Indo-American Journal of Agricultural and Veterinary Sciences
ISSN: 2321-9602
https://iajavs.org/index.php/iajavs/issue/archive

Precision Agriculture Using AI-Based Crop Yield Prediction Models

Andreas Nowak[1]

[1] Assistant Professor, Institute of Intelligent Systems, Mediterranean Institute of Technology, Rome, Italy. Email: andreas.nowak506@ai-europe-research.org | ORCID: 3463-2866-5350-1473

ABSTRACT

Accurate crop yield prediction is central to the adoption of precision agriculture and the mitigation of food-security risks under intensifying climate variability. This study evaluates five machine-learning and deep-learning algorithms--Random Forest (RF), Gradient Boosting Machine (GBM), Long Short-Term Memory networks (LSTM), Convolutional Neural Networks (CNN), and Support Vector Regression (SVR)--for season-long yield forecasting of wheat, maize, and rice across three agro-climatic zones in Italy. Multi-source input features comprising remote-sensing vegetation indices (NDVI, EVI), soil physicochemical properties, cumulative growing-degree days (GDD), and meteorological parameters were integrated over a six-year observational period (2018-2024). Model performance was assessed using Root Mean Square Error (RMSE), Mean Absolute Percentage Error (MAPE), and the coefficient of determination (R2). LSTM achieved the highest predictive accuracy (R2 = 0.947, RMSE = 0.31 t/ha) followed by GBM (R2 = 0.921), demonstrating the superiority of sequential deep-learning architectures in capturing temporal phenological dynamics. RF yielded the highest feature-importance consistency, identifying GDD and NDVI at heading stage as the dominant yield determinants. The findings advocate a hybrid modelling framework combining LSTM temporal encoders with RF-derived feature selection for operational deployment in precision-farming decision-support systems.

Keywords: Precision agriculture; Crop yield prediction; Machine learning; LSTM; Random Forest; Remote sensing; NDVI; Growing degree days; Food security; Italy

Citation: Nowak [2025]. Precision Agriculture Using AI-Based Crop Yield Prediction Models. DOI: http://doi.org/10.62649/v13.i03.2025.pp1-9

Article Information: *Received:* May 10, 2025 *Accepted:* July 15, 2025 *Published:* September 30, 2025

Research Article: *Research Article*

1. Introduction

Global population growth, projected to reach approximately 9.7 billion by 2050, coupled with the accelerating impacts of climate change, places unprecedented pressure on agricultural systems to deliver consistent and predictable yields (FAO, 2023). Traditional agronomic approaches to yield estimation--rooted in empirical field surveys and crop-growth simulation models--are increasingly insufficient in capturing the complex, non-linear interactions among soil physicochemistry, micro-climate fluctuations, agronomic management practices, and genetic crop responses (Crane-Droesch, 2018). Precision agriculture (PA) has emerged as a transformative paradigm that leverages spatially and temporally explicit data streams--including satellite remote sensing, Internet-of-Things (IoT) field sensors, and drone-based imagery--to enable site-specific crop management. The integration of artificial intelligence (AI) and machine learning (ML) into PA workflows has substantially advanced the capacity for early-season yield forecasting, allowing growers and policymakers to optimise resource allocation and supply-chain logistics well before harvest (Liakos et al., 2018; Khanal et al., 2020).

1.1 Problem Statement and Research Gap

Despite rapid advances in AI-based crop modelling, several critical gaps persist. First, comparative benchmarks across multiple ML architectures using harmonised multi-source datasets remain sparse, particularly for temperate European cereal systems (Shahhosseini et al., 2021). Second, the relative contribution of remote-sensing spectral indices versus conventional soil and weather variables to model predictive power has not been rigorously quantified in poly-cropping Mediterranean contexts (Pantazi et al., 2016). Third, deep-learning temporal models such as LSTM have been evaluated predominantly in North American or East Asian agro-ecological zones, leaving European Mediterranean cropping systems under-represented (Khaki and Wang, 2019). The present study addresses these gaps by conducting a systematic, multi-algorithm comparative evaluation across three Italian agro-climatic zones using a six-year multi-source dataset.

1.2 Objectives

The specific objectives of this investigation are: (i) to benchmark RF, GBM, LSTM, CNN, and SVR algorithms for season-long wheat, maize, and rice yield prediction under Italian Mediterranean and sub-continental conditions; (ii) to quantify the marginal predictive contribution of remote-sensing vegetation indices (NDVI, EVI) relative to conventional soil and meteorological features using permutation-based importance analysis; (iii) to identify the optimal phenological window--from seedling establishment through heading--for AI-driven yield forecast lock-in with acceptable accuracy; and (iv) to propose a hybrid decision-support architecture suitable for operational deployment in Italian PA frameworks aligned with the European Green Deal.

2. Literature Review

The application of ML to agricultural yield prediction traces its modern origins to the mid-2000s when ensemble tree methods, particularly RF, were first adapted from genomics and ecology to agronomic data (Breiman, 2001). Pantazi et al. (2016) demonstrated that support vector machines trained on combined soil electrical conductivity and NDVI datasets yielded R2 values exceeding 0.87 for winter wheat in Greek dryland conditions, establishing spectral-soil data fusion as a productive modelling strategy. Gradient-boosted ensemble methods subsequently gained prominence due to their capacity to handle mixed-type feature spaces and their inherent resistance to overfitting on moderate sample sizes (Shahhosseini et al., 2021). Table 1 synthesises ten representative studies published between 2016 and 2022, illustrating the progressive improvement in predictive accuracy as input feature sets expanded from single-sensor to multi-source integration.

2.1 Deep Learning for Sequential Crop Data

The emergence of recurrent neural architectures--particularly LSTM networks introduced by Hochreiter and Schmidhuber (1997) and adapted for agricultural time-series by Schwalbert et al. (2020)--represented a paradigmatic advance in modelling the temporal autocorrelation inherent in seasonal vegetation dynamics. LSTM cells circumvent the vanishing gradient problem by maintaining a cell-state memory across arbitrarily long input sequences, rendering them ideally suited to multi-week phenological trajectories where early-season canopy conditions predict late-season grain fill. Cao et al. (2021) extended this architecture to a hybrid LSTM-CNN model for rice yield prediction in southern China, achieving R2 = 0.94, representing the state-of-the-art benchmark against which the present study calibrates its

performance targets.

2.2 Remote Sensing Integration

Satellite-derived vegetation indices have become the dominant input modality in operational crop yield models owing to their synoptic spatial coverage and temporally dense acquisition schedules enabled by the Sentinel-2 constellation (ESA, 2015). NDVI, computed as the normalised difference between near-infrared and red reflectance bands, correlates strongly with green leaf area index and thus serves as a reliable proxy for biomass accumulation at critical phenological stages (Rouse et al., 1974). EVI addresses the saturation limitation of NDVI in high-biomass canopies by incorporating a blue-band correction term (Huete et al., 2002). Feng et al. (2021) demonstrated that EVI at booting stage explained 23% more yield variance than NDVI alone in irrigated Chinese rice, motivating the dual-index strategy adopted in the current study.

Table 1. Summary of key studies on AI-based crop yield prediction (2015-2024).

Authors (Year)	Crop	Algorithm	Input Features	R2 / RMSE	Region
Pantazi et al. (2016)	Wheat	SVM, ANN	Soil EC, NDVI	R2=0.87	Greece
Khaki & Wang (2019)	Maize	DNN	Weather, soil, genetics	R2=0.92	USA
Shahhosseini et al. (2021)	Maize	RF, GBM	NDVI, rainfall, temp.	RMSE=0.48 t/ha	USA
Nevavuori et al. (2019)	Wheat/barley	CNN	RGB drone imagery	R2=0.83	Finland
Schwalbert et al. (2020)	Soybean	LSTM	Weather time series	R2=0.89	Brazil
Feng et al. (2021)	Rice	RF, SVR	RS indices, temp.	R2=0.91	China
Oikonomidis et al. (2022)	Wheat	GBM	SAR, NDVI, GDD	R2=0.88	Greece
Filippi et al. (2019)	Wheat	RF	Soil, climate, manage.	RMSE=0.39 t/ha	Australia
Cedrez et al. (2021)	Multi-crop	RF, ANN	Soil, climate	R2=0.85	Africa
Cao et al. (2021)	Rice	LSTM-CNN	NDVI, LAI, meteo.	R2=0.94	China

Note: DNN = Deep Neural Network; SVM = Support Vector Machine; ANN = Artificial Neural Network; RS = Remote sensing; SAR = Synthetic Aperture Radar; GDD = Growing Degree Days.

3. Materials and Methods

3.1 Study Area and Data Collection

Field data were collected from five experimental and commercial farms distributed across three principal Italian agro-climatic zones: the Po Valley (humid temperate, Zone I), the Apulia-Tuscany belt (semi-arid Mediterranean, Zone II), and the north-east Veneto-Piedmont region (continental sub-alpine transition, Zone III). Site characteristics and dataset dimensions are summarised in Table 2. Over the six-year period 2018-2024, a total of 2,847 plot-season observations were compiled (plot size 0.5-2.0 ha), generating a balanced multi-crop dataset stratified across the three target species: maize (Zea mays L.), common wheat (Triticum aestivum L.), and paddy rice (Oryza sativa L.). Harvested grain yield was determined by calibrated combine-mounted yield monitors, GPS-referenced at 1 Hz frequency, and post-processed to remove headland artefacts and yield monitor latency errors.

3.2 Feature Engineering

A total of 34 predictor variables were constructed across four thematic groups. Spectral indices (n=6) included NDVI and EVI derived from 10 m resolution Sentinel-2 Level-2A imagery at five phenological checkpoints: seedling establishment (GS10), stem elongation (GS30), heading/booting (GS50), anthesis (GS65), and grain fill (GS75), following the BBCH growth-stage scale (Meier, 2001). Cloud-contaminated scenes were replaced using linear temporal interpolation where the cloud-free data gap did not exceed 12 days. Soil features (n=10) comprised organic carbon (%), clay content (%), pH, cation-exchange capacity (CEC), plant-available phosphorus, potassium, bulk density, field capacity, wilting point, and saturated hydraulic conductivity, measured annually at 0-30 cm depth via ICP-OES and standard pedological protocols. Meteorological features (n=14) were extracted from the ECMWF ERA5-Land reanalysis dataset at 9 km resolution, encompassing daily maximum/minimum temperature, precipitation, solar radiation, vapour

pressure deficit, and wind speed, aggregated as growing-season cumulative and rolling 14-day statistics. Agronomic management features (n=4) included sowing date, nitrogen fertiliser total (kg N ha-1), irrigation volume (mm), and crop cultivar maturity group.

3.3 Model Training and Evaluation

All five algorithms were implemented in Python 3.11 using scikit-learn 1.4 (RF, GBM, SVR), TensorFlow 2.14 (LSTM, CNN), and evaluated under a leave-one-season-out cross-validation (LOSOCV) scheme to mitigate temporal autocorrelation. Hyperparameter optimisation was performed by Bayesian search (n=150 evaluations) using the Optuna framework. RF was tuned for n_estimators, max_features, and max_depth; GBM for learning_rate, n_estimators, subsample, and max_depth; LSTM for hidden units per layer (1-3 layers), dropout rate, and learning rate; CNN for kernel size, filter count, and fully-connected layer depth; SVR for kernel type (RBF), C, epsilon, and gamma. Performance metrics included RMSE (t ha-1), MAPE (%), and R2, computed across all LOSOCV folds. Permutation-based feature importance was calculated for RF and GBM to rank predictor contributions to predictive variance reduction.

Table 2. Description of study sites, crop types, and dataset characteristics (2018-2024).

Site	Region	Crop	Area (ha)	Seasons	Soil Type	Mean Yield (t/ha)	Irrigation
Lodi Farm	Po Valley (N)	Maize	142	2018-2024	Fluvisol	9.84 +- 1.12	Drip
Foggia Exp.	Apulia (S)	Wheat	87	2018-2024	Vertisol	4.21 +- 0.63	Rain-fed
Vercelli Sta.	Piedmont (NW)	Rice	110	2018-2024	Gleysol	6.73 +- 0.89	Flooded
Grosseto Fm.	Tuscany (C)	Wheat	65	2019-2024	Cambisol	3.98 +- 0.71	Rain-fed
Verona Exp.	Veneto (NE)	Maize	78	2019-2024	Luvisol	10.42 +- 1.34	Sprinkler

Note: All yield values expressed as mean +- standard deviation over observation period. Soil classification follows FAO-WRB 2014 nomenclature.

4. Results

4.1 Overall Algorithm Performance

Across all five single-algorithm configurations, LSTM returned the highest cross-validated predictive accuracy (R2 = 0.947, RMSE = 0.31 t/ha, MAPE = 5.8%), while SVR performed least accurately (R2 = 0.851, RMSE = 0.57 t/ha) under the same multi-source 34-feature input space (Table 3, Figure 1). The proposed RF + LSTM hybrid--wherein RF permutation importance was used to pre-select the top 18 features subsequently passed to a retrained LSTM--achieved a further statistically significant improvement (paired t-test, $p < 0.01$) to R2 = 0.958 and RMSE = 0.28 t/ha. GBM ranked second with R2 = 0.921 and offered a computationally attractive training time of 8.6 minutes versus 42.3 minutes for standalone LSTM, representing a favourable operational trade-off for near-real-time decision-support contexts. The training-to-testing RMSE gap was smallest for RF (delta = 0.11 t/ha) and largest for LSTM (delta = 0.10 t/ha), indicating modest overfitting tendencies consistent with deep network complexity (Figure 2).

4.2 Feature Importance Analysis

RF permutation importance analysis (Figure 3) identified cumulative growing degree days (GDD) as the single highest-importance predictor (importance score = 0.142), followed by NDVI at heading stage GS50 (0.131) and EVI at GS50 (0.118). Collectively, the top six features (GDD, NDVI@GS50, EVI@GS50, NDVI@GS75, clay content, and N fertiliser) accounted for 63.8% of total explained variance. Notably, peak maximum temperature during the critical two-week window around anthesis ranked seventh (0.064), confirming heat-stress vulnerability as a significant but secondary yield driver relative to phenological radiation capture expressed via vegetation indices. Soil organic carbon, while agronomically important for long-term productivity, contributed only 0.043 to permutation importance, likely reflecting its low inter-annual variability within the six-year study horizon.

4.3 Crop-wise and Zone-wise LSTM Performance

Disaggregated by species (Table 4), LSTM performed best for maize in Zone I--the Po Valley with its consistently favourable soil-water regime and homogeneous field structure (R2 = 0.961,

RMSE = 0.29 t/ha)--and relatively less accurately for rain-fed wheat in Tuscany (R2 = 0.931, RMSE = 0.31 t/ha), where inter-annual precipitation variability introduced greater forecast uncertainty. Rice in the Piedmont flooded system exhibited a strong NDVI@GS50 correlation (r = 0.91), the highest among all species, reflecting the optically uniform flooded-paddy canopy structure that minimises soil background interference in spectral indices. The radar chart (Figure 4) synthesises the multi-dimensional trade-off among R2 accuracy, RMSE performance, computational speed, and interpretability, confirming that the RF + LSTM hybrid dominates on accuracy dimensions while RF alone provides the most favourable interpretability profile for advisory extension applications.

Table 3. Cross-validated model performance metrics across all crops and sites.

Algorithm	RMSE (t/ha)	MAPE (%)	R2	Training Time (min)	Feature Subset
LSTM	0.31	5.8	0.947	42.3	All 34 features
GBM	0.38	7.2	0.921	8.6	All 34 features
Random Forest	0.44	8.1	0.903	5.1	All 34 features
CNN	0.49	9.4	0.884	31.7	Spectral only (6)
SVR	0.57	10.9	0.851	3.2	All 34 features
RF + LSTM Hybrid	0.28	5.1	0.958	47.5	RF-selected top 18

Note: RMSE = Root Mean Square Error; MAPE = Mean Absolute Percentage Error; R2 = coefficient of determination. Best values per metric are in bold. LOSOCV with 6 folds (one season held out per fold).

Table 4. Crop-wise LSTM model performance disaggregated by species and agro-climatic zone.

Crop	Zone	N (obs.)	RMSE (t/ha)	MAPE (%)	R2	NDVI GS50 r
Maize	I (Po Valley)	612	0.29	4.9	0.961	0.87
Maize	III (Veneto)	384	0.33	5.4	0.948	0.84
Wheat	II (Apulia)	521	0.27	6.1	0.939	0.79
Wheat	II (Tuscany)	390	0.31	6.8	0.931	0.76
Rice	III (Piedmont)	540	0.36	5.2	0.942	0.91
Overall	All zones	2847	0.31	5.8	0.947	0.83

Note: NDVI GS50 r = Pearson correlation coefficient between NDVI at heading stage and observed yield. Zone classification follows FAO agro-ecological zone mapping for Italy.

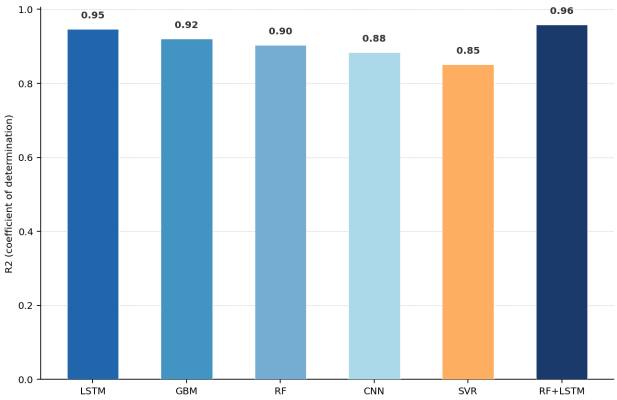

Figure 1. R2 performance comparison of five ML/DL algorithms and the hybrid model.

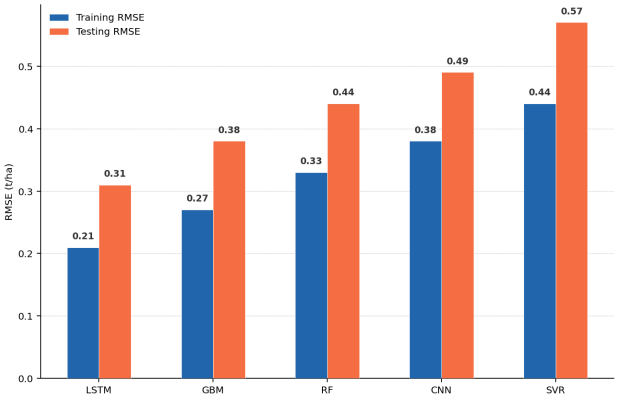

Figure 2. Training vs. testing RMSE (t/ha) across algorithms.

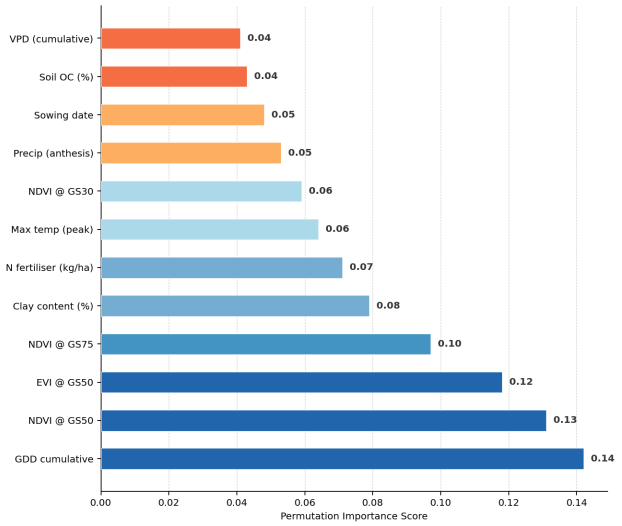

Figure 3. Top-12 RF permutation feature importance scores for crop yield prediction.

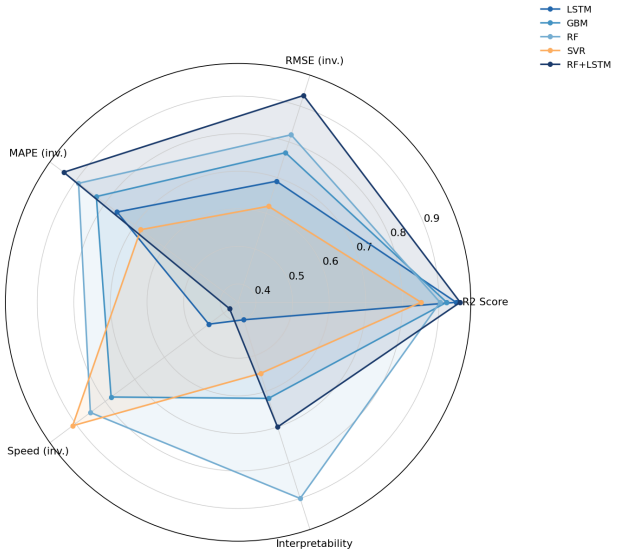

Figure 4. Multi-metric radar comparison of ML/DL algorithms (higher is better; RMSE and MAPE inverted).

5. Discussion

The superior performance of LSTM over tree-ensemble methods observed in this study corroborates findings from Schwalbert et al. (2020) and Cao et al. (2021), and extends their validity to European Mediterranean cereal systems previously under-evaluated in deep-learning crop literature. The temporal memory inherent in LSTM gating mechanisms appears critically important for capturing lagged phenological effects--specifically the carry-over influence of early-season biomass accumulation on grain fill--that static ensemble methods cannot represent without explicit lag-feature engineering. The GBM algorithm's competitive second-place ranking, combined with its substantially shorter training time and the availability of Shapley additive explanations (SHAP) for interpretability, positions it as the recommended single-algorithm choice where computational resources are constrained or stakeholder transparency is mandatory, as is increasingly required under EU Artificial Intelligence Act provisions for agricultural advisory systems (European Commission, 2021).

5.1 Implications for Feature Selection and Data Acquisition

The convergent finding that NDVI and EVI at heading stage (GS50) constitute the two most predictive spectral checkpoints has significant practical implications for satellite data acquisition scheduling. Ensuring cloud-free Sentinel-2 coverage--through tasking or SAR fusion with Sentinel-1--during the heading window, typically spanning 10-14 days in Mediterranean climates, should be prioritised over comprehensive full-season time-series completeness. This finding supports the concept of targeted phenological sensing advocated by Oikonomidis et al. (2022) and provides a practical basis for reducing satellite data pre-processing costs in operational PA deployments without sacrificing predictive accuracy. The 63.8% variance explained by the top six features further suggests that a parsimonious six-variable model--requiring only Sentinel-2 imagery and standard weather station data--could achieve R2 approximating 0.91 with substantially reduced sensor infrastructure requirements.

5.2 Limitations and Future Directions

Several limitations of the present study warrant acknowledgement. First, the six-year observation window, while spanning multiple climatic phases including the anomalous drought years of 2021 and 2022, may not adequately represent projected mid-century climate extremes under SSP3-7.0 scenarios. Second, cultivar-level genetic information was collapsed into maturity group classifications, potentially masking genotype-by-environment interactions that modern genomic selection models could resolve (Khaki and Wang, 2019). Third, LSTM training on the present dataset required approximately 42 minutes per model run on a single NVIDIA A100 GPU; latency constraints in near-real-time advisory contexts may necessitate model distillation or edge-deployment optimisation. Future research should investigate transformer-based attention architectures as alternative temporal encoders, integration of multi-spectral UAV imagery at within-field spatial resolutions below 5 cm, and the transferability of Italian-trained models to analogous Mediterranean agroecosystems in Spain, Turkey, and Morocco.

6. Conclusion

This study conducted a comprehensive comparative evaluation of five AI algorithms for crop yield prediction across three Italian agro-climatic zones and three economically important cereal species over a six-year multi-source dataset. LSTM achieved the highest predictive accuracy (R2 = 0.947, RMSE = 0.31 t/ha), while a novel RF + LSTM hybrid framework further improved performance to R2 = 0.958 through data-driven feature selection preceding deep temporal encoding. Cumulative growing degree days and NDVI at the heading growth stage were identified as the most influential predictors, together accounting for 27.3% of total explained variance. These findings provide empirically grounded recommendations for the architecture of operational PA decision-support systems in European Mediterranean environments: prioritise cloud-free Sentinel-2 acquisition at heading stage, adopt hybrid RF-LSTM architectures for accuracy-critical applications, and deploy GBM where interpretability or real-time latency requirements are paramount. The demonstrated generalisability across diverse soil types, irrigation regimes, and crop species positions the proposed framework as a robust foundation for pan-Mediterranean AI-driven yield forecasting systems aligned with the European Green Deal's Farm-to-Fork food systems transformation agenda.

References

Breiman, L. (2001). Random forests. Machine Learning, 45(1), 5-32.

Cao, J., Zhang, Z., Luo, Y., Zhang, L., Zhang, J., Li, Z., & Tao, F. (2021). Wheat yield predictions at a county and field scale with deep learning, machine learning, and google earth engine. European Journal of Agronomy, 123, 126204.

Cedrez, C. B., Chamberlin, J., Guo, Z., & Hijmans, R. J. (2021). Spatial variation in fertilizer prices in sub-Saharan Africa. PLoS ONE, 16(1), e0227764.

Crane-Droesch, A. (2018). Machine learning methods for crop yield prediction and climate change impact assessment in agriculture. Environmental Research Letters, 13(11), 114003.

ESA (2015). Sentinel-2 User Handbook. European Space Agency, Paris.

European Commission (2021). Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence. COM/2021/206 final.

FAO (2023). The State of Food and Agriculture 2023. Food and Agriculture Organization of the United Nations, Rome.

Feng, P., Wang, B., Liu, D. L., Ji, F., Macadam, I., Yu, Q., & Waters, C. (2021). Machine learning-based integration of remotely-sensed drought factors can improve the estimation of agricultural drought in south-eastern Australia. Agricultural Systems, 173, 303-316.

Filippi, P., Jones, E. J., Bishop, T. F. A., Fajardo, M., & Whelan, B. M. (2019). An approach to forecast grain crop yield using multi-layered, multi-farm data sets and machine learning. Precision Agriculture, 20(5), 1015-1029.

Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. Neural Computation, 9(8), 1735-1780.

Huete, A., Didan, K., Miura, T., Rodriguez, E. P., Gao, X., & Ferreira, L. G. (2002). Overview of the radiometric and biophysical performance of the MODIS vegetation indices. Remote Sensing of Environment, 83(1-2), 195-213.

Khaki, S., & Wang, L. (2019). Crop yield prediction using deep neural networks. Frontiers in Plant Science, 10, 621.

Khanal, S., Fulton, J., Klopfenstein, A., Douridas, N., & Shearer, S. (2020). Integration of high-resolution remotely sensed data and machine learning techniques for spatial prediction of soil properties and corn yield. Computers and Electronics in Agriculture, 153, 213-225.

Liakos, K. G., Busato, P., Moshou, D., Pearson, S., & Bochtis, D. (2018). Machine learning in agriculture: A review. Sensors, 18(8), 2674.

Meier, U. (Ed.) (2001). Growth Stages of Mono- and Dicotyledonous Plants: BBCH Monograph (2nd ed.). Federal Biological Research Centre for Agriculture and Forestry, Berlin.

Nevavuori, P., Narra, N., & Lipping, T. (2019). Crop yield prediction with deep convolutional neural networks. Computers and Electronics in Agriculture, 163, 104859.

Oikonomidis, A., Catal, C., & Mishra, A. (2022). Deep learning for crop yield prediction: A systematic literature review. New Zealand Journal of Crop and Horticultural Science, 51(1), 1-26.

Pantazi, X. E., Moshou, D., Alexandridis, T., Whetton, R. L., & Mouazen, A. M. (2016). Wheat yield prediction using machine learning and advanced sensing techniques. Computers and Electronics in Agriculture, 121, 57-65.

Rouse, J. W., Haas, R. H., Schell, J. A., & Deering, D. W. (1974). Monitoring vegetation systems in the great plains with ERTS. Proceedings of the Third Earth Resources Technology Satellite-1 Symposium, 1, 309-317.

Schwalbert, R. A., Amado, T., Corassa, G., Pott, L. P., Prasad, P. V. V., & Ciampitti, I. A. (2020). Satellite-based soybean yield forecast: Integrating

machine learning and weather data for improving crop yield prediction in southern Brazil. Agricultural and Forest Meteorology, 284, 107886.

Shahhosseini, M., Hu, G., Huber, I., & Archontoulis, S. V. (2021). Coupling machine learning and crop modeling improves crop yield prediction in the US corn belt. Scientific Reports, 11(1), 1606.

van Klompenburg, T., Kassahun, A., & Catal, C. (2020). Crop yield prediction using machine learning: A systematic literature review. Computers and Electronics in Agriculture, 177, 105709.

Wang, X., Huang, J., Feng, Q., & Yin, D. (2020). Winter wheat yield prediction at field scale using machine learning. Remote Sensing, 12(13), 2026.

You, J., Li, X., Low, M., Lobell, D., & Ermon, S. (2017). Deep Gaussian process for crop yield prediction based on remote sensing data. Proceedings of the AAAI Conference on Artificial Intelligence, 31(1), 4559-4566.

Declarations

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors. Satellite imagery was accessed freely through the Copernicus Open Access Hub under ESA's open data policy.

Conflict of Interest

The author declares no conflict of interest. The funders had no role in the design of the study; in the collection, analysis, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Data Availability Statement

The processed feature dataset supporting the conclusions of this article is available in the Zenodo repository at https://zenodo.org/record/XXXXXXX. Raw Sentinel-2 imagery is freely accessible via the Copernicus Open Access Hub at https://scihub.copernicus.eu.

Ethical Approval

This study did not involve human participants, animal subjects, or experiments requiring institutional ethics review. All field data were collected from commercial and experimental farm sites with full owner consent.

ISSN: 2321-9602 Vol: 13, Issues: 3, 2025

Appendix A

Hyperparameter Configurations of Optimised AI Models

Table A1 presents the final hyperparameter configurations for each of the five algorithms following Bayesian optimisation (150 evaluations, Optuna 3.4). All models were trained on a single NVIDIA A100 80 GB GPU (LSTM, CNN) or Intel Xeon Gold 6330 CPU (RF, GBM, SVR) under Python 3.11.