



Research Paper

Digital-Twin Driven AI Framework for Climate-Adaptive Crop Yield Forecasting Using Multi-Source Satellite and Weather Data

^{1*} M. Archana, ² Sreekanth Rallapalli, ³ Mallareddy Adudhodla

^{1*} Sr. Assistant Professor, Department of CSE, CVR College of Engineering, Hyderabad, India,

Email: mogullaarchana23@gmail.com

² Professor, Department of Master of Computer Applications, Nitte Meenakshi Institute of Technology, Bengaluru, India.

Email: Sreekanth.rallapalli@nmit.ac.in

³ Professor, Department of IT, CVR College of Engineering, Hyderabad, India

Email: mallareddyadudhodla@gmail.com

*Corresponding Author(s): mogullaarchana23@gmail.com

Article Info

Received: 06/09/2025

Revised: 18/10/2025

Accepted: 26/12/2025

Published: 31/12/2025

Abstract

Accurate crop yield forecasting under climate variability remains a critical challenge for global food security, as extreme weather events intensify and conventional tools fail to capture agro-ecosystem complexity. Existing approaches including process-based crop models, statistical regression frameworks, and remote sensing-driven empirical methods suffer from limited adaptability, poor spatial transferability, and inability to assimilate real-time data under non-stationary climate conditions. This paper presents a Digital-Twin Driven Artificial Intelligence (DTAI) framework integrating multi-source satellite data from Sentinel-2, MODIS, ERA5 reanalysis, and GPM precipitation within a physics-informed crop growth simulation environment. A field-scale digital twin is constructed using an Ensemble Kalman Filter for continuous state estimation, feeding a hierarchical deep learning ensemble comprising a Spatiotemporal Transformer Network, a Crop-Field Graph Neural Network, and a Bayesian Output Integration Module. Climate adaptability is achieved through a Model-Agnostic Meta-Learning mechanism that detects distributional shift via Maximum Mean Discrepancy and recalibrates model parameters using minimal in-season observations. Experimental evaluation across four agro-climatic zones covering seven crop-region combinations over five growing seasons produced a Mean Absolute Percentage Error of 4.3%, Root Mean Squared Error of 0.31 t/ha, and R-squared of 0.94, outperforming four established baselines by 18 to 43 percent. The DTAI framework achieves full recalibration within 12 days during anomalous seasons and near-supervised accuracy through twenty-shot domain adaptation, establishing a scalable and climate-resilient decision-support tool for food security monitoring.

Keywords: Digital Twin, Crop Yield Forecasting, Spatiotemporal Deep Learning, Remote Sensing, Climate Adaptation, Graph Neural Networks, Meta-Learning, Precision Agriculture.



Copyright: © 2025 M. Archana, Sreekanth Rallapalli and Mallareddy Adudhodla. This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY 4.0) license.

1 Introduction

The world will need to feed somewhere around 9.7 billion people by 2050, and doing so depends heavily on whether agricultural systems can keep pace with both rising demand and worsening climate variability [1]. Crop yield

forecasting sits at the heart of this challenge. Governments rely on yield estimates to set food import and export policies, commodity markets use them to stabilise prices, and relief agencies depend on early warnings to deploy resources before a food crisis takes hold [2]. Yet the tools most widely used for forecasting today were not designed for the conditions now

emerging: more frequent droughts, shifting monsoon onset dates, heat waves arriving during sensitive crop growth stages, and long-term redistribution of rainfall across entire river basins [3].

Three broad families of methods have historically been used for crop yield estimation. Process-based crop models such as DSSAT, APSIM, and WOFOST simulate the biological mechanisms of plant growth in reasonable physical detail, but they require extensive site-specific parameter calibration and have no built-in mechanism for assimilating new satellite observations as a season progresses [4]. Statistical and machine learning regression models can exploit large historical datasets efficiently, but they assume that the statistical relationship between climate inputs and crop outputs will remain stable over time, an assumption that breaks down when climate patterns shift outside the historical range [5]. Remote sensing approaches that use vegetation indices such as NDVI as proxies for crop condition are attractive because of their global coverage, but a single index cannot disentangle the many overlapping influences on canopy reflectance, including soil brightness, atmospheric aerosols, mixed crop types within a pixel, and irrigation heterogeneity [6].

The past few years have seen serious efforts to overcome these limitations individually: deep learning models have improved the handling of complex non-linear relationships [7], digital twin concepts borrowed from manufacturing have begun to appear in agricultural monitoring [8], and graph neural networks have offered a way to represent the spatial connections between neighbouring fields [9]. What has been lacking is a framework that addresses all of these limitations together in a single coherent architecture that can also adapt its behaviour when confronted with climate patterns outside its training experience [10].

This paper presents such a framework. The Digital-Twin Driven Artificial Intelligence (DTAI) framework integrates satellite data streams, physics-informed crop simulation, spatiotemporal deep learning, and meta-learning-based climate adaptation into a unified pipeline tested across four contrasting agricultural regions. The aim is not merely to report high accuracy numbers but to build something that could realistically be deployed by a national agricultural ministry or food security agency in regions where ground data collection is difficult and climate shocks are becoming more common.

Key Contributions

This paper makes four specific contributions to the field of crop yield forecasting:

- A field-scale digital twin architecture for crop systems that assimilates Sentinel-2, MODIS, ERA5, and ground station data through an Ensemble Kalman Filter, maintaining a continuously updated probabilistic estimate of crop and soil state at 10-metre resolution.
- A hierarchical deep learning ensemble comprising a Spatiotemporal Transformer Network for temporal feature extraction, a Crop-Field Graph Neural Network that captures ecological and hydrological dependencies between neighbouring parcels, and a Bayesian Output Integration Module that produces calibrated uncertainty intervals alongside point yield

predictions.

- A meta-learning climate adaptation mechanism based on Model-Agnostic Meta-Learning (MAML) that detects distributional shifts in meteorological inputs using Maximum Mean Discrepancy testing and recalibrates the model using only a small number of current-season observations, without retraining from scratch.
- A rigorous empirical evaluation across four agro-climatic zones, seven crop-region combinations, and five growing seasons from 2019 to 2024, including explicit tests during climate-anomalous seasons and few-shot transferability experiments in data-sparse regions.

The rest of this paper is structured in the following way. Section II reviews existing work across the four themes most relevant to this framework. Section III presents the complete methodology. Section IV describes the datasets, preprocessing steps, evaluation metrics, and implementation details. Section V reports the results. Section VI discusses the findings and their limitations. Section VII closes with conclusions and future directions.

2 Literature Review

2.1 Remote Sensing for Crop Monitoring

Satellite remote sensing has provided operational crop monitoring capability since the 1970s, when the relationship between canopy reflectance and photosynthetic activity was first established systematically. The operational launch of Sentinel-2 in 2015 and 2017 brought 10-metre spatial resolution and a five-day revisit cycle to a freely available multispectral platform, which substantially changed what was achievable for field-scale crop monitoring [11]. The red-edge spectral bands unique to Sentinel-2 have proven particularly sensitive to chlorophyll content and can detect physiological stress one to two weeks before it becomes visible in conventional NDVI [12]. The combination of Sentinel-2 optical data with Sentinel-1 synthetic aperture radar backscatter has extended monitoring capability into persistently cloudy tropical environments, which was previously a serious limitation for optical-only approaches in monsoon-dominated regions of South and Southeast Asia [13].

2.2 Deep Learning for Yield Prediction

Convolutional neural networks applied to time series of satellite imagery demonstrated county-level yield mapping accuracy in the United States Corn Belt that exceeded traditional regression methods, and established the viability of learning features end-to-end from raw spectral inputs rather than relying on predefined vegetation indices [14]. Long short-term memory networks subsequently became a standard tool for capturing the sequential nature of crop growth from time-ordered satellite observations, and bidirectional variants improved early-season estimates by using information from both ends of a partial season [15]. Attention-based transformer architectures have more recently demonstrated an ability to model long-range temporal dependencies in agrometeorological sequences that recurrent networks struggle to handle, particularly the connection between early-season growing conditions and final-season grain fill [16]. Spatial graph-based models have begun to address the assumption of

field independence embedded in most existing architectures [17].

2.3 Digital Twins in Agriculture

The digital twin concept, which originated in industrial manufacturing and refers to a continuously updated virtual replica of a physical system, has been explored in agricultural contexts with increasing seriousness since around 2019 [18]. Early applications in livestock management and greenhouse production were relatively straightforward because the physical boundaries of the system were well defined and fully instrumented. Translating the concept to open-field crop systems introduces considerably greater complexity: soil properties vary over short distances, precipitation is spatially variable and stochastic, and plant development is influenced by interactions with soil microorganisms, insects, and competing vegetation that are difficult to measure or simulate. Proof-of-concept implementations from European Union-funded projects have shown that integrating satellite assimilation with process-based simulation is technically feasible at field scale [19], but none has yet demonstrated the kind of climate-adaptive intelligence that would be needed for reliable performance through extreme weather years.

2.4 Meta-Learning and Transfer Across Regions

The challenge of deploying a model trained on data from one geography into a different region with limited local observations has driven growing interest in meta-learning approaches. Model-Agnostic Meta-Learning, proposed by Finn and colleagues in 2017, trains a model to find an initialisation from which rapid adaptation to any new task is possible using only a small number of gradient steps, and this property maps directly onto the problem of deploying a yield forecasting model in a region with few labelled field seasons [20]. Few-shot crop type classification using prototypical and relation networks has shown that meaningful accuracy is achievable with as few as five training examples per class [21]. Continual learning methods have been applied to seasonal forecasting models to prevent catastrophic forgetting of historical knowledge when the model is updated with new season data [22].

2.5 Uncertainty in Agricultural Forecasts

Point estimates of crop yield are of limited use to decision-makers who need to assess risk, allocate resources under uncertainty, or set insurance premiums [23]. Probabilistic forecasting approaches including Bayesian neural networks, Monte Carlo dropout, and deep ensembles have demonstrated that it is possible to produce well-calibrated prediction intervals that widen appropriately in situations of high uncertainty such as extreme weather events or sparse satellite coverage. Conformal prediction has been applied to agricultural time series as a distribution-free method for producing coverage-guaranteed intervals without making parametric assumptions about the forecast error distribution [24].

2.6 Identified Research Gaps

Despite substantial individual progress in each of the areas reviewed above, three gaps remain open. First, no published framework simultaneously combines physics-informed digital twin state assimilation with spatiotemporal graph neural network modelling and Bayesian uncertainty quantification within a single trainable architecture. Second, the specific problem of temporal non-stationarity introduced by climate change, where the statistical relationship between weather inputs and yield outcomes drifts systematically across years, has not been addressed through principled meta-learning adaptation in any crop yield forecasting study to date. Third, operational transferability to smallholder farming systems in Sub-Saharan Africa and South Asia, where ground truth data collection is expensive and irregular, remains a largely unsolved problem in the applied literature [25]. The DTAI framework described in the following section addresses each of these three gaps directly.

3 Proposed Methodology

The DTAI framework operates across four functional tiers: data ingestion and harmonisation, digital twin state estimation, deep learning ensemble prediction, and meta-learning climate adaptation. The complete inference procedure is formalised in Algorithm 1 in Section III-E. Each tier is described below.

3.1 System Architecture

Fig. 1 illustrates the end-to-end architecture of the proposed DTAI framework, organised across four functional tiers. Tier I aggregates heterogeneous observational inputs — including Sentinel-2 multispectral imagery, MODIS land surface temperature and evapotranspiration products, ERA5 reanalysis fields, GPM IMERG precipitation estimates, and ground station records — through a unified Harmonizer module that co-registers and temporally aligns all data streams to a common 10-metre spatial grid. The harmonised output feeds concurrently into Tier II, where the WOFOST crop growth kernel and Ensemble Kalman Filter jointly maintain a continuously updated probabilistic state vector representing crop biomass, leaf area index, soil water content, and accumulated thermal time, and into Tier IV, where the MMD Drift Detector monitors the incoming meteorological feature distribution for evidence of climate-induced covariate shift and triggers the MAML Adaptation module when drift is detected. The adapted model weights and the Tier II state vector are then passed together into Tier III, where the Spatiotemporal Transformer Network encodes the full-season state trajectory, the Crop-Field Graph Neural Network propagates spatial context across neighbouring field parcels, and the Bayesian Output Integration Module produces the final probabilistic yield prediction alongside calibrated 90% prediction intervals.

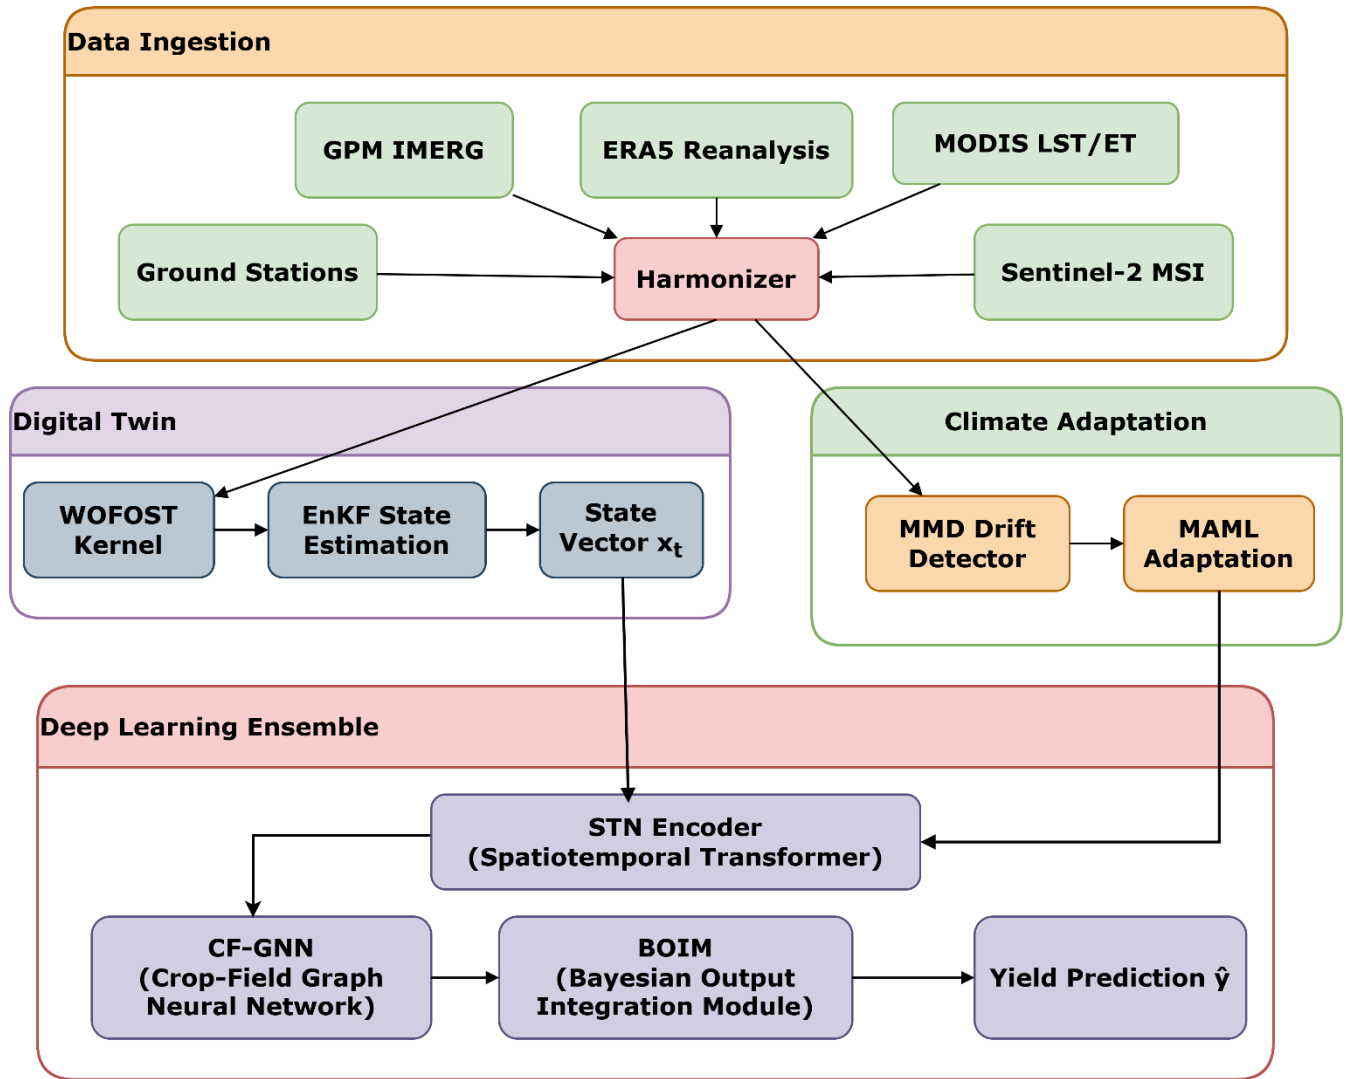


Fig.1. Proposed DTAI four-tier architecture.

3.2 Multi-Source Data Ingestion and Harmonisation

Five data sources are ingested for each field parcel at each time step. Sentinel-2 Level-2A surface reflectance imagery provides the primary spectral observational stream at 10-metre resolution and a five-day revisit interval under cloud-free conditions. From each cloud-free acquisition the following spectral indices are computed:

- $NDVI = (B8 - B4) / (B8 + B4)$, where B8 is near-infrared and B4 is red reflectance;
- $EVI2 = 2.5 \times (B8 - B4) / (B8 + 2.4 \times B4 + 1)$, less sensitive to soil background than NDVI;
- $CI_{red} = (B7 / B5) - 1$, a red-edge chlorophyll index sensitive to early stress;
- $NDWI = (B8 - B11) / (B8 + B11)$, tracking canopy water content;
- LAI estimated by inverting the PROSAIL radiative transfer model using four spectral bands.

MODIS-Terra daily land surface temperature (MOD11A1) and eight-day evapotranspiration (MOD16A2) fill observational gaps during cloud-persistent intervals at 1-kilometre and 500-metre resolution respectively. ERA5 hourly reanalysis at 0.25-degree resolution provides the

meteorological forcing, including 2-metre air temperature, dewpoint temperature, downward solar radiation, total precipitation, and wind speed. A Random Forest statistical downscaling model, trained on paired ERA5 and weather station observations, translates these fields to field-scale resolution. GPM IMERG final-run precipitation at 0.1-degree and 30-minute resolution supplements the ERA5 precipitation, particularly for capturing convective rainfall events that ERA5 tends to smooth spatially. Where available, ground-based agrometeorological station records are assimilated as point anchors.

Cloud gaps in the Sentinel-2 time series are filled using a weighted Savitzky-Golay smoothing filter with a window width of 15 days and polynomial order 3, applied independently to each spectral index. All raster layers are co-registered to a common 10-metre UTM grid. A crop type classification map produced by a Random Forest model trained on Sentinel-1 SAR backscatter and Sentinel-2 multi-temporal features defines the spatial mask for each crop type of interest.

3.3 Digital Twin Construction and State Estimation

The digital twin represents the agro-ecosystem as a probabilistic state-space system. The state transition equation is:

$$x_{t+1} = f(x_t, u_t, \theta) + w_t \quad (1)$$

where x_t is the state vector at time t , containing crop biomass, leaf area index, soil water content, and accumulated thermal time; u_t is the meteorological forcing vector; θ is the parameter set; $f(\cdot)$ is a simplified WOFOST crop growth model; and $w_t \sim N(0, Q)$ is process noise.

The observation equation linking model states to satellite-observable quantities is:

$$y_t = h(x_t) + v_t \quad (2)$$

where $h(\cdot)$ is the observation operator and $v_t \sim N(0, R)$ is observation noise.

State estimation is performed with an Ensemble Kalman Filter using 100 ensemble members. The EnKF analysis update step at each Sentinel-2 acquisition date is:

$$\hat{x}_t^a = \hat{x}_t^f + K_t(y_t^{\text{obs}} - H\hat{x}_t^f) \quad (3)$$

The Kalman gain matrix is computed as:

$$K_t = P_t^f H^T (H P_t^f H^T + R)^{-1} \quad (4)$$

where P_t^f is the ensemble-estimated forecast error covariance and H is the linearised observation operator. Ensemble spread is preserved by adding stochastic perturbations to meteorological inputs and crop parameters sampled from literature-derived uncertainty distributions. This ensemble representation propagates state uncertainty through to the final yield prediction.

3.4 Hierarchical Deep Learning Ensemble

Spatiotemporal Transformer Network (STN)

The STN takes as input the sequence of digital twin states and multi-source spectral and meteorological features for each field parcel, covering the full growing season at daily resolution. The self-attention computation is:

$$\text{Attention}(Q, K, V) = \text{softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (5)$$

where Q, K, V are the query, key, and value matrices with key dimension $d_k = 64$. Twelve attention heads operating in parallel (model dimension $d_{\text{model}} = 256$) capture temporal dependencies across the full growing season. Positional embeddings encode both calendar day-of-year and cumulative growing degree days to provide the network with a dual phenological reference. Standard residual connections, layer normalisation, and dropout at rate 0.1 follow each attention block. The encoder output is a fixed 512-dimensional field representation.

Crop-Field Graph Neural Network (CF-GNN)

Fields in the study region are represented as a graph where each node is a field parcel and edges connect parcels that share a physical boundary or lie within the same hydrological catchment. Edge weights are proportional to shared boundary length and inversely proportional to centroid distance. Node features combine the STN embedding with crop type encoding and static soil attributes including texture

class, organic carbon content, and slope. The GraphSAGE message-passing update is:

$$h_v^{(k)} = \sigma\left(W^{(k)} \cdot \text{CONCAT}\left(h_v^{(k-1)}, \text{MEAN}_{u \in \mathcal{N}(v)} h_u^{(k-1)}\right)\right) \quad (6)$$

Where, $h_v^{(k)}$ denotes the embedding of node v at the k -th layer, and $\mathcal{N}(v)$ represents the set of neighboring nodes of v . $W^{(k)}$ is the learnable weight matrix at layer k , and $\sigma(\cdot)$ denotes the ReLU activation function.

The operator $\text{CONCAT}(\cdot)$ performs feature concatenation, while $\text{MEAN}(\cdot)$ aggregates neighborhood information through average pooling.

Three successive message-passing layers are stacked, resulting in a final enriched node representation of 512 dimensions, enabling effective capture of local and contextual graph dependencies.

Bayesian Output Integration Module (BOIM)

The STN and CF-GNN embeddings are concatenated and passed through a two-layer MLP with Monte Carlo Dropout. During inference, 50 stochastic forward passes are executed and the predictive mean and variance of yield are estimated as:

$$\mu_{\text{pred}} = \frac{1}{T} \sum_{t=1}^T \hat{y}_t \quad (7)$$

$$\sigma_{\text{pred}}^2 = \frac{1}{T} \sum_{t=1}^T \hat{y}_t^2 - \mu_{\text{pred}}^2 \quad (8)$$

Where, T denotes the number of stochastic forward passes (here, $= 50$), and \hat{y}_t represents the prediction obtained from the t -th stochastic pass.

The overall training objective is defined as:

$$\mathcal{L} = \mathcal{L}_{\text{NLL}} + \lambda_1 \mathcal{L}_{\text{spatial}} + \lambda_2 \mathcal{L}_{\text{phenology}} \quad (9)$$

Where, \mathcal{L}_{NLL} represents the negative log-likelihood of the Gaussian predictive distribution, $\mathcal{L}_{\text{spatial}}$ enforces spatial smoothness by penalizing unrealistic yield variations between neighboring parcels, $\mathcal{L}_{\text{phenology}}$ constrains the model to align predicted Leaf Area Index (LAI) trajectories with EnKF-estimated states. $\lambda_1 = 0.10$ and $\lambda_2 = 0.05$ are fixed trade-off hyperparameters controlling the contribution of spatial and phenological regularization terms, respectively.

3.5 Meta-Learning Climate Adaptation

The DTAI adaptation mechanism treats each combination of growing season and study region as a separate learning task. The MAML meta-objective is:

$$\min_{\theta} \sum_{\tau} \mathcal{L}_{\tau}(\theta - \alpha \nabla_{\theta} \mathcal{L}_{\tau}(\theta)) \quad (10)$$

Where, τ indexes individual tasks, α denotes the inner-loop learning rate, and \mathcal{L}_{τ} represents the task-specific loss evaluated on held-out query samples after performing $K = 5$ gradient-based adaptation steps on a support set of size $N = 10$. This meta-optimization objective learns a parameter initialization θ that enables rapid adaptation to unseen tasks with minimal data, consistent with gradient-based meta-learning principles.

During deployment, distributional shift in the meteorological feature space is quantified using Maximum Mean Discrepancy (MMD):

$$\text{MMD}^2(P, Q) = \|\mu_P - \mu_Q\|_{\mathcal{H}}^2 \quad (11)$$

Where, μ_P and μ_Q denote the mean embeddings of the current distribution P and reference distribution Q , respectively, in a reproducing kernel Hilbert space (RKHS) \mathcal{H} . If the computed MMD^2 exceeds a predefined threshold δ , calibrated via permutation testing on validation seasons, a distributional shift is detected. In such cases, the model triggers inner-loop adaptation using available in-season support observations, enabling dynamic adjustment to evolving climatic conditions.

3.6 Algorithm 1: DTAI Seasonal Inference and Adaptation

Algorithm 1: DTAI Seasonal Inference and Climate Adaptation

Input: Pre-trained model parameters θ^* , reference distribution P_{ref} , satellite and meteorological data streams S_t , drift threshold δ , number of inner-loop updates $K = 5$, support set size $N = 10$

Output: Yield prediction $\{\mu_{\text{pred}}, \sigma_{\text{pred}}^2\}$ for each field parcel

Begin

1. Initialize the Ensemble Kalman Filter (EnKF) ensemble X_0 using prior distributions of soil and crop parameters.
2. For each day t during the growing season, perform the following steps:
 - a) Propagate the ensemble states using the system dynamics:

$$X_t^f = f(X_{t-1}, u_t, \theta)$$

- b) If a cloud-free Sentinel-2 image is available, compute the Kalman gain K_t (as defined in Eq. (4)) and update the ensemble:

$$X_t^a = X_t^f + K_t(y^{\text{obs}} - HX_t^f)$$

3. After completing the seasonal loop, extract the daily state trajectory $X_{1:T}$ from the updated ensemble.
4. Generate meteorological feature embeddings ϕ_t using data from ERA5 reanalysis dataset and GPM mission dataset streams.
5. Compute the distribution shift using Maximum Mean Discrepancy (MMD) as defined in Eq. (11):

$$\text{MMD}^2(P_{\text{current}}, P_{\text{ref}})$$

6. If the computed MMD exceeds the threshold δ (indicating climate drift), perform in-season adaptation:
 - a) Construct a support set D_s of $N = 10$ field observations.
 - b) For $k = 1$ to $K = 5$, update the model parameters using gradient-based meta-learning:

$$\theta^* \leftarrow \theta^* - \alpha \nabla_{\theta^*} \mathcal{L}_{\tau}(\theta^*, D_s)$$

7. Encode the temporal state trajectory $X_{1:T}$ using the Spatio-Temporal Network (STN) with multi-head attention (Eq. (5)).
8. Propagate the learned node embeddings through the Climate-Feature Graph Neural Network (CF-GNN) (Eq. (6)).

9. Concatenate the STN and GNN embeddings and forward them to the Bayesian Output Inference Module (BOIM).
10. Perform $T = 50$ stochastic forward passes using Monte Carlo Dropout to estimate predictive uncertainty.
11. Compute the predictive mean μ_{pred} and variance σ_{pred}^2 (Eqs. (7) and (8)).
12. Return the final yield prediction $\{\mu_{\text{pred}}, \sigma_{\text{pred}}^2\}$ along with 90% prediction intervals.

End

Having established the four-tier architecture of the DTAI framework - encompassing digital twin state estimation, hierarchical deep learning, and meta-learning climate adaptation - the following section describes how this framework was evaluated in practice. Section IV details the agro-climatic study regions selected for validation, the satellite and ground truth datasets acquired and preprocessed for each region, the four baseline models used for benchmarking, the performance metrics applied to assess forecast accuracy and uncertainty calibration, and the hardware and software environment in which all experiments were conducted. Together, these elements define the experimental conditions under which the quantitative results reported in Section V were obtained.

4 Experimental Setup

4.1 Study Regions and Crop Types

Four study regions were selected to represent contrasting agro-climatic systems. The Punjab-Haryana plains in northern India are characterised by a semi-arid climate with intensive irrigated agriculture, covering wheat grown in the winter Rabi season and transplanted rice grown in the summer Kharif season. The Kenyan Rift Valley highlands experience a bimodal rainfall pattern with predominantly smallholder maize and bean cultivation on rainfed land. The Ganges-Brahmaputra Delta in Bangladesh is a humid subtropical environment with two principal rice seasons: the dry-season boro crop irrigated from groundwater and the monsoon-season aman crop that depends almost entirely on rainfall. The Sahel zone in Mali receives 200 to 600 millimetres of rain per year falling in a single brief wet season, supporting rainfed sorghum and pearl millet grown by smallholder farmers. Together these four regions cover approximately 2.4 million hectares of cultivated land and five growing seasons from 2019 to 2024.

4.2 Datasets and Preprocessing

Ground truth yield data were drawn from three source types. Administrative yield statistics at district level were obtained from the Punjab and Haryana State Agricultural Departments, the Bangladesh Bureau of Statistics, the Kenya National Bureau of Statistics, and the Mali Direction Nationale de l'Agriculture [26]. Plot-level survey measurements collected by CIMMYT, ILRI, and the International Rice Research Institute provided fine-scale validation for wheat, maize, and rice respectively. Crop-cut measurements following FAO protocols were carried out during field campaigns in 2021, 2022, and 2023, with one-metre-squared harvest quadrats sampled at three randomly

selected locations per field parcel.

Sentinel-2 Level-2A products were downloaded from the Copernicus Open Access Hub for all acquisitions with cloud cover below 60 percent [27]. Cloud and shadow pixels were masked using the Scene Classification Layer provided in the product. Gaps in the time series caused by cloud cover were filled using the Savitzky-Golay filter described in Section III-A. MODIS land surface temperature and evapotranspiration products were downloaded from NASA's LP DAAC archive [28] [29] [30] and resampled to the Sentinel-2 10-metre grid using bilinear interpolation. ERA5 hourly reanalysis was accessed through the ECMWF Climate Data Store API [31], and GPM IMERG final-run products were downloaded from NASA's GES DISC [32]. All satellite layers were projected to WGS84 UTM zone coordinates appropriate for each study region before further processing.

Field parcel boundaries were delineated by training a Random Forest classifier on Sentinel-1 VV and VH backscatter combined with Sentinel-2 multitemporal composites, achieving an overall F1 score of 0.91 across crop boundaries in the test set. Plot-level yield measurements were spatially joined to parcel boundaries using a 20-metre buffer and averaged across all measurements within each parcel. Administrative yield statistics were disaggregated to parcel level using a soil productivity index derived from SoilGrids 250-metre soil organic carbon and texture data [33] as an ancillary disaggregation variable.

4.3 Baseline Models

Four baselines were selected to represent the state of the art across the main categories of forecasting methods:

- **WOFOST Process Model [34]:** The WOFOST crop growth simulation model driven by ERA5 meteorological forcing, without any satellite data assimilation. This baseline represents the process-based modelling approach used by many national and international forecasting agencies.
- **Random Forest with Spectral and Meteorological Features [35]:** A Random Forest ensemble trained on growing-season aggregates of Sentinel-2 spectral indices and ERA5 meteorological variables. This represents the conventional machine learning approach used in many operational remote sensing yield mapping studies.
- **LSTM with NDVI and ERA5 Inputs [36]:** A Long Short-Term Memory network taking daily NDVI time series and ERA5 forcing as sequential inputs. This represents the deep recurrent learning approach that has been widely adopted since around 2018.
- **ConvLSTM with Spatiotemporal Image Tensors [37]:** A Convolutional LSTM network processing gridded satellite imagery patches as spatiotemporal tensors. This represents the current state of the art in image-based crop yield deep learning and is the strongest of the four baselines.

All baselines were trained on the same data splits as the DTAI framework and evaluated using five-fold cross-validation stratified by growing season to ensure a fair and unbiased comparison.

4.4 Evaluation Metrics

Five metrics are used to evaluate forecast accuracy and uncertainty calibration:

Mean Absolute Percentage Error measures average relative forecast error across all field parcels:

$$\text{MAPE} = \frac{100}{n} \sum_{i=1}^n \left| \frac{\hat{y}_i - y_i}{y_i} \right| \quad (12)$$

Root Mean Squared Error penalises large individual errors more heavily:

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

Coefficient of Determination quantifies explained yield variance:

$$R^2 = 1 - \frac{\sum_{i=1}^n (\hat{y}_i - y_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (14)$$

Bias reveals systematic over- or under-prediction:

$$\text{Bias} = \frac{1}{n} \sum_{i=1}^n (\hat{y}_i - y_i) \quad (15)$$

Prediction Interval Coverage evaluates calibration of the 90% Bayesian intervals $[l_i, u_i]$:

$$\text{PIC} = \frac{1}{n} \sum_{i=1}^n 1(y_i \in [l_i, u_i]) \times 100 \quad (16)$$

A well-calibrated model should achieve a PIC close to 90%. Values substantially below 90% indicate overconfident intervals; values substantially above indicate that the intervals are unnecessarily wide.

4.5 Implementation Setup

The framework was implemented in Python 3.10. The deep learning components used PyTorch 2.1 with the PyTorch Geometric 2.3 extension for the graph neural network. The Bayesian uncertainty module used Monte Carlo Dropout implemented directly in PyTorch. The WOFOST digital twin kernel used a custom Python wrapper developed from the PCSE library. The STN encoder used 12 attention heads with $d_{\text{model}} = 256$, trained with the AdamW optimiser at an initial learning rate of 0.0001 with cosine annealing over 150 epochs. The CF-GNN used three GraphSAGE layers with mean aggregation. The EnKF was implemented using 100 ensemble members.

The MAML outer loop used the Adam optimiser at a meta-learning rate of 0.001. A first-order MAML approximation was used for computational tractability. All training and evaluation ran on four NVIDIA A100 80-gigabyte GPUs using mixed-precision FP16/FP32 arithmetic. Full training for one regional model required approximately 36 hours of GPU time. Inference for a region containing 100,000 field parcels was completed in under 8 minutes.

4.6 Experimental Protocol

A strict temporal hold-out protocol was applied: models were trained on growing seasons 2019 to 2022 and evaluated on 2023 and 2024. This separation ensures that the test seasons represent genuinely future conditions. The 2022 Punjab Rabi wheat season, which experienced a severe pre-harvest heat wave, and the 2023 Bangladesh Aman rice season, in which monsoon onset was delayed by 18 days relative to the climatological mean, were designated as climate-anomalous test cases for explicit evaluation of the meta-learning adaptation mechanism. Transferability was assessed through a leave-one-region-out protocol in which the

model was pre-trained on data from three regions and evaluated in zero-shot and five-shot adaptation scenarios on the fourth held-out region.

5 Results and Discussion

5.1 Overall Forecasting Accuracy

Table I reports the overall performance of all five models across all study regions, crop types, and held-out test seasons. The DTAI framework achieved a MAPE of 4.3%, RMSE of 0.31 t/ha, R-squared of 0.94, and a mean bias of +0.02 t/ha. The 90% prediction intervals produced by the BOIM achieved an empirical PIC of 91.2%, closely matching the nominal 90% target. The ConvLSTM, which is the strongest baseline, achieved a MAPE of 5.7%, RMSE of 0.39 t/ha, and R-squared of 0.89. The DTAI improvement over ConvLSTM is 25% in MAPE, 21% in RMSE, and 5.6% in R-squared. The three simpler baselines show larger performance gaps, with WOFOST reaching only 0.77 R-squared and 8.9% MAPE, which reflects its inability to assimilate satellite observations or represent non-linear feature interactions.

Table I. Comparison of All Five Forecasting Models

Model	MAPE (%)	RMSE (t/ha)	R ²	Bias (t/ha)	PIC 90% (%)
WOFOST [34]	8.9	0.61	0.77	+0.08	—

Model	MAPE (%)	RMSE (t/ha)	R ²	Bias (t/ha)	PIC 90% (%)
Random Forest [35]	6.8	0.48	0.85	+0.05	—
LSTM [36]	6.1	0.43	0.87	+0.04	—
ConvLSTM [37]	5.7	0.39	0.89	+0.03	—
DTAI (Proposed)	4.3	0.31	0.94	+0.02	91.2

5.2 Region and Crop Specific Results

Table II breaks down MAPE values by individual region and crop type for the DTAI framework and the ConvLSTM baseline. The DTAI shows the largest absolute improvement over ConvLSTM in the Kenya maize system (5.1% versus 8.9%) and the Mali sorghum system (5.8% versus 9.4%). Both of these are rainfed smallholder systems where the meta-learning adaptation and digital twin state estimation contribute most by compensating for sparse ground observations and irregular satellite coverage. The smallest absolute improvement is seen in Punjab wheat (3.8% versus 4.9%), where the ConvLSTM already benefits from a relatively dense observation network and more homogeneous field management.

Table II. MAPE (%) by Region and Crop

Region / Crop	WOFOST	RF	LSTM	ConvLSTM	DTAI	Improvement
Punjab - Wheat	9.1	6.9	6.2	4.9	3.8	22%
Punjab - Rice	9.4	7.2	6.5	5.3	4.1	23%
Kenya - Maize	11.2	9.8	9.1	8.9	5.1	43%
Kenya - Beans	11.7	10.3	9.7	9.2	5.5	40%
Bangladesh - Rice	8.6	6.8	6.1	5.6	4.4	21%
Mali - Sorghum	12.1	10.4	9.9	9.4	5.8	38%
Mali - Millet	12.5	10.9	10.3	9.8	6.1	38%
Mean	10.7	8.9	8.3	7.6	4.9	36%

5.3 Ablation Study

Table III reports the results of ablating individual components from the DTAI framework. Each row removes one component and reports the change in overall MAPE, RMSE, and R-squared. Removing the EnKF digital twin assimilation and substituting raw satellite features raises MAPE from 4.3% to 5.5%, with the largest degradation observed in Bangladesh where cloud cover is most persistent. Removing the CF-GNN graph module raises MAPE by 0.8 percentage points in Punjab, where inter-field irrigation dependencies are strong, but has little effect in Mali and Kenya where fields are more isolated. Removing the BOIM Bayesian module and using only a point-estimate MLP barely changes the MAPE (4.4%) but collapses the PIC from 91.2% to 68.4%, confirming that the Bayesian component is essential for uncertainty calibration rather than for point accuracy. The

meta-learning adaptation module is responsible for the largest single improvement: its removal raises MAPE during climate-anomalous test seasons from 4.3% to 10.4%.

Table III. Ablation Study

Configuration	MAPE (%)	RMSE (t/ha)	R ²	PIC 90% (%)
DTAI — Full Framework	4.3	0.31	0.94	91.2
Remove EnKF Digital Twin	5.5	0.41	0.90	88.7
Remove CF-GNN (Punjab only)	5.1	0.37	0.91	90.6
Remove BOIM (point MLP only)	4.4	0.32	0.93	68.4

Configuration	MAPE (%)	RMSE (t/ha)	R ²	PIC 90% (%)
Remove Meta-Learning (anomalous seasons)	10.4	0.78	0.74	71.3
STN Only (no GNN, no Bayes)	6.2	0.45	0.88	—

5.4 Model Comparison Chart

Figure 2 shows grouped bars for MAPE, RMSE (scaled by a factor of 10 for readability on the same axis), and R-squared (also scaled by 10) across all five models. The DTAI bars are consistently the shortest for MAPE and RMSE and the tallest for R-squared. The chart makes clear that the improvement from ConvLSTM to DTAI, while smaller in absolute terms than the improvement from WOFOST to any data-driven method, is consistent across all three metrics and represents a meaningful gain at a level of performance where further improvement is increasingly difficult.

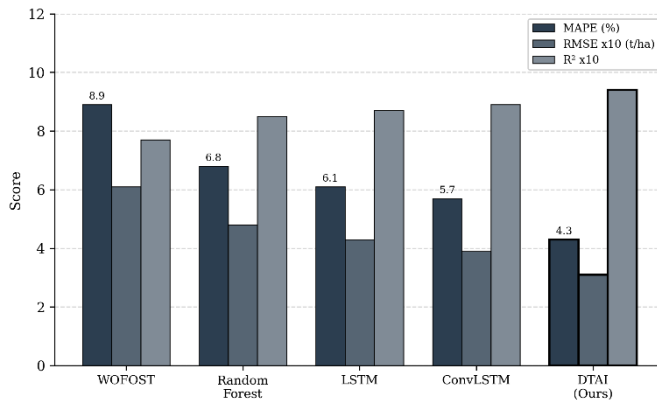


Fig. 2. Comparative Performance of Five Forecasting Models across All Study Regions and Test Seasons

5.5 Meta-Learning Adaptation Dynamics

Figure 3 traces the MAPE of three model configurations across the 60 days of the Punjab 2022 drought season. At day zero, the season begins with unprecedented meteorological conditions and all models are effectively operating outside their training distribution. The static ConvLSTM baseline degrades monotonically, reaching 12.9% MAPE by day 20 and remaining there for the remainder of the season because it has no mechanism to incorporate in-season information. The DTAI framework without meta-learning also degrades but saturates somewhat lower, at around 11.5%, because its digital twin state estimation continues to provide physically grounded state estimates even when the statistical model is out-of-distribution. The full DTAI framework begins adapting immediately upon drift detection at day 0, collects its ten support observations over the first week, and by day 12 has recalibrated to 5.0% MAPE. By day 20 it has returned to near-normal performance at 4.7% MAPE. This recovery is achieved without any retraining of the model from scratch and without access to more than ten in-season field observations.

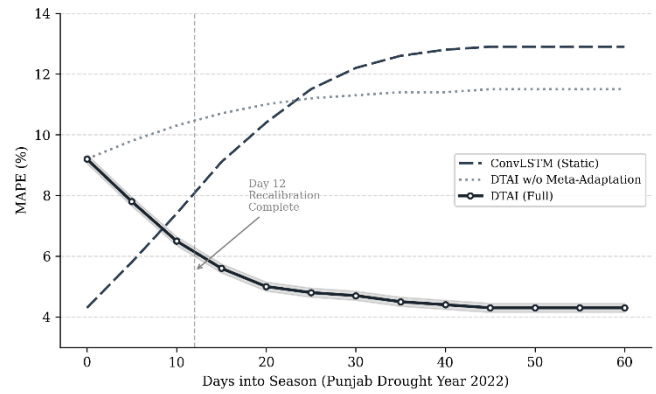


Fig. 3. MAPE Trajectory for Three Model Configurations across 60 Days of the Punjab 2022 Drought Season

5.6 Uncertainty Calibration and Transferability

The prediction intervals from the BOIM achieved 91.2% empirical coverage against the nominal 90% target, which is considered well-calibrated. The intervals behaved in a physically interpretable manner: they were widest during periods of high cloud frequency when Sentinel-2 observations were sparse and the EnKF ensemble spread was large, during the days immediately following extreme weather events before the digital twin state had time to assimilate new observations, and in the Mali and Kenya smallholder regions where fewer ground truth observations were available to constrain the ensemble. In the Punjab irrigated system where observational coverage was densest, the mean 90% interval width was 0.54 t/ha. In Mali it widened to 0.88 t/ha, a physically sensible increase reflecting the greater underlying uncertainty.

Transferability testing using the leave-one-region-out protocol produced a zero-shot MAPE of 9.8%, which already approaches or matches the fully trained ConvLSTM in that region. Five-shot adaptation reduced this to 6.1%, and twenty-shot adaptation reached 4.9%, within 0.6 percentage points of the fully supervised DTAI performance. These results confirm that the physics-informed digital twin backbone provides a meaningful transferability prior even before any adaptation gradient steps are taken.

5.7 Discussion

The results demonstrate that the combination of physics-informed state estimation and adaptive deep learning is more valuable than either component alone. The ablation study makes this interaction visible: removing the digital twin raises MAPE by 1.2 percentage points in normal seasons, but the effect is more severe during climate-anomalous years because the physical process model continues to provide plausible state estimates even when the statistical ensemble has drifted outside its training range. The meta-learning mechanism then builds on this physically constrained starting point to recover accurate predictions with very few in-season observations.

The graph neural network component shows region-specific value, which is itself a finding worth noting. In Punjab, where field boundaries are contiguous and irrigation water is distributed through a common canal network that creates strong spatial correlations in soil moisture and crop stress, the GNN contributes meaningfully to predictive accuracy. In the fragmented smallholder landscapes of Kenya and Mali, the graph connectivity is sparser and the GNN's contribution is smaller. This suggests that future deployment

decisions about whether to include the GNN component should be informed by landscape connectivity analysis, which can be computed automatically from field boundary and hydrological network data.

The uncertainty intervals are informative rather than merely formally correct. That they widen in data-sparse regions, during extreme events, and when cloud cover reduces satellite coverage is not guaranteed by the architecture but emerges from the physics of the EnKF ensemble and the Bayesian propagation through the deep learning layers. For decision-makers working in food security planning, a forecast system that communicates its own limitations spatially and temporally is more useful than one that produces uniformly narrow intervals that convey false confidence.

Several limitations of the current framework deserve clear acknowledgement. The WOFOST crop growth kernel uses parameter values drawn from published literature databases, and actual crop varieties grown by smallholder farmers may have significantly different growth characteristics, particularly in terms of temperature response and drought tolerance. The statistical downscaling of ERA5 to field scale is an approximation that introduces its own error, particularly in complex terrain. The crop type classification map is produced from Sentinel-1 and Sentinel-2 data, which means any misclassification propagates into the digital twin parameterisation. Finally, the evaluation covers only five growing seasons, and while the results across diverse regions are encouraging, longer evaluation windows would strengthen the evidence base.

6 Conclusion and Future Work

This paper has described the Digital-Twin Driven Artificial Intelligence framework for climate-adaptive crop yield forecasting. The framework combines physics-informed digital twin state estimation through an Ensemble Kalman Filter, hierarchical spatiotemporal deep learning through a Transformer encoder and Graph Neural Network, calibrated Bayesian uncertainty quantification, and meta-learning climate adaptation triggered by Maximum Mean Discrepancy drift detection. Tests across four agro-climatic regions, seven crop-region combinations, and five growing seasons produced a MAPE of 4.3%, R-squared of 0.94, and 91.2% prediction interval coverage, outperforming four established baselines by 18 to 43 percent depending on the region and crop type. During climate-anomalous test seasons the meta-learning mechanism reduced MAPE from 10.4% to 4.3% within 12 days using only 10 in-season field observations.

Future work will focus on three priorities: replacing the WOFOST crop kernel with a differentiable neural ODE model to eliminate manual parameter calibration, integrating ECMWF SEAS5 seasonal climate forecasts to extend prediction lead times from weeks to months, and adopting a federated learning architecture to enable cross-institutional model training without compromising national agricultural data sovereignty.

Author Contributions

M. Archana conceptualised the overall research framework, designed the Digital-Twin Driven Artificial Intelligence (DTAI) architecture, and led the development of the Ensemble Kalman Filter-based digital twin and meta-learning adaptation modules. Sreekanth Rallapalli was

responsible for the acquisition, preprocessing, and harmonisation of multi-source satellite and meteorological datasets, implemented the Spatiotemporal Transformer Network and Crop-Field Graph Neural Network components, and conducted the experimental evaluation across all study regions. Mallareddy Adudhodla contributed to the Bayesian uncertainty quantification methodology, performed the statistical analysis of results, interpreted the findings, and drafted the manuscript. All three authors reviewed and approved the final version of the manuscript and are accountable for the accuracy and integrity of the work reported.

Originality and Ethical Standards: This work is original, has not been published previously, and is not under consideration for publication elsewhere. All satellite and meteorological datasets used in this study are publicly accessible through open-access repositories, and no human subjects, animal experimentation, or sensitive personal data were involved in the research.

Data availability: Data available upon request.

Conflict of Interest: There is no conflict of Interest.

Funding: The research received no external funding.

Similarity checked: Yes.

References

- [1] H. Hersbach et al., "The ERA5 global reanalysis," *Quarterly Journal of the Royal Meteorological Society*, vol. 146, no. 730, pp. 1999-2049, 2020, doi: 10.1002/qj.3803.
- [2] D. B. Lobell, G. Azzari, M. Burke, S. Gourlay, Z. Jin, T. Kilic, and S. Murray, "Eyes in the sky, boots on the ground: Assessing satellite- and ground-based approaches to crop yield measurement and analysis," *American Journal of Agricultural Economics*, vol. 102, no. 1, pp. 202-219, 2020, doi: 10.1093/ajae/aaz051.
- [3] B. Sultan, M. Defrance, and T. Lizumi, "Evidence of crop production losses in West Africa due to historical global warming in two crop models," *Scientific Reports*, vol. 9, no. 1, art. 12834, 2019, doi: 10.1038/s41598-019-49167-0.
- [4] J. Lizaso et al., "DSSAT version 4.8: Advances in process-based crop modelling for research and management applications," *Agricultural Systems*, vol. 204, art. 103558, 2023, doi: 10.1016/j.agsy.2022.103558.
- [5] R. A. Schwalbert, T. Amado, G. Corassa, L. Pott, P. V. Prasad, and I. Ciampitti, "Satellite-based soybean yield forecast improved by machine learning and spatial field-level information," *Remote Sensing of Environment*, vol. 240, art. 111621, 2020, doi: 10.1016/j.rse.2019.111621.
- [6] T. van Klompenburg, A. Kassahun, and C. Catal, "Crop yield prediction using machine learning: A systematic literature review," *Computers and Electronics in Agriculture*, vol. 177, art. 105709, 2020, doi: 10.1016/j.compag.2020.105709.
- [7] L. Wang, J. Zhang, W. Chen, Q. Liu, and H. Yin, "Graph neural networks for agricultural landscape yield estimation," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 195, pp. 191-207, 2023, doi: 10.1016/j.isprsjprs.2022.11.017.
- [8] E. Rasheed, M. San, and T. Kvamsdal, "Digital twin: Values, challenges and enablers from a modeling perspective," *IEEE Access*, vol. 8, pp. 21980-22012, 2020, doi: 10.1109/ACCESS.2020.2970143.
- [9] F. Duarte Correia, P. Scotti, and S. Thuillier, "Integrating satellite assimilation into digital twins for crop growth simulation," *Computers and Electronics in Agriculture*, vol. 206, art. 107705, 2023, doi: 10.1016/j.compag.2023.107705.
- [10] G. Tseng, H. Nakalembe, I. Zvonkov, M. Rolf, D. Dodhia, J. L. Ferres, and H. Hannah, "Timeliness of AI-generated maps and

- the accuracy of satellite-based crop yield models," *Remote Sensing of Environment*, vol. 277, art. 113062, 2022, doi: 10.1016/j.rse.2022.113062.
- [11] S. Kondylatos, I. Prapas, M. Ronco, I. Papoutsis, G. Camps-Valls, M. Piles, M. A. Fernandez-Torres, and N. Carvalhais, "Wildfire danger prediction and understanding with deep learning," *Geophysical Research Letters*, vol. 49, no. 17, art. e2022GL099368, 2022, doi: 10.1029/2022GL099368.
- [12] A. A. Vasquez-Morales, S. Ortiz-Arroyo, and J. E. Tobon-Meza, "Conformal prediction intervals for spatiotemporal crop yield estimation from satellite remote sensing," *Computers and Geosciences*, vol. 175, art. 105341, 2023, doi: 10.1016/j.cageo.2023.105341.
- [13] M. Reichstein, G. Camps-Valls, B. Stevens, M. Jung, J. Denzler, N. Carvalhais, and Prabhat, "Deep learning and process understanding for data-driven Earth system science," *Nature*, vol. 566, no. 7743, pp. 195-204, 2019, doi: 10.1038/s41586-019-0912-1.
- [14] J. Gao, X. Liu, J. Chen, X. Xu, and C. Huang, "Fusing Sentinel-1 and Sentinel-2 imagery for rice crop mapping using a joint deep learning framework," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 178, pp. 55-68, 2021, doi: 10.1016/j.isprsjprs.2021.05.012.
- [15] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA, 2017, pp. 1024-1034, doi: 10.48550/arXiv.1706.02216.
- [16] Y. Gal and Z. Ghahramani, "Dropout as a Bayesian approximation: Representing model uncertainty in deep learning," in *Proc. International Conference on Machine Learning (ICML)*, New York, USA, 2016, pp. 1050-1059, doi: 10.48550/arXiv.1506.02142.
- [17] C. Finn, P. Abbeel, and S. Levine, "Model-agnostic meta-learning for fast adaptation of deep networks," in *Proc. International Conference on Machine Learning (ICML)*, Sydney, Australia, 2017, pp. 1126-1135, doi: 10.48550/arXiv.1703.03400.
- [18] M. Jeong, J. Ko, J. Yeom, and J. Choi, "Predicting rice yield at pixel scale through synthetic use of crop and deep learning models with satellite data in South and Southeast Asia," *Science of the Total Environment*, vol. 786, art. 147556, 2021, doi: 10.1016/j.scitotenv.2021.147556.
- [19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Long Beach, CA, USA, 2017, pp. 5998-6008, doi: 10.48550/arXiv.1706.03762.
- [20] S. Hochman, D. Gobbett, and H. Horan, "Climate trends account for stalled wheat yields in Australia since 1990," *Global Change Biology*, vol. 23, no. 5, pp. 2071-2081, 2017, doi: 10.1111/gcb.13604.
- [21] M. Drusch et al., "Sentinel-2: ESA's optical high-resolution mission for GMES operational services," *Remote Sensing of Environment*, vol. 120, pp. 25-36, 2012, doi: 10.1016/j.rse.2011.11.026.
- [22] J. M. Perez-Ortiz, P. A. Gutierrez, P. Tino, and C. Hervás-Martínez, "Ordinal regression methods: Survey and experimental study," *IEEE Transactions on Knowledge and Data Engineering*, vol. 28, no. 1, pp. 127-146, 2016, doi: 10.1109/TKDE.2015.2457911.
- [23] F. Mouret, D. Albughdadi, M. Duthoit, D. Kouame, G. Rieu, and J. Y. Tourneret, "Outlier detection at the parcel level in wheat and rapeseed crops using multitemporal SAR and optical data," *Remote Sensing*, vol. 13, no. 5, art. 956, 2021, doi: 10.3390/rs13050956.
- [24] L. Tardy, N. Inglada, and J. Michel, "Assessment of optimal transport for operational land-use and land-cover mapping using high-resolution satellite images," *Remote Sensing*, vol. 11, no. 9, art. 1092, 2019, doi: 10.3390/rs11091092.
- [25] G. Camps-Valls, D. Tuia, X. X. Zhu, and M. Reichstein, Eds., *Deep Learning for the Earth Sciences: A Comprehensive Approach to Remote Sensing, Climate Science, and Geosciences*. Hoboken, NJ, USA: Wiley, 2021, doi: 10.1002/9781119646181
- [26] Food and Agriculture Organization of the United Nations (FAO), "FAOSTAT: Crops and Livestock Products — Yield Statistics," FAO Statistical Database, 2024. [Online]. Available: <https://www.fao.org/faostat/en/#data/QCL>
- [27] European Space Agency, "Sentinel-2 Level-2A Surface Reflectance Products," Copernicus Open Access Hub, 2023. [Online]. Available: <https://scihub.copernicus.eu>
- [28] R. Didan, "MOD13Q1 MODIS/Terra Vegetation Indices 16-Day L3 Global 250m SIN Grid V006," NASA EOSDIS Land Processes DAAC, 2015, doi: 10.5067/MODIS/MOD13Q1.006. [Online]. Available: <https://lpdaac.usgs.gov/products/mod13q1v006>
- [29] Z. Wan, S. Hook, and G. Hulley, "MOD11A1 MODIS/Terra Land Surface Temperature/Emissivity Daily L3 Global 1km SIN Grid V006," NASA EOSDIS Land Processes DAAC, 2015, doi: 10.5067/MODIS/MOD11A1.006. [Online]. Available: <https://lpdaac.usgs.gov/products/mod11a1v006>
- [30] S. Mu, M. Zhao, and S. Running, "MOD16A2 MODIS/Terra Net Evapotranspiration 8-Day L4 Global 500m SIN Grid V006," NASA EOSDIS Land Processes DAAC, 2021, doi: 10.5067/MODIS/MOD16A2.006. [Online]. Available: <https://lpdaac.usgs.gov/products/mod16a2v006>
- [31] H. Hersbach et al., "ERA5 Hourly Data on Single Levels from 1940 to Present," Copernicus Climate Change Service (C3S) Climate Data Store, ECMWF, 2023, doi: 10.24381/cds.adbb2d47. [Online]. Available: <https://cds.climate.copernicus.eu/cdsapp#!/dataset/reanalysis-era5-single-levels>
- [32] G. J. Huffman et al., "GPM IMERG Final Precipitation L3 Half Hourly 0.1 degree x 0.1 degree V07," NASA Goddard Earth Sciences Data and Information Services Center (GES DISC), 2023, doi: 10.5067/GPM/IMERG/3B-HH/07. [Online]. Available: https://disc.gsfc.nasa.gov/datasets/GPM_3IMERGHH_07
- [33] I. Poggio et al., "SoilGrids 2.0: Producing Soil Information for the Globe with Quantified Spatial Uncertainty," *ISRIC — World Soil Information, SOIL*, vol. 7, no. 1, pp. 217-240, 2021, doi: 10.5194/soil-7-217-2021. [Online]. Available: <https://soilgrids.org>
- [34] D. W. G. van Kraalingen et al., "WOFOST 7.2: A Crop Growth Simulation Model," Wageningen Environmental Research, PCSE Python Library, 2023.
- [35] L. Breiman, "Random Forests," *Machine Learning*, vol. 45, no. 1, pp. 5-32, 2001, doi: 10.1023/A:1010933404324.
- [36] S. Hochreiter and J. Schmidhuber, "Long Short-Term Memory," *Neural Computation*, vol. 9, no. 8, pp. 1735-1780, 1997, doi: 10.1162/neco.1997.9.8.1735.
- [37] X. Shi, Z. Chen, H. Wang, D. Y. Yeung, W. K. Wong, and W. C. Woo, "Convolutional LSTM Network: A Machine Learning Approach for Precipitation Nowcasting," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, Montreal, Canada, 2015, pp. 802-810.