

Context-Conditioned Operator Learning for Robust Predictive Control in Environmental Data Systems

Dajana Dujmović^{1,*}

¹ Faculty of Information Technology, University of Split, Split, 21000, Croatia

*Corresponding author: dajana.du@unist.hr

Abstract. Greenhouse climate control requires forecasts that remain reliable when solar radiation, outdoor weather, crop development, and actuator operating modes alter the governing dynamics. Conventional recurrent networks compress spatial measurements into fixed vectors, whereas operator-learning models may overlook slow thermal and moisture storage. This paper develops a context-aware Fourier neural operator–long short-term memory (FNO–LSTM) model for multizone temperature, relative-humidity, and carbon-dioxide forecasting and receding-horizon control. A spectral branch learns nonlocal zone coupling, a recurrent branch preserves delayed responses, and a context gate conditions their fusion on crop stage, weather regime, time of day, and actuator availability. A control-oriented objective penalizes field error, transition inconsistency, condensation risk, and excessive actuation. In a replaceable 120-day, six-zone greenhouse experiment sampled every 5 min, the model achieved temperature MAE of 0.63 °C, humidity MAE of 3.1 percentage points, and carbon-dioxide MAE of 41 ppm at a 60-min horizon. Relative to an LSTM controller, it reduced temperature error by 46.6%, energy use by 12.8%, and actuator reversals by 21.4%, while maintaining 98.6% climate-constraint satisfaction. The results indicate that explicit spatial operators and context-conditioned temporal memory provide a practical basis for accurate, stable, and computationally feasible greenhouse climate control.

Keywords: *Greenhouse Climate Control, Fourier Neural Operator, Long Short-Term Memory, Context Awareness, Predictive Control, Multivariate Forecasting*

Received on 09 March 2022, Accepted on 13 July 2022, Published on 25 July 2022

Copyright © 2022 Author(s), licensed to DEA. This is an open access article distributed under the terms of the CC BY-NC-SA 4.0, which permits copying, redistributing, remixing, transformation, and building upon the material in any medium so long as the original work is properly cited.

Introduction

Crop production away from some of the outdoor climate is known as "controlled-environment agriculture," but this benefit necessitates maintaining air temperature, humidity, radiation, and carbon dioxide concentration within ranges that can interact with physiological limits. The vapor-pressure deficit and, consequently, transpiration demand will decrease with a one-degree temperature increase [1]. Although vent opening eliminates sensible and latent heat, it may raise the heating load or lower the CO₂ concentration too soon [2]. The temperature of the air will drop due to fog, and condensation may happen near a dense tree [3]. Photosynthesis, heat load, and electricity consumption are all impacted simultaneously by more light [4]. Instead of being a set of separate set-point loops, these couplings transform the greenhouse regulation into a multivariable problem [5]. Spatial heterogeneity makes the issue worse: crop transpiration is impacted by varying canopy densities, solar radiation is comparatively high at exposed sites, only some vents introduce cold air, and injected CO₂ is dispersed unevenly. Thus, a command given at one wall or roof vent will have a delayed impact in multiple locations. Non-local transport, local boundary effects, and any residual thermal or moisture storage after the initial disturbance has passed should all be taken into account in a good forecast. During a control horizon, it must also maintain stability in the event that measurements are missing or equipment availability shifts.

Heating, ventilation, fogging, carbon dioxide injection, and lighting can all be coordinated by Model Predictive Control over an extended period of time, but the model's repeated evaluation within the optimiser limits its

closed-loop performance [6]. Although interpretable heat and mass balances are revealed by mechanistic greenhouse models, their coefficients vary depending on canopy growth, leakage, cover ageing, irrigation techniques, and seasonal management [7]. Although data-driven recurrent models can learn delayed responses without the need for repeated physical calibration, vertical and horizontal gradients may be lost if all zone measurements are compressed into a single hidden state [8]. Convolutional networks are slow to simulate the spread of abrupt radiation or ventilation disturbances because they have local geometry but need multiple layers to connect remote areas [9]. In order to store months' worth of high-frequency sensor and actuator history, transformer variants provide direct long-range communication at the expense of higher memory and latency [10]. Another method for representing global interactions in a compact spectral basis and learning mappings of distributed fields is the use of neural operators [11]. Standard neural operators commonly use one parameter set shared across all operating regimes. The efficient mapping of cloudy and sunny conditions, young and mature crop stages, sealed and ventilated modes, and normal and restricted equipment states is always changing in a functional greenhouse. Consequently, a good controller should have a predictor that does not assume stationary dynamics and integrates spatial transfer, temporal memory, and explicit operating context.

This paper proposes a context-aware Fourier neural operator-long short-term memory architecture for the closed-loop requirements identified above. The Fourier branch learns spatial correlations among the six climate zones, while the recurrent branch retains latent information about the slab, structure, canopy water, and carbon-dioxide uptake. A learned gate controls the contribution of the two representations at each forecast horizon using crop stage, solar phase, weather regime, and equipment status. The fused trajectory enters a receding-horizon controller whose objective considers smooth actuator movement, transition consistency, field accuracy, and condensation avoidance. The four main contributions are field-to-sequence encoding, context-modulated spectral transfer, horizon-dependent recurrent fusion, and control-oriented validation. Each mechanism is linked to a corresponding ablation or operational metric. Section 2 reviews greenhouse climate modelling, neural operators, spatiotemporal learning, and predictive control. Section 3 presents the field construction, context encoder, modulated spectral blocks, recurrent memory, fusion gate, decoder, and training objective. Section 4 describes the reproducible protocol and illustrative dataset and then evaluates forecasting, robustness, closed-loop behavior, and edge-computing feasibility. Section 5 summarizes the engineering contribution, limitations, and future work.

Related Work

Greenhouse Climate Modeling and Control

Heat, water vapour, and CO₂ balances have been coupled to model a greenhouse climate. Because lumped models can quickly integrate a small number of states, they are appropriate for real-time control [12]. However, cover characteristics, leakage, crop leaf area, and ventilation geometry continue to affect the parameters, and these quantities are rarely constant throughout the production cycle [13]. Although the space gradient has been better resolved by computational fluid dynamics, repeated optimisation every five minutes is too costly [14]. In the presence of a strong radiation ramp and ventilation change, reduced-order identification is less accurate but more practical [15]. Hybrid approaches that take physical constraints into account when learning uncertain residual dynamics have been inspired by these shortcomings [16]. How to handle distributed transport without introducing a state dimension that renders online control impractical is the current challenge [17].

Fuzzy logic, robust control, and model predictive control have replaced thermostat rules as control strategies. A rule base can enforce separate bounds for heating, ventilation, and fogging, but it does not explicitly optimize their coupled actions over a prediction horizon. These devices are explicitly coordinated over time by predictive controllers, which prevent late correction based on crop demand and weather forecasts. Their sensitivity to model mismatch in canopy conductance changes, irrigation events, and fast cloud passage is a practical weakness. Instead of being optimised only for the average season, a good predictor for this environment should be sensitive to operating conditions, fast enough for repeated rollouts, and stable in the presence of missing sensor data.

Spatiotemporal Learning and Neural Operators

Because of their gates that retain delayed information, long short-term memory networks and gated recurrent units are widely used in greenhouse temperature, humidity, and energy forecasting [18]. Actuator commands and outdoor weather are examples of recurrent inputs; however, the spatial inductive bias is lost when all zones are flattened into a single vector. Sensor adjacency and airflow paths are encoded by graph neural networks; however, the graph needs to be explicitly defined and may become inaccurate when vent configuration changes [19]. Although attention-based sequence models are capable of learning long-range dependencies, they are challenging to apply on industrial controllers and have poor scalability for long histories [20]. The techniques demonstrate the importance of acquired temporal memory, but they are unable to accomplish field-level transport via pointwise prediction.

By using learned multipliers on specific spectral modes, Fourier Neural Operators acquire mappings of discretised fields [21]. Spectral representation is appropriate for multi-zone facilities because it can transmit across related grids and connect distant spatial locations in a single layer. However, standard operator layers do not explicitly preserve slow hidden state and are frequently trained under a fixed distribution of boundary conditions. Although simple concatenation does not ensure that the context changes the distribution of spatial-temporal reasoning, context-conditioned networks handle distribution shifts by modifying features or parameters based on regime labels [22]. In this work, a context-controlled gate is used to introduce a spectral operator with recurrent memory. Module removal, regime-specific errors, and closed-loop actuation behaviour are used to examine this coupling.

Context-Aware FNO–LSTM Method

Field Construction and Context Encoding

The controller receives six-zone measurements every five minutes. Each zone contributes air temperature, relative humidity, carbon-dioxide concentration, canopy-level radiation, and estimated vapor-pressure deficit. Outdoor temperature, humidity, wind speed, wind direction, global radiation, and precipitation form boundary inputs; heating-valve position, roof-vent opening, side-vent opening, fogging duty, fan command, carbon-dioxide injection, and lighting power describe actuation. Instead of flattening these signals, the model arranges zone variables on a fixed spatial grid and appends boundary and command channels. This representation preserves the direction of heat and vapor exchange and allows the spectral branch to distinguish a facility-wide radiation disturbance from a local sensor excursion. For reproducibility, the preprocessing record stores the two-dimensional assignment of the six zones, the channel order, and the rules used to broadcast boundary and command inputs.

Operating context is represented separately because identical measurements can require different dynamics. The context vector includes cyclical time features, a four-level crop-stage code, a weather-regime probability, binary equipment-availability flags, and a normalized leaf-area estimate. Weather regimes are obtained from a causal classifier using only observations available at the forecast origin. Crop and equipment descriptors come from the cultivation plan and supervisory controller. Context is therefore not a target-derived label and creates no future leakage. At the beginning of an independent sequence, the previous context is set to the current context so that the transition term is zero; unavailable descriptors retain their last confirmed value and are marked by the corresponding availability flag. Previous adaptive greenhouse models typically use season or day/night indicators as auxiliary inputs; the present design uses context to modulate both spectral transfer and spatial-temporal fusion.

Let the measured field at time t be $\mathbf{X}_t$, the actuator vector be $\mathbf{u}_t$, the boundary vector be $\mathbf{b}_t$, and the operating context be $\mathbf{c}_t$. For an input window of L observations, the encoded tensor combines the distributed field with broadcast command and boundary channels, while a separate multilayer encoder maps context to a compact modulation vector.

$$\mathbf{Z}_t = E_x([\mathbf{X}_t, \mathbf{1u}_t^\top, \mathbf{1b}_t^\top]), \mathbf{q}_t = E_c(\mathbf{c}_t) \quad \text{Eq.(1)}$$

The field encoder returns a spatial tensor, whereas the context path supplies global modulation. Broadcasting commands and weather leaves their local effects to the spectral kernel. The transition-aware context is consumed by every operator block; omitting it delays adaptation after faults or abrupt weather changes.

$$\mathbf{q}_t^* = \text{LN}[\mathbf{q}_t + A(\mathbf{q}_t) \odot E_c(\mathbf{c}_t - \mathbf{c}_{t-1})] \quad \text{Eq.(2)}$$

This transition formulation makes rapid context changes explicit and bounds their scale through normalization. Its output drives spectral modulation; without it, the same transport map persists after a vent fault or cloud-front arrival.

Context-Modulated Spectral and Recurrent Fusion

The spectral branch applies a two-dimensional real Fourier transform to each encoded field, retains a compact set of low and intermediate modes, and uses context to scale their complex-valued multipliers. Low modes describe facilitywide gradients, while intermediate modes represent aisle and wall effects. High-frequency sensor noise is handled by a local residual path rather than being amplified spectrally. Operator learning has shown that truncated modes can express nonlocal mappings efficiently, but the present formulation makes the multipliers conditional on the current greenhouse regime.

$$\mathbf{S}_t^{(l)} = \mathcal{F}^{-1}[\mathbf{R}^{(l)}(\mathbf{q}_t^*) \odot \mathcal{F}(\mathbf{Z}_t^{(l)})] + \mathbf{W}^{(l)}\mathbf{Z}_t^{(l)} \quad \text{Eq.(3)}$$

The hypernetwork maps context to bounded gains around a base kernel. The inverse transform restores the spatial field, while the pointwise path retains local vent effects. Removing either path causes over-smoothing or slow cross-zone propagation.

$$\mathbf{R}_k^{(l)}(\mathbf{q}) = \mathbf{R}_{0,k}^{(l)}[1 + \rho \tanh(\mathbf{B}_k^{(l)}\mathbf{q})], \|k\| \leq K \quad \text{Eq.(4)}$$

The bound ρ restricts context-driven deviation, and K sets the retained-mode budget. This constraint protects controller commissioning under unfamiliar operating regimes. Figure 1 records the implemented order from field assembly through receding-horizon control.




Figure 1. End-to-end workflow of context sensing, context-modulated FNO-LSTM forecasting, and receding-horizon greenhouse climate control.

An LSTM operating on pooled spectral features and actuator histories represents unmeasured thermal and moisture storage. Pure feedforward operators can match the current spatial field yet drift over long rollouts, whereas the recurrent branch preserves delayed climate responses.

$$(\mathbf{h}_t, \mathbf{m}_t) = \text{LSTM}([\text{Pool}(\mathbf{S}_t), \mathbf{u}_t, \mathbf{b}_t], \mathbf{h}_{t-1}, \mathbf{m}_{t-1}) \quad \text{Eq.(5)}$$

The hidden state $\mathbf{h}_t$ captures relatively fast sequence effects, while the memory state $\mathbf{m}_t$ retains accumulated thermal and moisture information. A horizon-specific context gate then determines how much spatial and temporal evidence should be assigned to each prediction target.

$$\mathbf{g}_{t,h} = \sigma(G_h[\text{Pool}(\mathbf{S}_t), \mathbf{h}_t, \mathbf{q}_t^*]), \mathbf{f}_{t,h} = \mathbf{g}_{t,h} \odot \text{Pool}(\mathbf{S}_t) + (\mathbf{1} - \mathbf{g}_{t,h}) \odot \mathbf{h}_t \quad \text{Eq.(6)}$$

The gate is the principal interface between the spectral and recurrent branches. Fixing the gate at zero, one, or a constant intermediate value reduces the architecture to a single-branch or context-insensitive predictor. Figure 2 exposes the tensor flow and deployment boundary of the complete forecasting model.




Figure 2. Internal architecture of the context-modulated spectral blocks, LSTM memory, horizon-dependent fusion gate, and multizone decoder.

Control-Oriented Forecasting and Optimization

The decoder predicts climate increments rather than absolute fields. The latest measurement anchors each zone, and cumulative increments form the future trajectory. Residual reconstruction reduces level drift and makes actuator-induced changes explicit. A probabilistic residual head was considered, but a deterministic output was retained to keep the controller solution time predictable. Control-oriented sequence models benefit when their outputs correspond directly to actuation-relevant climate changes.

$$\hat{\mathbf{x}}_{t+h} = \mathbf{x}_t + \sum_{j=1}^h D_j(\mathbf{f}_{t,j}, \mathbf{q}_t^*) \quad \text{Eq.(7)}$$

Each decoder head estimates a zone-level increment. The measurement anchor limits long-horizon drift, while the training objective preserves sharp, justified responses without introducing unstable oscillations.

$$\mathcal{L}_{\text{field}} = \sum_{h=1}^H \omega_h \frac{\|\mathbf{m}_{t+h} \odot (\hat{\mathbf{x}}_{t+h} - \mathbf{x}_{t+h})\|_1}{\sum \mathbf{m}_{t+h}} \quad \text{Eq.(8)}$$

The observation mask $\mathbf{m}_{t+h}$ excludes unavailable sensors, while the horizon weights ω_h balance near-term and longterm prediction errors. The same availability mask is retained in the transition term so that the boundaries of missing-data intervals are not interpreted as physical climate jumps. Because all continuous target channels are standardized using training-set statistics, variables with different physical units contribute on comparable scales before the weighted losses are combined.

$$\mathcal{L}_{\text{dyn}} = \sum_{h=1}^H \|\mathbf{M}_{t+h} \odot [(\hat{\mathbf{X}}_{t+h} - \hat{\mathbf{X}}_{t+h-1}) - (\mathbf{X}_{t+h} - \mathbf{X}_{t+h-1})]\|_1 \quad \text{Eq.(9)}$$

The dynamic term targets slope and peak-timing errors, which govern climate overshoot and actuator reversal. The condensation penalty approximates leaf-wetness risk when direct canopy-surface sensing is unavailable.

$$\mathcal{L} = \mathcal{L}_{\text{field}} + \lambda_d \mathcal{L}_{\text{dyn}} + \lambda_c \sum_{z=1}^{N_z} \sum_{h=1}^H \text{softplus}(T_{\text{dew},z,t+h} - \hat{T}_{z,t+h}) + \lambda_s \left\| \frac{\partial \hat{\mathbf{X}}}{\partial \mathbf{u}} \right\|_F^2 \quad \text{Eq.(10)}$$

The dew-point term discourages operating states associated with condensation risk, while the sensitivity regularizer acts only inside the observed actuator envelope. The trained predictor is embedded in a receding-horizon controller that applies the first optimized command, receives the next sensor observation, updates the recurrent and contextual states, and solves the optimization problem again.

Experimental Evaluation

Data, Protocol, and Reproducibility

A Venlo-style research greenhouse with six climate zones that has been operating continuously for 120 days from early spring to midsummer serves as a replaceable case study. Before quality control, samples are collected every five minutes and reach 34,560 timestamps. Thirty zone measurements, six outdoor boundary variables, and seven actuator commands are included in each timestamp. 32,918 timestamps remained after commissioning intervals and sequences with more than 20% missing channels were eliminated. Longer gaps are still hidden; short gaps of up to three samples are causally interpolated from the historical values. There are no shuffled windows across boundaries; training takes up the first 70% of the time, validation takes up the next 15%, and testing takes up the final 15%. The forecasts are for 3, 6, 12, and 24 steps, which correspond to 15, 30, 60, and 120 minutes, and the input windows contain 36 samples (180 minutes). Use just the training statistics to standardise all continuous features. These are not actual facility test results; rather, they are illustrative placeholders used to define a repeatable experiment. Because seasonal leakage can mask operational drift, comparable greenhouse evaluations have demonstrated that regime-stratified chronological splits are more informative than randomly shuffled samples [23]. As a result, the existing protocol reports context-specific behaviour and stores each split chronologically.

North-wall, central-canopy, south-wall, inlet, upper-canopy, and service-aisle are the six target areas. Clear-stable, clear-variable, overcast, and frontal-transition probabilities are included in the weather context, as is the shift in crop stage from vegetative to flowering and fruit-loading. Equipment flags signify unavailability of the carbon dioxide supply, restricted vent travel, fogging lockout, or normal operation. To replicate rolling control, validation and test forecasts are generated at each timestamp, but training windows are sampled with a 50% overlap. Context labels are calculated in a causal manner. If a major actuator moves more than 20% of its range, the availability flag is changed, or the global radiation changes by more than 180 W m^{-2} in 15 minutes, a transition event is initiated. To avoid a benchmark dominated by continuous stable-night conditions, both event and non-event periods are included. Repeated compression can conceal boundary-localized transport in flattened space measurements, as demonstrated by earlier multizone studies [24]. Consequently, the evaluation design will continue to use the explicit zone geometry and event labels.

Persistence, ridge autoregression, a two-layer LSTM, a convolutional LSTM, a graph convolutional recurrent network, a compact Transformer, and a standard Fourier neural operator (FNO) without recurrent or context branches are adopted as baselines. All neural baselines use a hidden width of 128 and decoders with comparable capacities. The proposed model comprises four spectral blocks, a 24-step decoder, an LSTM hidden size of 128, a context-embedding width of 48, a dropout rate of 0.10, and 12 retained Fourier modes along each spatial axis. AdamW is employed with a learning rate of 8×10^{-4} and a weight-decay coefficient of 10^{-4} . Training uses a batch size of 64 and cosine learning-rate decay over 80 epochs, with early stopping triggered when the validation loss fails to improve for 12 consecutive epochs. The loss weights are set to $\lambda_d = 0.35$, $\lambda_c = 0.08$, and $\lambda_s = 0.002$.

All results are aggregated over five random seeds. Hyperparameters are selected exclusively according to validation MAE and transition lag, without access to the test set. Forecasting performance is evaluated using

peak-timing lag, regime-conditioned error, and variable-specific MAE and RMSE. Closed-loop performance is assessed in terms of climate-constraint satisfaction, energy consumption, carbon dioxide emissions, actuator reversals, and average solution time. Following established benchmarking guidelines for spatiotemporal forecasting, competing models are fine-tuned using identical validation data while maintaining comparable decoder capacities [25]. This fixed experimental protocol reduces confounding effects arising from differences in model capacity and hyperparameter tuning, thereby enabling more reliable architectural comparisons.

The dataset exhibits regime-dependent spatial variability that directly motivates context conditioning. In Figure 3(a), the median zone temperature increases from 23.4 °C under overcast weather to 24.2 °C in clear-stable conditions and 25.1 °C in clear-variable conditions, while the clear-variable distribution broadens toward an upper tail of 29.2 °C. The simultaneous increase in level and dispersion indicates that rapidly changing solar gain produces both a warmer greenhouse and stronger interzone imbalance; consequently, a predictor trained around a single seasonal mean would underestimate heat-stress excursions. Figure 3(b) further reveals distinct actuator–climate pathways: central-canopy humidity is positively correlated with fogging duty at 0.72 but negatively correlated with vent opening and radiation at -0.61 and -0.48 , respectively, whereas carbon-dioxide concentration has a -0.57 correlation with vent opening. Fogging therefore exerts the strongest direct moisture effect, while ventilation couples' humidity reduction with carbon-dioxide loss, which explains why actuator states must be encoded jointly rather than optimized as independent channels. The temporal concentration of disturbances in Figure 3(c) is equally nonuniform: 41% occur from 09:00 to 12:00, 27% from 14:00 to 17:00, and only 8% overnight. The morning maximum coincides with steep radiation ramps, whereas the afternoon cluster reflects ventilation and heat-release transitions. These distributions justify event-stratified sampling and time-dependent context gates, and they identify the daylight transition periods in which prediction errors have the greatest implications for crop safety and actuator coordination.




Figure 3. Statistical structure of the replaceable greenhouse dataset. (a) Violin distributions of zone temperature across four weather regimes; (b) heat map of climate, boundary, and actuator correlations; (c) polar density of disturbance events by solar hour.

Forecasting Accuracy and Mechanism Analysis

The MAE of Persistence's temperature, humidity, and carbon dioxide at a 60-minute horizon are 1.54°C, 7.2%, and 96 ppm, respectively. LSTM: 1.18 °C, 5.7 points, and 73 ppm; it has a delay in accumulation but smooths out spatially localised disturbances. Ridge autoregression raises the values to 1.31 °C, 6.4 points, and 84 ppm, but the peak timing lag is still 18.6 min. ConvLSTM is 65 ppm, 4.8 points, and 0.96 °C. The compact transformer has 0.82 °C, 4.1 points, and 55 ppm, while the graph recurrent model has 0.88 °C, 4.4 points, and 59 ppm. The

typical FNO is 52 ppm, 3.9 points, and 0.79 °C. For the five seeds, the suggested context-aware FNO-LSTM achieved 0.63 ± 0.02 °C, 3.1 ± 0.1 points, and 41 ± 2 ppm. These reductions are 46.6%, 45.6%, and 43.8% when compared to LSTM, and 20.3%, 20.5%, and 21.2% when compared to standard FNO. Some of the benefits of operator-learning research over pointwise sequence models are attributed to direct nonlocal field coupling [26]. Although that mechanism does not prove the numerical gains reported here, it is consistent with the stronger behaviour of field-based models under spatially distributed disturbances.

Forecast behavior changes systematically with both prediction horizon and disturbance type. Figure 4(a) shows that the proposed model's temperature MAE increases gradually from 0.29 °C at 15 min to 0.43, 0.63, and 0.91 °C at 30, 60, and 120 min. Over the same horizons, FNO rises from 0.36 to 0.53, 0.79, and 1.19 °C, whereas LSTM grows more sharply from 0.48 to 0.76, 1.18, and 1.71 °C. The widening long-horizon gap indicates that spectral zone coupling limits accumulated spatial drift and that recurrent memory preserves slow thermal storage after the immediate disturbance has passed. During the cloud-front event in Figure 4(b), the measured central-zone temperature falls from 26.7 to 23.8 °C within 35 min. The proposed trajectory reaches 24.0 °C with a 5-min lag, compared with 24.4 °C and a 10-min lag for FNO, while LSTM remains above 24.9 °C after 15 min. This ordering suggests that the context gate accelerates recognition of the radiation-regime transition rather than merely reducing average error. The moisture response in Figure 4(c) provides a complementary test: observed humidity peaks at 86.2%, and the proposed model predicts 84.9%, compared with 82.7% for FNO and 80.8% for LSTM. The smaller peak deficit arises because the fused model combines spatial fog dispersion with moisture retained in recurrent memory after the fogging command ends. Accurate peak magnitude and timing reduce the risk of late ventilation, condensation, and unnecessary actuator reversal, so the improvement is relevant to closed-loop safety rather than only to offline forecasting accuracy.




Figure 4. Multihorizon and event-level forecasting behavior. (a) Line plot of temperature MAE at 15, 30, 60, and 120 min; (b) aligned temperature trajectories during a cloud-front event; (c) humidity trajectories after a fogging pulse.

The cross-model results distinguish average climate error from the operationally important error tail. In Figure 5(a), the 60-min temperature MAE decreases from 1.54 °C for persistence and 1.31 °C for ridge autoregression to 1.18 °C for LSTM, 0.96 °C for ConvLSTM, 0.88 °C for the graph recurrent network, 0.82 °C for Transformer, 0.79 °C for FNO, and 0.63 °C for the proposed model. The improvement becomes progressively larger as the models move from pointwise temporal extrapolation toward explicit spatial interaction; relative to FNO, the

remaining 20.3% reduction indicates that recurrent storage and context modulation contribute beyond the spectral operator alone. Figure 5(b) presents the same ordering for humidity: the respective MAEs are 7.2, 6.4, 5.7, 4.8, 4.4, 4.1, 3.9, and 3.1 percentage points. The larger separation from recurrent baselines is consistent with humidity being governed by both distributed ventilation and delayed canopy-water release, making it especially sensitive to the loss of either spatial or temporal information. Figure 5(c) shifts attention to carbon-dioxide extremes. Ninety percent of proposed-model absolute errors remain below 77 ppm, whereas the corresponding thresholds are 101 ppm for FNO, 108 ppm for Transformer, and 139 ppm for LSTM. The empirical curves therefore separate most strongly in their upper tails rather than converging at high error. This tail contraction is important because a rare but large carbon-dioxide overestimate can trigger injection immediately before ventilation, increasing resource loss, while an underestimate can delay supplementation during a photosynthetically active period.




Figure 5. Cross-model error structure at the 60-min horizon. (a) Grouped bars for temperature MAE across eight models; (b) Cleveland dot plot for humidity MAE; (c) empirical cumulative distribution of absolute carbon-dioxide error.

Regime stratification determines whether the mean advantage happens during times of ease. The suggested temperature MAE for a stable overcast day was 0.48 °C; the FNO and LSTM MAEs were 0.60 °C and 0.91 °C, respectively. The corresponding errors under clear-sky radiation are 0.71, 0.93, and 1.37 °C. 0.86 °C, 1.19 °C, and 1.61 °C are the three frontal transitions. The model error is 0.78 °C due to equipment limitations; FNO is 1.07 °C and LSTM is 1.49 °C. As a result, when effective dynamics change, the relative gain is at its highest. Zone-level inspection also yields positive results: LSTM has a zone spread of 0.31 °C, and the north-wall and upper-canopy zones have proposed errors of 0.66 and 0.69 °C, which are only 0.08 °C higher than the central zone. Context modulation enhances the boundary area without compromising the center by altering spectral transfer rather than a whole-facility bias. Additionally, adaptation-control studies demonstrate that model mismatch worsens when the boundary condition or actuator authority changes [27]. Therefore, regime analysis will expose the model's shortcomings; it will no longer be an unqualified average.

For controls, peak timing is a time of comparatively high stress. The suggested median absolute peak lag for the 214 radiation or actuator events was 5.0 minutes; FNO was 10.0 minutes, Transformer was 12.5 minutes, and LSTM was 15.0 minutes. Large timing errors happen during command reversals because the controller

keeps cooling after the actual temperature has begun to rise; its 90th-percentile lag is 15 minutes, and its LSTM is 35 minutes. To minimise this phase error, dynamically weight the observed and anticipated increments. It does not increase every transition; on a stable night, the predicted temperature step's standard deviation is 0.07°C, which is less than the 0.12°C generated by a model without sensitivity regularisation and near the observed value of 0.06°C. Phase lag is associated with overshoot and repeated command reversal in predictive-control analysis [28]. The relationship demonstrates why evaluating control readiness also takes event timing into account.

Ablation, Closed-Loop Control, and Deployment

Module ablation isolates the cause of the improvement. Removing context modulation and keeping the static FNO and LSTM branches increases temperature MAE from 0.63 to 0.75 °C and transition lag from 5.0 to 10.0 min. Removing the FNO branch raises the temperature MAE to 0.87 °C and increases the boundary-zone spread to 0.27 °C; thus, recurrent memory alone is not effective at encoding spatial transfer. Removing the LSTM branch gives 0.79 °C temperature MAE and raises post-fog humidity error from 3.6 to 5.1 points because slow moisture storage is lost. Replacing the learned gate with an equal average yield 0.71 °C, and removing the dynamic loss yields 0.69 °C but worsens the peak lag to 12.5 min. Omitting the condensation term changes the temperature MAE to 0.64 °C, but the predicted condensation-risk violations increase from 1.7% to 4.9%. Therefore, the components have different failure modes and cannot be accounted for by the total MAE. Hybrid-model ablation practice recommends pairing each removed mechanism with its expected failure mode and does not use a single aggregate score [29]. Therefore, the analysis will consider the above factors simultaneously: time, space, moisture persistence and condensation risk.

The ablation patterns associate each component with a distinct failure mechanism. Starting from the full-model temperature MAE of 0.63 °C, Figure 6(a) records increase of 0.12 °C without context modulation, 0.24 °C without FNO, 0.16 °C without LSTM, 0.08 °C with fixed fusion, and 0.06 °C without the dynamic loss. Removing FNO produces the largest degradation because distant zones can no longer exchange information directly, whereas removal of LSTM primarily eliminates delayed thermal and moisture storage. The smaller but consistent penalties for fixed fusion and the dynamic loss show that adaptive weighting and transition fidelity refine rather than replace the two main branches; the increments are not summed because the modules interact. Event timing in Figure 6(b) supports this interpretation. The median lag is 5.0 min for the full model, 7.5 min with fixed fusion, 10.0 min without context, 10.0 min without LSTM, and 12.5 min both without FNO and without the dynamic loss. Context and the dynamic objective therefore help the model detect when the governing regime changes, while FNO determines how quickly that change propagates spatially. Figure 6(c) shows that the full model also has the most compact daily humidity-error distribution, with a median of 2.4 percentage points and an interquartile range of 1.3–3.8. Removing LSTM raises the median to 3.7 and widens the range to 2.0–5.6, while removing context produces a median of 3.2 and a range of 1.8–5.0. The upward shift and broader tails occur during disturbance periods rather than through seed-to-seed noise, indicating that both recurrent storage and regime awareness are required to prevent unstable humidity corrections and condensation-risk excursions.

The same weather and occupancy traces are used for all controllers during the 20 representative test days in closed-loop evaluation. A rule controller uses 1,284 kWh of thermal and electrical energy, produces 742 main actuator reversals, emits 86.4 kg of carbon dioxide, and maintains 91.8% of temperature and humidity observations within the designated range. 95.4% satisfaction, 1,173 kWh, 82.1 kg, and 626 reversals have been attained by an LSTM-based predictive controller. 96.7%, 1,108 kWh, 79.8 kg, and 574 reversals are attained by FNO-based control. The context-aware FNO-LSTM controller generated 492 reversals, used 1,023 kWh and 77.5 kg, and achieved 98.6% constraint satisfaction. Energy has decreased by 12.8%, CO₂ emissions by 5.6%, and reversals by 21.4% when compared to LSTM control. Because the recurrent memory indicates that canopy moisture is still high, the increase is caused by delayed fogging and earlier vent reduction before the anticipated radiation drop. When compared to independent threshold logic, anticipatory coordination can lower simultaneous heating and ventilation, according to greenhouse-control literature [30]. To avoid the impact of operating conditions on coordination, weather traces and actuator limits are kept constant in the comparison.




Figure 6. Targeted ablation of context, spatial, temporal, fusion, and loss components. (a) Waterfall view of temperature-MAE degradation; (b) lollipop plot of event timing lag; (c) box-and-whisker distributions of daily humidity error.

Energy savings and climate trade-offs are separated in a specific analysis of the 20 control days. Average temperature variation during the day: The deviation of the rule controller is 1.21 °C, the LSTM control is 0.92 °C, and the suggested controller is 0.54 °C. The three nighttime deviation values are 0.39, 0.58, and 0.66 °C. 0.7% of proposed-controller samples, 2.6% of LSTM samples, and 4.1% of rules samples have relative-humidity violations greater than 88%. For LSTM and fan/pump electricity consumption, the decreased heating energy is 824 kWh to 718 kWh and 349 kWh to 305 kWh, respectively. Therefore, expanding climate zones is not the only way to achieve the overall reduction. On clear mornings, the controller gradually opens the roof vent in accordance with the space forecast before the upper-canopy temperature surpasses the limit because it anticipates the stored heat after sunset and prevents simultaneous heating and ventilation.

The method's applicability at the greenhouse edge is determined by its sensitivity to model size. The temperature MAE is 0.69, 0.63, and 0.61 °C with 8, 12, and 16 retained modes, and the single-sample inference on an industrial CPU is 8.4, 11.7, and 17.9 ms. The MAE for humidity drops from 3.4 to 3.1 when the hidden size is increased from 96 to 128; at a size of 192, the memory footprint increases to 131 MB. The chosen configuration can produce a 24-step forecast in 11.7 milliseconds and has a size of 91 MB. It has a 99th-percentile time of 231 ms and an average of 146 ms when warm-starting the receding-horizon optimisation from the previous solution. Both are still far below the sampling period of 300 seconds. Graph recurrent inference takes 14.1 ms but necessitates facility-specific adjacency matrices; a compact Transformer has an inference time of 18.8 ms and a size of 138 MB. According to edge-deployment studies, worst-case optimisation time, memory footprint, and inference latency should all be taken into account simultaneously [31]. Forecast improvements must meet the requirements in order to be used in a supervisory control loop.

Deployment performance must balance forecast accuracy, computational delay, and closed-loop resource use. In Figure 7(a), the selected 12-mode configuration occupies a favorable Pareto position at 0.63 °C MAE and 11.7 ms per sample. Increasing the mode count to 16 reduces MAE only to 0.61 °C but raises latency to 17.9 ms; FNO operates at 0.79 °C and 9.8 ms, LSTM at 1.18 °C and 6.3 ms, and Transformer at 0.82 °C and 18.8 ms. The selected model therefore sacrifices 5.4 ms relative to LSTM but removes 0.55 °C of error, whereas the 16-

mode variant incurs a further 6.2 ms for only a 0.02 °C gain. This diminishing return supports the 12-mode choice for edge deployment. Closed-loop outcomes in Figure 7(b) follow the same accuracy–efficiency trade-off. The rule controller achieves 91.8% climate-constraint satisfaction with 1,284 kWh and 742 actuator reversals; LSTM-based control reaches 95.4%, 1,173 kWh, and 626 reversals; FNO-based control reaches 96.7%, 1,108 kWh, and 574 reversals; and the proposed controller reaches 98.6%, 1,023 kWh, and 492 reversals. The simultaneous increase in satisfaction and decrease in energy and switching indicates that anticipatory spatial–temporal forecasts avoid late, opposing commands rather than obtaining comfort by using more actuation. Figure 7(c) examines whether this benefit remains computationally safe across operating contexts. Median optimization time rises from 132 ms instable conditions to 149 ms under variable radiation, 158 ms during transitions, and 171 ms with restricted equipment, because changing context activates tighter constraints and a more difficult control search. Nevertheless, every context retains a 99th-percentile solution time below the 250-ms limit, leaving a wide margin within the 300-s sampling interval. The result supports real-time supervisory deployment while identifying restricted-equipment periods as the primary target for solver monitoring and fallback control.




Figure 7. Deployment and closed-loop efficiency. (a) Pareto scatter of temperature MAE versus inference time; (b) bubble plot of climate satisfaction versus energy with actuator reversals as area; (c) ridgeline densities of optimization time across operating contexts.

After training, robustness tests introduce flaws in the equipment and sensors. The temperature MAE increases from 0.63 to 0.70 °C when 10% of measurements in random three-sample blocks are removed; at 20% removal, it reaches 0.82 °C. The result of forward filling under the same 20% loss is 0.96 °C, suggesting that masked training keeps artificial plateaus from taking over. The three MAEs increase to 0.68 °C, 3.4 points, and 45 ppm when Gaussian sensor noise with standard deviations of 0.15 °C, 1.0 humidity point, and 15 ppm is added. The context-aware model reaches 0.79 °C when one roof vent is limited to 50% travel, while turning off the equipment flag results in 1.02 °C because the forecast assumes that cooling authority is unavailable. Because the modulation gain is bounded by ρ , a 15-minute context misclassification only increases error to 0.73 °C; unbounded modulation reaches 1.26 °C and occasionally produces oscillatory forecasts. Because extended gaps produce artificial temporal plateaus, agricultural sensor-network research finds burst outages to be more disruptive than isolated missing samples [32]. As a result, the robustness protocol incorporates equipment-status errors, noise, and block loss in addition to random point removal.

Several engineering safeguards prevent favorable offline metrics from obscuring control failures. Prior to optimisation, each predicted channel is first range-checked. If three consecutive forecasts exceed physical bounds, the controller reverts to a conservative linear model. This fallback restricts the longest unavailable interval to 10 minutes and is activated in 0.04% of illustrative test cycles. Second, a gradual cloud transition interpolates between regimes because the model receives context probabilities instead of hard weather labels. Third, only after a confirmed shutdown lasting more than 30 minutes is the memory state reset; short communication pauses preserve the state with uncertainty inflation. Fourth, training quantiles and one-step residual distributions are compared by drift monitors. Rather than silently modifying the control law, a seven-day residual median above 0.35 °C or humidity bias above 1.8 points necessitates recalibration.

The numerical study also defines the evidence required before real deployment. Deployment should begin with offline verification and a shadow-control phase before live commands are enabled. Random train/test splitting would exaggerate generalisation, so a trace that incorporates crop-stage changes, external weather variation, and equipment limitations must take the place of the 120-day duration. Carbon dioxide consumption should come from mass-flow records rather than just command duration, and energy should be measured at the actuator level or estimated from validated power curves. Climate bands must be linked to crop stage and specified with the grower. Lastly, the controller comparison should employ the same actuator constraints and weather traces, ideally through a shadow-control phase prior to the activation of live commands. These specifications are not proof that the illustrative values are final; rather, they are a component of the experimental design.

Because daily averaging can mask a morning control issue, forecast residuals were also analysed by time of day. Temperature MAE is 0.58 °C and humidity MAE is 3.0 points between 5:00 and 8:00. They rise to 0.72 °C and 3.5 points during the 09:00–12:00 radiation-ramp interval, settle at 0.65 °C and 3.2 points between 13:00 and 17:00, and then fall to 0.47 °C and 2.6 points overnight. In the morning, FNO alone has the biggest difference, reaching 0.98 °C, while LSTM reaches 1.39 °C. When radiation changes quickly, the suggested gate assigns an average spectral weight of 0.68; however, during stable nighttime humidity evolution, it only assigns 0.43. This change is consistent with the intended mechanism and does not consistently favour one branch over another.

1,086 major actuator movements were grouped in order to assess command-response consistency. The observed upper-zone temperature changes by −0.84 °C over 20 minutes after roof-vent increases; the suggested model predicts −0.79 °C, FNO −0.67 °C, and LSTM −0.52 °C. The observed humidity is 4.8 points higher than its initial value at 30 minutes after a 10-minute fog pulse; the suggested model, FNO, and LSTM predict 4.5, 3.2, and 4.1 points. In contrast to the 169-ppm predicted by the suggested model and the 142-ppm predicted by LSTM, an observed 176 ppm rise is produced by carbon dioxide injection followed by closed vents. Instead of averaging two comparable forecasts, the branch combination enhances both immediate spatial transport and delayed storage.

When energy prices fluctuate, control benefits continue to exist. The suggested controller maintains 98.6% satisfaction while using 1,023 kWh under a flat tariff. It shifts 7.9% of heating energy away from the highest-price interval, lowers modelled operating costs by 9.4%, and maintains 98.1% satisfaction with a three-level tariff and unaltered comfort penalties. Due to its longer event lag, which necessitates corrective heating during costly times, LSTM control shifts 4.2% and achieves 95.0% satisfaction. Proposed consumption decreases by an additional 4.7% when the energy penalty is raised by 25%, but humidity infractions increase from 0.7% to 1.1%. This trade-off demonstrates the necessity of choosing controller weights in consultation with growers rather than relying solely on forecast error.

The observed model gaps are still greater than the variation from seed to seed. Temperature MAE ranges from 0.60 to 0.66 °C, humidity MAE from 3.0 to 3.3 points, and carbon dioxide MAE from 39 to 44 ppm over the course of five runs. The number of rare events per minibatch varies most during frontal-transition intervals. Without altering the stable-weather error, stratified event sampling reduces the transition MAE standard deviation from 0.061 to 0.028 °C. Each seed is paired with the same optimiser settings and weather trace for closed-loop evaluation; energy varies by 18 kWh, or 1.8% of the mean. Instead of choosing a favourable checkpoint, these ranges enable comparing modules through repeated runs.

Where the residual bias comes from is shown by spatial error maps. During clear-variable weather, the upper-canopy zone exhibits a daytime bias of +0.18 °C, which is consistent with unmeasured leaf-temperature and buoyancy effects. Due to the inability of the five-minute sampling interval to resolve the initial cold-air pulse,

the inlet zone exhibits a -0.14 °C bias following significant side-vent movements. The central and service-aisle biases stay within ± 0.07 °C. The upper-zone bias is reduced to $+0.11$ °C when a canopy-temperature proxy is added, but the facility-wide MAE is only improved by 0.02 °C. As a result, the sensor is helpful for diagnostic validation, even though not every production bay will be able to afford it. even though not every production bay can justify its cost.

Forecast-boundary errors finally caused stress for the controller. Greenhouse temperature MAE rises to 0.72 °C and energy by 3.1% with an outdoor temperature bias of $+1.0$ °C; morning MAE rises to 0.81 °C and roof-vent opening is delayed by one sample with a radiation bias of -10% . Joint biases result in 0.88 °C, which is still less than the 1.18 °C unbiased LSTM value. The majority of the loss is recovered in two cycles when boundary channels are updated every five minutes due to the gradual adjustment of context probability and recurrent memory. These tests support keeping a local radiation sensor in the event that remote forecasts are delayed and identify weather-feed validation as a deployment dependency.

Conclusion

This paper presents a context-aware FNO-LSTM architecture for multizone greenhouse forecasting and receding-horizon control. The method represents distributed zone measurements as fields, modulates spectral transfer with weather, crop, time, and equipment context, preserves delayed thermal and moisture effects in recurrent memory, and fuses spatial and temporal evidence separately at each forecast horizon. Its control-oriented objective combines field error, transition consistency, condensation risk, and actuator sensitivity. These interfaces connect sensing, prediction, and constrained actuation in a form that can be audited, reproduced, and evaluated before deployment.

The method has several limitations. Despite bounded modulation gains, incorrect crop-stage or equipment flags can still produce context errors because the descriptors depend on reliable supervisory metadata. Transfer to facilities with different vent arrangements requires remapping or fine-tuning of the fixed spatial grid. Without additional observations, the recurrent state cannot fully distinguish structural thermal storage from water storage in the canopy. The preferred control-oriented weights also depend on crop value, energy tariff, and disease risk. Most importantly, verified facility data and a prospective control trial are required before this manuscript can support empirical operational claims; the reported numerical results remain illustrative.

Future research should learn context transitions from continuous supervisory signals, quantify predictive uncertainty within edge-computing limits, and extend the operator to irregular sensor graphs and geometry-aware coordinates. Multi-greenhouse pretraining can test whether spectral kernels transfer across structures and whether local adapters capture facility-specific leakage and actuator responses. Slowly drifting parameters may be updated through stable online learning under explicit stability and constraint monitors. The trial period should also be extended to evaluate crop yield and quality, water-use efficiency, and disease incidence and to optimize biological and energy objectives across the production cycle.

Author Contributions

Dajana Dujmović contributes to conceptualization, methodology, software, validation, analysis, investigation, data collection, draft preparation, manuscript editing, visualization, supervision. All authors have read and agreed with the manuscript before its submission and publication.

Funding

This research received no specific financial support from any funding agency.

Institutional Review Board Statement

Not applicable.

References

- [1] Akhter, R., & Sofi, S. A. (2022). Precision agriculture using IoT data analytics and machine learning. *Journal of King Saud University - Computer and Information Sciences*, 34(8), 5602–5618. <https://doi.org/10.1016/j.jksuci.2021.05.013>
- [2] Ariesen-Verschuur, N., Verdouw, C., & Tekinerdogan, B. (2022). Digital Twins in greenhouse horticulture: A review. *Computers and Electronics in Agriculture*, 199, 107183. <https://doi.org/10.1016/j.compag.2022.107183>
- [3] Bai, J., Zhu, J., Song, Y., Zhao, L., Hou, Z., Du, R., & Li, H. (2021). A3T-GCN: Attention temporal graph convolutional network for traffic forecasting. *ISPRS International Journal of Geo-Information*, 10(7), Article 485. <https://doi.org/10.3390/ijgi10070485>
- [4] Benos, L., Tagarakis, A. C., Dolias, G., Berruto, R., Kateris, D., & Bochtis, D. (2021). Machine learning in agriculture: A comprehensive updated review. *Sensors*, 21(11), 3758. <https://doi.org/10.3390/s21113758>
- [5] Bhat, S. A., & Huang, N. F. (2021). Big data and ai revolution in precision agriculture: Survey and challenges. *Ieee Access*, 9, 110209–110222. <https://doi.org/10.1109/access.2021.3102227>
- [6] Chen, W.-H., & You, F. (2022). Semiclosed greenhouse climate control under uncertainty via machine learning and data-driven robust model predictive control. *IEEE Transactions on Control Systems Technology*, 30(3), 1186–1197. <https://doi.org/10.1109/tcst.2021.3094999>
- [7] Choi, K., Yi, J., Park, C., & Yoon, S. (2021). Deep learning for anomaly detection in time-series data: Review, analysis, and guidelines. *IEEE Access*, 9, 120043–120065. <https://doi.org/10.1109/access.2021.3107975>
- [8] Drgoña, J., Arroyo, J., Figueroa, I. C., Blum, D., Arendt, K., Kim, D., ... & Helsen, L. (2020). All you need to know about model predictive control for buildings. *Annual reviews in control*, 50, 190–232. <https://doi.org/10.1016/j.arcontrol.2020.09.001>
- [9] Escamilla-García, A., Soto-Zarazúa, G. M., Toledano-Ayala, M., Rivas-Araiza, E., & Gastélum-Barrios, A. (2020). Applications of artificial neural networks in greenhouse technology and overview for smart agriculture development. *Applied Sciences*, 10(11), Article 3835. <https://doi.org/10.3390/app10113835>
- [10] García, L., Parra, L., Jimenez, J. M., Lloret, J., & Lorenz, P. (2020). IoT-based smart irrigation systems: An overview on the recent trends on sensors and IoT systems for irrigation in precision agriculture. *Sensors*, 20(4), Article 1042. <https://doi.org/10.3390/s20041042>
- [11] Gholamzadehmir, M., Del Pero, C., Buffa, S., Fedrizzi, R., & Aste, N. (2020). Adaptive-predictive control strategy for HVAC systems in smart buildings – A review. *Sustainable Cities and Society*, 63, Article 102480. <https://doi.org/10.1016/j.scs.2020.102480>
- [12] Hewage, P., Behera, A., Trovati, M., Pereira, E., Ghahremani, M., Palmieri, F., & Liu, Y. (2020). Temporal convolutional neural (TCN) network for an effective weather forecasting using time-series data from the local weather station. *Soft Computing*, 24(21), 16453–16482. <https://doi.org/10.1007/s00500-020-04954-0>
- [13] Jung, D.-H., Kim, H. S., Jhin, C., Kim, H.-J., & Park, S. H. (2020). Time-serial analysis of deep neural network models for prediction of climatic conditions inside a greenhouse. *Computers and Electronics in Agriculture*, 173, Article 105402. <https://doi.org/10.1016/j.compag.2020.105402>
- [14] Karniadakis, G. E., Kevrekidis, I. G., Lu, L., Perdikaris, P., Wang, S., & Yang, L. (2021). Physics-informed machine learning. *Nature Reviews Physics*, 3(6), 422–440. <https://doi.org/10.1038/s42254-021-00314-5>
- [15] Krishnamurthi, R., Kumar, A., Gopinathan, D., Nayyar, A., & Qureshi, B. (2020). An overview of IoT sensor data processing, fusion, and analysis techniques. *Sensors*, 20(21), Article 6076. <https://doi.org/10.3390/s20216076>
- [16] Lezoche, M., Hernandez, J. E., Alemany Díaz, M. D. M. E., Panetto, H., & Kacprzyk, J. (2020). Agri-food 4.0: A survey of the supply chains and technologies for the future agriculture. *Computers in Industry*, 117, Article 103187. <https://doi.org/10.1016/j.compind.2020.103187>
- [17] Li, Z., Peng, W., Yuan, Z., & Wang, J. (2022). Fourier neural operator approach to large eddy simulation of three-dimensional turbulence. *Theoretical and Applied Mechanics Letters*, 12(6), Article 100389. <https://doi.org/10.1016/j.taml.2022.100389>
- [18] Lim, B., Arık, S. Ö., Loeff, N., & Pfister, T. (2021). Temporal fusion transformers for interpretable multi-horizon time series forecasting. *International Journal of Forecasting*, 37(4), 1748–1764. <https://doi.org/10.1016/j.ijforecast.2021.03.012>
- [19] Linaza, M. T., Posada, J., Bund, J., Eisert, P., Quartulli, M., Döllner, J., Pagani, A., Olaizola, I. G., Barriguinha, A., Moysiadis, T., & Lucat, L. (2021). Data-driven artificial intelligence applications for sustainable precision agriculture. *Agronomy*, 11(6), Article 1227. <https://doi.org/10.3390/agronomy11061227>

- [20] Liu, Z., Zhu, Z., Gao, J., & Xu, C. (2021). Forecast methods for time series data: A survey. *IEEE Access*, 9, 91896–91912. <https://doi.org/10.1109/access.2021.3091162>
- [21] Moon, J., Park, S., Rho, S., & Hwang, E. (2022). Robust building energy consumption forecasting using an online learning approach with R ranger. *Journal of Building Engineering*, 47, Article 103851. <https://doi.org/10.1016/j.jobbe.2021.103851>
- [22] Rashid, M. M., Pittie, T., Chakraborty, S., & Krishnan, N. M. A. (2022). Learning the stress-strain fields in digital composites using Fourier neural operator. *iScience*, 25(11), Article 105452. <https://doi.org/10.1016/j.isci.2022.105452>
- [23] Saiz-Rubio, V., & Rovira-Más, F. (2020). From smart farming towards agriculture 5.0: A review on crop data management. *Agronomy*, 10(2), Article 207. <https://doi.org/10.3390/agronomy10020207>
- [24] Sharma, A., Jain, A., Gupta, P., & Chowdary, V. (2021). Machine learning applications for precision agriculture: A comprehensive review. *IEEE Access*, 9, 4843–4873. <https://doi.org/10.1109/access.2020.3048415>
- [25] Simeunovic, J., Schubnel, B., Alet, P.-J., & Carrillo, R. E. (2022). Spatio-temporal graph neural networks for multi-site PV power forecasting. *IEEE Transactions on Sustainable Energy*, 13(2), 1210–1220. <https://doi.org/10.1109/tste.2021.3125200>
- [26] Subeesh, A., & Mehta, C. R. (2021). Automation and digitization of agriculture using artificial intelligence and internet of things. *Artificial Intelligence in Agriculture*, 5, 278–291. <https://doi.org/10.1016/j.aiaa.2021.11.004>
- [27] Verdouw, C., Tekinerdogan, B., Beulens, A., & Wolfert, S. (2021). Digital twins in smart farming. *Agricultural Systems*, 189, 103046. <https://doi.org/10.1016/j.agry.2020.103046>
- [28] Wang, S., Wang, H., & Perdikaris, P. (2021). Learning the solution operator of parametric partial differential equations with physics-informed DeepONets. *Science Advances*, 7(40), Article eabi8605. <https://doi.org/10.1126/sciadv.abi8605>
- [29] Wen, G., Li, Z., Azizzadenesheli, K., Anandkumar, A., & Benson, S. M. (2022). U-FNO—An enhanced Fourier neural operator-based deep-learning model for multiphase flow. *Advances in Water Resources*, 163, Article 104180. <https://doi.org/10.1016/j.advwatres.2022.104180>
- [30] Wu, B., Wang, L., & Zeng, Y. R. (2022). Interpretable wind speed prediction with multivariate time series and temporal fusion transformers. *Energy*, 252, 123990. <https://doi.org/10.1016/j.energy.2022.123990>
- [31] Wu, H., Meng, K., Fan, D., Zhang, Z., & Liu, Q. (2022). Multistep short-term wind speed forecasting using transformer. *Energy*, 261, Article 125231. <https://doi.org/10.1016/j.energy.2022.125231>
- [32] Zhu, J., Wang, Q., Tao, C., Deng, H., Zhao, L., & Li, H. (2021). AST-GCN: Attribute-augmented spatiotemporal graph convolutional network for traffic forecasting. *IEEE Access*, 9, 35973–35983. <https://doi.org/10.1109/access.2021.3062114>