-out folds were formed. In each fold, both cultivar sequences of one light-temperature treatment were excluded from training and evaluated together, so every destructive sample was predicted by a model that had not seen its treatment. Point, Interpolation and DPMF used identical settings and training procedures and differed only in how biomass supervision was organised in time (Section 3.3.2). VegSyst was kept as an external process-based reference in the sequence-level plots and was not treated as a matched-architecture baseline.

All state models were trained with AdamW (learning rate 0.006, weight decay 10^{-4}) for 300 epochs, using a cosine-annealing learning-rate schedule and global gradient-norm clipping at 5.0. Each update used one padded full batch containing all training sequences of the fold, without early stopping.

4.1.2 Daily-Flux Protocol

Four configurations were compared on the same 154 Reference-compartment dates. Independent comprises four single-output ExtraTrees models and is the zero-sharing reference. Fixed shared comprises one four-output ExtraTrees model. Global fusion combines the Independent and Shared predictions with task-specific, time-invariant coefficients. DPMF applies phenology-conditioned stage-specific coefficients together with the Independent-protection rule of Section 3.4.2.

AICU, Automatoes, Digilog and IUACAAS formed the four training compartments, TheAutomators the validation compartment, and Reference was reserved for independent testing. For each seed, leave-one-compartment-out prediction across the four training compartments generated 616 out-of-fold rows (4×154), from which the global and stage-specific fusion coefficients were fitted. The final Independent and Shared predictors were then refitted on the combined training and validation compartments (770 rows, 5×154), and the learned fusion coefficients were applied unchanged to the 154 valid Reference dates.

The Reference targets were not used for ExtraTrees fitting, feature-imputation statistics, fusion-coefficient fitting or hyperparameter selection. The Reference phenology coordinate used the complete Reference temperature record for full-season GDD10 normalisation but no Reference water or nutrient targets. The reported DPMF results do not use sequential target-bias assimilation.

Predictions were first averaged across the five seeds for every sample or date, and RMSE and MAE were then calculated from the ensemble prediction:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i)^2}, \quad (13)$$

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i|. \quad (14)$$

Where a seed-level RMSE (mean $\pm$ SD) is reported, it was computed separately from the five individual seed RMSE values using the sample standard deviation.

To measure how much single-task accuracy is retained when the four outputs are organised within one framework, Independent is used as the zero-sharing reference. For a joint method m , the equal-weight mean relative RMSE gap is

$$G_m = \frac{1}{4} \sum_{j \in \{W, N, P, K\}} \frac{\text{RMSE}_{m,j} - \text{RMSE}_{\text{ind},j}}{\text{RMSE}_{\text{ind},j}} \times 100\%. \quad (15)$$

A smaller G_m means that the unified method stays closer to the task-specific Independent reference.

The implemented components and fixed experimental settings are summarised in Table 1.




Figure 2. Representative sampling-date total-dry-matter reconstruction under two contrasting temperatures. (a) Orange cultivar at 20.8 °C and DLI 17.5 mol m⁻² d⁻¹. (b) Orange cultivar at 29.1 °C and DLI 17.5 mol m⁻² d⁻¹. Open circles denote observed total dry matter at destructive sampling dates; coloured markers denote predictions from DPMF, Interpolation, Point and VegSyst.




Figure 3. Leave-one-treatment-out evaluation of total-dry-matter reconstruction. (a) Agreement between DPMF predictions and observations for 64 held-out destructive samples; the dashed line is the 1:1 line, and marker shapes distinguish the seven held-out light-temperature treatments. (b) Paired ensemble RMSE of DPMF, Interpolation and Point for each of the seven held-out treatments; lower values indicate better reconstruction.

4.2 Experimental Results

4.2.1 Multi-timescale Supervision and Crop-State Reconstruction

Figure 2 shows sampling-date reconstructions for two contrasting temperatures. Under the 20.8 °C treatment, DPMF followed the rise in observed total dry matter at the later sampling dates more closely than Interpolation and Point (Figure 2(a)). Under the 29.1 °C treatment, DPMF was not the closest prediction at the final sampling date (Figure 2(b)). Reconstructions for all 14 cultivar-treatment sequences

are given in Figure A1 (Appendix A).

Over the complete held-out test set, DPMF achieved an ensemble RMSE of 132.821 g m⁻² and an ensemble MAE of 65.246 g m⁻² on 64 destructive samples (Figure 3(a)), and the reconstructed state increased consistently with observed total dry matter across the full measurement range. Treatment-wise errors are shown in Figure 3(b). DPMF gave the lowest paired RMSE in four of the seven held-out treatments (D1, D2, D2a and D4), whereas Interpolation was lowest in D3 and D3a. In D2a, DPMF reached 35.083 g m⁻²,




Figure 4. Ensemble RMSE of the four modelling strategies for daily net (a) water, (b) N, (c) P and (d) K fluxes on the 154 Reference test dates. All strategies used the same dates, response definitions, historical inputs and update protocol; lower values are better.

compared with 52.563 g m^{-2} for Interpolation and 58.672 g m^{-2} for Point.

On the complete common test set, Interpolation and Point obtained ensemble RMSE values of 141.428 and 140.263 g m^{-2} , respectively, whereas DPMF obtained 132.821 g m^{-2} (Table 2). Multi-timescale supervision therefore reduced RMSE by 6.09% relative to Interpolation and by 5.31% relative to Point.

4.2.2 Phenology-Conditioned Fusion and Single-Task Accuracy

DPMF improved all four daily resource estimates relative to Fixed shared under the matched AGC2 protocol (Figure 4; Table 3). Water RMSE decreased from 0.6371 to 0.6101 mm d^{-1} (4.24%), and N RMSE from 2.7277 to $2.5819 \text{ kg ha}^{-1} \text{ d}^{-1}$ (5.34%). P RMSE decreased from 0.5907 to $0.5759 \text{ kg ha}^{-1} \text{ d}^{-1}$ and K RMSE from 5.6531 to $5.6153 \text{ kg ha}^{-1} \text{ d}^{-1}$, reductions

of 2.51% and 0.67%, respectively (Figure 4(a-d)).

Table 4 summarises how much single-task accuracy each unified strategy retained. The equal-weight mean relative RMSE increase over Independent was 5.28% for Fixed shared, 2.89% for Global fusion and 1.91% for DPMF; DPMF thus reduced the mean increase by 3.37 and 0.98 percentage points relative to Fixed shared and Global fusion, respectively. The advantage was not uniform across tasks. For water and N, DPMF was within 1.6% of Independent; for P, Global fusion reproduced the Independent result and was marginally better than DPMF; and for K, all three unified strategies remained more than 5% above Independent.

Figure 5 shows the daily flux trajectories over the Reference test period, with the ensemble predictions of Fixed shared, Global fusion, Independent and DPMF




Figure 5. Daily time series of observed and predicted net fluxes over the Reference test period: (a) water, (b) N, (c) P and (d) K. The black line is the reference flux on each of the 154 test dates; coloured lines are five-seed ensemble predictions from Fixed shared, Global fusion, Independent and DPMF. Water is in mm d^{-1} , and N, P and K are in elemental $\text{kg ha}^{-1} \text{d}^{-1}$.

plotted against the observed day-to-day variation on the same 154 dates. For water, DPMF followed the seasonal increase and short-term variability and had a lower full-period RMSE than Fixed shared and Global fusion (Figure 5(a)). The N, P and K trajectories (Figure 5(b-d)) show when the task-specific RMSE differences in Figure 4 and Table 3 arise.

4.3 Discussion

4.3.1 Multi-timescale Supervision Reconstructs Sparse Crop States

The state experiment shows that the way supervision is organised affects crop-state reconstruction. Point-only training anchors the model at destructive sampling dates but says little about the change between them. Interpolation produces a dense trajectory, but its intermediate labels are constructed, not observed. DPMF keeps the measured points and adds two further kinds of recorded evidence: dry-matter change between adjacent samples and cumulative harvested-fruit production. The state loss sets the

biomass level, the interval loss constrains growth over each sampling interval, and the fruit loss ties the latent trajectory to reproductive output. Together they lowered both RMSE and MAE on the pooled held-out treatments. The gain was, however, treatment-dependent: DPMF was best in four of seven treatments, and Interpolation was more accurate in the two D3 treatments. Identifying which conditions favour interval and fruit supervision requires more treatments than are available in the present dataset.

4.3.2 Phenology-Conditioned Fusion Retains Task-Specific Accuracy

The flux experiment asks a different question: how to organise four resource tasks within one framework without losing the accuracy of specialised models. Independent models are the reference because each output is optimised separately. Fixed shared modelling imposes one cross-task representation over the whole production cycle, Global fusion adds task-specific but constant coefficients, and DPMF lets the balance change with crop development.

Table 1. Implemented DPMF components and fixed experimental settings.

Component	Setting
State temporal model	Single-layer GRU (19 inputs, 32 hidden units) with phenology-gated process-increment correction and LAI readout
Common state objective	State loss + $0.4 \times$ LAI loss + $0.01 \times$ process-residual regularisation
DPMF state additions	$0.5 \times$ interval loss + $0.35 \times$ cumulative-fruit loss
Flux base predictors	Four single-output ExtraTrees models and one four-output ExtraTrees model
Phenology representation	Six normalised Gaussian stage bases; path-specific thermal-time coordinates and parameters
Fusion learning	Stage-wise NNLS on OOF predictions from four training compartments, with sum-to-one normalisation and Independent protection
Random seeds	11, 23, 37, 53 and 71
Principal estimate	Five-seed prediction ensemble formed before metric calculation
State training	AdamW, learning rate 0.006, weight decay 10^{-4} , 300 epochs, cosine annealing, gradient clipping at 5.0

Table 2. Overall state-reconstruction performance on 64 held-out destructive samples (g m^{-2}).

Method	Ens. RMSE	Ens. MAE	Seed RMSE (mean $\pm$ SD)
DPMF	132.821	65.246	134.238 $\pm$ 3.316
Interpolation	141.428	72.193	142.163 $\pm$ 2.239
Point	140.263	70.950	140.512 $\pm$ 1.560

Ens.: five-seed ensemble. Bold marks the proposed method.

DPMF reduced the equal-weight mean relative RMSE gap to 1.91%, against 2.89% for Global fusion and 5.28% for Fixed shared, so it came closest overall to the four Independent predictors. The gain came mainly from water and N; for P, the time-invariant mixture was already equivalent to Independent, and for K none of the unified strategies closed the gap. The per-task improvements over Fixed shared show that stage-dependent weighting is better than direct multi-output sharing, while the remaining K gap

Table 3. Daily net-flux performance on 154 Reference-compartment dates under the matched protocol.

Task	Method	Ens. RMSE	Ens. MAE	Seed RMSE
Water	Fixed shared	0.6371	0.3978	0.6379 ± 0.0042
	Global fusion	0.6183	0.3823	0.6190 ± 0.0075
	Independent	0.6010	0.3766	0.6018 ± 0.0027
	DPMF	0.6101	0.3868	0.6116 ± 0.0067
N	Fixed shared	2.7277	1.9602	2.7297 ± 0.0346
	Global fusion	2.6454	1.8689	2.6498 ± 0.0409
	Independent	2.5663	1.7888	2.5693 ± 0.0306
	DPMF	2.5819	1.8029	2.5846 ± 0.0212
P	Fixed shared	0.5907	0.4540	0.5917 ± 0.0073
	Global fusion	0.5745	0.4329	0.5759 ± 0.0022
	Independent	0.5745	0.4329	0.5759 ± 0.0022
	DPMF	0.5759	0.4360	0.5772 ± 0.0027
K	Fixed shared	5.6531	3.6566	5.6565 ± 0.0491
	Global fusion	5.6322	3.6392	5.6365 ± 0.0683
	Independent	5.3336	3.3666	5.3412 ± 0.0930
	DPMF	5.6153	3.6126	5.6202 ± 0.0337

Units: water mm d^{-1} ; N, P, K $\text{kg ha}^{-1} \text{d}^{-1}$. Seed RMSE: mean $\pm$ SD over five seeds. Bold marks the proposed method; Independent is the zero-sharing reference.

Table 4. Equal-weight mean relative RMSE increase of each unified strategy over the Independent reference across the four flux tasks.

Method	Mean relative RMSE increase (%)
Fixed shared	5.28
Global fusion	2.89
DPMF	1.91

indicates that sharing adds little for this resource.

4.3.3 Relationship to Process-Based Water and Nutrient Modelling

Process-based greenhouse models describe crop growth and resource use through explicit physiological relations. FAO-56 provides the reference evapotranspiration framework underpinning crop water demand [18]. TOMSIM and its evaluation describe dry-matter production and organ partitioning in tomato [23, 24], and the Wageningen crop models offer broader mechanistic frameworks for crop simulation [25]. HortSyst simulates dry matter, N uptake, LAI and transpiration for greenhouse tomato [19], with uncertainty quantified across the full N, P, K and water variable set [20]. GreenLight

integrates greenhouse climate, energy and crop production [26]. The VegSyst family has been validated for crop growth, evapotranspiration and N uptake for greenhouse tomato [27], extended to model N, P and K uptake [31], and built into decision-support systems for water and nutrient management of tomato [28, 29]; later versions have been applied under prescriptive-corrective management and across different regions, substrates and soilless systems [32–36]. Models for deficit irrigation of greenhouse tomato have also been developed and validated [22]. Broader reviews of dry matter production and growth modelling for greenhouse tomatoes [30], and data assimilation approaches using Kalman-family filters to update state estimates of greenhouse tomato growth models from in-season observations [21], further extend the applicability of these process-based frameworks. Building on this landscape, DPMF works on the observation and task-organisation layers around such models: multi-timescale supervision lets sparse biomass, interval growth and cumulative harvest constrain one temporal representation, and phenology-conditioned fusion uses development as an explicit coordinate for weighting water and nutrient predictions. Crop development therefore does not enter only as an input feature; it also controls how task-specific and shared estimates are combined. The outputs, one crop-state trajectory and four daily resource series, keep their original units and temporal supports.

4.3.4 Implications and Limitations

In an exploratory replay, the same DPMF predictions were passed through alternative feedback rules to show how predicted demand translates into simulated root-zone water and nutrient trajectories (Figure B1, Appendix B). A scenario-based simulation further examined the sensitivity of downstream outcomes to imposed changes in resource demand and initial root-zone conditions (Figure C1, Appendix C). In these simulations the DPMF prediction-control configuration kept simulated resource-deficit exposure low while operating close to the reference input level. These results illustrate a possible downstream use and are not an independent validation of closed-loop greenhouse control.

The daily-flux evaluation has three further limitations. First, the predictors include one-day-lagged observed water, N, P and K fluxes and contemporaneous greenhouse, root-zone and drainage variables, so the task is daily net-flux estimation rather than

prospective forecasting. Second, because the reference flux is defined as supply minus drainage, the contemporaneous drainage predictor shares a component with the target; the reported accuracies are therefore conditional on observed same-day drainage. Third, the phenology coordinate is normalised by the maximum cumulative GDD10 of the complete compartment record, which uses full-season temperature covariates, although no test-compartment water or nutrient targets enter this normalisation. Prospective deployment will require strictly online covariates and causal phenology normalisation, and closed-loop fertigation trials will be needed to confirm the simulated benefits.

5 Conclusion

DPMF couples crop-state reconstruction with phenology-conditioned daily resource prediction. Multi-timescale supervision reduced total-dry-matter RMSE to 132.821 g m^{-2} on 64 held-out destructive samples, 6.09% and 5.31% lower than Interpolation and Point supervision, although the gain varied between treatments. Phenology-conditioned fusion reduced daily net-flux RMSE relative to Fixed shared modelling by 4.24%, 5.34%, 2.51% and 0.67% for water, N, P and K, and gave the smallest equal-weight mean RMSE gap to the Independent reference (1.91%, against 2.89% for Global fusion and 5.28% for Fixed shared), with most of the benefit for water and N. Future work will test the framework with strictly online inputs and in closed-loop fertigation experiments.

Data Availability Statement

The WUR dwarf-tomato dataset used in this study is publicly available from 4TU.ResearchData at <https://doi.org/10.4121/fa102772-32db-4b30-bace-12f2016722ce.v1>. The AGC2 greenhouse dataset is publicly available from 4TU.ResearchData at <https://doi.org/10.4121/uuid:88d22c60-21b3-4ea8-90db-20249a5be2a7>. Processed data and code are available from the corresponding authors upon reasonable request.

Funding

This work was supported by the National Major Agricultural Science and Technology Project under Grant 0304.

Conflicts of Interest

The authors declare no conflicts of interest.

AI Use Statement

The authors declare that no generative AI was used in the preparation of this manuscript.

Ethical Approval and Consent to Participate

Not applicable. This study analysed publicly available plant-cultivation datasets and did not involve human participants or animals.

References

- [1] Shamshiri, R., Kalantari, F., Ting, K. C., Thorp, K. R., Hameed, I. A., Weltzien, C., ... & Shad, Z. M. (2018). Advances in greenhouse automation and controlled environment agriculture: A transition to plant factories and urban agriculture. *International Journal of Agricultural and Biological Engineering*, 11(1), 1-22. [CrossRef]
- [2] Van Klompenburg, T., Kassahun, A., & Catal, C. (2020). Crop yield prediction using machine learning: A systematic literature review. *Computers and electronics in agriculture*, 177, 105709. [CrossRef]
- [3] Kamilaris, A., & Prenafeta-Boldú, F. X. (2018). Deep learning in agriculture: A survey. *Computers and electronics in agriculture*, 147, 70-90. [CrossRef]
- [4] Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2018). Machine learning in agriculture: A review. *Sensors*, 18(8), 2674. [CrossRef]
- [5] Lim, B., & Zohren, S. (2021). Time-series forecasting with deep learning: a survey. *Philosophical transactions of the royal society a: mathematical, physical and engineering sciences*, 379(2194). [CrossRef]
- [6] Cho, K., Van Merriënboer, B., Gulçehre, Ç., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014, October). Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)* (pp. 1724-1734). [CrossRef]
- [7] Caruana, R. (1997). Multitask learning. *Machine Learning*, 28(1), 41-75. [CrossRef]
- [8] Evgeniou, T., & Pontil, M. (2004, August). Regularized multi-task learning. In *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining* (pp. 109-117). [CrossRef]
- [9] Kendall, A., Gal, Y., & Cipolla, R. (2018). Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (pp. 7482-7491). [CrossRef]
- [10] Zhang, Y., & Yang, Q. (2021). A survey on multi-task learning. *IEEE transactions on knowledge and data engineering*, 34(12), 5586-5609. [CrossRef]
- [11] Ma, J., Zhao, Z., Yi, X., Chen, J., Hong, L., & Chi, E. H. (2018). Modeling task relationships in multi-task learning with multi-gate mixture-of-experts. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 1930-1939). [CrossRef]
- [12] Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5-32. [CrossRef]
- [13] Geurts, P., Ernst, D., & Wehenkel, L. (2006). Extremely randomized trees. *Machine learning*, 63(1), 3-42. [CrossRef]
- [14] Hemming, S., de Zwart, F., Elings, A., Petropoulou, A., & Righini, I. (2020). Cherry tomato production in intelligent greenhouses—Sensors and AI for control of climate, irrigation, crop yield, and quality. *Sensors*, 20(22), 6430. [CrossRef]
- [15] Hemming, S., de Zwart, F., Elings, A., Petropoulou, A., & Righini, I. (2020). Autonomous greenhouse challenge, (2019). 4TU. *ResearchData Data Set*, 678. [CrossRef]
- [16] Maree, S. C., Zhang, P., van Marrewijk, B. M., de Zwart, F., Bijlaard, M., & Hemming, S. (2025). Autonomous greenhouse cultivation of dwarf tomato: Performance evaluation of intelligent algorithms for multiple-sensor feedback. *Sensors*, 25(14), 4321. [CrossRef]
- [17] Maree, S., Zhang, P., van Marrewijk, B., de Zwart, F., Bijlaard, M., & Hemming, S. (2024). 4th Autonomous Greenhouse Challenge: Dwarf Tomato Timeseries and Images. 4TU. *ResearchData*. [CrossRef]
- [18] Allen, R. G., Pereira, L. S., Raes, D., & Smith, M. (1998). Crop evapotranspiration-Guidelines for computing crop water requirements-FAO Irrigation and drainage paper 56. *Fao, rome*, 300(9), D05109. <https://www.fao.org/3/X0490E/X0490E00.htm>
- [19] Martínez-Ruiz, A., López-Cruz, I. L., Ruiz-García, A., Pineda-Pineda, J., & Prado-Hernández, J. V. (2019). HortSyst: A dynamic model to predict growth, nitrogen uptake, and transpiration of greenhouse tomatoes. *Chilean journal of agricultural research*, 79(1), 89-102. [CrossRef]
- [20] Martínez-Ruiz, A., López-Cruz, I. L., Ruiz-García, A., Pineda-Pineda, J., Sánchez-García, P., & Mendoza-Pérez, C. (2021). Uncertainty analysis of the HORTSYST model applied to fertigated tomatoes cultivated in a hydroponic greenhouse system. *Spanish Journal of Agricultural Research*, 19(3), e0802-e0802. [CrossRef]
- [21] de Oliveira, M. P. G., Zorzeto-Cesar, T. Q., de Faissol Attux, R. R., & Rodrigues, L. H. A. (2025). Leveraging data from plant monitoring into crop models. *Information Processing in Agriculture*, 12(3), 408-429. [CrossRef]
- [22] Locatelli, S., Barrera, W., Jr., Verdi, L., Nicoletto, C., Dalla Marta, A., & Maucieri, C. (2024). Modelling the response of tomato on deficit irrigation under greenhouse conditions. *Scientia Horticulturae*, 326,

112770. [[CrossRef](#)]
- [23] Heuvelink, E. (1995). Dry matter production in a tomato crop: Measurements and simulation. *Annals of Botany*, 75(4), 369-379. [[CrossRef](#)]
- [24] Heuvelink, E. (1999). Evaluation of a dynamic simulation model for tomato crop growth and development. *Annals of Botany*, 83(4), 413-422. [[CrossRef](#)]
- [25] van Ittersum, M. K., Leffelaar, P. A., Van Keulen, H., Kropff, M. J., Bastiaans, L., & Goudriaan, J. (2003). On approaches and applications of the Wageningen crop models. *European journal of agronomy*, 18(3-4), 201-234. [[CrossRef](#)]
- [26] Katzin, D., van Mourik, S., Kempkes, F., & van Henten, E. J. (2020). GreenLight—An open source model for greenhouses with supplemental lighting: Evaluation of heat requirements under LED and HPS lamps. *Biosystems Engineering*, 194, 61-81. [[CrossRef](#)]
- [27] Gallardo, M., Fernández, M. D., Giménez, C., Padilla, F. M., & Thompson, R. B. (2016). Revised VegSyst model to calculate dry matter production, critical N uptake and ETc of several vegetable species grown in Mediterranean greenhouses. *Agricultural systems*, 146, 30-43. [[CrossRef](#)]
- [28] Gallardo, M., Thompson, R. B., Giménez, C., Padilla, F. M., & Stöckle, C. O. (2014). Prototype decision support system based on the VegSyst simulation model to calculate crop N and water requirements for tomato under plastic cover. *Irrigation Science*, 32(3), 237-253. [[CrossRef](#)]
- [29] Gallardo, M., Elia, A., & Thompson, R. B. (2020). Decision support systems and models for aiding irrigation and nutrient management of vegetable crops. *Agricultural Water Management*, 240, 106209. [[CrossRef](#)]
- [30] Higashide, T. (2022). Review of dry matter production and growth modelling to improve the yield of greenhouse tomatoes. *The Horticulture Journal*, 91(3), 247-266. [[CrossRef](#)]
- [31] Gallardo, M., Cuartero, J., de la Torre, L. A., Padilla, F. M., Segura, M. L., & Thompson, R. B. (2021). Modelling nitrogen, phosphorus, potassium, calcium and magnesium uptake, and uptake concentration, of greenhouse tomato with the VegSyst model. *Scientia Horticulturae*, 279, 109862. [[CrossRef](#)]
- [32] Gallardo, M., Peña-Fleitas, M. T., Padilla, F. M., Cedeño, J., & Thompson, R. B. (2023). Prescriptive-corrective irrigation and macronutrient management in greenhouse soil-grown tomato using the vegsyst-DSS v2 decision support tool. *Horticulturae*, 9(10), 1128. [[CrossRef](#)]
- [33] Berrueta, C., Grasso, R., García, C., Thompson, R. B., & Gallardo, M. (2023). Use of the VegSyst model to simulate seasonal dry matter production, N and K uptake and evapotranspiration in greenhouse soil-grown tomato in Uruguay. *Agricultural Water Management*, 286, 108395. [[CrossRef](#)]
- [34] Carmassi, G., Cialli, S., Cela, F., Romero, E. B., Gallardo, M., & Incrocci, L. (2025). Use of the VegSyst v3 model to simulate seasonal dry matter production, uptake of nutrients, and evapotranspiration in a greenhouse soilless tomato crop in Tuscany. *Agricultural Water Management*, 314, 109508. [[CrossRef](#)]
- [35] Gallardo, M., Cedeño, J. M., Magán, J. J., Fernández, M. D., & Thompson, R. B. (2025). Adaptation of VegSyst-DSS for N, P and K recommendations for grafted tomato grown in perlite in Mediterranean greenhouses. *Agricultural Water Management*, 310, 109351. [[CrossRef](#)]
- [36] Faliagka, S., Naounoulis, I., Gallardo, M., & Katsoulas, N. (2026). Adaptation of the VegSyst model to predict crop nutrient uptake and water needs for precise soilless crop fertigation in greenhouses. *Irrigation Science*, 44(3), 66. [[CrossRef](#)]
- [37] Jin, X., Kumar, L., Li, Z., Feng, H., Xu, X., Yang, G., & Wang, J. (2018). A review of data assimilation of remote sensing and crop models. *European journal of agronomy*, 92, 141-152. [[CrossRef](#)]
- [38] Zhou, X., Liu, Q., Katzin, D., Qian, T., Heuvelink, E., & Marcelis, L. F. M. (2023). 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. [[CrossRef](#)]
- [39] Gong, L., Yu, M., Cutsuridis, V., Kollias, S., & Pearson, S. (2022). A novel model fusion approach for greenhouse crop yield prediction. *Horticulturae*, 9(1), 5. [[CrossRef](#)]
- [40] Florakis, K., Trevezas, S., & Letort, V. (2023). Predicting tomato water consumption in a hydroponic greenhouse: contribution of light interception models. *Frontiers in Plant Science*, 14, 1264915. [[CrossRef](#)]
- [41] Simm, J., De Abril, I. M., & Sugiyama, M. (2014). Tree-based ensemble multi-task learning method for classification and regression. *IEICE TRANSACTIONS on Information and Systems*, 97(6), 1677-1681. [[CrossRef](#)]
- [42] Ge, J., Zhao, L., Yu, Z., Liu, H., Zhang, L., Gong, X., & Sun, H. (2022). Prediction of greenhouse tomato crop evapotranspiration using XGBoost machine learning model. *Plants*, 11(15), 1923. [[CrossRef](#)]
- [43] Gong, L., Yu, M., Jiang, S., Cutsuridis, V., & Pearson, S. (2021). Deep learning based prediction on greenhouse crop yield combined TCN and RNN. *Sensors*, 21(13), 4537. [[CrossRef](#)]
- [44] Lin, Y. S., Fang, S. L., Kang, L., Chen, C. C., Yao, M. H., & Kuo, B. J. (2024). Combining recurrent neural network and sigmoid growth models for short-term temperature forecasting and tomato growth prediction in a plastic greenhouse. *Horticulturae*, 10(3), 230. [[CrossRef](#)]

Appendix

A. Full Cultivar-Treatment Reconstruction Results

Figure A1 gives the sampling-date reconstruction results for all 14 cultivar-treatment sequences of the leave-one-treatment-out experiment, i.e., the two cultivars (Orange and Red) under seven light-temperature treatments. Panels (a) and (i) are the two cases shown in Figure 2. The panels illustrate treatment- and cultivar-dependent variation in reconstruction accuracy; no single method was best under all conditions.

B. Simulated Root-Zone Dynamics Under Alternative Control Rules

To examine how predicted demand could translate into root-zone dynamics, the same five-seed ensemble DPMF predictions were replayed under three control rules: Direct, Balance and DPMF control (Figure B1). Direct supplies resources according to the predicted demand; Balance adds feedback from the simulated root-zone inventory; and DPMF control combines inventory feedback with a prediction window and constraints on action changes. The three simulations thus differ in the control rule, not in the prediction model. The replay covered 154 simulation steps; because dates with missing records were excluded, the step axis represents ordered valid observations rather than a continuous calendar. Feedback-based control changed the timing of supply and shortened low-inventory periods relative to Direct control. For water, Direct control produced extended periods of low or depleted storage, whereas Balance and DPMF control generally kept storage higher, including during the marked depletion around steps 78-81 (Figure B1(a)). The nutrient inventories behaved similarly: Direct control approached depletion for N and stayed at relatively low P and K levels in several periods, whereas the feedback-based rules generally kept larger stocks (Figure B1(b)). These are exploratory outcomes from a simplified root-zone reservoir model, not an independent validation of closed-loop control.

C. Scenario-Based Evaluation of Resource-Management Strategies

A scenario-based simulation examined how alternative prediction-control strategies behaved under changes in resource demand and initial root-zone conditions. Seven strategy configurations were evaluated under six predefined scenarios, giving 42 simulation cases of 154 steps each. Normal used unit multipliers for water, N, P and K and an initial reservoir level of 55% of capacity. High ET (evapotranspiration) used multipliers of 1.25, 1.05, 1.00 and 1.10 for water, N, P and K; Low light used 0.75, 0.85, 0.85 and 0.85 with an initial reservoir level of 65%; Low initial kept unit multipliers but lowered the initial level to 25%; Long-cycle used 1.08, 1.10, 1.08 and 1.12; and High yield used 1.10, 1.20, 1.15 and 1.25. Unless stated otherwise, the initial reservoir level was 55%, and all multipliers were applied over the full 154-step replay. These scenarios are controlled sensitivity settings, not observed greenhouse experiments or explicit simulations of the processes implied by their names.

The evaluated configurations were DPMF (DPMF predictions with feedback- and prediction-window-based control), Independent + Smooth (independent task predictions with smoothed control), Quota (stage-based water and nutrient quotas), ET rule (evapotranspiration-based water prediction with nutrient quotas and threshold feedback), VegSyst (VegSyst predictions with direct control), Forecast (VegSyst predictions with forecast-based control) and Combination (element-wise median of candidate predictions with balance control). Modelled yield retention was used only as a simulation-derived summary of resource-deficit exposure. Mean unmet-demand fractions for water, N, P and K were weighted by 0.45, 0.25, 0.08 and 0.22, respectively, and converted to percentage retention as $100 \exp(-0.24E)$, where E is the weighted exposure. A value of 100% is the modelled no-deficit reference and is not measured yield.

Across the six scenarios, DPMF produced 1-2 stress steps per simulation (mean 1.50), whereas Independent + Smooth produced 3-9 (mean 5.83) (Figure C1(a)). The corresponding modelled yield retention for DPMF ranged from 99.983% to 99.993% (mean 99.988%) (Figure C1(b)). DPMF had a mean normalised input ratio of 1.004 and a mean modelled yield retention of 99.988%, against 0.993 and 99.928% for Independent + Smooth (Figure C1(c)). Within this exploratory reservoir simulation, the DPMF prediction-control configuration therefore kept modelled deficit exposure low across the imposed perturbations while operating near the reference input level.






Figure A1. Sampling-date reconstructions for all 14 cultivar-treatment sequences. Within each of the seven held-out light-temperature treatments, the Orange cultivar is shown first and the Red cultivar second. Open circles denote observed total dry matter at destructive sampling dates; coloured markers denote predictions from DPMF, Interpolation, Point and VegSys. Horizontal axis: days after sowing (DAS); vertical axis: total dry matter (g m^{-2}).



(a)



(b)

Figure B1. Simulated root-zone water and nutrient dynamics under alternative control rules driven by the same DPMF ensemble predictions. (a) Irrigation actions and simulated available root-zone water storage under Direct, Balance and DPMF control. (b) Simulated available root-zone N, P and K stocks under the same rules. Horizontal dotted, dashed and dash-dotted lines mark 25%, 55% and 85% of the predefined reservoir capacities (low-inventory diagnostic level, initial/feedback target and overflow threshold). Water storage is in mm, and nutrient stocks are in elemental kg ha^{-1} .




Figure C1. Exploratory scenario-based evaluation of resource-management strategies. (a) Number of stress steps for each strategy-scenario combination; a stress step is a simulation step in which the unmet-demand fraction exceeded 10% for at least one resource (lower is better). (b) Modelled yield retention derived from simulated unmet-demand exposure (higher means lower modelled yield loss); this is a simulation indicator, not measured yield. (c) Mean water/N/P/K input-to-reference ratio versus modelled yield retention for the 42 strategy-scenario combinations; the lower panel enlarges DPMF and Independent + Smooth. The vertical line at a ratio of 1 marks equality between mean normalised input and reference demand. Colours identify strategies and marker shapes identify scenarios.



Chenlei Fang is a Ph.D. candidate in Computer Science and Technology at Northwestern Polytechnical University, China. His research covers large language models, reinforcement learning and representation learning, with interests in multimodal perception, resource-constrained autonomous decision-making and applications of artificial intelligence in precision agriculture, particularly crop-state modelling and intelligent agricultural management. (Email: chenleifang@mail.nwpu.edu.cn)



Yisheng Miao is a Research Professor at the Information Technology Research Center, Beijing Academy of Agriculture and Forestry Sciences. His research interests include the development and application of artificial intelligence in agriculture, particularly agricultural multimodal large models, intelligent vegetable production and autonomous farming systems. (Email: Miaoys@nercita.org.cn)



Youlin Cui is a Senior Agronomist at the Information Technology Research Center, Beijing Academy of Agriculture and Forestry Sciences. His research interests include unmanned vegetable production, intelligent agricultural machinery and agricultural informatisation, particularly the mechanisation and intelligentisation of vegetable production, unmanned vegetable farms, unmanned harvesting equipment, vegetable data collection, and intelligent management and control. (Email: cuiyl@nercita.org.cn)



Chunjiang Zhao is an Academician of the Chinese Academy of Engineering, a Professor and a Doctoral Supervisor. He is Director of the National Engineering Research Center for Information Technology in Agriculture. His research focuses on agricultural information technology, precision agriculture, and intelligent agricultural systems and equipment, including agricultural decision-support systems, intelligent agricultural machinery, precision irrigation and fertilisation, and intelligent greenhouse management. (Email: zhaocj@nercita.org.cn)