

Green Intelligence-Driven Explainable and Energy-Efficient AI for Climate-Resilient Crop Yield Prediction

Sanjay SINGLA

Post Doc Research Scholar, Multimedia University, Cyberjaya, Malaysia.

Department of Computer Science & Engineering, University Institute of Engineering,
Chandigarh University, Gharuan, Punjab, India

dr.ssinglacs@gmail.com

Chinnaiyan SENTHILPARI

Centre for Advanced Devices and Systems, Faculty of Artificial Intelligence and Engineering
(FAIE), Multimedia University, Cyberjaya, Malaysia

c.senthilpari@mmu.edu.my (corresponding author)

Chu-Liang LEE

Centre for Advanced Devices and Systems, Faculty of Artificial Intelligence and Engineering
(FAIE), Multimedia University, Cyberjaya, Malaysia

cllee@mmu.edu.my

Abstract— Despite the proposed solutions of models accounting for temporal generalization, climate-stress robustness, and the computational efficiency of forecasts, they are still under-explored for reliable forecasts of crop yields. The present study presents GreenzakaNet (GCN), a compact climate-aware neural regression model with a 390-dimensional representation of the predictor using 15 Agricultural and Climate predictors and a unique 9-dimensional Climate-feature pathway with the feature-attention and residual blocks. The model was tested against nine baseline models with a strict chronological protocol: during the validation period (from 2001 to 2007, 14944 samples), the model was compared with known models, while during the unseen test period (from 2008 to 2014, 76878 samples), the model was tested. On validation, GCN achieved the highest performance ($R^2 = 0.953329$; RMSE = 73,080.49 hg/ha), followed by Heavy MLP ($R^2 = 0.795811$) and XGBoost ($R^2 = 0.667444$). During the unseen test period, the R^2 value for GCN was 0.689650, and the climate-stress evaluation ranged from 0.564 to 0.799. Global feature attribution was given by SHAP analysis. GCN has 571,585 parameters, a 6.59 MB checkpoint, and an inference latency of 0.0723 ms/sample, which is suitable for all the objectives of predictive, robust, explainable, and efficient agricultural forecasting for climate resilience.

Keywords— Crop yield prediction; climate-resilient AI; explainable AI (XAI); SHAP; Green AI; energy-efficient deep learning; chronological generalization; climate-stress robustness; lightweight neural networks; precision agriculture.

1. Introduction

As climate variability grows to a greater threat to agricultural productivity, accurate and timely forecasts of crop yields are needed for planning food security and managing resources for climate adaptation. The

incorporation of a wide range of agricultural, climatic, and environmental predictors and more sophisticated architectures have led to better yield prediction via machine learning and deep learning. The three major gaps to sit realistically with the deployment are, however.

First, many reported results are based on random train/test splits, and these do not reflect the forecasting conditions under which models are required to predict for the future, unseen years. An out-of-year evaluation study showed significant degradation under such conditions [13]. Secondly, the evaluation of behaviour under climate and temperature stress conditions is still not adequately represented, and robustness and interpretability are not always considered as part of predictive evaluation. Thirdly, computation and energy consumption are not always considered together with predictive performance, and it remains unclear to what extent the use of resources is warranted when considering the gains in accuracy in resource-constrained agricultural applications.

As such, previous studies tend to focus on optimizing one of the three goals (accuracy, interpretability or efficiency) without examining the impact that one goal has on the other. Practically, one question that arises is whether the gains achieved by the predictions will hold up with unknown years and climate stress, and what is the computational cost?

This study introduces the compact climate-aware neural regression model, GreenzakaNet (GCN), which is tested in a systematic framework with multiple evaluation metrics addressing predictive accuracy, strict chronological generalization, climate- and temperature-stress robustness, SHAP-based explainability, and Green AI efficiency. GCN is evaluated against nine baseline methods (linear, ensemble, boosting, and neural) with the same evaluation protocol as is commonly used in the literature, where the validation period (2001–2007) is separated from the strictly unseen chronological test period (2008–2014).

The contributions of the principal are:

- A compact architecture focused on climate-aware architecture and that integrates in real-time agricultural predictors with a dedicated climate-feature pathway.
- A detailed explanation of how the behavior was generalized and when.
- Assessment of robustness in climate and temperature stress conditions.
- Global feature attribution using SHAP.
- Resource-efficiency assessment based on the inference efficiency measured, and the size of the model, the number of parameters, and the training time.

These elements collectively tackle the pressing need for climate-resilient agricultural AI, combining predictive performance, generalization, robustness, interpretability, and computational efficiency.

2. Literature Review

Classical machine learning methods like Random Forest, XGBoost, and ensemble regressors are adopted in the latest studies, where it is applied to agricultural and meteorological data to make predictions about crops yield. Lionel et al. [1] and Patil et al. [2] showed that amongst several crop types, the Random Forest, gradient boosting and Gaussian Process models performed best, whereas, Gupta et al. [3] and Johri et al. [4] reported high accuracy performance of models trained using temperature, rainfall and soil-derived

features such as using Random Forest, XGBoost and LSTM. In addition to the generic weather input, physiologically-informative predictors have been helpful in improving the prediction of evaporation, thus Naghdyzadegan Jahromi et al. [5] used all-satellite-derived wheat yield related metrics and Joshi et al. [6] took advantage of CHI and gradient boosting. Building on this line, Sahu et al. [7] optimized the balance of drought event within their training data and enhanced the prediction of wheat and barley over Indian and German agro climatic regions.

These baselines have been gradually replaced by deep learning and hybrid architectures. Jeong et al. [8] fused deep neural networks and process-based crop models for rice; Mirhoseini Nejad et al. [10] used a hybrid of 3DCNN, ConvLSTM, and Vision Transformer to preserve the spatial and temporal global context of MS images. Jeong et al. [8] combined a deep neural network with a process-based crop model for rice and Mirhoseini Nejad et al. [10] integrated 3D-CNN, ConvLSTM and Vision Transformer to capture spatial, temporal and global context on MS images. Helber et al. [11] highlighted the modeling capabilities at both the operational and the subfield scale, without needing any prior crop-specific training, while Menon et al. [12] tested a deep regression approach with multi-temporal satellite data in wheat, corn and soybean fields. The above-mentioned studies suggest that in many cases adding architectural complexity produces tangible, albeit small, improvements upon classical ensembles without compromising the performance duration.

However, there is limited ability to generalize outside the training distribution. Temporal cross validation, predicting growing seasons that were not used during the training phase, significantly reduced corn-yield performance compared to random splits, cautioning against the random partitioning for operational forecasting as done by Desloires et al. [13]. A parallel finding was presented by Khan et al. [14] who found that there were inconsistency in the generalization of transfer learning across crops and climatic zones. As noted earlier, if architectural complexity is added to networks, it is not quantitatively better for all major measures of robustness; Zhou et al. [15] showed that a light-weight custom network trained using a handcrafted datasets could outperform deep learning baselines in terms of accuracy while using a much lighter computational load, especially in water-stressed environments.

At the same time, the priority of explainability has emerged as a complementary priority. A naturalistic LSTM was designed by Mateo-Sanchis et al. [16] whose weights make insight into phenological drivers of yield possible. Najjar et al. [17] examined the distinction quality in the explanation of subfield predictions for different crops and countries, revealing that the quality of the explanation depends on the spatial scale and model architectures. Energy and computational cost have not been comparatively discussed as much: for example, to reduce the carbon footprint of agricultural forecasting, Elbasi et al. [18] proposed a Green AI framework with the focus on being lightweight, compressing models, and using edge computing. Energy economy has been treated up to now as a separate goal instead of being part of the accuracy and robustness assessment.

All taken together, these studies demonstrate some of the incremental progress in predictive accuracy, the depth of the available climate-informed features, and an emerging but still intermittent drive for interpretability and efficiency. A limited number of studies consider all three factors (chronological generalization, climatic stress robustness, explainability, and computational efficiency) together; most of studies just focus on a single one, maintaining the other two as secondary or not reported. Whether accuracy advantages from architecturally complex design exist in the face of unknown climatic years and stress remains unclear, as does the cost of computation. GreenzakaNet (GCN) is proposed to jointly assess these

two relatively under-explored points by considering the predictive accuracy, strict chronological generalization, climate-stress robustness, SHAP-based explainability and energy efficiency.

3. Methodology

3.1 Research Framework

The problem of predicting crop yields is a supervised regression problem from agricultural and climate-related variables to a continuous yield. The framework consists of a series of steps that involve preprocessing, chronological assessment, training of the GreenzakaNet (GCN), robustness, explainability and Green AI analysis as illustrated in Fig. 1. Ten models are benchmarked including Mean Baseline, Linear Regression, Ridge, Random Forest, Extra Trees, HistGradientBoosting, XGBoost, LightGBM, Heavy MLP and GreenzakaNet (Proposed V3).

Fig. 1: Workflow Architecture of GreenzakaNet (GCN)




Fig. 1. The Green Intelligence framework for preprocessing, chronological evaluation, training (GreenzakaNet), robustness and explainability and efficiency analysis.

3.2 Data and Preprocessing

The dataset contains 15 original predictors that are geographic, crop, temporal, agricultural, irrigation and climate related. Numerical features were imputed with median values and StandardScaler was applied, and categorical features were imputed with most frequent values and one-hot encoding was applied. The target values Yield_hg_per_ha were log transformed with $\log(1+y)$. The all preprocessing was only trained on the training data and not applied to the validation or test sets in any way. The final representation had 390 features.

3.3 Chronological Evaluation




Fig. 2. Chronological evaluation strategy showing the validation period (2001–2007) and unseen chronological test period (2008–2014).

The validation period was 2001–2007, and the period of the unseen chronological test was 2008–2014 (76,878 samples), as shown in Fig. 2. This split over time is evaluating towards future observations, not random hold-out performance.

3.4 GreenzakaNet (GCN)

The 390-dimensional predictor representation is combined with a 9-dimensional climate vector (vector representing the climate) in GreenzakaNet. Attention and residual blocks handle predictor information, climate gating and interaction add climate effects. The representations are a 448 dimensional fusion, a 160 dimensional latent representation, and a single output of yield. There are 571,585 parameters in the model, and the model has a checkpoint of 6.59 MB.

$$\hat{y} = f(X, C; \theta) \quad (1)$$

where X, C, θ & $\hat{y}$ represents the representation for predictor, the climate feature vector, the learned parameters and the predicted yield.

Training was performed with AdamW, initial learning rate 0.0005, batch size 4096, mixed precision training and warmup–cosine learning rate decay. The training was set up for 100 epochs of training, with early stopping of 15 epochs, with two Tesla T4 GPUs and data parallel training.

3.5 Evaluation, Explainability and Robustness.

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

$$MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i| \quad (3)$$

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2} \quad (4)$$

R^2 , RMSE and MAE were used as performance measures. Feature attribution was conducted globally with 64 background samples, 128 explanation samples, and 390 features with SHAP GradientExplainer. During the SHAP analysis, no retraining was done. The chronological test period was used to assess climate-stress robustness using distribution shifts related to climate.

3.6 Green AI Efficiency

The efficiency of the computation was evaluated based on the number of parameters, the model size, the training time, the inference time, throughput and latency on Tesla T4 GPUs. An estimated energy proxy was used to represent energy, rather than either direct power or emissions instrumentation. Pareto analysis was also used to compare the predictive performance and resource requirements for the 10 benchmark models.

4. Results and Discussion

4.1 Overall Predictive Performance

Table 1 shows that the model that performed the best was GreenzakaNet (GCN) with $R^2 = 0.953329$ and RMSE = 73,080.49 hg/ha. It exceeded Heavy MLP ($R^2 = 0.795811$) and XGBoost ($R^2 = 0.667444$) by 0.157518 and 0.285885 R^2 points, respectively. The comparative validation R^2 performance is shown in Fig. 3.

Table 1. Validation performance of the ten models.

Model	R^2	RMSE (hg/ha)	MAE (hg/ha)	sMAPE (%)
GreenzakaNet (GCN)	0.953329	73,080.49	NR	NR
Heavy MLP	0.795811	152,860.30	24,616.59	23.7132
XGBoost	0.667444	195,079.01	40,255.82	39.8655
Extra Trees	0.607462	211,942.98	28,409.04	25.5144
Random Forest	0.530723	231,735.56	33,002.74	29.3141
LightGBM	0.396306	262,837.01	40,270.74	38.7345
Linear Regression	0.328902	277,122.16	45,675.18	44.4048
Ridge Regression	0.304980	282,017.96	45,698.69	44.4877
HistGradientBoosting	0.292248	284,589.47	46,223.29	44.0792
Mean Baseline	-0.038613	344,750.71	82,950.29	95.9117




Fig. 3. The ten models evaluated have been compared in terms of validation R^2 .

4.2 Chronological Generalization

In Table 2, on the completely unseen 2008–2014 test period, GCN achieved $R^2 = 0.689650$, $RMSE = 254,075.43$ hg/ha, $MAE = 35,777.43$ hg/ha, and $sMAPE = 34.2682\%$ over 76,878 samples. The R^2 value of the model between training and validation set was -0.263679 showing the model had a distribution shift over time but would be able to predict in an out-of-sample scenario. The annual R^2 values were 0.9373, 0.8977, 0.7546, 0.7249, 0.5750, 0.5493, and 0.6719 for 2008–2014, respectively, as shown in Fig. 4. The validation $RMSE$ and MAE comparison of the 10 benchmark models is shown in Fig. 5.





Fig. 4. Temporal generalization of GreenzakaNet for the unseen time period 2008-2014 by year.

Table 2. GreenzakaNet validation and chronological-test performance.

Metric	Validation	Chronological Test
Period	2001–2007	2008–2014
Samples	75,446	76,878
R ²	0.953329	0.689650
RMSE (hg/ha)	73,080.49	254,075.43
MAE (hg/ha)	NR	35,777.43
sMAPE (%)	NR	34.2682
ΔR^2	—	−0.263679

**Fig. 5.** Comparison of validation RMSE and MAE by the 10 models.

4.3 Climate- and Temperature-Stress Robustness

76,878 chronological-test predictions were kept for stress analysis purposes in Table 3. The climate-stress R² ranged between 0.798877 under low stress to 0.641190 under extreme stress while the temperature-stress R² ranged between 0.564331 to 0.780176.

Table 3. Stress-wise robustness of GreenzakaNet.

Stress	Climate R ²	Climate MAE	Temperature R ²	Temperature MAE
Low	0.798877	32,381.08	0.564331	31,622.05
Moderate	0.670157	34,504.59	0.780176	30,184.02
High	0.744169	34,666.72	0.715688	38,945.30
Extreme	0.641190	41,568.40	0.700277	42,378.25

The RMSE for climate stress increased by 232,727.06 hg/ha when it went from low stress to extreme stress, as shown in Figure 6.




Fig. 6. The resistance of GreenzakaNet to climatic and temperature-stress.

4.4 Explainability

In Table 4, these features and 128 explanation samples were used for the SHAP analysis. The highest mean absolute SHAP contributions were found for Item Code (0.102193), Arable land (0.079911), Item: Sesame seed (0.078135), Total area equipped for irrigation (0.070531) and Agricultural area (≈0.0702), as depicted in Fig. 7. The Item Code is not a causal agronomic factor, but an attribute that is attached to the identifier.

Table 4. GreenzakaNet's top 5 feature contributions to the world.

Rank	Feature	Mean Absolute SHAP Value
1	Item Code	0.102193
2	Arable land	0.079911
3	Item: Sesame seed	0.078135
4	Total area equipped for irrigation	0.070531
5	Agricultural area	≈0.0702




Fig. 7. Explainability analysis of GreenzakaNet (GCN): (a) Global feature importance based on mean absolute SHAP values (b) SHAP contribution landscape across explanation samples.

4.5 Green AI and Computational Efficiency

The model has 571,585 trainable parameters and a 6.59 MB checkpoint size. The mean measured value of the inference was obtained in 5.5588 s, at a sampling frequency of 13830.08 samples/s and a sampling time of 0.072306 ms/sample. The previous estimate of 0.0966 s is not included in the final report in Table 5.

Table 5. Computational properties of GreenzakaNet.

Metric	Value
Trainable parameters	571,585
Checkpoint size	6.59 MB
Test inference time	5.5588 s
Throughput	13,830.08 samples/s
Latency	0.072306 ms/sample

The common validation-benchmark latency basis is used for the inference Pareto comparison, and therefore the separately measured 5.5588-s chronological-test inference is not replaced, as shown in Figure 8.




Fig. 8. The benchmark-reported inference latency versus R².

The training-energy analysis is reported as an energy proxy (an estimation) instead of electrical energy measured directly, as shown in figure 9.




Fig. 9. The correlation between R² and estimated training-energy consumption.

4.6 Statistical Robustness

The chronological-test predictions were subjected to a 2000-repetition bootstrap analysis. On the log1p scale, the bootstrap estimates were $R^2 = 0.858977$ with a 95% CI of [0.855233, 0.862714], RMSE = 0.530980 [0.523634, 0.538563], MAE = 0.364603 [0.362020, 0.367221], and sMAPE = 3.554094 [3.520643, 3.588327]. These are scale-transformed statistics and are reported independently of the raw-scale test statistics.

4.7 Overall Findings

Overall, GCN outperformed the other validation models, maintaining $R^2 = 0.689650$ on the test set of unseen chronological data points; was valid under climate stress and temperature stress; gave feature-level attribution; and had a compact architecture with 571,585 parameters and measured inference latency of 0.072306 ms per sample. The results are consistent with the proposed model's predictive, robustness, interpretability, and computational goals.

5. Conclusion and Future Work

The result of the model evaluation showed that with and RMSE of 73,080.49 hg/ha, GreenzakaNet (GCN) was the best validated model in the 10 models evaluated. Over the elusive timeframe 2008 to today, it showed noticeable chronological generalization, even with a measurable time series gap; over this time period it showed no change at all. The model exhibited good predictive performance in both the climate and temperature-stress scenarios, and the feature-level attribution is provided by the SHAP analysis. It's 571,585-parameter architecture, 6.59 MB checkpoint and measured inference latency of 0.072306 ms/sample further illustrate computational efficiency.

Future work will involve examining how shifts in temporal distribution can be addressed (either adaptively or continuously over time) and the robustness of systems in extreme climate, as well as how to minimize the need to rely on identifier like variables by using more suitable features. It further ensures the practical applicability and generalizability of the proposed framework at the energy side of the hardware level and other areas via direct hardware-level energy measurements and more-wide cross-region validation and cross-dataset validation.

References

- [1] B. M. Lionel, R. Musabe, O. Gatera, and C. Twizere, "A comparative study of machine learning models in predicting crop yield," *Discover Agriculture*, vol. 3, no. 1, art. 151, 2025, doi: 10.1007/s44279-025-00335-z.
- [2] Y. Patil, H. Ramachandran, S. Sundararajan, and P. Srideviponmalar, "Comparative analysis of machine learning models for crop yield prediction across multiple crop types," *SN Computer Science*, vol. 6, art. 64, 2025, doi: 10.1007/s42979-024-03602-w.

- [3] P. Gupta, S. Jha, C. Tharani, and V. S. Kirthika Devi, "Crop yield prediction based on various environmental factors using machine learning," in *Proc. 2024 4th Int. Conf. Sustainable Expert Systems (ICSES)*, 2024, doi: 10.1109/ICSES63445.2024.10763249.
- [4] M. Johri, Y. Mangal, N. Parashar, A. Ashish, and Kumud, "Optimizing crop yield prediction using machine learning and deep learning approach," in *Proc. 2024 7th Int. Conf. Contemporary Computing and Informatics (IC3I)*, 2024, doi: 10.1109/IC3I61595.2024.10828647.
- [5] M. Naghdyzadegan Jahromi, S. Zand-Parsa, F. Razzaghi, S. Jamshidi, S. Didari, A. Doosthosseini, and H. R. Pourghasemi, "Developing machine learning models for wheat yield prediction using ground-based data, satellite-based actual evapotranspiration and vegetation indices," *European Journal of Agronomy*, vol. 146, art. 126820, 2023, doi: 10.1016/j.eja.2023.126820.
- [6] A. Joshi, B. Pradhan, S. Chakraborty, and M. D. Behera, "Winter wheat yield prediction in the conterminous United States using solar-induced chlorophyll fluorescence data and XGBoost and random forest algorithm," *Ecological Informatics*, vol. 77, art. 102194, 2023, doi: 10.1016/j.ecoinf.2023.102194.
- [7] H. Sahu, P. K. Garg, S. Vijay, and A. Dasgupta, "Predictive drivers and transferability of multi-scale machine learning based crop yield prediction under drought across European and Asian climates," *Agricultural Water Management*, vol. 327, art. 110254, 2026, doi: 10.1016/j.agwat.2026.110254.
- [8] S. Jeong, J. Ko, J.-O. Ban, T. Shin, and J.-M. Yeom, "Deep learning-enhanced remote sensing-integrated crop modeling for rice yield prediction," *Ecological Informatics*, vol. 84, art. 102886, 2024, doi: 10.1016/j.ecoinf.2024.102886.
- [9] J. Madhuri, M. Indiramma, and N. Nagarathna, "M-Bi-GRU-CNN: a hybrid deep learning model with optimized feature selection for enhanced crop yield prediction," *Multimedia Tools and Applications*, vol. 84, pp. 39787–39811, 2025, doi: 10.1007/s11042-025-20747-9.
- [10] S. M. Mirhoseini Nejad, D. Abbasi-Moghadam, and A. Sharifi, "ConvLSTM–ViT: a deep neural network for crop yield prediction using earth observations and remotely sensed data," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 17, pp. 17489–17502, 2024, doi: 10.1109/JSTARS.2024.3464411.
- [11] P. Helber, B. Bischke, P. Habelitz, C. Sanchez, D. Pathak, M. Miranda, H. Najjar, F. Mena, J. Siddamsetty, D. Arenas, M. Vollmer, M. Charfuelan, M. Nuske, and A. Dengel, "Crop yield prediction: an operational approach to crop yield modeling on field and subfield level with machine learning models," in *Proc. IGARSS 2023 – 2023 IEEE Int. Geoscience and Remote Sensing Symp.*, 2023, pp. 2763–2766, doi: 10.1109/IGARSS52108.2023.10283302.
- [12] A. S. Menon, J. Aravinth, R. Sankaran, and P. Kiran, "Deep learning based farm-level crop yield prediction using multi-temporal satellite data for complex engineering application," *Smart Agricultural Technology*, vol. 12, art. 101562, 2025, doi: 10.1016/j.atech.2025.101562.
- [13] J. Desloires, D. Ienco, and A. Botrel, "Out-of-year corn yield prediction at field-scale using Sentinel-2 satellite imagery and machine learning methods," *Computers and Electronics in Agriculture*, vol. 209, art. 107807, 2023, doi: 10.1016/j.compag.2023.107807.

- [14] S. N. Khan, D. Li, and M. Maimaitijiang, "Using gross primary production data and deep transfer learning for crop yield prediction in the US Corn Belt," *International Journal of Applied Earth Observation and Geoinformation*, vol. 131, art. 103965, 2024, doi: 10.1016/j.jag.2024.103965.
- [15] Y. Zhou, S. Ma, H. Zhang, and S. Aakur, "Enhancing corn yield prediction: optimizing data quality or model complexity?" *Smart Agricultural Technology*, vol. 9, art. 100671, 2024, doi: 10.1016/j.atech.2024.100671.
- [16] A. Mateo-Sanchis, J. E. Adsuara, M. Piles, J. Muñoz-Marí, A. Pérez-Suay, and G. Camps-Valls, "Interpretable long-short term memory networks for crop yield estimation," *IEEE Geoscience and Remote Sensing Letters*, vol. 20, pp. 1–5, 2023, doi: 10.1109/LGRS.2023.3244064.
- [17] H. Najjar, M. Miranda, M. Nuske, R. Roscher, and A. Dengel, "Explainability of subfield level crop yield prediction using remote sensing," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 18, pp. 4141–4161, 2025, doi: 10.1109/JSTARS.2025.3528068.
- [18] E. Elbasi, Y. I. Alzoubi, A. E. Topcu, and M. Nadeem, "Green AI for smart agriculture: energy-efficient predictive models for crop yield and resource management," *IEEE Access*, vol. 13, pp. 204924–204953, 2025, doi: 10.1109/ACCESS.2025.3637027.